Redbooks
ibm.com/redbooks

IBM Power Systems Virtual Server Guide for IBM i

Henry Vo
Gayathri Gopalakrishnan
Diego Kesselman
Rohit Chauhan
Ash Giddings
Luis Eduardo Silva Viera
Tim Simon



Cloud

IBM Power

IBM

Redbooks



IBM Redbooks

Power Virtual Server for IBM i

April 2025

SG24-8513-01

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

Second Edition (April 2025)

This edition applies to Version
IBM i 7.5
IBM i Access Client Solutions 1.1.9.6
Bus4i version 7.1
IBM Power Virtual Server October 2024

© Copyright International Business Machines Corporation 2025. All rights reserved.
Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices ix
Trademarks x

Preface xi
Authors xi
Now you can become a published author, too! xiii
Comments welcome xiii
Stay connected to IBM Redbooks xiii

Summary of changes xv
January 2025, Second Edition xv

Chapter 1. Introduction to IBM Power Virtual Server 1
1.1 IBM Power Virtual Server overview 2
1.1.1 Potential consumers 3
1.1.2 IBM Power Virtual Server Regions 3
1.2 Benefits of running IBM i on IBM Power Virtual Server 5
1.2.1 Availability 5
1.2.2 Performance 5
1.2.3 Backup and recovery 6
1.3 Storage and storage tiers 6
1.4 Power Virtual Servers networking 8
1.4.1 Public network 9
1.4.2 Private network 9
1.5 Creating a Power Virtual Server instance 9
1.5.1 IBM Cloud Dashboard 10
1.5.2 IBM Cloud catalog 11
1.5.3 Power Virtual Server service 13
1.5.4 Power Virtual Server instance 16
1.5.5 Fields in the Virtual Servers Instances page 17
1.5.6 VSI ready to use 18

Chapter 2. Migrating an IBM i LPAR to IBM Power Virtual Server 21
2.1 IBM i data migration to Power Virtual Server 22
2.1.1 Other challenges in cloud migrations 22
2.1.2 Migration options 23
2.2 Migration with partial saves and restores 23
2.3 Backup and Restore migration 26
2.3.1 Considerations for backup and restore migrations 27
2.3.2 Full system backup and restore by using native commands 28
2.3.3 Full-system backup from IBM i Source by using BRMS and IBM Cloud Storage and Restore on IBM Power Virtual Server 47
2.4 Migrating IBM i to Power Virtual Server using PowerVC 61
2.5 Other migration options in IBM Cloud catalog 62
2.6 Migrating IBM i by using IBM i Migrate While Active 66
2.6.1 Migration use cases 67
2.6.2 Anatomy of a migration 67
2.6.3 Data synchronization jobs 69
2.7 Migration example and use case 69

© Copyright IBM Corp. 2025.

2.7.1 Use case overview 70
2.7.2 Scope 71
2.7.3 Architectural decisions 73
2.7.4 Architecture overview 78

Chapter 3. Networking Solutions for Power Virtual Server 81
3.1 Power Virtual Server virtual private network connectivity introduction 82
3.2 VPN in cloud environments 82
3.2.1 Significance of VPN in cloud environments 83
3.2.2 Core components of VPN connectivity for IBM Power Systems Virtual Servers . 84
3.2.3 Role of security and data encryption in VPN connections 86
3.2.4 Scenarios that use VPN for initial setups and scaling environments 88
3.3 IBM Power Virtual Server Network overview 90
3.3.1 Key aspects of the IBM Power Virtual Server network 90
3.3.2 Detailed description of the IBM Power Virtual Server network structure 91
3.4 Integration with IBM Cloud and other cloud services 94
3.4.1 Seamless hybrid cloud connectivity 94
3.4.2 Direct Link for high-performance, low-latency connections 95
3.4.3 Integration with IBM Cloud Services 95
3.4.4 Multi-Cloud Connectivity Using IBM Cloud Transit Gateway 96
3.4.5 Interoperability with Open-Source and Containerized Environments 96
3.4.6 API-driven automation and integration 96
3.4.7 Data and application mobility across environments 97
3.4.8 Network topology: Subnets, routing, and IP addressing 97
3.5 Security measures incorporated in Power Virtual Server networks 101
3.6 Comparing IBM Power Virtual Server network features to similar cloud offerings . . . 105
3.7 IBM Power Virtual Server Network Scenarios 109
3.7.1 Non-production proof-of-concept scenarios 109
3.7.2 Production Scenarios 111

Chapter 4. Backup for IBM i on IBM Power Virtual Server 117
4.1 Backup and restore for IBM Power Virtual Server instances 118
4.2 IBM BRMS 119
4.2.1 IBM BRMS overview 119
4.2.2 IBM Cloud Storage Solutions 119
4.2.3 Cloud solutions for i characteristic 121
4.2.4 IBM BRMS turn-key 122
4.2.5 Backing up the IBM i system to the cloud 123
4.2.6 Full-system backups from the cloud 126
4.2.7 Recovering the IBM i system from the cloud 131
4.2.8 Full-system recovery from the cloud using IBM i as an NFS server 132
4.3 Creating object-level backups 150
4.3.1 Introduction 150
4.3.2 Copying files to the cloud 151
4.3.3 Copying files from the cloud 152
4.3.4 Backup and restore by using save files 153
4.3.5 Creating backups by using image catalog and IBM COS 155
4.3.6 Saving and restoring IBM i objects to IBM COS 157
4.3.7 Transferring a save library to the cloud by using IBM BRMS 161
4.3.8 Automatically transferring media to IBM COS 162
4.4 FalconStor overview 165
4.4.1 Installation and Configuration Prerequisites for FalconStor Storsafe 165
4.4.2 Deployment Scenario with IBM Cloud 166

4.4.3 IBM i migration example that uses FalconStor StorSafe 172
4.4.4 Sizing and ordering FalconStor StorSafe 181
4.4.5 Configuring FalconStor StoreSafe by using StorSight 181
4.5 Power Virtual Instance 187
4.5.1 Snapshot 187
4.5.2 Cloning a volume 188
4.5.3 Restoring a snapshot 188
4.5.4 Snapshot operations by using IBM Cloud CLI 189
4.5.5 Use cases 192

Chapter 5. Hints and tips for IBM i deployments on IBM Power Virtual Server 195
5.1 Connecting to an IBM i virtual machine 196
5.1.1 Remote access to IBM i by using tunneling 199
5.1.2 IBM i 5250 console through LAN adapter 203
5.2 Using snapshots on IBM i instances 206
5.2.1 Uses for snapshots 206
5.2.2 Taking snapshots by installing and using IBM Cloud CLI 206
5.3 Power Virtual Server with VPC Landing Zone 208
5.3.1 Standard variation 209
5.3.2 Extend Power Virtual Server with VPC landing zone - standard variation 210
5.3.3 Quickstart variation 211
5.3.4 Import variation 213

Chapter 6. Disaster Recovery with IBM Power Virtual Server 215
6.1 Disaster recovery by using system-level replication 216
6.1.1 Solution components 216
6.1.2 Solution requirements 217
6.2 OS level replication with PowerHA SystemMirror for i geographic mirroring 242
6.2.1 Prerequisites 243
6.2.2 Planning 243
6.2.3 Implementation of PowerHA SystemMirror for i geographic mirroring in the Power Virtual Server environment 249
6.2.4 Switching a geographic mirroring environment 258
6.2.5 Troubleshooting 262
6.3 Logical replication use case with Bus4i 262
6.3.1 Introduction 262
6.3.2 Prerequisites 263
6.3.3 Planning 263
6.3.4 Implementation of Bus4i for IBM i Logical Replication 270
6.3.5 Switching a Bus4i environment between two IBM i virtual machines in the IBM Power Virtual Server Cloud 279
6.3.6 BUS4i ISB Disaster Recovery 291
6.4 Disaster Recovery with storage replication 295
6.4.1 Data replication 295
6.4.2 Benefits of Global Replication Service 295

Appendix A. IBM i licensing considerations 297
A.1 IBM i Subscription Term Licensing 298
A.2 IBM i ISV Licensing in Power Virtual Server 298
A.3 IBM i Support and Power Virtual Server 299
A.4 Other Licensed Program Products 299
A.5 Passport Advantages software 303
A.6 FAQs about IBM i on IBM Power Systems Virtual Server 303

Abbreviations and acronyms 307

Related publications 309
IBM Redbooks 309
Online resources 309
Help from IBM 311

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.

© Copyright IBM Corp. 2025.

Trademarks

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

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

AIX®
Aspera®
DB2®
Db2®
FlashCopy®
IBM®
IBM Cloud®
IBM Cloud for Financial Services®
IBM Cloud Pak®
IBM Cloud Satellite®
IBM FlashSystem®
IBM Garage™
IBM Security®
IBM Watson®
Passport Advantage®
POWER®
Power9®
PowerHA®
PowerVM®
Rational®
Redbooks®
Redbooks (logo) ®
Satellite™
System z®
SystemMirror®
Tivoli®
WebSphere®

ITIL is a Registered Trade Mark of AXELOS Limited.

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.

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

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

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

Preface

This IBM® Redbooks® publication delivers a how-to usage content perspective that describes deployment, networking, and data management tasks on the IBM Power® Systems Virtual Server by using sample scenarios.

During the content development, the team used available documentation, IBM Power Systems Virtual Server environment, and other software and hardware resources to document the following information:

- IBM Power Systems Virtual Server networking
- Data management deployment scenarios
- Migrations use case scenarios
- Backups case scenarios
- Disaster recovery case scenarios

This book addresses topics for IT architects, IT specialists, developers, sellers, and anyone who wants to implement and manage workloads in the IBM Power Systems Virtual Server. This publication also describes transferring the how-to-skills to the technical teams, and solution guidance to the sales team.

This book compliments the documentation that is available in IBM Documentation and aligns with the educational materials that are provided by IBM Garage™ for Systems Technical Education.

Authors

This book was produced by a working at IBM Redbooks.

Henry Vo IBM Redbooks Project Lead. He joined IBM right after graduating from University of Texas at Dallas 2012 with MIS (Management Information System) major. He shared his technical expertise in business problem solving, risk/root-cause analyze, writing technical plan for business. Within IBM, he hold multiple roles such as Project management, Project lead, ST/FT/ETE Test, Back-End Developer, DOL agent for NY.

Tim Simon is a Redbooks Project Leader in Tulsa, Oklahoma, USA. He has over 40 years of experience with IBM primarily in a technical sales role working with customers to help them create IBM solutions to solve their business problems. He holds a BS degree in Math from Towson University in Maryland. He has worked with many IBM products and has extensive experience creating customer solutions using IBM Power, IBM Storage, and IBM System z® throughout his career.

Rohit Chauhan is a Senior Technical Specialist with expertise in IBM i architecture, working at Tietoevry Tech Services, Stavanger, Norway, an IBM Business Partner and also one of the biggest IT service provider in Nordics. He has over 12 years of experience working on the IBM Power platform with design, planning, and implementation of IBM i infrastructure including High availability/Disaster recovery solutions for many customers during this tenure. Before his current role, Rohit worked for clients in Singapore and U.A.E in the technical leadership and security role on IBM Power domain. He possesses rich corporate experience in architecting solution design, implementations, and system administration. He is also a member of Common Europe Norway with strong focus on IBM i platform and security. He is

© Copyright IBM Corp. 2025.

also recognized as IBM Advocate, Influencer and Contributor for 2024 through IBM Rising Champions Advocacy Badge program. He holds a Bachelor's degree in Information Technology. He is an IBM certified technical expert and also holds ITIL CDS certificate. His areas of expertise include overall IBM i, IBM Hardware Management Console (HMC), Security enhancements, IBM PowerHA®, Systems Performance analysis and tuning, BRMS, External storage, Power VM, and providing solutions to the customers on the IBM i platform.

Ash Giddings is an accomplished IT professional with over 30 years of experience specializing in IBM midrange servers. As the Product Manager at Maxava, a leading provider of disaster recovery, high availability, and cloud-based monitoring solutions for IBM i and beyond, Ash uses his extensive expertise to drive innovation and deliver robust solutions. Recognized as an IBM Champion, Ash is renowned for his deep technical knowledge, strategic insights, and commitment to the IBM i community. Ash is an IBM Champion and specializes in systems management, HA/DR and cloud computing. He is a regular speaker at IBM i user group events around the globe.

Gayathri Gopalakrishnan With over 24 years of extensive experience in technical solution architecture and IT consulting, Gayathri Gopalakrishnan is a results-oriented IT Architect known for leading impactful projects from conception through implementation. She brings exceptional expertise in managing, designing, developing, and testing innovative solutions that align with critical business objectives. Gayathri is recognized for her strategic vision and leadership in applying cutting-edge technical solutions to solve complex challenges across industries. Her deep understanding of both the business and technical landscapes enables her to effectively prioritize tasks, align teams, and deliver results that exceed expectations. She is adept at transforming business requirements into robust, scalable technical frameworks that drive measurable success.

Diego Kesselman is general manager at Esselware Soluciones, an IBM Business Partner in Mèxico. The company provides IBM Power Systems services and solutions to customers in Mexico and Latin America, with a focus on managed services, hardware architecture, virtualization, performance analysis and tuning, software development, and integration on IBM i and AIX® environments. He is also an IBM i and IBM Power Systems specialist with 15 years of experience. A former programmer, he uses Open Source tools on IBM i and tests new tools on the platform. He also scripts by using OpenSSH and BASH. Diego published several projects on GitHub about IBM i Open Source and participates in different communities about IBM i and Power Systems.

Luis Eduardo Silva Viera is an IBM i Expert IT specialist and consultant who worked for IBM in Latin America since 1999 for 20 years, developing expertise as IBM i Consultant in the design, planning, and implementation of external storage, high availability, Disaster Recovery (DR), backup and recovery, and performance review solutions for small and large customers of the IBM i platform in countries in South, Central, and North America. Since 2019, he Works for the globally operating n-Komm group of companies that are based in Karlsruhe, Germany, and focuses on the company's IBM i services for its customers in Germany, Switzerland, and Austria. Starting in 2021, he also supports the group's worldwide activities with the T.S.P. Bus4i logical replication product for IBM i, and focuses in the high availability, DR, and migration projects for IBM i customers. Luis also has experience teaching technical courses on IBM i in different countries.

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

Dan Sundt
IBM i Power as-a-Service Product Manager

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 books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

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

  ibm.com/redbooks

- Send your comments in an email to:

  redbooks@us.ibm.com

- Mail your comments to:

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

Stay connected to IBM Redbooks

- Find us on LinkedIn:

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

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

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

- Stay current on recent Redbooks publications with RSS Feeds:

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

Summary of changes

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

January 2025, Second Edition

This revision includes the following new and changed information.

New information

- Migration while active
- IBM Cloud® Network
- FalconStor Virtual Tape Library
- Logical replication use case with Bus4i

Changed information

- Deleted outdated information
- Updated new workplace for IBM Cloud
- Added support for IBM Power10

© Copyright IBM Corp. 2025.

1

Introduction to IBM Power Virtual Server

The IBM Power Systems Virtual Server is an IBM Power Systems service offering that is hosted by IBM data centers. It delivers enterprise-class computing capabilities with the flexibility of a hybrid cloud deployment. In addition to the public cloud Power Virtual Service offering, IBM also provides Power Virtual Server Private Cloud, an on-premises version of the Power Virtual Server offering for clients that require more control over the location of their workloads and data.

This chapter provides a look at the benefits of running your AIX, IBM i, or Linux on Power workloads on the cloud and discussing some of the use cases that can be effectively run on Power Virtual Server.

The chapter also provides a short tutorial on creating a Power Virtual Server instance.

This chapter contains the following topics:

- "IBM Power Virtual Server overview"
- “Benefits of running IBM i on IBM Power Virtual Server”
- “Power Virtual Servers networking”
- "Creating a Power Virtual Server instance"

© Copyright IBM Corp. 2025.

1.1 IBM Power Virtual Server overview

IBM Power Virtual Server is a cloud-based infrastructure service that allows businesses to run workloads on IBM Power hardware. It combines the performance and reliability of IBM's Power architecture with the flexibility and scalability of a virtualized environment. Key features include:

- Flexibility. Users can scale resources up or down based on demand, making it suitable for various workloads, including enterprise applications and databases.
- Performance. Using IBM POWER® processors, the service is designed to support business-critical workloads like high-performance databases.
- Hybrid Cloud Integration. It supports seamless integration with on-premises systems and other cloud services, facilitating a hybrid cloud approach.
- Security. Built-in security features help protect sensitive data and comply with industry standards.
- Managed Services. IBM provides management options, allowing users to focus on their applications rather than infrastructure.

IBM Power Virtual Server is targeted at organizations looking for robust, enterprise-grade cloud solutions that can support demanding applications. IBM Power Virtual Server resources are located in IBM data centers with dedicated networking and storage area network attached Fibre Channel storage. At the time of writing, 21 data centers that are located around the world host Power Virtual Server, you can choose one of the data centers that are nearest to your data center. The list of data center regions is provided in section 1.1.2, “IBM Power Virtual Server Regions” on page 3. IBM Power clients who rely on private cloud infrastructure can now quickly and economically extend their Power IT resources to the cloud.

In the data centers, the Power Virtual Server are separated from the rest of the IBM Cloud servers with separate networks and direct attached storage. The internal networks are fenced for security, but offer connectivity options to IBM Cloud infrastructure or private cloud environments on your premises. This infrastructure design enables Power Virtual Server to maintain key enterprise software certification and support as the Power Virtual Server architecture is identical to certified private cloud infrastructure.

Power Virtual Server is an infrastructure as a service offering (IaaS) where there are no up-front costs for deploying resources and resources are paid for based on usage. Power Virtual Server uses a monthly billing rate that includes the licenses for the AIX, IBM i, or Linux operating systems. The monthly billing rate is prorated by the hour based on the resources that are deployed to the Power Virtual Server instance for the month. When you create the Power Virtual Server instance, you can see the total cost for your configuration based on the options that you specify. You can quickly identify what configuration options provide you with the best value for your business needs. There is also an option for a bring your own Linux image. In this case the customer is responsible for acquiring the subscription and paying the Linux distributor for licenses used.

You can configure and customize the following options when you create a Power Virtual Server:

- Number of virtual server instances
- Number of cores
- Amount of memory
- Data volume size and type
- Network interfaces

Configuration, billing, and management are all handled by IBM Cloud through a cloud-based graphical user interface or through defined APIs.

For those customers that want to move to a cloud-like environment but with a preference for on-premises hardware, IBM offers IBM Power Virtual Server Private Cloud. This on-premises based solution enables customers who potentially have sensitive data concerns to implement their Power based applications in an on-premises cloud solution. The solution provides the same cloud capabilities and management as the public cloud based version of the IBM Power Virtual Server without placing their data in a public cloud infrastructure.

1.1.1 Potential consumers

Potential consumers of the Power Virtual Server service offering are IT administrators, Managed Service providers/Cloud Service providers (MSPs/CSPs), independent software vendors (ISVs), application developers, and those businesses looking to move from a traditional capital expenditure approach for IBM Power hardware, to operational expenditure.

IT administrators

The people who manage infrastructure technology are interested in migrating to the cloud to speed up time for value their Power workloads, shift capital expense to operating expense, and improve business resilience and scalability. It allows clients to manage a truly hybrid environment with flexible burst environments for spikes in usage, development and test environments, or production workloads. The flexibility of quickly implementing new hardware for a project speeds time to results and reduces costs, providing significant benefits to the enterprise.

Managed Service providers and Cloud Service providers

Service providers that are interested in expanding the level of service can take advantage of the flexibility of Power Virtual Server to offer their clients enhanced offerings. Many MSPs have client workloads running on Power, and they can provide more services around the cloud.

Independent software vendors

Companies that sell software as a service can take this capability to deploy infrastructure to host their software on the cloud for client basis.

Application developers

Companies with Power servers feature subject matter experts in AIX, IBM i, and Linux development can take advantage of this offering, developing mission-critical applications in the cloud and then deploying either in the cloud or on-premises. Reducing the time to implement new development and test environments increases the productivity of the developers.

Businesses changing from capital expenditures to operational expenditures

Companies that have traditionally made capital expenditure IBM Power hardware purchases that are looking to move to an operational expenditure consumption model can do so with this offering. Having the flexibility to pay only for resources as they are used can provide financial benefits.

1.1.2 IBM Power Virtual Server Regions

IBM Power Virtual Server infrastructure is located in IBM data centers. Currently the IBM Power Virtual Server offering is available in 21 data centers around the globe and this

number continues to expand. In addition, Power Virtual Server Private Cloud provides a private data center option that is installed on your premises as well. For more information on Power Virtual Server Private Cloud, see Introduction to IBM Power Virtual Server Private Cloud.

Data center locations are provided in both North and South America. Figure 1-1 displays regions and zones in the Americas.

Location	VPC Region	Power Virtual Server region	VPC Zone	Power Virtual Server zone	Classic Infrastructure
Dallas	us-south	us-south	us-south-1 us-south-2 us-south-3	dal10 dal12 us-south	dal10 dal12 dal13
Sao Paulo	br-sao	sao	br-sao-1 br-sao-2 br-sao-3	sao01 sao04 x	sao01 sao04 x
Toronto	ca-tor	tor, mon	ca-tor-1 ca-tor-2 ca-tor-3	tor01 [1] mon01 x	tor01 mon01 x
Washington DC	us-east	us-east	us-east-1 us-east-2 us-east-3	us-east wdc06 wdc07	wdc04 wdc06 wdc07

Figure 1-1 Power Virtual Server regions and zones in Americas

Additional data centers are available in Europe. Figure 1-2 is a list of the regions and zones in Europe.

Location	VPC Region	Power Virtual Server region	VPC Zone	Power Virtual Server zone	Classic Infrastructure
Frankfurt	eu-de	eu-de	eu-de-1 eu-de-2 eu-de-3	x eu-de-1 eu-de-2	x fra04 fra05
London	eu-gb	lon	eu-gb-1 eu-gb-2 eu-gb-3	lon04 x lon06	lon04 x lon06
Madrid	eu-es	mad	eu-es-1 eu-es-2 eu-es-3	mad02 mad04 x	mad02 mad04 x

Figure 1-2 Power Virtual Server regions and zones in Europe

If you need service in the Asia Pacific regions, data centers are available as shown in Figure 1-3.

Location	VPC Region	Power Virtual Server region	VPC Zone	Power Virtual Server zone	Classic Infrastructure
Sydney	au-syd	syd	au-syd-1 au-syd-2 au-syd-3	x syd04 syd05	x syd04 syd05
Tokyo	jp-tok	tok	jp-tok-1 jp-tok-2 jp-tok-3	x tok04 x	x tok04 x
Osaka	jp-osa	osa	jp-osa-1 jp-osa-2 jp-osa-3	osa21 x x	osa21 x x
Chennai	x	X	x	che01 [2]	che01

Figure 1-3 Power Virtual Server regions and zones in Asia Pacific

1.2 Benefits of running IBM i on IBM Power Virtual Server

There are many benefits to moving your IBM i workloads to IBM Power Virtual Server instances. Some of these benefits are:

- Cloud based pricing – only pay for what you use.
- Quick, self-service definition of Power resources – faster setup of the environment can lead to meeting customer requirements more quickly.
- Moving from a capital expense model to an operating expense model – allow capital to be spent on more important business imperatives.

There are additional benefits that provide significant advantages when the workloads are run in an IBM Power Systems Virtual Server environment compared to building the environment yourself.

1.2.1 Availability

IBM Power Virtual Server is designed to support a range of environments, from development/test environments to critical business applications requiring high availability and disaster recovery. The following availability features are built into the offering:

- Live Partition Mobility support

  The IBM i instance is migrated seamlessly to a new server with Live Partition Mobility at no cost when the original server is being maintained or is failing. This issue results in no charge high availability, which is the default behavior but can be tuned.
- Data protection and replication
  - Instant snapshots can be taken from your IBM i instances.
  - Clones can be taken from disks with just a couple of clicks.
  - Instances can be exported to an OVA and saved for later use or migrated to a new region.
- Disaster recovery solutions available
  - Data replication by using storage replication is available by using Global Replication Services (GRS).
  - PowerHA is available for automated failover.
  - Other third-party IBM i replication services are also supported.

1.2.2 Performance

Power Virtual Server is designed to meet the performance requirements of any environment. Based on IBM Power9® and Power10 servers, your workloads benefit from the latest generations of IBM Power technology and the environment is designed to provide excellent performance. Consider the following:

- Power9 and Power10 processors provide significant improvement in performance over older generations of IBM POWER processor based server.
- Power Virtual Server uses SAN attached high-performance storage provided by IBM FlashSystem® servers with SSD drives.
  - Choose from a variety of IOPs-based storage pools to meet the performance needs of your application.
- Network speed is fast:
  - Internet connection can scale 1 Gbps–10 Gbps.

– Local networking is redundant, supports Jumbo Frames and link aggregation, and the link speed is over 10 Gbps.

Most of these features are not evident to the user.

1.2.3 Backup and recovery

Power Virtual Server provides several options for backing up and restoring data to your LPARs.

- An iSCSI Virtual Tape Library offering is available that uses the fast networking facility.
- IBM Cloud Object Storage can be accessed from the internet or privately, with a cost-effective offer.
  - IBM Cloud Object Storage is fast and can be accessed with native or open source software.
- BRMS or IBM Cloud Storage Solution for i does not need to be installed. Instead, select the license feature when creating an instance and the standard image includes the product with a basic configuration.

1.3 Storage and storage tiers

Power Virtual Server uses direct attached SAN storage to provide a high-performance storage environment. The storage infrastructure is designed with high availability in mind, providing multiple connections between the storage and servers with no single point of failure. In addition, Power Virtual Server provides multiple methods of replicating storage in your environment to support high availability between data centers and disaster recovery options.

When you define your Power Virtual Server instance (LPAR, you define the boot image and the storage tier that you want the boot image volume to be created. You can create additional data volumes when you define your LPAR or you can attach them to the instance at a later time. For each storage volume, you define the storage tier for that volume.

IBM Power Virtual Server offers you the option to select an I/O operation per second (IOPS) based storage option according to your requirement. Flexible IOPS is a tier-less storage offering that removes the notion of disk type and replace it with a storage pool. Each of the storage pools supports multiple storage tiers. The storage tiers are based on different IOPS levels. Figure 1-4 shows the supported storage tiers with corresponding IOPS.

Tier level	IOPS	Performance
Tier 0	25 IOPS/GB	A 100-GB volume receives 2500 IOPS. This is 2.5x faster than tier 1 and 8.3x faster than tier 3.
Tier 1	10 IOPS/GB	A 100-GB volume receives 1000 IOPS. This is 3.3x faster than tier 3.
Tier 3	3 IOPS/GB	A 100-GB volume receives 300 IOPS.
Fixed IOPS	5000 IOPS regardless of size	A 100-GB volume receives 5000 IOPS.

Figure 1-4 Tier and IOPS mapping

Important: The use of fixed IOPS is limited to volumes with a size of 200 GB or less, which is the break even size with Tier 0 (200 GB @ 25 IOPS/GB = 5000 IOPS).

Benefits of flexible IOPS

Flexible IOPS offers multiple IOPS levels (tiers) of storage to choose from. Each IOPS level has its own pricing, performance, and capability.

By using flexible IOPS, you can perform the following tasks:

- Create a volume (or multiple volumes) and specify the IOPS level that you want.
- Change the volume level of IOPS of existing volumes while staying within the same storage pool.
- Clone one or more volumes to your choice of IOPS level. If any of the volumes is greater than 200 GB, then you cannot have a target tier as Fixed IOPS.
- Deploy the boot volume of a virtual server instance in any of the supported IOPS levels. Additional data volumes that are attached to the new virtual server instance can have different IOPS levels from that of the boot volume of an instance.
- Import an image from IBM Cloud Object Storage to any of the supported IOPS levels.

Note: Best practice for an image import is to use the default IOPS level (Tier 3). While deploying the image you can choose the IOPS level for your boot volume during virtual server instance deployment, the boot volume's IOPS level does not need to match the IOPS level of the image.

Selecting a storage tier

You can use flexible IOPS to select a tier that you need for your boot volume and data volume.

Boot volume

- When you are creating a virtual server instance, you can define the boot volume by performing the following steps:
- Select your Operating system.
- Select or clear the Configure for Epic workloads indicator.
- Select your Image.
- Select from Tier 0, Tier 1, Tier 3, or Fixed IOPS.
- For Storage pool as Figure 1-5, select from Auto-select pool, Affinity, Anti-affinity.

VM storage pool affinity setting	Action
Auto-select pool	Power Virtual Server determines the best storage pool available for you.
Affinity	The storage pool must be similar to the storage pool of affinity object that you choose.
Anti-affinity	Storage pool makes the storage pool of affinity object that you choose as an exception.

Figure 1-5 Storage pool affinity setting

Note: All volumes that are created during VM provisioning are created on the same storage pool as the boot volume irrespective of their tier selection.

Data volume

To create a volume, complete the following steps:

- Enter a unique name.
- Enter the desired size of the volume.
- Enter the quantity of volumes that you need.
- Select from Tier 0, Tier 1, Tier 3, or Fixed IOPS.

During VM provisioning, if you create additional data volumes to attach to the new virtual server instances then these data volumes can be created on any of the supported storage tiers. All these additional data volumes are located in the same storage pool where the boot volume of the virtual server instance resides.

Limitations of flexible IOPS

Flexible IOPs includes the following limitations:

- Snapshot data cannot be changed from one tier to another. All volumes of a snapshot must reside in the same storage pool.
- Any volume that has a storage type of tier 0, tier 1, or tier 3 and the volume size is greater than 200 GB then the option to change to Fixed IOPS is not allowed.

1.4 Power Virtual Servers networking

IBM Power Virtual Server supports various network scenarios that address diverse needs, from low-cost, flexible POC setups to robust, high-performance production environments. By selecting the right scenario, organizations can optimize their network architectures, achieve secure and reliable connectivity, and support business requirements, whether it's for a single cloud, hybrid, or multi-cloud deployment. These scenarios are foundational for organizations seeking the scalability and flexibility of the cloud while maintaining control, security, and performance for their critical workloads.

IBM Power Virtual Servers have separate networking hardware:

- Cisco Nexus9000 N9K-C9364C (Spine 10 G)
- Cisco Nexus9000 9348GC-FXP (Leaf 1 G)
- Cisco Nexus9000 93180YC-FX (Leaf 25 G)
- Cisco UCS - APIC controller
- Cisco ASR1001-HX Routers
- Avocent ACS8032DAC-400

You can use Power Virtual Server Cloud Connections to connect your Power Virtual Server instances to IBM Cloud resources on IBM Cloud classic network and Virtual Private Cloud (VPC) infrastructures. Power Virtual Server Cloud Connection creates a Direct Link (2.0) Connect instance to connect your Power Virtual Server instances to the IBM Cloud resources within your account. For cross-account connectivity, use IBM Transit Gateway to interconnect your Power Virtual Server to the IBM Cloud classic and Virtual Private Cloud (VPC) infrastructures. The speed and reliability of the Direct Link connection extends your Power Virtual Server network to the IBM Cloud network and offers more consistent and higher-throughput connectivity, while keeping the network traffic within IBM Cloud.

When you create a Power Virtual Server, you can select a private or public network interface to connect your Power Virtual Server instance to your on-premises networks.

1.4.1 Public network

The public network option provides a method to connect to a Power Virtual Server instance:

- IBM configures the network environment to enable a secure public network connection from the internet to the Power Virtual Server instance.
- Connectivity is implemented by using an IBM Cloud Virtual Router Appliance (VRA) and a Direct Link Connect connection.
- The public network is protected by a firewall and supports the following secure network protocols:
  - SSH
  - HTTPS
  - Ping
  - IBM i 5250 terminal emulation with SSL (port 992)

1.4.2 Private network

You can use the private network option so that your Power Virtual Server instance can access existing IBM Cloud resources, such as IBM Cloud Bare Metal Servers, Kubernetes containers, and Cloud Object Storage The private networks uses a Direct Link Connect connection to connect to your IBM Cloud account network and resources. A private network is required for communication between different Power Virtual Server instances.

For more detail on the types of networks, see Chapter 3 - Chapter 3, “Networking Solutions for Power Virtual Server” on page 81.

1.5 Creating a Power Virtual Server instance

Before creating a virtual server, you must understand the difference in terminology between a Power Systems Virtual Server service and a Power Systems Virtual Server instance. You can think of the Power Systems Virtual Server service as a container for all Power Systems Virtual Server instances at a specific geographic region.

The Power Systems Virtual Server service is available from the Resource list in the Power Systems Virtual Server user interface. The service can contain multiple Power Systems Virtual Server instances. For example, you can have two Power Systems Virtual Server services: one in Dallas, Texas, US, and another Washington, DC, US. Each service can contain multiple Power Systems Virtual Server instances.

Before you create your first Power Virtual Server instance, log in to IBM Cloud. See Figure 1-6 on page 10.

If you do not have an account, see Create an IBM Cloud acount to register an account.

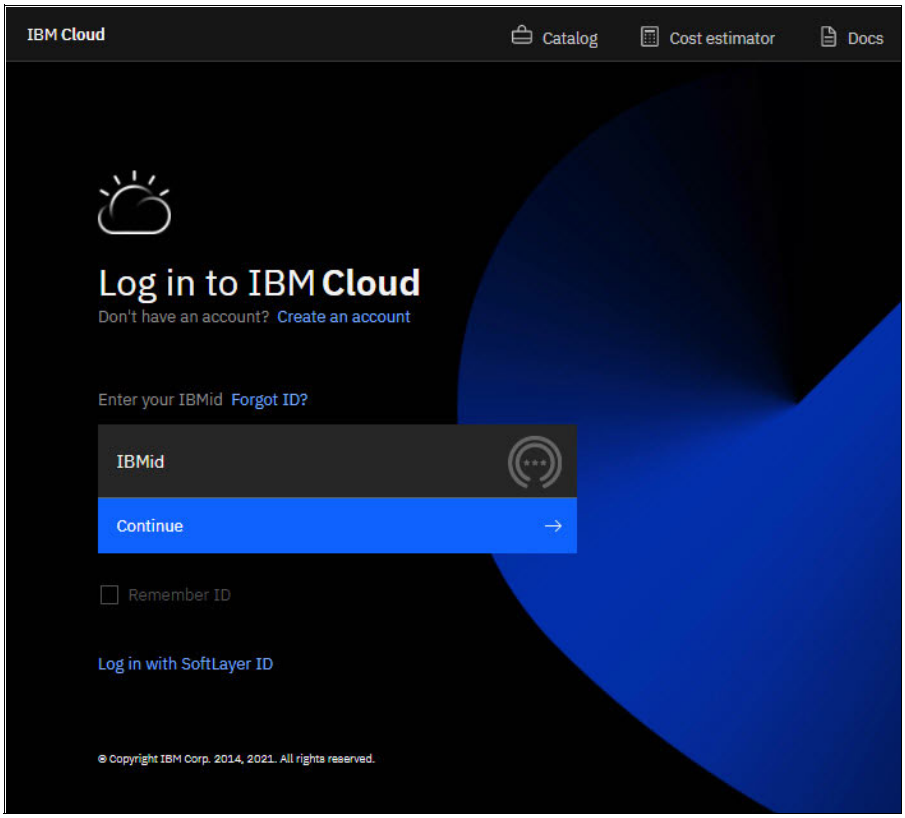

Figure 1-6 Log in to IBM Cloud (https://cloud.ibm.com/login)

1.5.1 IBM Cloud Dashboard

The Dashboard involves arranged parts and details and points to applicable documentation. In addition, many menus, links, and icons help you fully manage your environment, your Power Virtual Server instances, and associated resources. See Figure 1-7 on page 11.

Dashboards are customizable. You can create a dashboard that displays information that is relevant to you.

Note: The dashboards that you create can be scoped to specific resources. Also, you can share the dashboards with users in your account, so you can group resources for particular projects or teams.

You can securely authenticate users, control access to Power Virtual Server resources with resource groups, and use access groups to allow access to specific resources for a set of users. This service is based on an Identity and Access Management (IAM) mechanism, which provides all user and resource management in the IBM Cloud.

You can assign IAM authorizations based on the following criteria:

- Individual users
- Access groups
- Specific types of resources
- Resource groups

For more information about IAM, see Identity and access management (IAM) services.

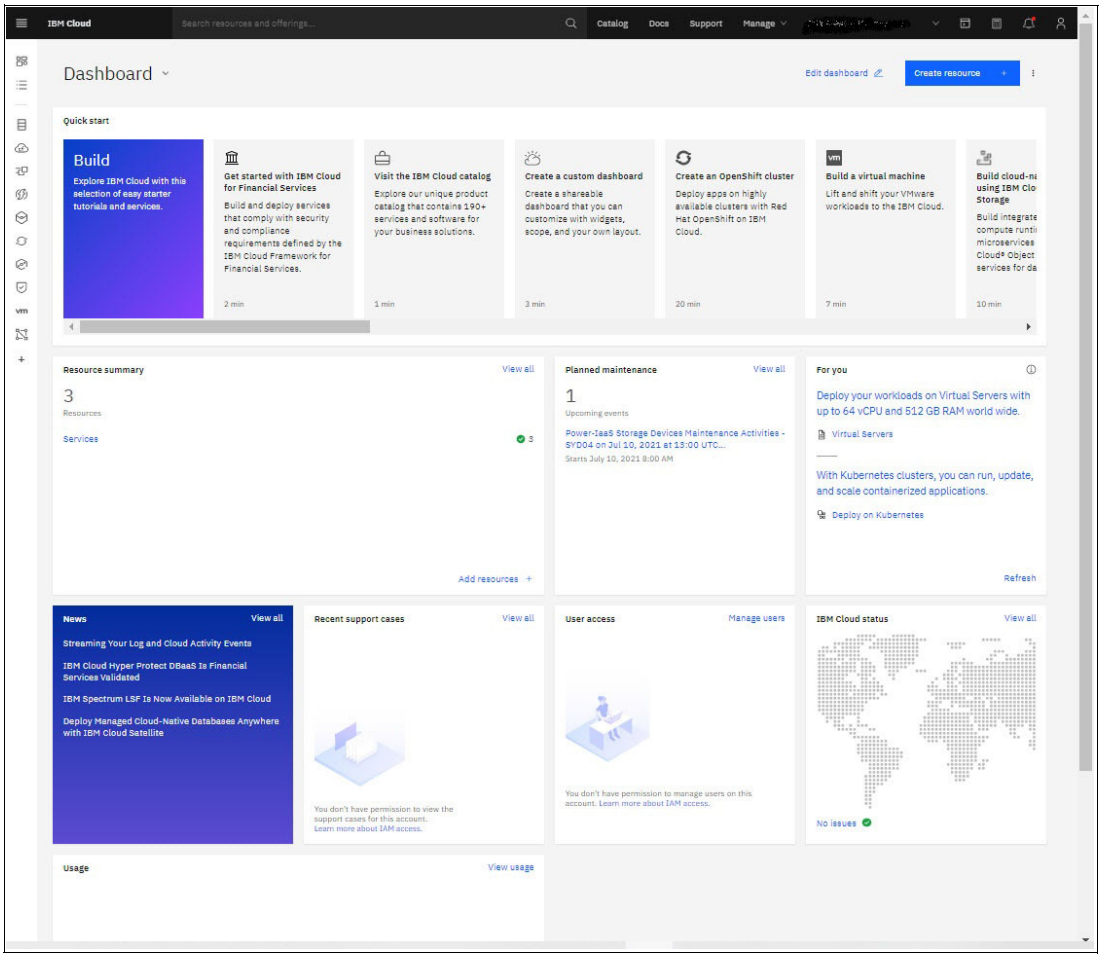

Figure 1-7 IBM Cloud Dashboard

1.5.2 IBM Cloud catalog

Select **Catalog** to open the IBM Cloud catalog, which is an inventory of all offerings that are provided by IBM Cloud. See Figure 1-8 on page 12.

The catalog provides a stack with various products, including the following examples:

- Computing
- Containers
- Storage
- Networking
- End-to-end developer solutions for application development
- Testing and deployment
- Security management services
- Traditional and open source databases
- Cloud-native services
- AI / Machine Learning

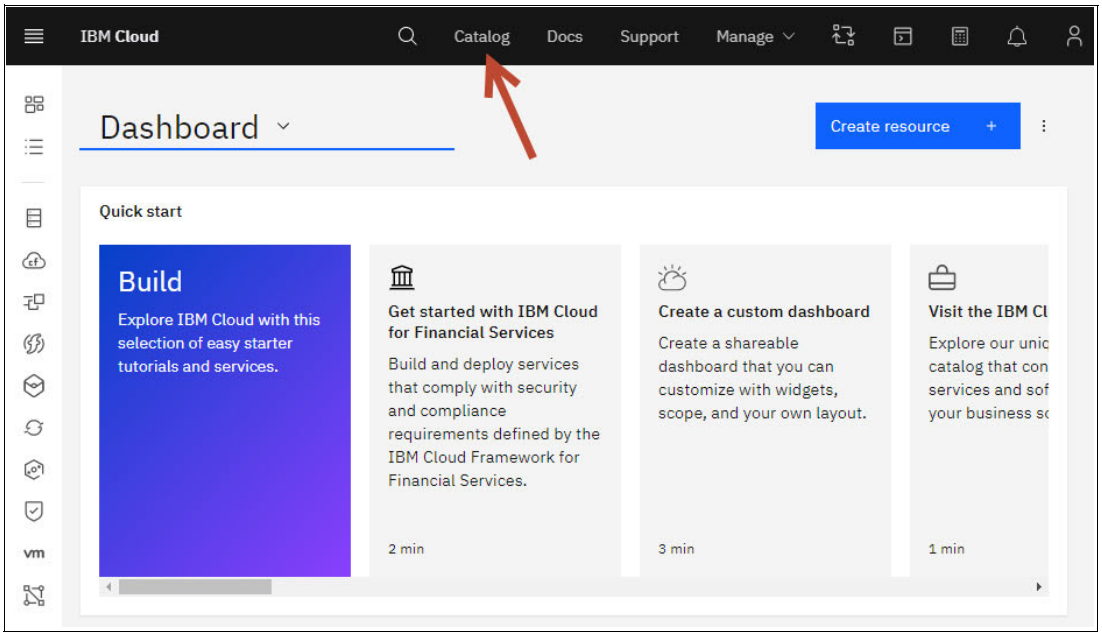

Figure 1-8 IBM Cloud catalog

The offering relies on services, software, and consulting. You can share the catalog, as shown in Figure 1-4.

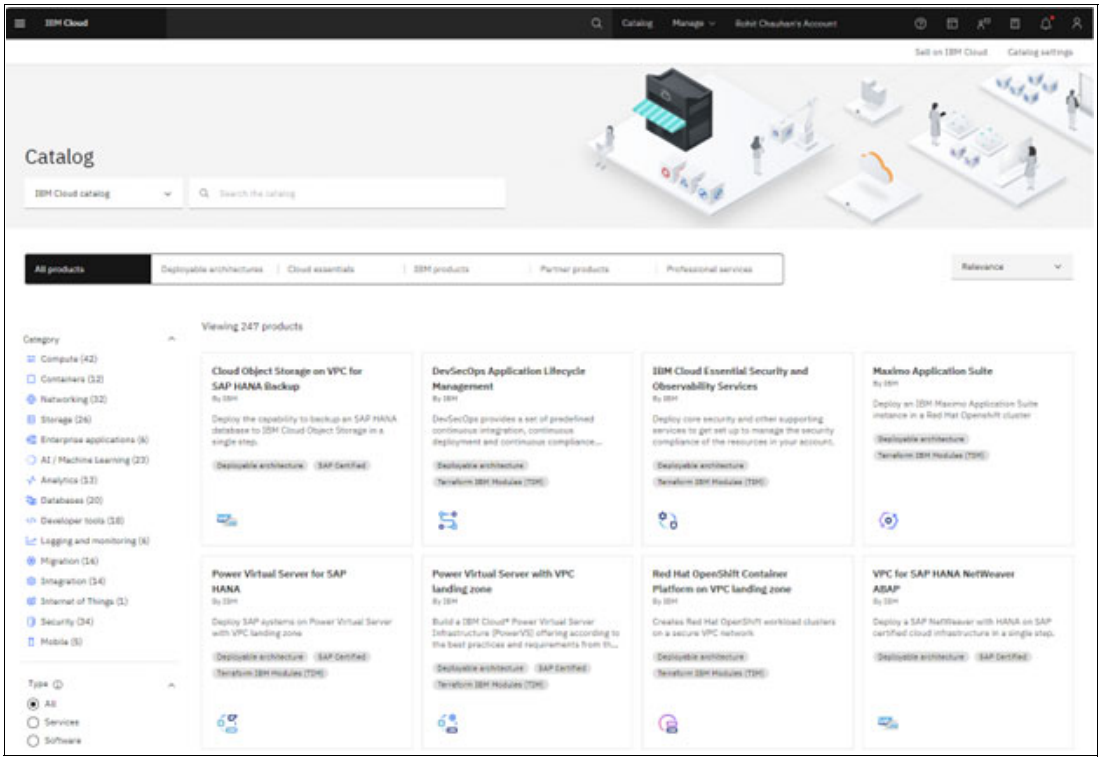

Figure 1-9 IBM Cloud Catalog Search

The catalog includes the following sections:

- Services

  It contains a portfolio of managed services for infrastructure, developer tools, and more to build your applications on the public cloud.

- Software

  List of software solutions that take advantage of a simplified installation process.

- Deployable Architectures

  A deployable architecture involves components, modules, and dependencies in a way that allows for seamless deployment and helps make it easier for developers and operations teams to quickly deploy new features and updates to the system without requiring extensive manual intervention.

- Hub

  Three different services that can share trust data and can use managed service and connected IBM IOT.

- Professional Services

  To get help, from technical to strategic and more, from IBM and its network of partners. You can search the catalog as shown in Figure 1-10

.I

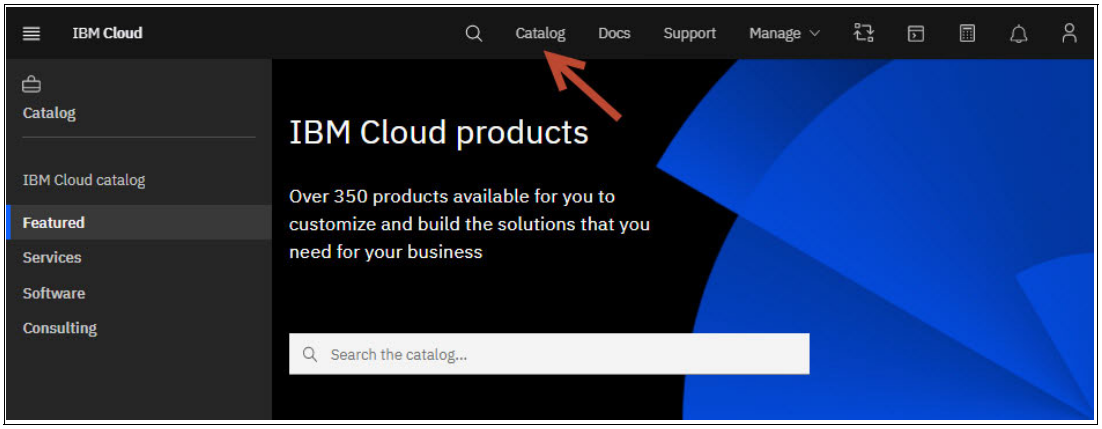

Figure 1-10 IBM Cloud catalog search

1.5.3 Power Virtual Server service

To create a Power Virtual Server service, do the following steps:

1. Search for Power Virtual Server in the catalog's search box. See Figure 1-11 on page 14.
2. After you select **Workspace for Power Systems Virtual Server,** a new window opens.
3. Select **Create a workspace** and specify a name for your service and choose where you want to deploy your Power Virtual Server instances. Refer to Figure 1-12 on page 14.

   After you click **Create**, the Resource List page opens, which contains a list of your account resources. Use this page to view and manage your platform and infrastructure resources in your IBM Cloud.

> **Note:** Another way to access the Resource List page is to click the Navigation menu in the upper-left, and then click **Resource List** in the Dashboard menu. See Figure 1-13.

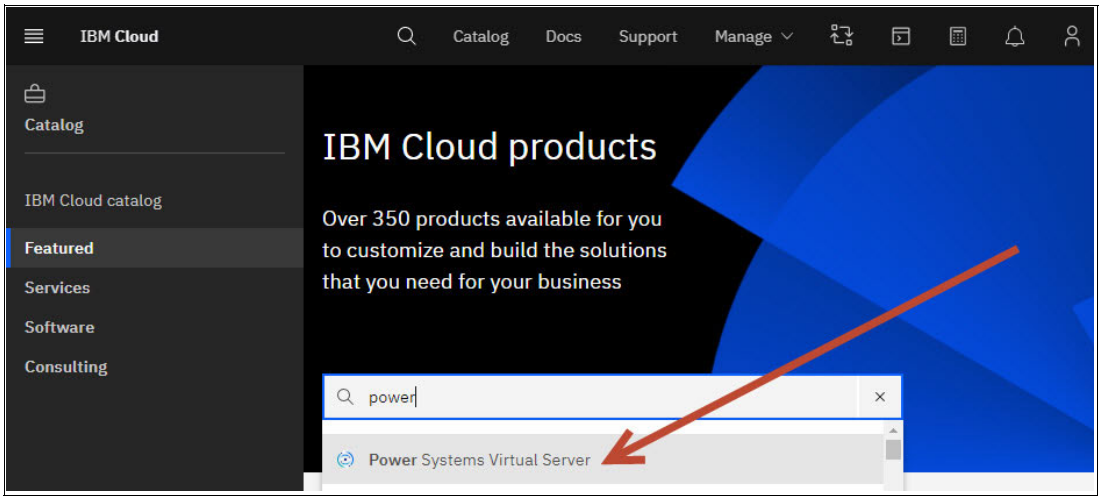

Figure 1-11 Search for Power Virtual Server

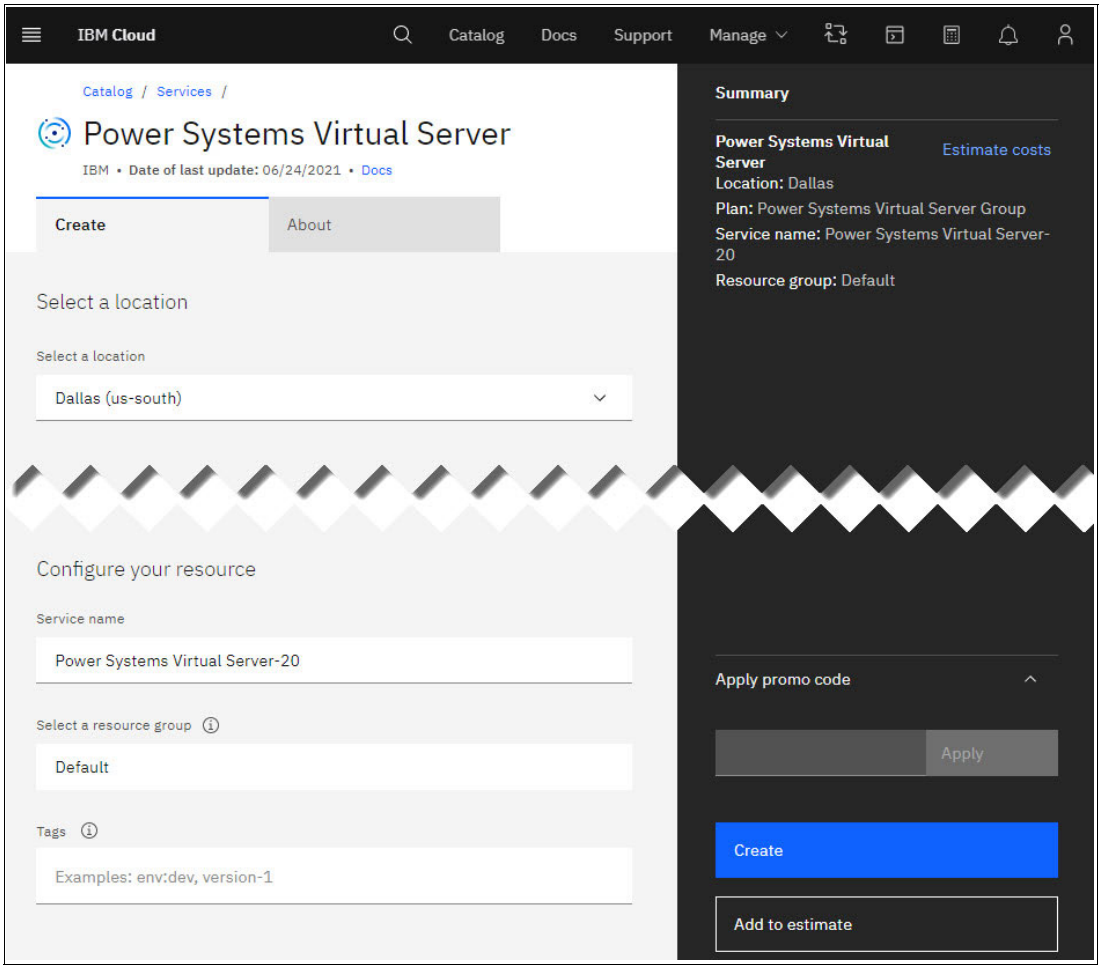

Figure 1-12 Service name and location

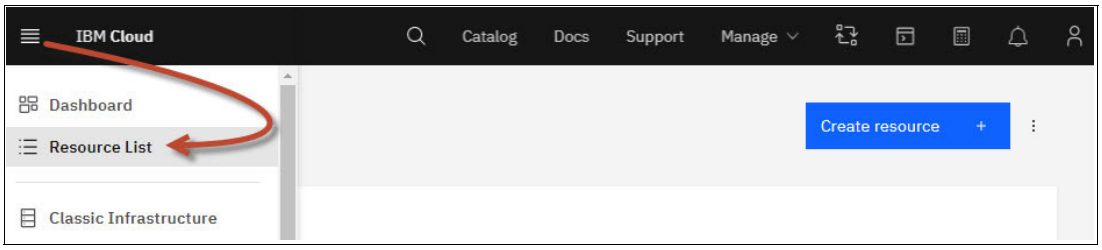

Figure 1-13 Resource List on Dashboard menu

4. Search for resources from anywhere in the IBM Cloud by entering the resource or tag in the search field from the menu bar.

   Figure 1-14 shows the result after Power Virtual Server is created. All Power Virtual Servers are under **Services and software** resources.

   Each resource is displayed in its row, and an Actions icon is included at the end of the row. Click the **Actions** icon to start, stop, rename, or delete a resource.

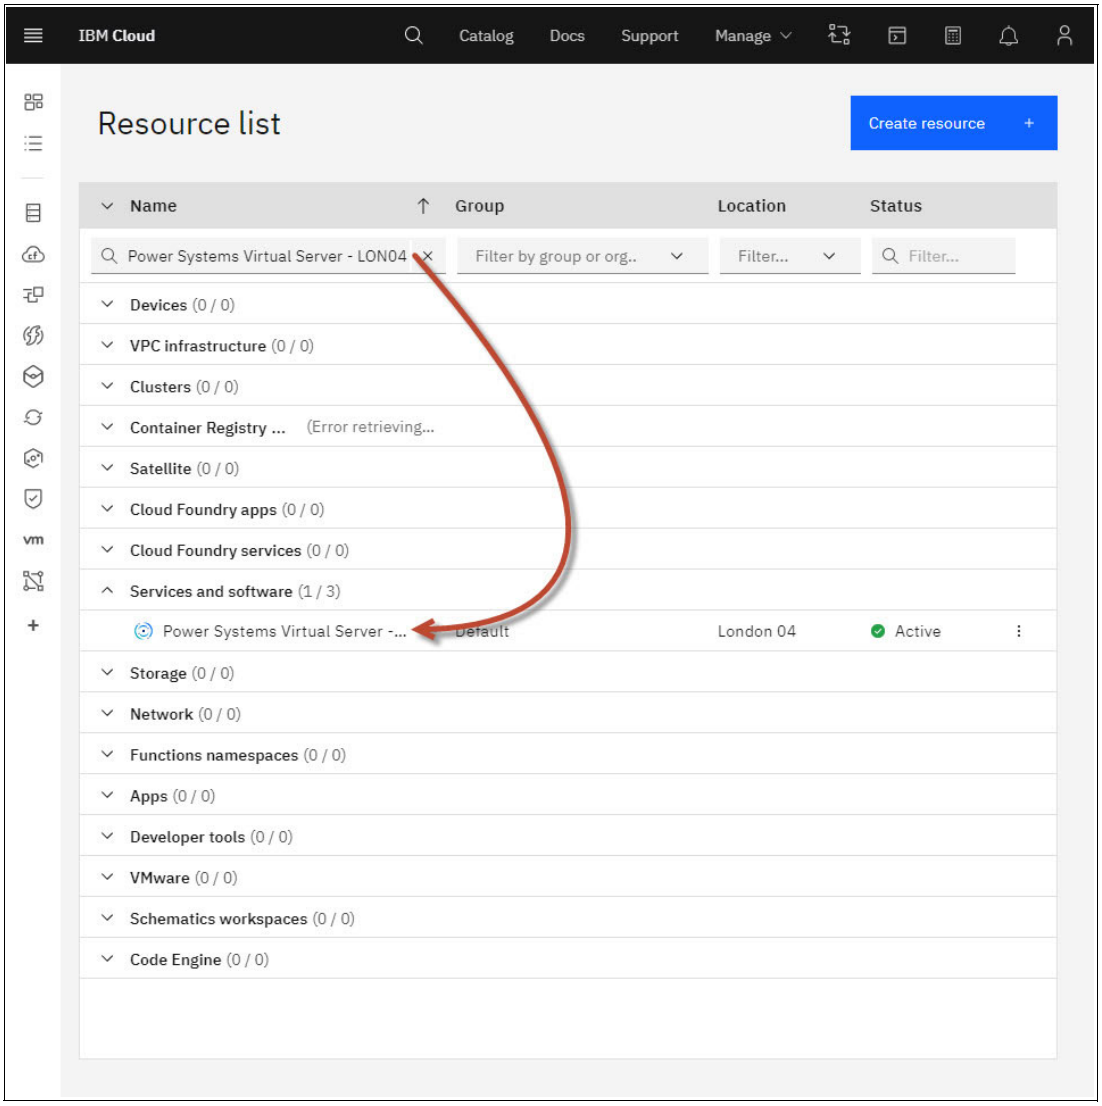

Figure 1-14 Search for resources in the IBM Cloud

1.5.4 Power Virtual Server instance

You can work with your resources in various ways from your resource list. To directly manage the Power Virtual Server service, click the name of the resource to go to the resource details page as shown in Figure 1-15.

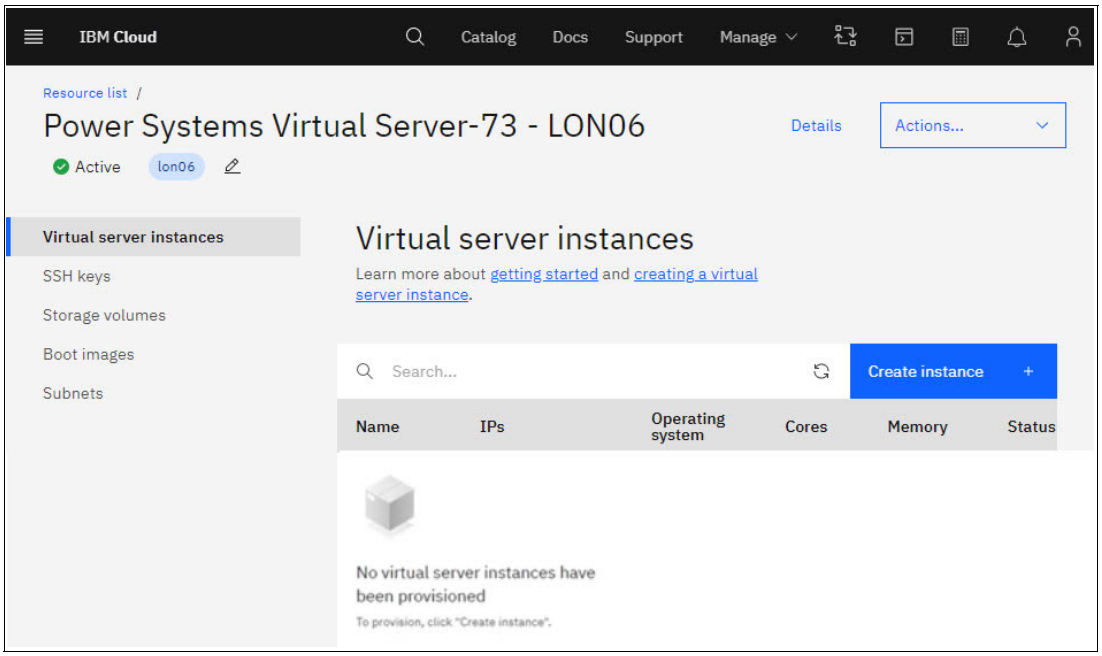

Figure 1-15 Virtual servers instances section

Click the **Create instance** button to create a Power Virtual Server instance (VSI), for example an LPAR or VM, then complete the required fields under the Virtual servers instances section.

The total due per month is dynamically updated in the Summary panel based on your selections as shown in Figure 1-16 on page 17. You can create a cost-effective Power Virtual Server instance that satisfies your business needs.

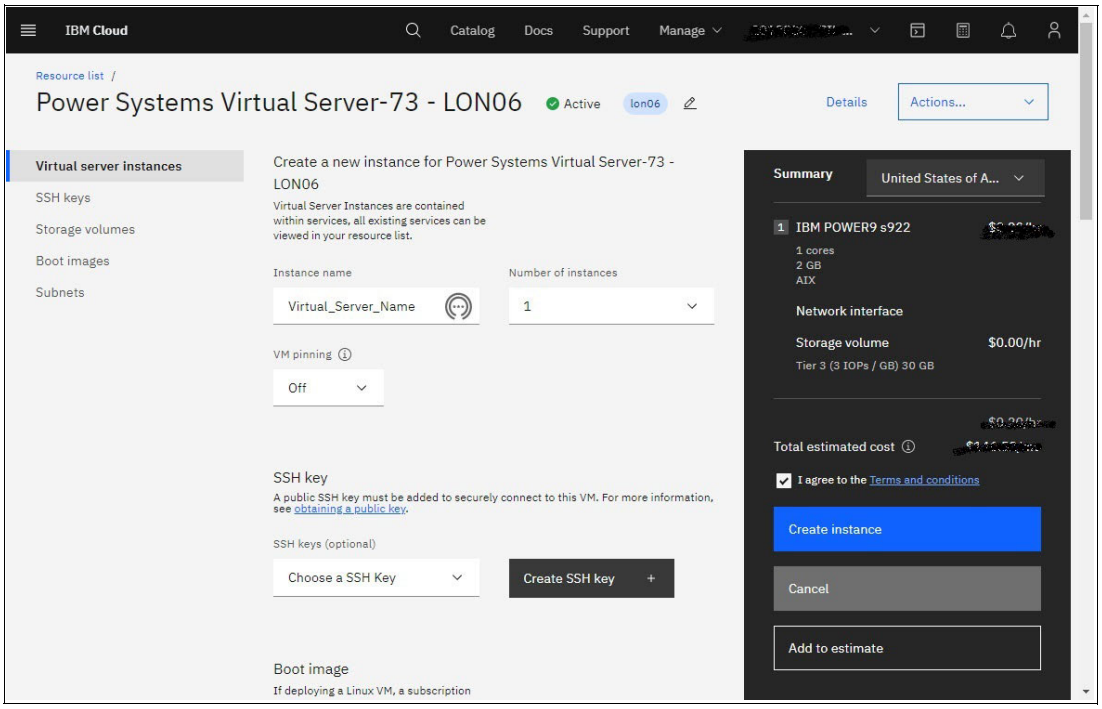

Figure 1-16 Virtual servers instances section and summary cost

1.5.5 Fields in the Virtual Servers Instances page

After you specify the instance, enter the required data in the remaining fields.

Number of instances

Specify the number of instances that you want to create for the Power Virtual Server. If you specify more than one instance, other options are available, such as hosting all instances on the same server or not and VM pinning. You can choose to soft pin or hard pin a VM to the host where it is running. When you soft pin a VM for high availability, PowerVC automatically migrates the VM back to the original host after the VM returns to its operating state. The hard pin option restricts the movement of the VM during remote restart, automated remote restart, Dynamic Resource Optimizer, and live partition migration.

SSH key

Choose an existing SSH key or create one to connect to your Power Virtual Server securely.

Machine type

Specify the machine type. The machine type that you select determines the number of maximum cores and maximum memory that is available.

Cores

There is a core-to-vCPU ratio of 1:1. For shared processors, fractions of cores round up to the nearest whole number. For example, 1.25 cores equal 2 vCPUs.

Memory

Select the amount of memory for the Power Virtual Server.

Boot image

Click **Boot image** to select boot images from a group of stock images or the list of images in your catalog.

Attached volumes

You can either create a new data volume or attach an existing one that you defined in your account.

Network interfaces

At least one private or public network is required. Network interfaces are created by adding a public network, private network, or both. When you add an existing private network, you can choose a specific IP address or have one auto-assigned.

1.5.6 VSI ready to use

After you click **Create an instance**, some messages are displayed about your instance, network, and disk space creation. Then, the main VSI page is displayed. See Figure 1-17. Note the status sequence: Build, Warning, and then Active.

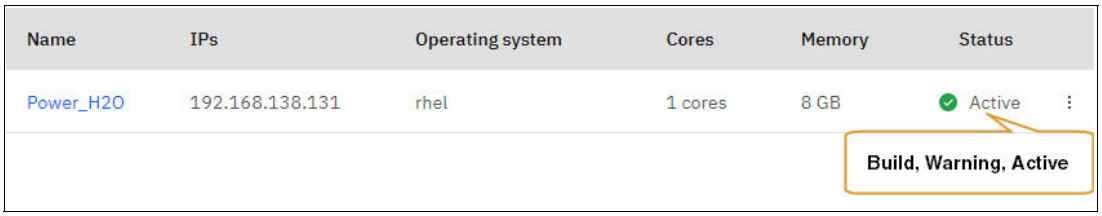

Figure 1-17 VSI status

When the VSI is running, you see an **Actions** icon included at the end of the row. When you click the icon, a pull-down menu is displayed from which you can select to **Immediate shutdown**, **Restart**, **Open console**, or **Delete** your instance as shown in Figure 1-18. If you click the **Open console**, a new window with the login prompt opens.

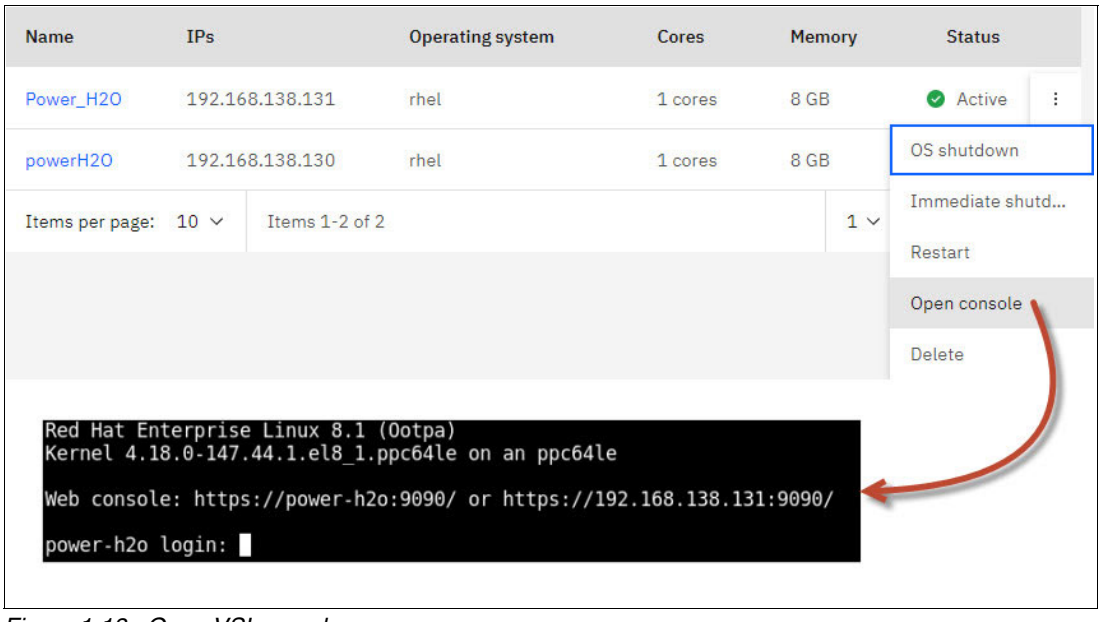

Figure 1-18 Open VSI console

The following chapters in this publication expand and show more details about managing the Virtual Server instances and the associated resources.

2

Migrating an IBM i LPAR to IBM Power Virtual Server

This chapter explains how to migrate IBM i LPARs from customer data centers to virtual machines (VMs) on IBM Power Virtual Server.

IBM i clients that want to use the benefits of IBM Power Virtual Server often want to move existing workloads to Power Virtual Server. The process typically involves saving the system on-premises, transferring the save to Power Virtual Server, and performing a restore.

A well-planned migration strategy is crucial to the project's success. Because every customer has unique business needs, the right strategy drives the migration project forward. This chapter discusses the various tools and techniques for transferring IBM i workloads, including the requirements, challenges, and advantages of each method.

This chapter includes the following topics:

- “IBM i data migration to Power Virtual Server”
- “Migration with partial saves and restores”
- “Backup and Restore migration”
- “Migrating IBM i to Power Virtual Server using PowerVC”
- “Other migration options in IBM Cloud catalog”
- “Migrating IBM i by using IBM i Migrate While Active”
- “Migration example and use case”

© Copyright IBM Corp. 2025.

2.1 IBM i data migration to Power Virtual Server

Chapter 1, “Introduction to IBM Power Virtual Server” on page 1 described how to create a Power Virtual Server instance to run IBM i workloads. To create a new instance, you can log in to the IBM i instance to set up and customize your LPAR. However, to move an existing IBM i workload to Power Virtual Server, you must move the LPAR configuration setup, programs, business logic, and associated data to the newly defined instances.

When you move workloads to a new system, you must consider business rules, technical resources, and the environment. In addition, IBM Cloud Power Virtual Servers have specific requirements and challenges to consider because you are using a self-service capability within a multi-tenant cloud environment.

Consider the following differences between the Power Virtual Server instance and an IBM i LPAR that is running on on-premises hardware:

- There is no access to the hardware layer, so no HMC or Storage direct access is provided, and hardware replication cannot be used.
- There is no access to the data center, so a physical tape cannot be sent to perform a restore.
- No physical tape is available for backup or restore operations.
- The system is in a remote facility, and communications can have some latency, depending on your location and link type.
- Most IBM i backup devices use SAS or Fibre Channel connections with proven reliability and fast transfer rates, but network backups are slow, and even backups to disk can seem slow compared to physical or virtual tape.
- Fewer third-party backup software solutions are available for IBM i than in other environments, and only a small portion of them use the network to transfer data.

2.1.1 Other challenges in cloud migrations

Sometimes, you must address the following business and technical constraints:

- There is not enough space for backing up to disk.
- Your communications are not fast enough to transfer data as you need.
- You use an outdated or unsupported operating system releases.
- You have third-party software with serial number licensing and validation.

Most IBM i customers use physical media for backup and restore, and their data recovery strategies include this media as a main option, even when a replication mechanism is in place. In the Power Virtual Server instance, the use of physical media is not an option, and the backup and restore process must be modified to use a supported option.

When data must be moved without disrupting the business operation, replication methods or a migration software setup can be used to reduce the backup window. PowerHA and third-party products provide solutions to transfer information gradually and synchronize systems while they are running. Figure 2-1 on page 23 illustrates how data mirroring can be used in a migration.

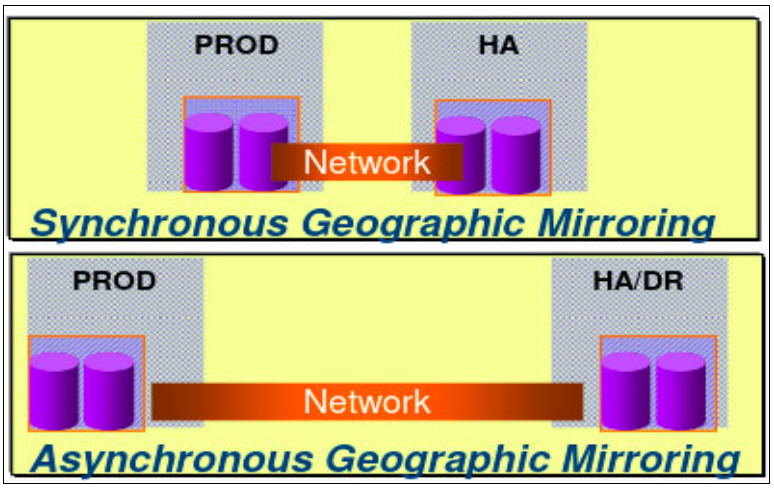

Figure 2-1 From PowerHA IBM SystemMirror® for i

Many resources are available that can be used to migrate data from on-premises systems to IBM Power Systems Virtual Servers. Therefore, planning is crucial to the migration process.

2.1.2 Migration options

Migrating an IBM i LPAR to Power Virtual Server requires the following steps:

1. Create the new Power Virtual Server instance
2. Save the current LPAR information in a format that can be transported or migrated to Power Virtual Server
3. Restore the LPAR information, including the IBM i settings, any programs installed, and the application data

Any migration from an on-premises LPAR to Power Virtual Server requires some downtime for users. Choosing the appropriate migration scenario is important to match the length of the downtime to your business requirements. One of the major considerations in planning your migration is the amount of data to be migrated and the amount of time that is available to do the migration. Larger LPARs with a large amount of data require more planning than smaller LPARs. For non-critical LPARS with a tolerance for longer outages for the migration, there are many options available. For critical LPARs that can tolerate only a short outage, choose a migration option that uses data replication and synchronization. Consider an IBM i migration tool like IBM PowerHA, MIMIX by Assure, or IBM i Migrate While Active to meet your business requirements.

If you need assistance with planning and implementing a migration of your IBM i workloads to Power Virtual Server, IBM Expert Labs has a standard offering available to assist you. For more information, see IBM Expert Labs Migrate Services for IBM Power or contact IBM Expert Labs or your IBM sales representative.

2.2 Migration with partial saves and restores

Migration with partial saves and restores consists of using multiple save files, ISO images with virtual optical media, or virtual tape images with which a system can be backed up in parts. These files are then transferred to the new Power Virtual Server instance by using FTP, SFTP, SCP, or any other method.

The lack of HMC access forces users to use restart mechanisms and requires the creation of an empty VM instance with a working IBM i operating system. Partial restores can be a good option to migrate LPARs when the operating system images that are provided on IBM Cloud fit a user's needs. This scenario can be easier to deploy because the target system is already running.

By using a working IP address, communications can be checked. Stock images frequently include the latest re-saves and PTFs. Most of the software required to run workloads in IBM Power Virtual Server is in place:

- BRMS with basic policies.
- IBM i Cloud Storage Solution (5733ICC) running with certificates.
- SSL certificates are ready for IBM Access.
- YUM is installed.

Figure 2-2 shows a high-level view of the process.

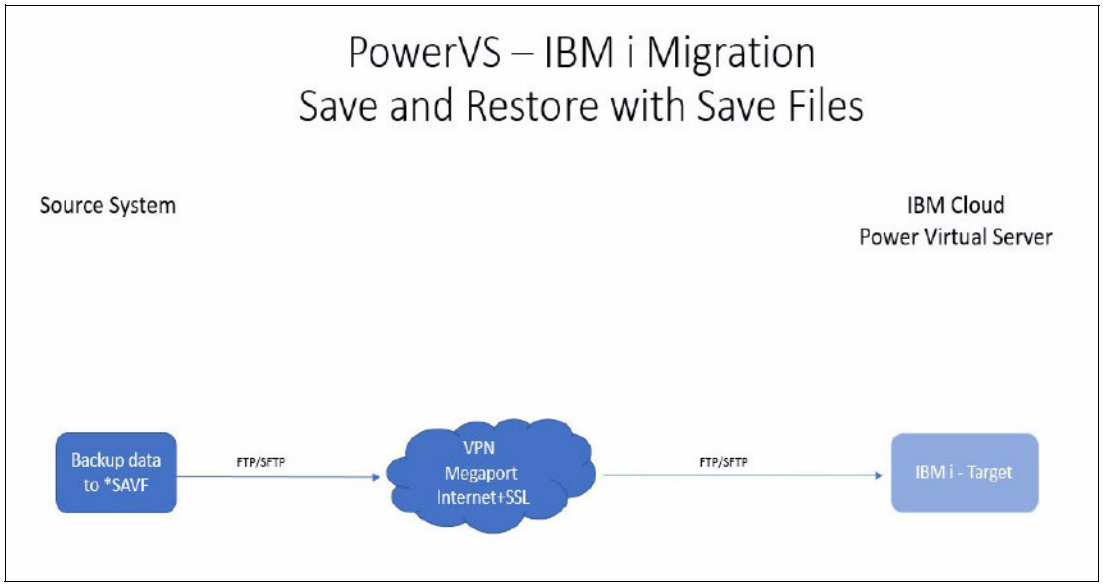

Figure 2-2 Save and restore by using save files

A sample scenario that uses a backup and restore scenario includes the following steps:

1. Create the Power Virtual Server service in one location (in this example, Dallas 13).
2. Create an IBM i instance or LPAR with an existing image and add the needed resources with at least six disks, 0.25 CPUs, and 8 GB of memory so that the VM runs smoothly. This system requires more space for backup files when restoring.

Note: A proper sizing is recommended. Request a sizing from IBM or your IBM Business Partner.

3. Communications must be configured between the local network and IBM Cloud so that the system that was created from the source data center can be reached.
4. To allow object restore, some changes by using the target server console might be required. See Example 2-1 on page 25.

Example 2-1 Changes to allow object restore

```
CHGSYSVAL SYSVAL(QCTLSBSD) VALUE('QCTL')
CHGSYSVAL SYSVAL(QALWOBJRST) VALUE(*ALL)
CHGSYSVAL SYSVAL(QALWUSRDMN) VALUE(*ALL)
CHGSYSVAL SYSVAL(QAUTOVRT) VALUE(*NOMAX)
CHGSYSVAL SYSVAL(QSECURITY) VALUE('40')
CHGSYSVAL SYSVAL(QJOBMSGQFL) VALUE(*PRTWRAP)
CHGSYSVAL SYSVAL(QJOBMSGQMX) VALUE(64)
CHGSYSVAL SYSVAL(QJOBMSGQSZ) VALUE(16384)
CHGSYSVAL SYSVAL(QJOBMSGQTL) VALUE(16384)
CHGSYSVAL SYSVAL(QQRYDEGREE) VALUE(*OPTIMIZE)
CHGTCPA TCPKEEPALV(20) TCPRCVBUF(1048576) TCPSNDBUF(1048576)
CHGSHRPOOL POOL(*BASE) ACTLVL(4000) PAGING(*CALC)
```

Note: For some stock images, enable System Value changes from DST/SST.

5. Start ***FTP** or ***SSHD** servers to allow file transfer as shown in Example 2-2.

Example 2-2 Start services

```
STRTCPSVR *SSHD
STRTCPSVR *FTP
```

6. On the source system, enough disk space is needed so that the system can be backed up to disk by using Save Files or IMAGE CATALOG (optical or tape). In this example, Save Files are used. A control language (CL) with all the commands is recommended, as shown in Example 2-3.

Example 2-3 CL commands for saving data

```
CRTLIB BKPIBMI
CRTSAVF BKPIBMI/SAVSECDTA1
SAVSECDTA DEV(*SAVF) SAVF(BKPIBMI/SAVSECDTA1)
CRTSAVF BKPIBMI/SAVCFG1
SAVCFG DEV(*SAVF) SAVF(BKPIBMI/SAVCFG1)
CRTSAVF BKPIBMI/MYLIB1
SAVLIB LIB(MYLIB1) DEV(*SAVF) SAVF(BKPIBMI/MYLIB1) SAVACT(*LIB) SAVACTWAIT(30)
SAVACTMSGQ(*WRKSTN)
```

7. In this step, all libraries are migrated and saved as shown in Example 2-4.

Example 2-4 Libraries to be migrated

```
CRTSAVF BKPIBMI/SAVDLO1
SAVDLO DLO(*ALL) DEV(*SAVF) SAVF(BKPIBMI/SAVDLO1)
CRTSAVF BKPIBMI/SAVIFS1
SAV DEV('/QSYS.LIB/BKPIBMI.LIB/SAVIFS1.FILE') OBJ(('/home/'))
SAVACT(*YES)SAVACTOPT(*ALL)
```

   All files that must be migrated are saved on IFS.

8. Upload the Save Files to the target system.
9. On the target system, create the temporary library and the empty Save Files on it.
10. Transfer backups to the target LPAR by using FTP, FTPS, or SFTP. In this example, FTP is used.

11. End the subsystems with the following command:

```
ENDSBS *ALL *IMMED
```

12. Restore user profiles and security with the following command:

```
RSTUSRPRF DEV(*SAVF) SAVF(BKPIBMI/SAVSECDTA) ALWOBJDIF(*ALL)
OMITUSRPRF(QSECOFR)
```

13. Restore system configuration:

```
RSTCFG OBJ(*ALL) DEV(*SAVF) SAVF(BKPIBMI/SAVCFG1) SRM(*NONE) ALWOBJDIF(*ALL)
```

14. Restore all your libraries, one by one:

```
RSTLIB SAVLIB(MYLIB1) DEV(*SAVF) SAVF(BKPIBMI/MYLIB1) MBROPT(*ALL)
ALWOBJDIF(*ALL)
```

15. Because Save Files are used for this example, deferred mode cannot be used. Therefore, the previous step must be repeated to avoid errors with logical files.

16. Restore the documents:

```
RSTDLO DLO(*ALL) DEV(*SAVF) SAVF(BKPIBMI/SAVDLO1) ALWOBJDIF(*ALL) SAVASP(*ANY)
```

17. Restore folders and files from IFS, which can be one or more save files to restore:

```
RST DEV('/QSYS.LIB/BKPIBMI.LIB/SAVIFS1.FILE') CRTPRNDIR(*YES)
```

18. Restore the authorities:

```
RSTAUT *ALL
```

19. Review the messages. If an error occurs and the restore is unsuccessful, attempt the restore again.

20. Check the devices, line descriptions, and any configuration that affects a correct system start.

21. Adjust the net name:

```
CHGNETA SYSNAME(MYIBMI) LCLCPNAME(MYIBMI) LCLLOCNAME(MYIBMI) DTACPR(*ALLOW)
DTACPRINM(*REQUEST) NETSERVER((MYIBMI *ANY)) NWSDOMAIN(MYIBMI) ALWVRTAPPN(*YES)
ALWHPRTWR(*YES)
```

22. IPL the target system:

```
PWRDWNSYS *IMMED RESTART(*YES) CONFIRM(*NO)
```

Check and validate the new system. If the system passes validation, the new server can be published.

2.3 Backup and Restore migration

Most IBM i customers use physical tapes for backups of day-to-day operations. Customers also use tapes for migrating workloads to a new system. Tapes are the main migration method when moving data (and the operating system) from the system to a new box.

However, physical tapes cannot be used with cloud environments. Therefore, alternative methods must be used, such as saving data to disk and by using the network for transfer. An additional option is available by using FalconStor StorSafe, which is a virtual tape library solution available in the IBM Cloud Catalog for Power Virtual Server environments.

Using FalconStor StorSafe as a migration option involves installing FalconStor StorSafe in the customer location and connecting it to the FalconStor StorSafe instance in the Power Virtual Server. Data is backed up to FalconStor StorSafe in the customer location and replicated to

the Power Virtual Server instance where it can be restored to the Power Virtual Server instance. This is described in section 4.4.3, “IBM i migration example that uses FalconStor StorSafe” on page 172

Important: It is important to mention that you can save your data to Physical tape and convert the media to Image Catalog Virtual Tape cartridges in a different system with enough space and the adequate communications to reach Cloud Object Storage or your target system.

2.3.1 Considerations for backup and restore migrations

Consider the following aspects of using backup and restore for migration:

- When data is moved to the cloud, data must be uploaded by using the network. Sometimes, the target systems can be accessed, and the data is stored in an intermediate “bucket” at the target system. IBM Cloud Object Storage can be a useful environment for this intermediate storage. Because the amount of data that is being migrated affects the time that it takes to transfer to a target environment, compression can reduce the amount of data to transfer.
- Virtual media can be very large, depending on the amount of data that must be transferred. You can use multiple virtual media devices to enable parallel file transfers, which can improve transfer speeds and reduce the time that is required to transfer the data. Consider the size of your virtual media devices to enable this parallel connectivity.
- Media types often use UDF and ISO formats because they can be combined with virtual optical devices. Also, their size can be 4.7 GB if you want to create a physical media for your own purposes or as large as 1 TB in the versions of IBM i that are available at the time of writing.

  For more information, see Image Catalog and Virtual Media - Maximum Capacities and Limitations.
- There are multiple ways to manage your backups within the IBM i environment. The following backup options are available on IBM i:
  - Native commands. SAV, SAVLIB, SAVDLO, SAVCFG, SAVSECDTA, and menu GO SAVE.
  - Backup and Recovery Media Services for i (BRMS). A complete backup solution for IBM i environment with many features to automate your backup and restore process.
- IBM Cloud provides Cloud Object Storage for storing information. Data can be uploaded to Cloud Object Storage by using SSL, VPN, or dedicated communications to transfer backups.
- IBM Aspera® software improves data transfers with IBM Cloud Object Storage and improves the upload of backups by using multiple streams. You can run IBM Aspera client from a PC or side server to enable faster file transfers to IBM Cloud.
- When the amount of data to be transferred is too large to use the network, IBM Cloud can be used to send a Massive Data Migration device. This portable storage device can save up to 120 TB and then this data can be uploaded directly to the IBM Cloud network, which avoids direct file transfers from IBM data centers.
- When your system is up and running, plan a data backup solution. See Chapter 4, “Backup for IBM i on IBM Power Virtual Server” on page 117 for more information on solutions to back up data in Power Virtual Server instances.

2.3.2 Full system backup and restore by using native commands

Full system backup saves the operating system's configurations, security, and data to media to perform a full system restore from bare metal on a new system.

This process is the most traditional migration method. Most customers are confident when performing a full system backup by using physical tape devices or Virtual Tape Libraries (VTL) because this process is part of the Disaster Recovery procedure.

IBM i can be installed from this special backup, starting with License Machine Code (SLIC) and then the operating system. However, because cloud environments cannot use the physical media that is required for a "D" IPL, a different procedure must be used.

On IBM Power Virtual Servers, the VM cannot be started by using network installation parameters, and license keys and IP addresses are needed to identify the correct Ethernet adapter to use when restoring the system. Therefore, a starting point must be created by using a new empty instance that uses the target operating system.

At the time of this writing, the following versions of IBM i are supported in IBM Cloud Power Virtual Server service:

- IBM i versions 7.1, 7.2, 7.3 and 7.4 are supported on Power8.
- IBM i versions 7.2, 7.3, 7.4, and 7.5 are supported on Power9.
- IBM i versions 7.3, 7.4, and 7.5 are supported on Power10. Power 10 9028 supports IBM i versions 7.4 and 7.5 only.

When a restore must be done from this virtual media, special requirements that are based on cloud characteristics are important. The iSCSI VTL for IBM i or an IBM i NFS server that is combined with TFTP service can be used so that the system can use BOOTP and load the SLIC that is necessary to IPL.

In this section. a temporary IBM i VM with NFS and TFTP is used to install the system from the full system backup media. This special restore process works with Virtual Optical media, and the backups must be done on this media.

Some specific requirements must be met to perform this migration, and some of them need time to be provisioned. Therefore, good planning is essential for success.

The following Bill of Materials and services were used:

- Target system. IBM Power Virtual Server with IBM i and enough resources for data and workloads.
- Network installation server. IBM Power Virtual Server with IBM i. At least 1 core and 32 GB, and enough disk space for backup images. For more information, see "Setting up IBM i network installation server with NFS Server and NFS Client" on page 29.
- Communications between source and target system. Usually, a Power Edge Router (PER), Transit Gateway between Power Virtual Servers and IBM Cloud VPC environment, and a VPN gateway to route traffic between VLANs and enable an IPsec tunnel to the on-premises environment are required.
- Optional. Space in Cloud Object Storage to store the backup.
- An NFS server on the source environment to back up data. Any version of Linux can be used, such as Red Hat Enterprise Linux, CentOS, Fedora, SUSE, and Debian. The example that is presented here was made by using Ubuntu Server.
- A private network (subnet) on the IBM Power Virtual Server workspace to connect both servers

- Source System. Enough available disk for backup.

Tip: When disk space is not available on the source system, use an NFS server or a Mass Data Migration device for backup purposes.

Setting up IBM i network installation server with NFS Server and NFS Client

Complete the following steps to set up IBM i network installation server with NFS server and NFS client:

1. Provision an IBM i VSI in the target Power Systems Virtual Server location to be an NFS Server. Run the IBM i NFS Server on Version 7.2 or later with current PTFs.
2. To use virtual optical images through an NFS server, the IBM i NFS client must meet the following requirements:
   - The IBM i includes a Version 4 Internet Protocol (IP) address.
   - During setup, the shared NFS server directory is mounted over a directory on the IBM i client.
   - An IBM i service tools server or a LAN console connection is configured by using a Version 4 IP address.
   - A 632B-003 virtual optical device is defined that uses the IP address of the NFS server.

Note: The IP address of both the IBM i server and the IBM i service tools server (LAN console connection) can be the same.

Setup process

Figure 2-3 shows how data flows in this type of migration process.

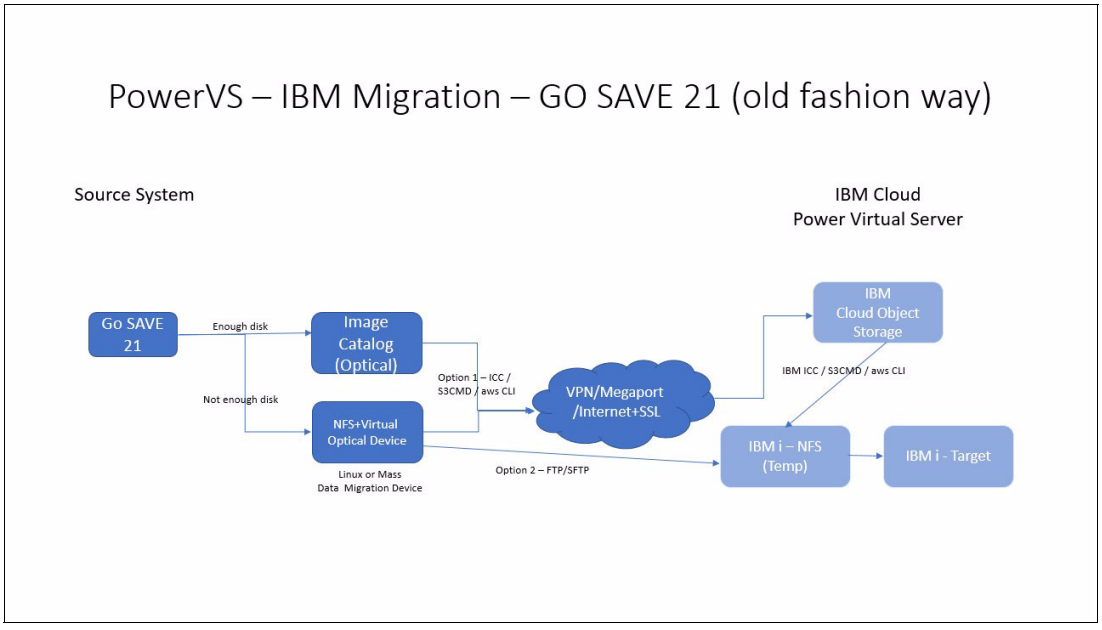

Figure 2-3 IBM i Migration that uses GO SAVE option 21

As shown in Figure 2-3, different mechanisms are based on resources and the migration strategy that is used.

The process includes the following steps:

1. Back up IBM i by using GO SAVE menu Option 21 and save to the NFS server disk by using the Image Catalog with Virtual Optical media.
2. Transfer data without compression to IBM Cloud Object Storage by using the internet and SSL.
3. Download data from IBM Cloud Object Storage to IBM i NFS Server by using IP address `192.168.80.12`.
4. Install IBM i on the target server. The IP address is `192.168.80.11`. The Service Tools Server LAN Adapter IP address is `192.168.80.21`.

Backup process

It is assumed that not enough space is available on the IBM i source server to be migrated. When the free disk space is less than 50%, a full system backup can be performed to an NFS disk.

Consider the following points:

- A working Linux partition with enough space is needed to accommodate the backup images. In this example, an Ubuntu server is used, but any distribution can be used. It is a good practice to use more than two physical cores and SSDs or SAN-attached storage when possible.
- The NFS directions are based on instructions that are found in IBM i save and restore using virtual optical images on an NFS server.
- IBM i can back up to an NFS server from V7R2 and later. Previous versions can also back up in this way, but with limitations. The test environment was validated by using V7R2 and V7R4.
- Install the latest cumulative PTF package, Technology Refresh, and Group PTFs that are related to TCP/IP, Backup and Restore, and IBM Db2®.
- A `cloud-init` program must also be installed to communicate with IBM Cloud Power Virtual Server installation process. For more information, see the following resources:
  - Cloud-Init Support for IBM i
  - RPM pile for IBM i releases in standard support
- Linux and IBM i must be connected by a network with a speed of 1 Gbps or faster.

Complete the following steps:

1. Install the following tools on your Linux partition by using YUM or IBM Access Client Solutions for i (see Example 2-5):
   - Python3
   - NFS Server
   - awscli (AWS Client works with most S3 servers)
   - pigz

Example 2-5 Install Linux tools

```
sudo apt install python3 python3-pip nfs-kernel-server awscli pigz nfs-common
```

2. Set up an NFS server on the Linux system. In Example 2-6, the source network is `IBMi04` with an IP address of `192.168.50.244` and Linux `192.168.50.22`.

Example 2-6 NFS directory creation

```
#Create NFS work directory
mkdir /NFS
chmod 755 /NFS
mkdir /NFS/IBMi04
chmod 777 /NFS/IBMi04
# Now edit the /etc/exportfs file
sudo vi /etc/exportfs
# Add the following line and save the file with wq!
/NFS    192.168.50.0/24(rw,sync,no_subtree_check)
```

3. Restart the NFS service on Linux as shown in Example 2-7.

Example 2-7 Restart NFS

```
# Restart NFS sevice
sudo systemctl restart nfs-kernel-server
# Export shares
sudo exportfs -ra
# Validate our shares
showmount --exports localhost
```

4. At IBMi04, set the Service Tools Server from SST or DST. For more information, see the following resources:
   - Configuring the service tools server using SST
   - Configure a service tools server for DST for the virtual optical device to use

   Consider the following points:
   - The Line Description that is associated with the IP address must be in the same network as the Linux NFS server:

     ```
     CFGTCP
     Option 1
     F11
     ```

   - This command in the following example shows the network interface and the Line Description name:

     ```
     ETH01:
     DSPLIND ETH01
     ```

5. Start Service Tools (SST): **`STRSST`**
6. Select **Option 8 Work with Service tools Server Security and Devices** (see Figure 2-4 on page 32).

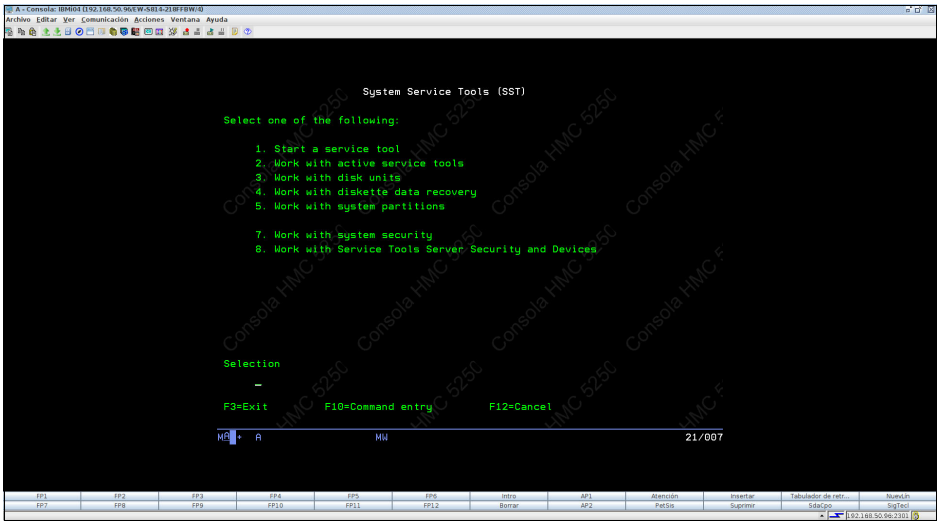

Figure 2-4 SST Main Menu

7. Press **F13** to Select STS LAN adapter, as shown in Figure 2-5.

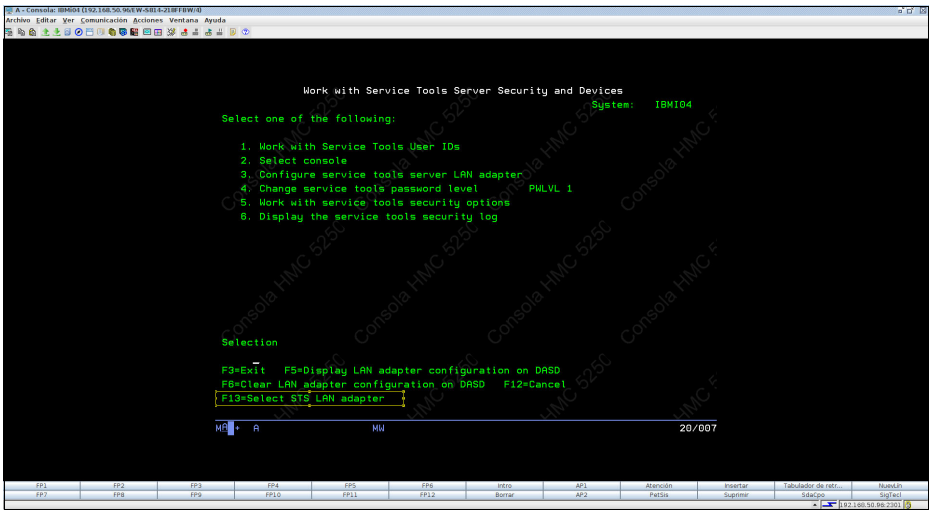

Figure 2-5 SST Work with Service Tools Server Security and Devices

8. Enter 1, which is the same resource that was found with **DSPLIND** and press **Enter**, as shown in Figure 2-6 on page 33.

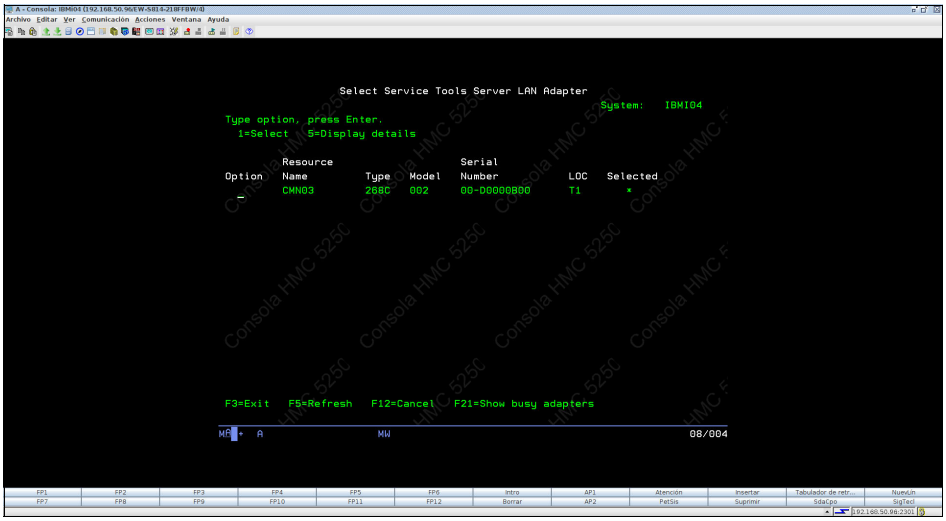

Figure 2-6 Select Service Tools Server LAN Adapter

9. As shown in Figure 2-7, set the hostname, address (which is different from the main interface), gateway address, and subnet mask.
10. “Store” the LAN adapter by pressing **F7**, deactivate the LAN adapter by pressing **F13**, and Activate the LAN adapter by pressing **F14**.
11. Press **F3** to exit Service Tools.



Figure 2-7 Configure Service Tools Server LAN Adapter menu

12. Restart the system to take the changes:

```
PWRDWNSYS *IMMED RESTART(*YES) CONFIRM(*NO)
```

13. Log in to the system again and add an entry to the host table:

```
ADDTCPHTE INTNETADR('192.168.50.22') HOSTNAME((NFSSERVER))
```

14. Mount the NFS share and start creating the virtual media, as shown in Example 2-8 on page 34.

Note: It is important to know the amount of data to be backed up so that enough virtual media can be created. It is not possible to create the images during the backup process.

Example 2-8 Mount resource

```
# Mount the resource
MOUNT TYPE(*NFS) MFS('192.168.50.22:/NFS/IBMi04') MNTOVRDIR('/NFS')

# Create the empty files - 4 x 10GB virtual optical disks
CALL QP2TERM

cd /NFS/IBMi04
dd if=/dev/zero of=IMAGE01.ISO bs=1M count=10240
dd if=/dev/zero of=IMAGE02.ISO bs=1M count=10240
dd if=/dev/zero of=IMAGE03.ISO bs=1M count=10240
dd if=/dev/zero of=IMAGE04.ISO bs=1M count=10240

# Now we need to create the VOLUME_LIST file entries to tell our device how to
change media when the mounted ISO gets full.

echo 'IMAGE01.ISO W'  > VOLUME_LIST
echo 'IMAGE02.ISO W' >> VOLUME_LIST
echo 'IMAGE03.ISO W' >> VOLUME_LIST
echo 'IMAGE04.ISO W' >> VOLUME_LIST
```

15. Create a device description that is linked to the Service Tools Server LAN adapter, as shown in Example 2-9.

Example 2-9 Device creation

```
# Create the device
CRTDEVOPT DEVD(SAV2NFS) RSRCNAME(*VRT) LCLINTNETA(*SRVLAN)
RMTINTNETA('192.168.50.22') NETIMGDIR('/NFS')
# Vary on the device
VRYCFG CFGOBJ(SAV2NFS) CFGTYPE(*DEV) STATUS(*ON)
```

16. Initialize the virtual optical media, as shown in Example 2-10.

Example 2-10 Loading images to catalog

```
LODIMGCLGE IMGCLG(*DEV) IMGCLGIDX(4) DEV(sav2nfs) INZOPT NEWVOL(IBMi04_4)
DEV(sav2nfs) CHECK(*NO)
LODIMGCLGE IMGCLG(*DEV) IMGCLGIDX(3) DEV(sav2nfs) INZOPT NEWVOL(IBMi04_3)
DEV(sav2nfs) CHECK(*NO)
LODIMGCLGE IMGCLG(*DEV) IMGCLGIDX(2) DEV(sav2nfs) INZOPT NEWVOL(IBMi04_2)
DEV(sav2nfs) CHECK(*NO)
LODIMGCLGE IMGCLG(*DEV) IMGCLGIDX(1) DEV(sav2nfs) INZOPT NEWVOL(IBMi04_1)
DEV(sav2nfs) CHECK(*NO)
```

17. Run the cloud-init process:

```
CALL PGM(QSYS/QAENGCHG) PARM(*ENABLECI)
```

18. Call the SAVE menu and run option 21 for a full system:

```
ENDTCP
ENDSBS *ALL *IMMED
```

19. Click **Save** → **Option 21** and press Enter.

    On the main window, use the `SAV2NFS` device name.

20. After the backup is done, deactivate the device:

    ```
    VRYCFG CFGOBJ(SAV2NFS) CFGTYPE(*DEV) STATUS(*OFF)
    ```

21. For example purposes, uploads are made to IBM Cloud Object Storage. However, you can use FTP or SFTP to upload your files from your NFS Linux system to the NFS IBM i server:

    – Create a bucket and configure the **`awscli`**. For more information, see Using the AWS CLI.

    – Create your Object Storage resource and bucket, as shown in Figure 2-8, Figure 2-9, Figure 2-10 on page 36, Figure 2-11 on page 36, Figure 2-12 on page 37, and Figure 2-13 on page 37.

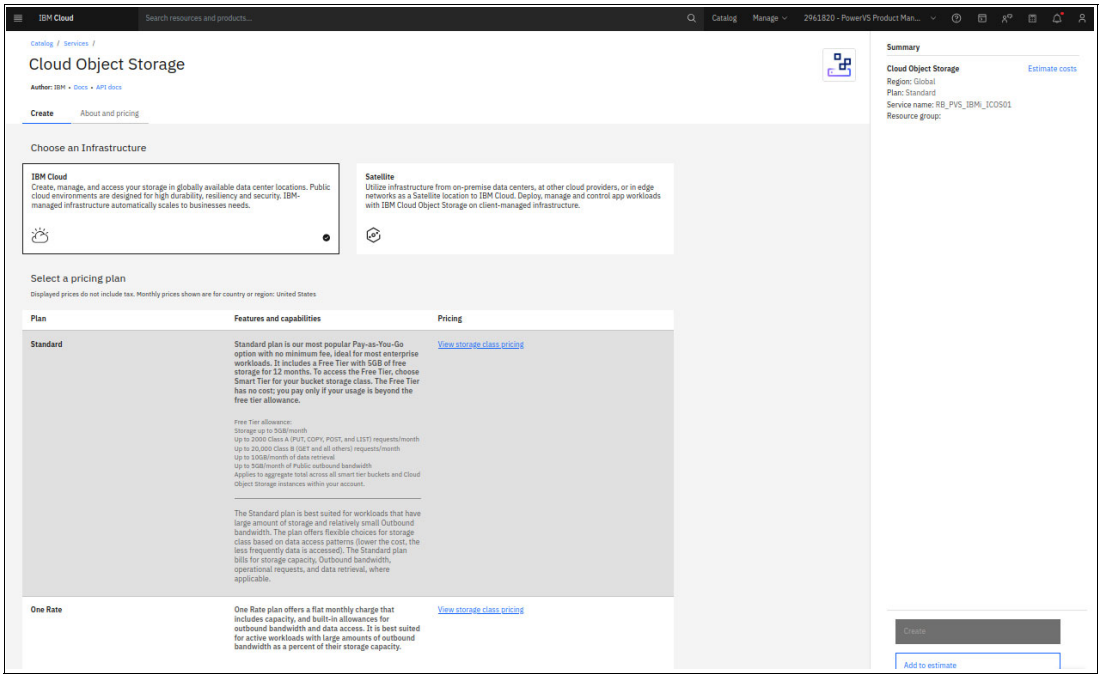

Figure 2-8 Select the resource type: Lite or Standard and create the resource

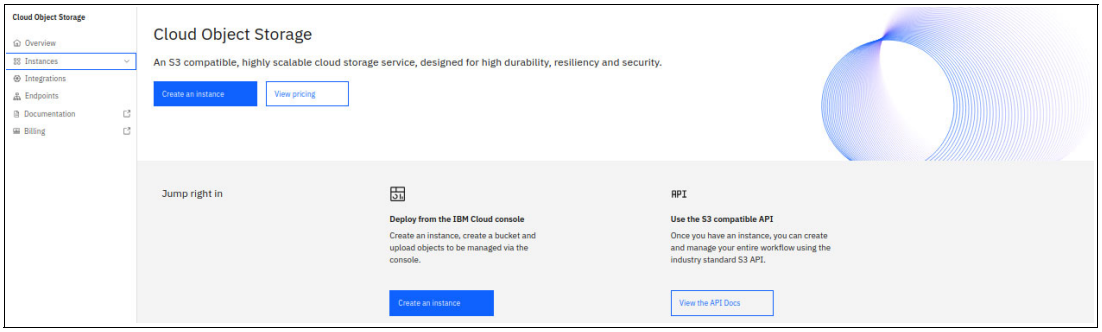

Figure 2-9 Cloud Object Storage main window

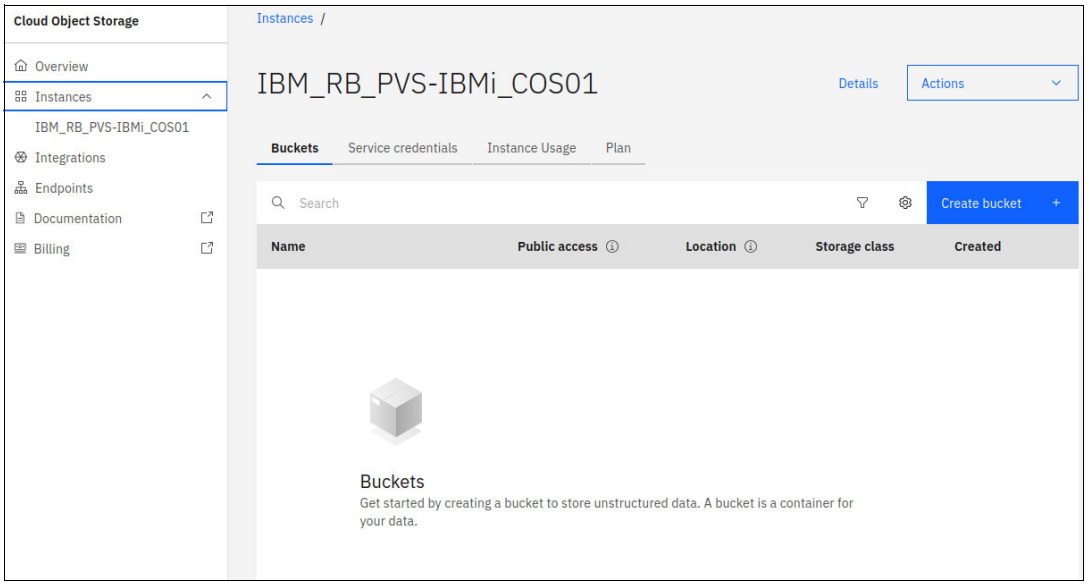

Figure 2-10 Select and Create the bucket

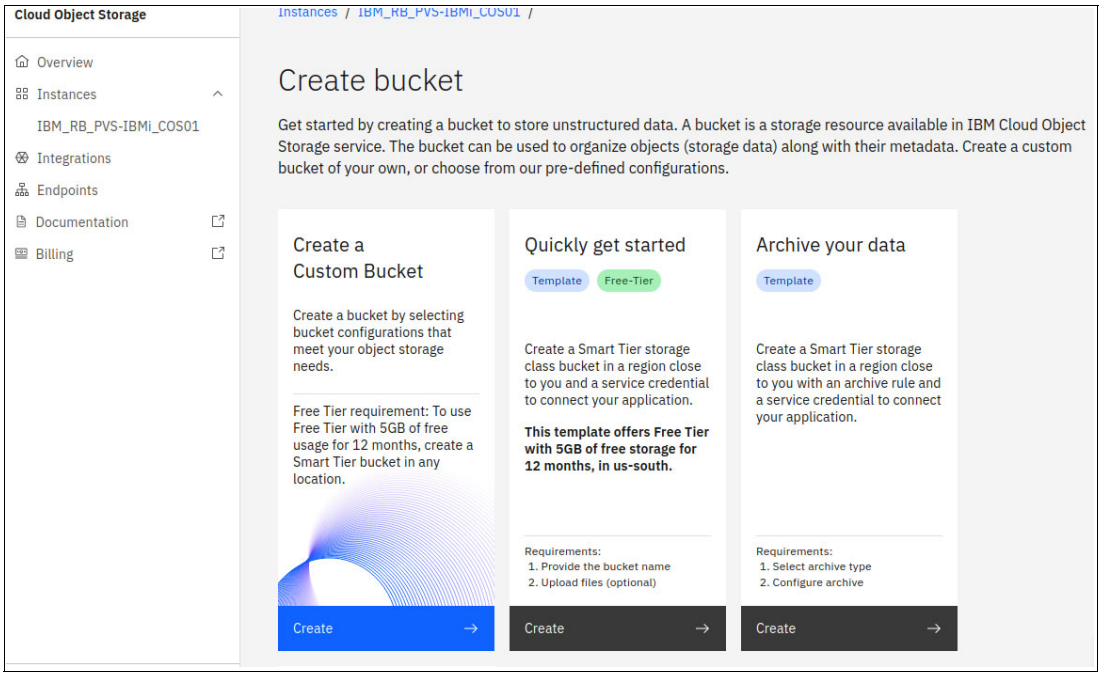

Figure 2-11 Create your bucket

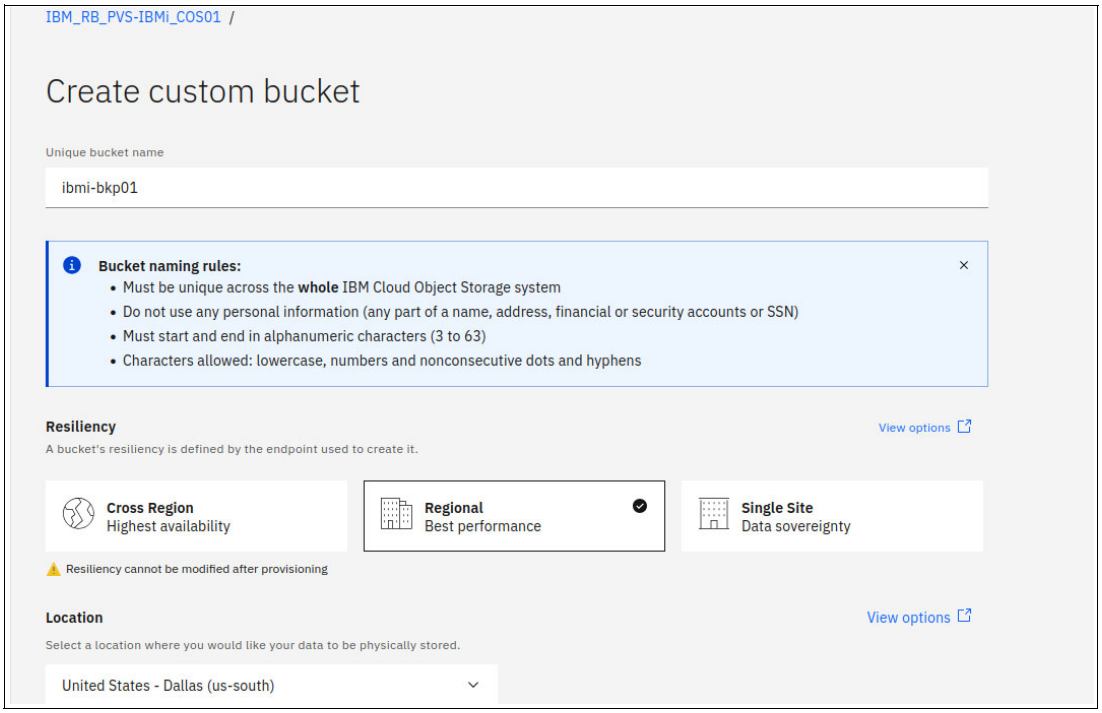

Figure 2-12 Select bucket configuration

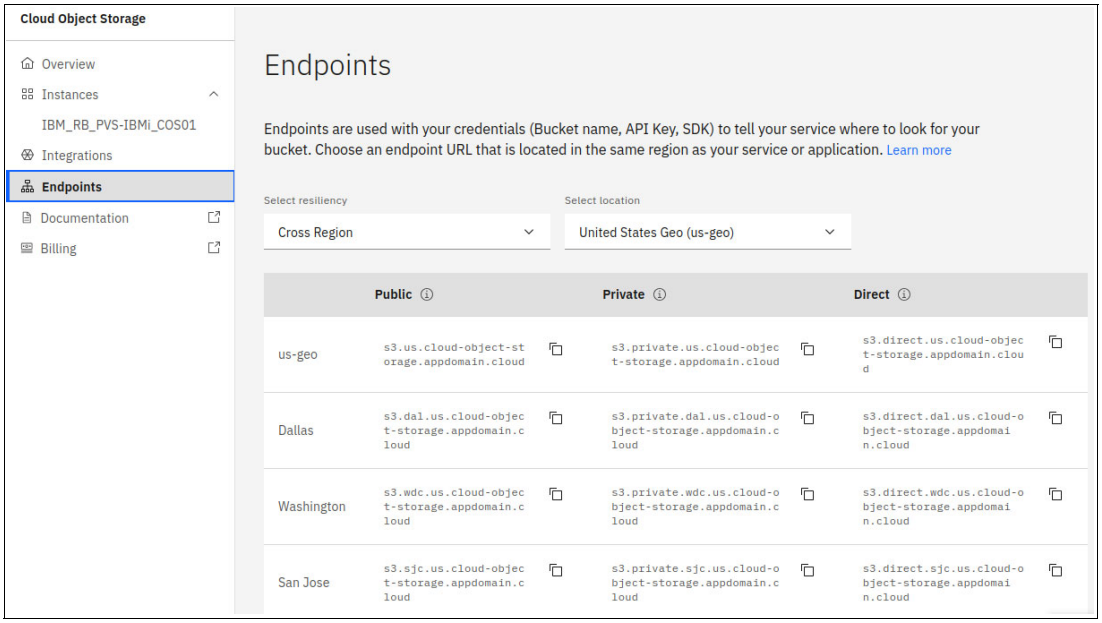

Figure 2-13 The endpoint option shows URLs

22. Take note of the endpoint URI. When connecting through the VPN, select **Private**. When connecting to the internet, use **Public** (see Figure 2-14, Figure 2-15, Figure 2-16, and Figure 2-17 on page 39).

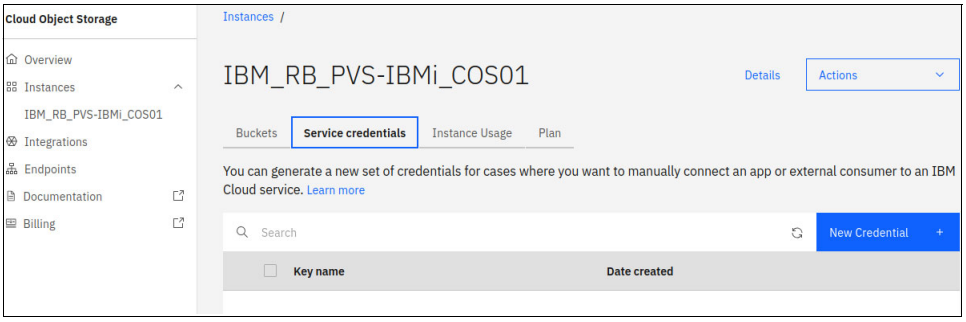

Figure 2-14 Service Credentials - Create Credential



Figure 2-15 Create Credential - Advanced Option - Include HMAC Credential

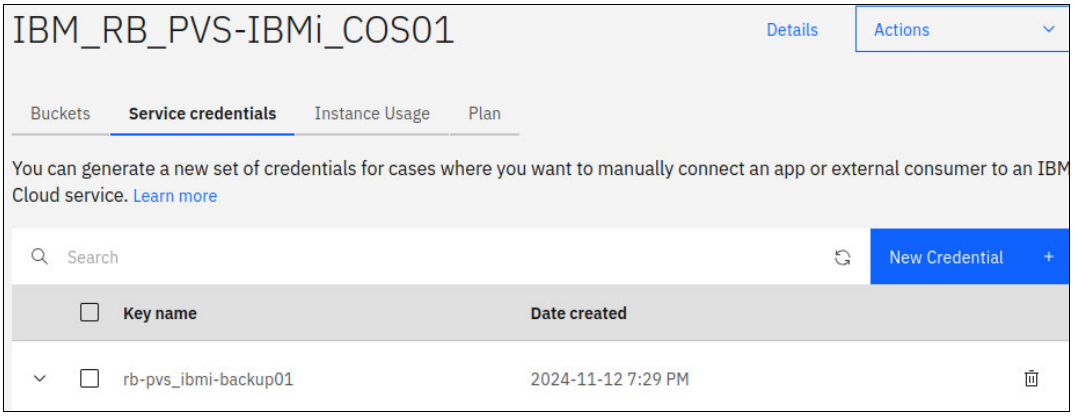

Figure 2-16 List new credential



Figure 2-17 Expand credential details

Take note of the Access Key and the Secret Key.

23. Return to Linux.

24. AWS configure:

    – Use your credentials to log in.
    – For this example, the region is `us-south` and the default output format is **None**.

25. Compress and upload the image files. To avoid the use of more disk space, compress the files as they are being uploaded, as shown in Example 2-11.

Example 2-11 Compressing and uploading image files

```
cd /NFS/IBMi04
cat IMAGE01.ISO | pigz -9 -p40 -c | aws
--endpoint-url=https://s3.us-south.cloud-object-storage.appdomain.cloud s3 cp -
s3://ibmi-bkp01/IMAGE01.ISO.gz

cat IMAGE02.ISO | pigz -9 -p40 -c | aws
--endpoint-url=https://s3.us-south.south.cloud-object-storage.appdomain.cloud s3 cp -
s3://ibmi-bkp01/IMAGE02.ISO.gz

cat IMAGE03.ISO | pigz -9 -p40 -c | aws
--endpoint-url=https://s3.us-south.cloud-object-storage.appdomain.cloud s3 cp -
s3://ibmi-bkp01/IMAGE03.ISO.gz

cat IMAGE04.ISO | pigz -9 -p40 -c | aws
--endpoint-url=https://s3.us-south.cloud-object-storage.appdomain.cloud s3 cp -
s3://ibmi-bkp01/IMAGE04.ISO.gz
```

26. After the process finishes, verify whether all files were uploaded to the Cloud Object Storage portal, as shown in Figure 2-18 on page 40.

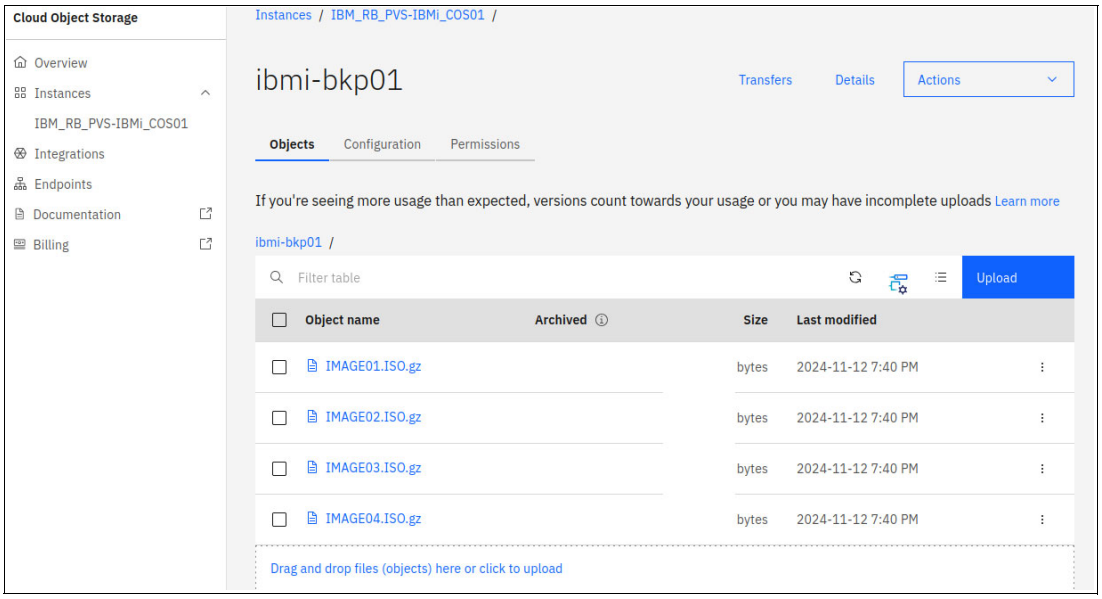

Figure 2-18 View objects in a bucket

27. In the command line, use `awscli` and the **`ls`** command, as shown in Example 2-12.

Example 2-12 Another way by using Linux commands

```
aws --endpoint-url=https://s3.us-south.cloud-object-storage.appdomain.cloud s3 ls
s3://ibmi-bkp01/
```

28. Return to the IBM Cloud Power Virtual Server environment.
29. Prepare the NFS server and the target system, both running IBM i.
30. The following requirements must be met to perform a full system restore because the system must be restarted from the “A” or “B” side (*not* “D” IPL), and point to an NFS server to load the LIC:
    - One TEMPORARY NFS instance is available to serve the images that are running IBM i V7R2, V7R3 or V7R4.
    - The latest PTFs are on the TEMPORARY NFS instance.
    - Enough disk space is available on the TEMPORARY NFS instance
    - The target IBM i system is on the same network that is connected to the TEMPORARY NFS instance with enough resources to restore the backup. Open a Support ticket to enable the VLAN and connect the systems.
    - Communications are in place from the TEMPORARY NFS instance to the IBM Cloud Object Storage or the place data was saved. Consider the following points:
      - The backups can also be downloaded by using the internet and the public access on IBM Cloud Object Storage. However, charges might be incurred on a public connection. Internet access from IBM Power Virtual Server instances has no public access cost.

Restore process

Complete the following steps to restore the instance:

1. On the TEMPORARY NFS instance update YUM, as shown in Example 2-13.

Example 2-13 Updating YUM on PASE

```
CALL QP2TERM
/QOpenSys/pkgs/bin/yum -y update
```

2. Install the following tools by using the ACS or command line, as shown in Example 2-14:

   `pigz gzip gunzip python3 python3-pip`

Example 2-14 Install other tools from Access Client solutions

```
CALL QP2TERM

PATH=$PATH:/QOpenSys/pkgs/bin
export PATH

yum -y install pigz compress gunzip python3 python3-pip

pip3 install s3cmd
pip3 install awscli

# Then you can configure awscli
aws configure

# Paste the Secret Key, Access Key and select the region.
```

3. Download the compressed images, as shown in Example 2-15.

Example 2-15 Downloading compressed image files

```
CALL QP2TERM
mkdir /NFS
mkdir /NFS/RESTORE21
chmod 755 /NFS
chmod 777 /NFS/RESTORE21

aws --endpoint-url=https://s3.us-south.cloud-object-storage.appdomain.cloud s3 cp
s3://ibmi-bkp01/IMAGE01.ISO.gz /NFS/RESTORE21
aws --endpoint-url=https://s3.us-south.cloud-object-storage.appdomain.cloud s3 cp
s3://ibmi-bkp01/IMAGE02.ISO.gz /NFS/RESTORE21
aws --endpoint-url=https://s3.us-south.cloud-object-storage.appdomain.cloud s3 cp
s3://ibmi-bkp01/IMAGE03.ISO.gz /NFS/RESTORE21
aws --endpoint-url=https://s3.us-south.cloud-object-storage.appdomain.cloud s3 cp
s3://ibmi-bkp01/IMAGE04.ISO.gz /NFS/RESTORE21
```

4. Decompress the images, as shown in Example 2-16.

Example 2-16 Decompressing image files

```
/QOpenSys/pkgs/bin/gunzip /NFS/RESTORE21/IMAGE01.ISO.gz
/QOpenSys/pkgs/bin/gunzip /NFS/RESTORE21/IMAGE02.ISO.gz
/QOpenSys/pkgs/bin/gunzip /NFS/RESTORE21/IMAGE03.ISO.gz
/QOpenSys/pkgs/bin/gunzip /NFS/RESTORE21/IMAGE04.ISO.gz
```

5. Press **F3**. Complete the following steps in QCMD:
   a. Create an Image Catalog and add our virtual media:

      ```
      CRTIMGCLG IMGCLG(RESTORE21) DIR('/NFS/RESTORE21') CRTDIR(*NO)
      ADDVRTVOL(*DIR) IMGTYPE(*ISO)
      ```

   b. Create and make available a Virtual Optical Device:

      ```
      CRTDEVOPT DEVD(NFSOPT) RSRCNAME(*VRT) LCLINTNETA(*N) VRYCFG CFGOBJ(NFSOPT)
      CFGTYPE(*DEV) STATUS(*ON)
      ```

   c. Load the Image Catalog with the device:

      ```
      LODIMGCLG IMGCLG(RESTORE21) DEV(NFSOPT)
      ```

   d. Verify the IMAGE Catalog (some files are extracted for the remote IPL process):

      ```
      VFYIMGCLG IMGCLG(RESTORE21) TYPE(*LIC) SORT(*YES) NFSSHR(*YES)
      ```

   e. Change authorities:

      ```
      CHGAUT OBJ('/NFS/RESTORE21') USER(*PUBLIC) DTAAUT(*RWX) SUBTREE(*ALL)
      CHGAUT OBJ('/NFS/RESTORE21') USER(QTFTP) DTAAUT(*RX) SUBTREE(*ALL)
      ```

   f. Export NFS Share as read-only:

      ```
      CHGNFSEXP OPTIONS('-i -o ro') DIR('/NFS/RESTORE21')
      ```

   g. Start NFS servers:

      ```
      STRNFSSVR *ALL
      ```

   h. Change the TFTP server attributes:

      ```
      CHGTFTPA AUTOSTART(*YES) ALTSRCDIR('/NFS/RESTORE21/BOOTP')
      ```

   i. Restart the TFTP server:

      ```
      ENDTCPSVR SERVER(*TFTP)
      STRTCPSVR SERVER(*TFTP)
      ```

Complete the following steps to set up the environment in the TARGET instance:

1. Configure Service Tools Server LAN Adapter by using a different IP address on the same private network segment and adapter (IP address is `192.168.80.21`). Use the guidelines that were described in the previous backup process.
2. Add the NFS instance to the host table:

   ```
   ADDTCPHTE INTNETADR('192.168.80.12') HOSTNAME((IBMiNFS))
   ```

3. Mount NFS SHARE on the target instance to confirm whether the files are visible:

   ```
   MOUNT TYPE(*NFS) MFS('192.168.80.12:/NFS') MNTOVRDIR('/NFS')
   WRKLNK '/NFS/RESTORE21/*'
   ```

4. Create the virtual optical device:

   ```
   CRTDEVOPT DEVD(NFSRESTORE) RSRCNAME(*VRT) LCLINTNETA(*SRVLAN)
   RMTINTNETA('192.168.80.12') NETIMGDIR('/NFS/RESTORE21')
   VRYCFG CFGOBJ(NFSRESTORE) CFGTYPE(*DEV) STATUS(*ON)
   ```

5. IPL the instance so that the changes can take effect.
6. Record the following information (Figure 2-19 on page 43, Figure 2-20 on page 43, and Figure 2-21 on page 44):
   – TCP/IP information: IP address, gateway, and hostnames

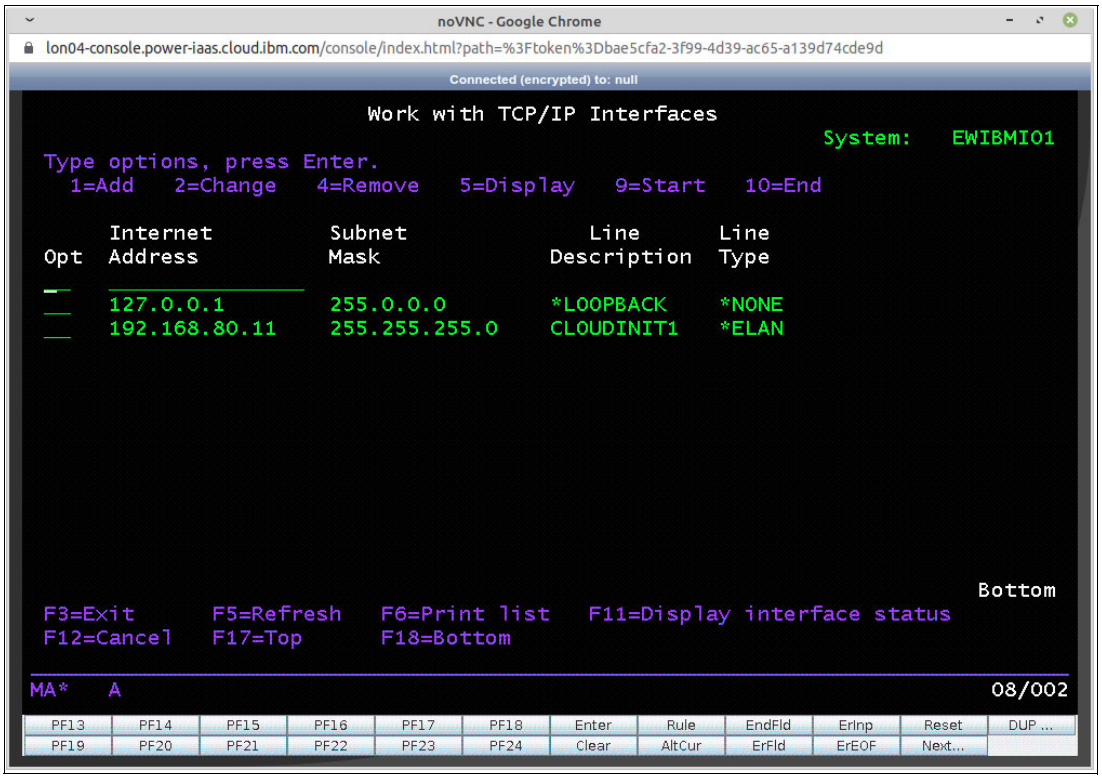

Figure 2-19 Get IP address and associated line description

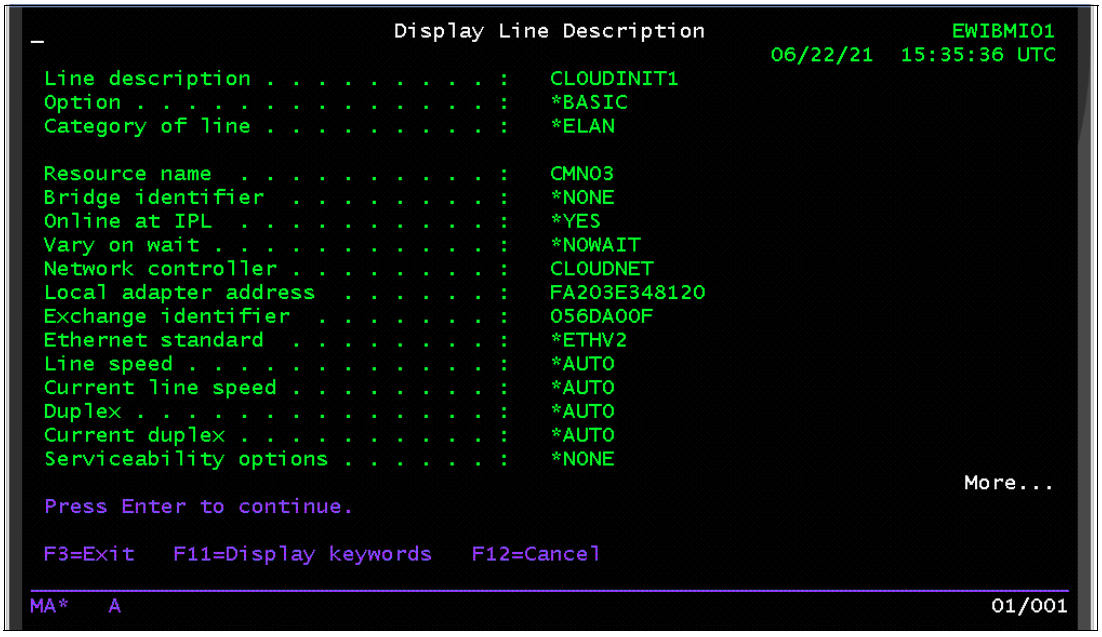

Figure 2-20 Display line description details and get the Resource Name

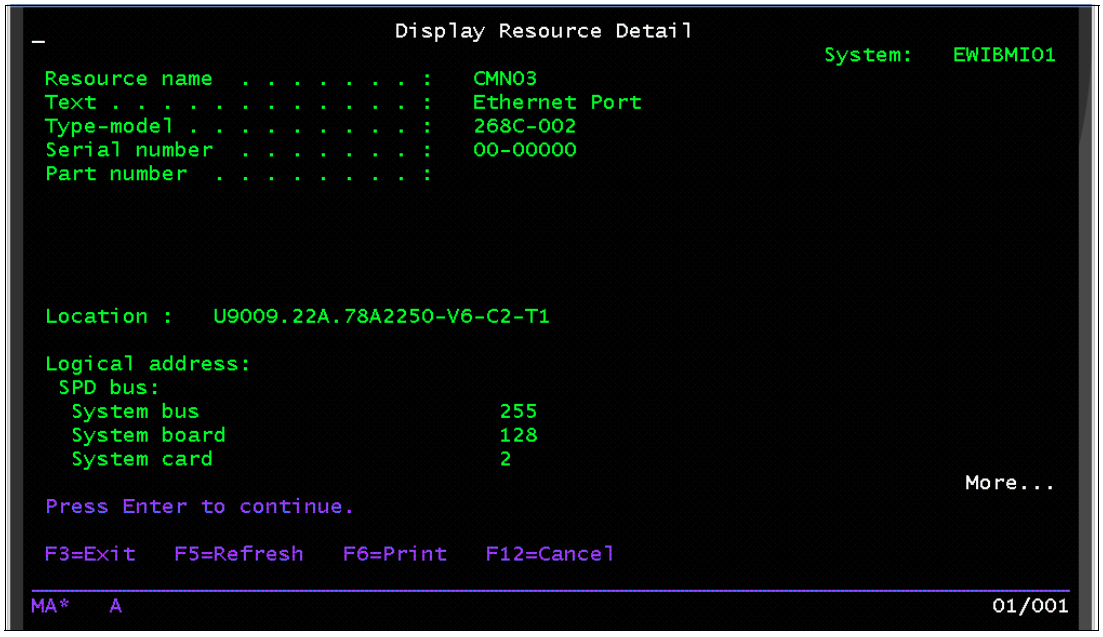

*Figure 2-21 Display Resource Detail by using WRKHDWRSC *CMN option 7*

- License keys. Back up the licenses to a file and transfer them to the NFS server:

```
DSPLICKEY OUTPUT(*LICKEYFILE) LICKEYFILE(QGPL/LIC01)
CRTSAVG QGPL/LIC
SAVOBJ OBJ(LIC01) LIB(QGPL) DEV(*SAVF) SAVF(QGPL/LIC)
```

7. Transfer the license keys to the NFS instance.
8. Start the network installation (see Figure 2-22):

```
STRNETINS DEV(NFSRESTORE) OPTION(*LIC) KEYLCKMOD(*MANUAL)
```

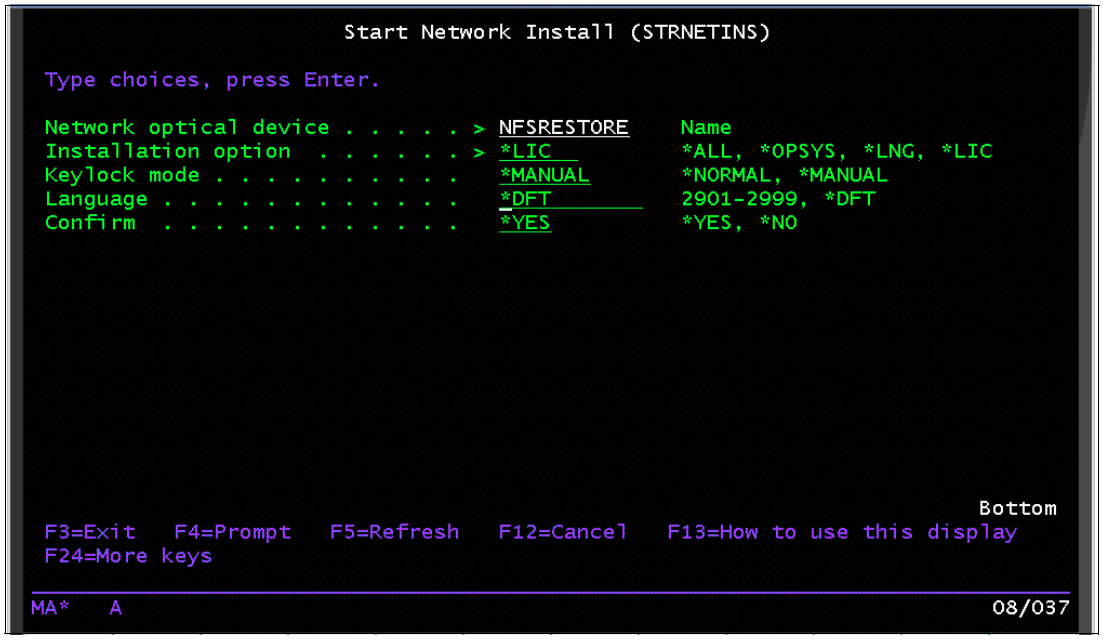

Figure 2-22 STRNETINS command

9. Click **Finish** or confirm the command at TARGET. The IPL process starts (a connection must be established to the console for manual operation).
10. This process can take several minutes until the DST window is displayed.
11. From the DST window, complete the following steps to install the License Machine Code (SLIC):
    a. On the Install License Internal Code Menu, choose **Option 2**, and confirm by pressing **F10**.

       The process to initialize the load source disk begins.
    b. Select **Install Operating System**. In this example, select the installation device type **Option 5 - Network device**.
    c. Press **F10**. The Network Device - Configuration menu is displayed (see Figure 2-23).

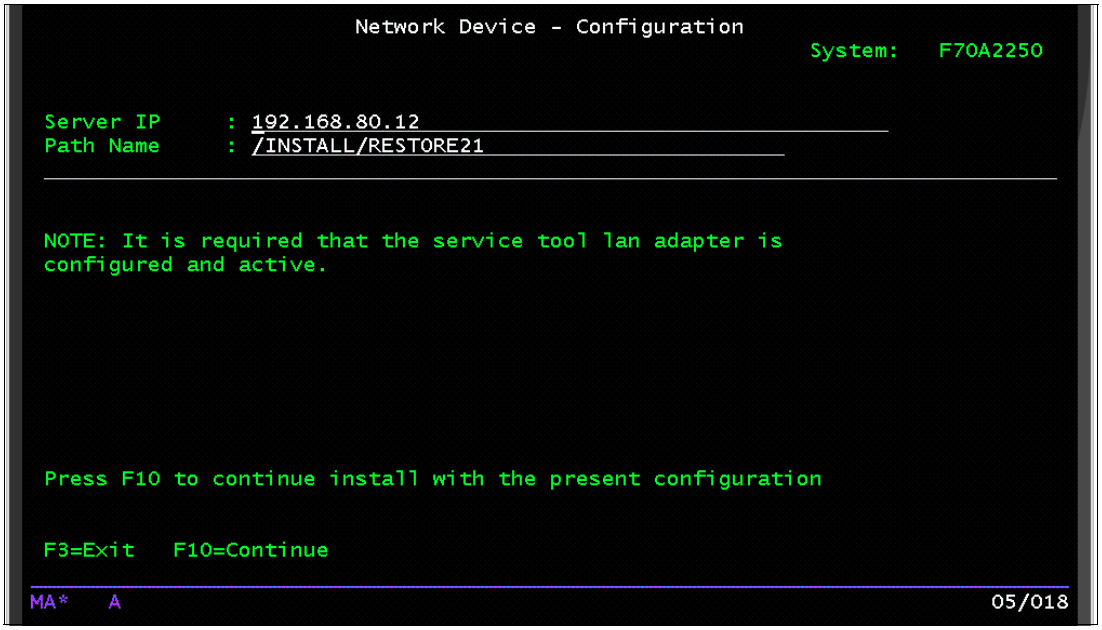

Figure 2-23 Confirm your NFS instance values and continue with F10

    d. Confirm the installation of the operating system and press **Enter**.
    e. Select a language group from the Select a Language Group Menu, (in this example, **2924** is selected).
    f. If disks must be added, use the Add All Disk Units to the System menu. In the example, option 1 is selected: **Keep the current disk configuration**.
    g. The sign-on window opens. Log in by using the `QSECOFR` user profile. The IPL options are displayed. Enter your choice (see Figure 2-24 on page 46).



Figure 2-24 Set IPL Options as shown in this image.

h. Set the Major System Options:

- Enable automatic configuration = Yes
- Device configuration naming = *NORMAL
- - Default special environment = *NONE

The IBM i Installation status progress is displayed.

i. After you install the operating system, you can sign on. However, the password must be changed to continue the sign-on request.

12. Continue with the setup and restore process:

a. Check the resource name for your network adapter:

```
WRKHDWRSC *CMN
```

b. If the adapters changed, delete them and then re-create them by using the correct resource:

```
CRTLINETH LIND(ETH01) RSRCNAME(CMN03)
CRTLINETH LIND(ETH02 RSRCNAME(CMN04)
```

c. Create a Virtual Optical Device on Service Tool Server LAN adapter (CMN04) again:

```
CRTDEVOPT DEVD(NFSRESTORE) RSRCNAME(*VRT) LCLINTNETA(*SRVLAN)
RMTINTNETA('192.168.80.12') NETIMGDIR('/NFS/RESTORE21')
```

d. Vary on the Virtual Optical Device:

```
VRYCFG CFGOBJ(NFSRESTORE) CFGTYPE(*DEV) STATUS(*ON)
```

e. Check the NFS Image Catalog:

```
WRKIMGCLGE IMGCLG(*DEV) DEV(NFSRESTORE)
```

f. Continue with the restore process:

```
GO RESTORE
Option 21
```

g. Use NFSRESTORE as the device name.

After the restore process is done, complete the following steps to check the messages and system:

1. Check the software products:

   ```
   GO LICPGM
   Option 10
   ```

2. Change system values to fit the new environment:

   ```
   WRKSYSVAL
   ```

3. Change the network name:

   ```
   CHGNETA SYSNAME(IBMI04) LCLCPNAME(IBMI04) LCLLOCNAME(IBMI04) DTACPR(*ALLOW)
   DTACPRINM(*REQUEST) NETSERVER((IBMI04 *ANY)) NWSDOMAIN(IBMI04) ALWVRTAPPN(*YES)
   ALWHPRTWR(*YES)
   ```

4. Change the network configuration by using the previously saved values:

   ```
   CFGTCP
   ```

5. On the NFS instance, start the FTP service. It is stopped by default:

   ```
   STRTCPSVR *FTP
   ```

6. License keys should be injected with the “cloudinit” process.

7. Validate license keys:

   ```
   WRKLICINF OUTPUT(*PRINT)
   ```

8. Change the IPL mode to NORMAL from IBM Cloud Power Virtual Server portal.

9. Restart the TARGET instance:

   ```
   PWRDWNSYS *IMMED RESTART(*YES) CONFIRM(*NO)
   ```

The IBM i TARGET system is now running and the validation process can begin.

Alert: If a failure occurs during the LIC or operating system installation process, delete the instance and start again by installing the LIC and operating system.

2.3.3 Full-system backup from IBM i Source by using BRMS and IBM Cloud Storage and Restore on IBM Power Virtual Server

This section describes how to save and transfer an IBM i workload on-premises to Power Systems Virtual Server by using Backup, Recovery, and Media Services (BRMS) with the IBM Cloud Storage licensed program product (LPP) for IBM i.

Migrating IBM i to the Cloud using IBM Cloud Object Storage

Figure 2-25 on page 48 shows a migration scenario that uses IBM Cloud Object Storage. Note the following characteristics:

- It uses BRMS and either 5733-ICC or 5770-ICC to transfer VMs to IBM Cloud Object Storage.
- Requires 2x disk space on-premises.
- Requires reasonable network connection to IBM Cloud.

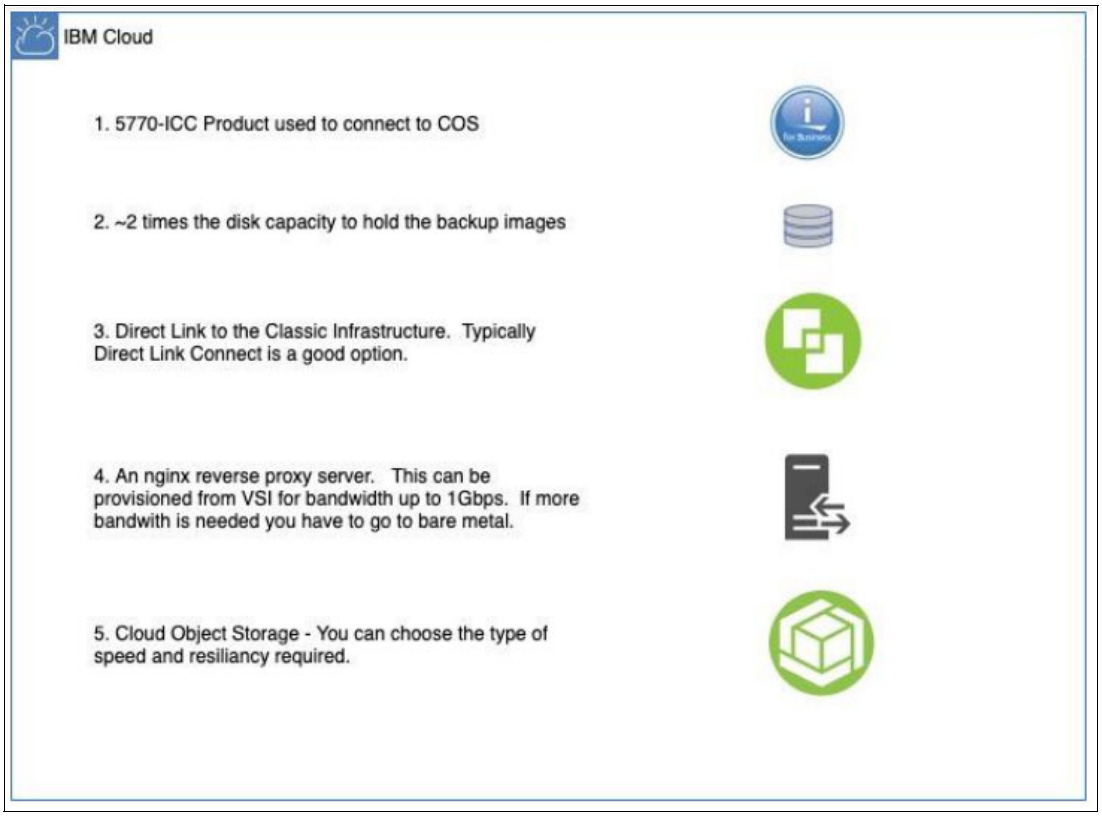

Figure 2-25 Migrating IBM i to the Cloud by using IBM Cloud Object Storage

Figure 2-26 shows the steps to retrieve your image by using IBM Cloud Storage Solutions for IBM i.

i

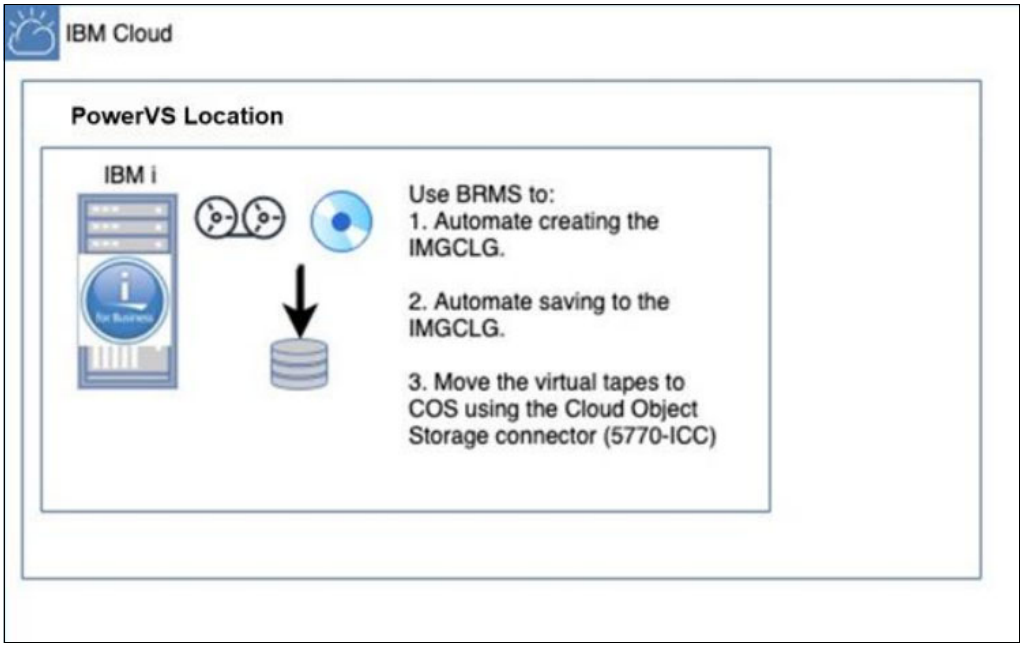

Figure 2-26 Get your BRMS Image Catalog on IBM Cloud Object Storage

Solution components and requirements

This section describes the solution components and requirements.

The following components must be configured on the IBM i Local Source:

- BRMS and IBM Cloud Storage software.
- IBM i System Minimum PTF levels.
- BRMS and IBM Cloud Storage PTFs.
- Cloud-Init for IBM i and PTFs.
- System is saved by using BRMS/IBM Cloud Storage to IBM Cloud Object Storage.
- BRMS Recovery Reports are created.

IBM i Source system requirements

Install the BRMS and IBM Cloud Storage software. Complete the following steps:

1. Install the following operating system options and BRMS:
   - 5770-SS1Option 18: Media and Storage Extensions
   - 5770-SS1Option 44: Encrypted Backup Enablement (optional)
   - 5770-BR1*BASE
   - 5770-BR1Option 1: Network feature (optional)
   - 5770-BR1Option 2: Advanced Functions feature (optional)
2. Install IBM Cloud Storage Solutions:
   - 5733ICC *BASE IBM Cloud Storage Solutions for i
   - 5733ICC Option 1: Cloud Storage
   - 5733ICC Option 2: Advanced
   - 5733ICC Option 3: Reserved – Option 3
   - 5733ICC Option 4: Reserved – Option 4
   - 5733ICC Option 5: Reserved – Option 5
   - 5733ICC Option 6: Reserved – Option 6
   - 5733ICC Option 7: Reserved – Option 7

IBM i Source System PTF minimum levels

Install the following program temporary fixes (PTFs) on your IBM i VM:

- PTFs for IBM i License Keys Post Deployment
- License Keys that are incorporated as part of the deployment of the LPAR.
  - IBM i 7.2 - 5770SS1 SI71091 (prerequisite SLIC PTFs: MF66395, MF66394, MF66391)
  - IBM i 7.3 - MF99207 (TR7) and SI69686
  - IBM i 7.4 - MF99301 (TR1) and SI70544

  For information, see Minimum PTF levels for IBM i.
- The following BRMS PTFs must be installed on a stand-alone BRMS system or on all of the systems in a BRMS network:
  - IBM i 7.3 - SI61153
  - IBM i 7.2 - SI61152
  - IBM i 7.1 - SI61151
- The SI73401 - Install the PTF Cloud Storage Solutions Proxy Support 5733ICC IBM Cloud Storage Solutions for i for Direct Link (DL) Reverse Proxy Server Support. Support Google Cloud Storage.

IBM i software requirements

For Cloud-Init Support for IBM i, this Cloud-Init requirement is installed after deployment. Any images that are imported to the cloud must have `cloud-init` installed.

If you are bringing your own IBM i custom SAVSYS, you must install these SAVSYS PTFs and the software that is required for `cloud-init`. For more information, see Cloud-Init Support for IBM i.

Installing cloud-init on IBM i

Specific LPPs and PTFs are needed for IBM i so that cloud-init can work. Complete the following steps to install and configure `cloud-init` on IBM i:

1. Download and install the required license programs:
   - 5770DG1 with *ALL
   - 5770SS1 with Option 30 and 33
   - 5770SC1 *BASE and option 1
2. Install the following rpm packages:
   - `python2-ibm_db-2.0.5.8`
   - `python2-six-1.10.0-1`
   - `python2-2.7.15-1`
   - `cloud-init-1.2-100`

Note: Consider the following points:

- The listed version of rpm packages is the lowest required version.
- The rpm packages can be installed by using yum. For more information, see yum for IBM i.

3. Install the following required PTFs and PTF groups:
   - IBM i 7.2: SF99713, MF61899 (permanently apply), SI62642, SI63163, MF65261, MF65218, and SI72162
   - IBM i 7.3: MF61942 (permanently apply), SI62642, SI63041, MF63830, and SI72161
   - IBM i 7.4: SI72142
4. Save and migrate your system.

Note: For IBM i version 7.2 and later on Power8, `cloud-init` is enabled automatically.

For versions earlier than IBM i 7.2, the process that is described in “Installing cloud-init on IBM i” on page 50 *must* be completed to run `cloud-init` on IBM i.

Enabling cloud-init on IBM i

To enable the cloud-init on an IBM i system manually, run this on the virtual machine, as shown in Example 2-17.

Example 2-17 Enable cloud-init on IBM i

```
CALL PGM(QSYS/QAENGCHG) PARM(*ENABLECI)
```

This command must be run with a user profile that includes the following access privileges:

- *ALLOBJ
- *AUDIT
- *IOSYSCFG

- *JOBCTL
- *SAVSYS
- *SECADM
- *SERVICE
- *SPLCTL

After you enable `cloud-init` on the IBM i server, power off the system so that it is ready to save and migrate.

Downloading the required software

Complete the following steps to download the required software. For more information, see Cloud-Init Support for IBM i.

1. Click **My Entitled Systems Support** (Figure 2-27).

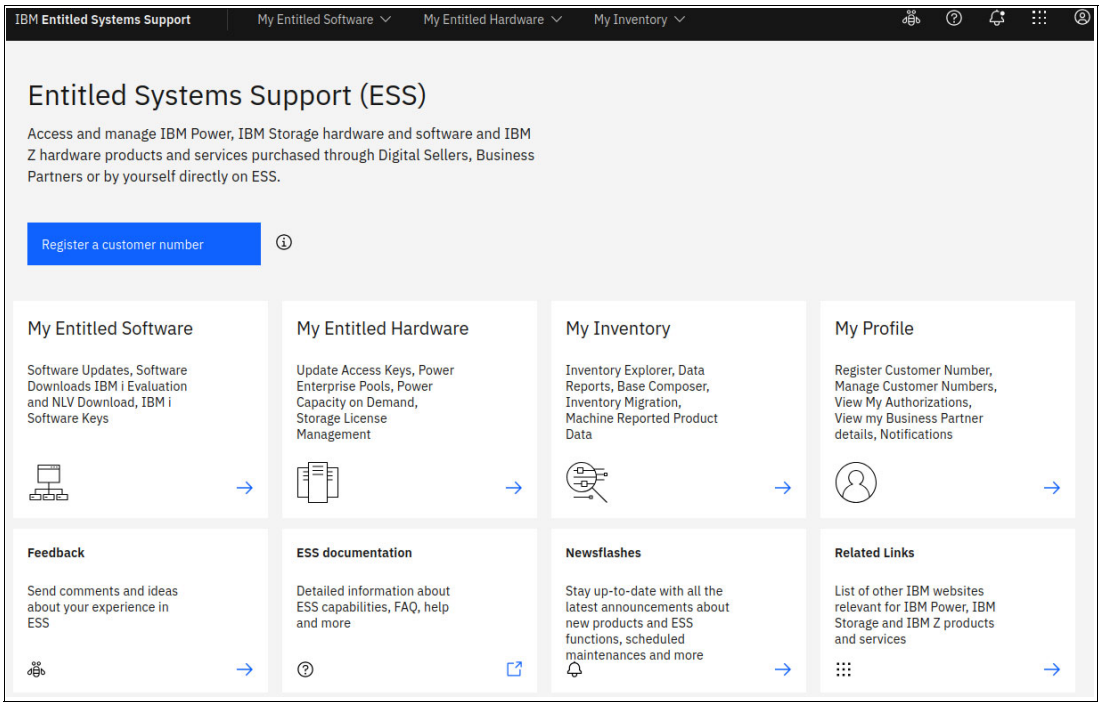

Figure 2-27 Entitled Systems Support (ESS) portal

2. Click **My Entitled Software -> IBM i evaluation**. (Figure 2-28). For more information, see Entitled Systems Support (ESS).

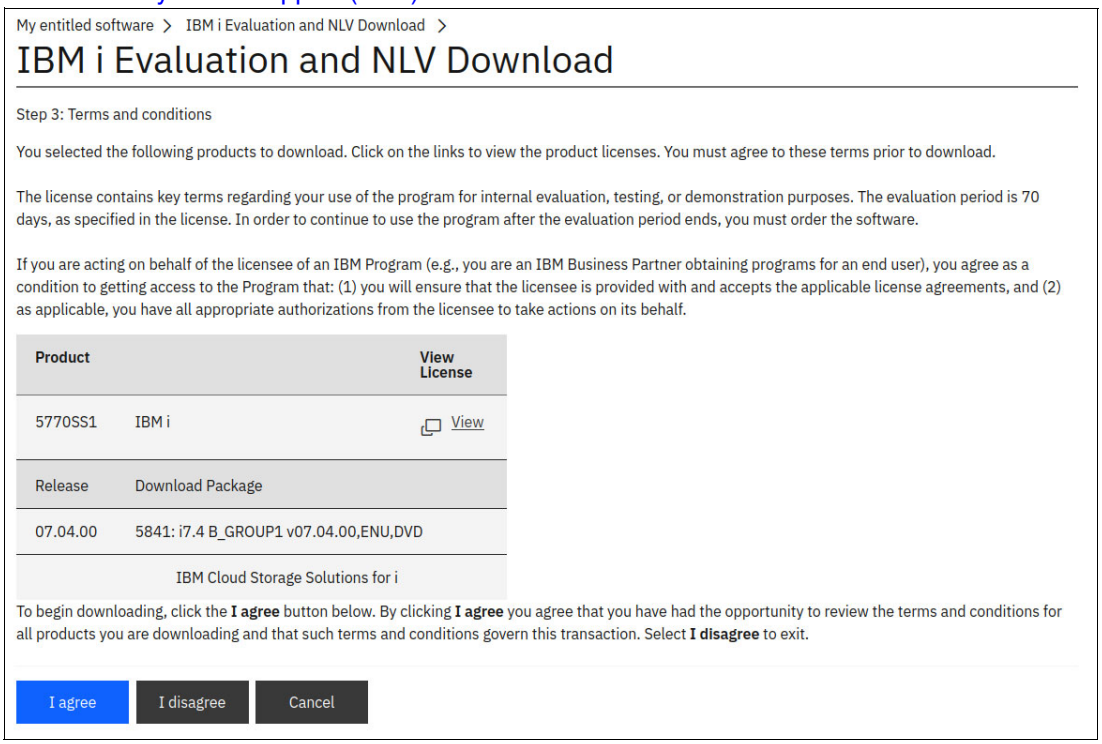

Figure 2-28 Software downloads option

3. Select IBM i product **5770-SS1** and click **Continue** (Figure 2-29).

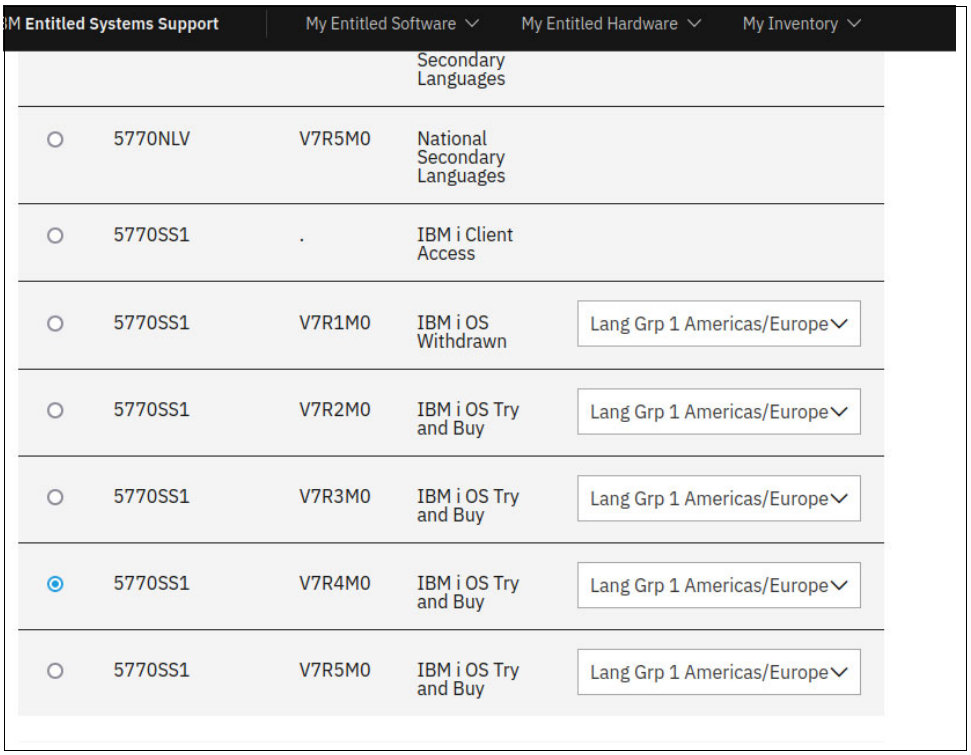

Figure 2-29 Selecting IBM i and release category

4. Expand and select the product and click **Continue** (Figure 2-30).

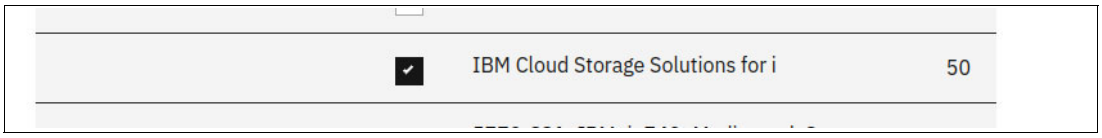

Figure 2-30 Select desire product

5. Agree to the evaluation terms and click **Continue** (Figure 2-31).

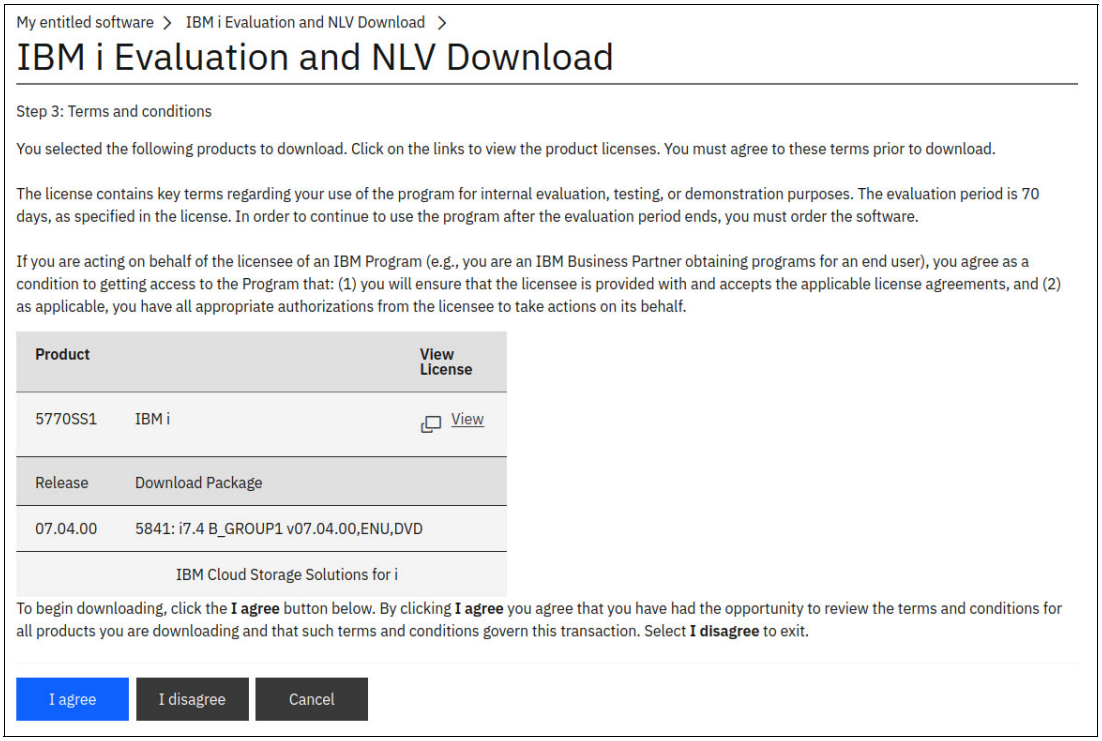

Figure 2-31 Evalucation and NV download

6. Select the new download (Figure 2-32 on page 54).

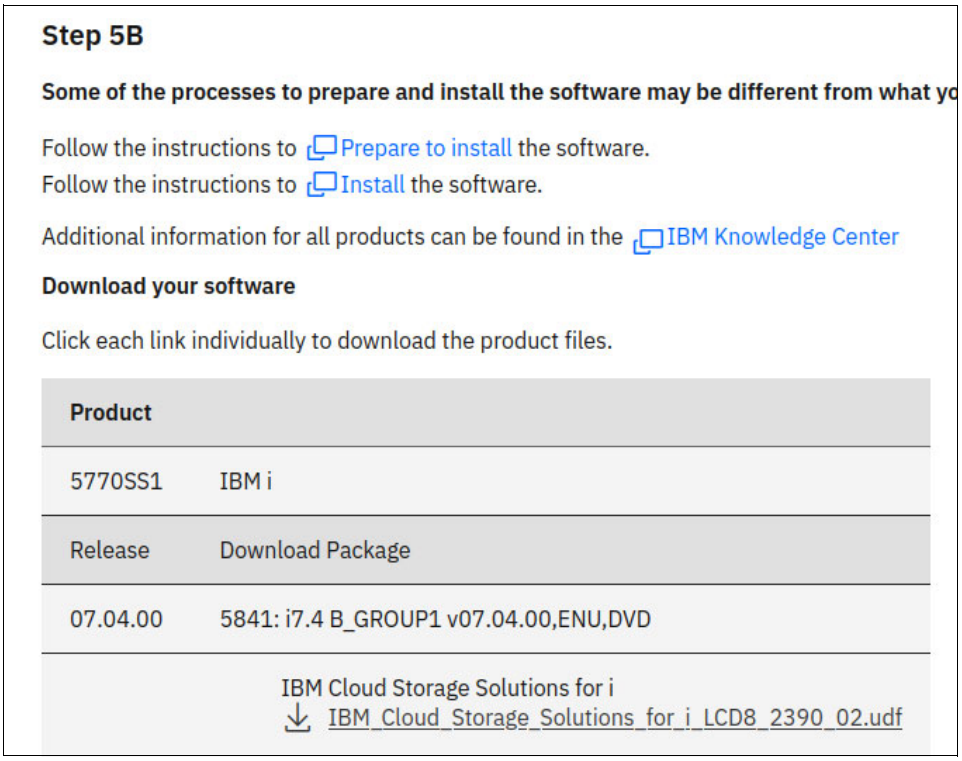

Figure 2-32 Select 5770-SS1

7. To use IBM Download Director, select **Download now**. To use HTTP, select **Click here to use HTTP.** (Figure 2-33)

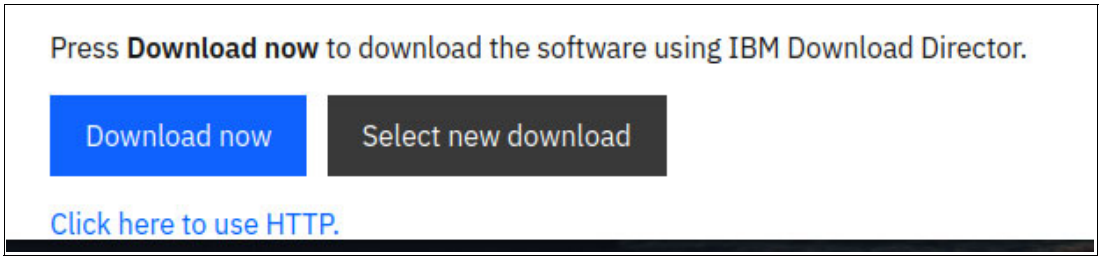

Figure 2-33 Entitled software downloads - Select language and press Continue

In the example system, the compressed file content was extracted and renamed to `5733ICC.udf`.

Preparing an optical image catalog to install software

This section describes how to create an image catalog and an image catalog entry, add an image catalog entry, and load the image catalog to prepare for an IBM i software upgrade. The steps use a virtual optical device in the example. The server setup requires an image catalog setup that is then shared with the client partition.

Complete the following steps:

1. Create a virtual optical device.

   To create a device description, enter the command that is shown in Example 2-18 on page 55.

Example 2-18 Create a device description

```
CRTDEVOPT DEVD(virtual-device-name) RSRCNAME(*VRT)+
 ONLINE(*YES) TEXT(text-description)
```

2. Vary on the virtual optical device, as shown in Example 2-19.

Example 2-19 Vary on the Virtual Optical Device

```
VRYCFG CFGOBJ(virtual-device-name) CFGTYPE(*DEV) STATUS(*ON)
```

3. Create an image catalog for the licensed programs that you want to install. The Create Image Catalog (**CRTIMGCLG**) command associates an image catalog with a target directory where the optical image files are loaded, as shown in Example 2-20.

Example 2-20 Create image catalog

```
CRTIMGCLG IMGCLG(catalog-name) DIR(catalog-path) CRTDIR(*YES)
TEXT(text-description)
```

4. If you downloaded your images into an image catalog directory, two quick methods are available to add all of the images simultaneously into your image catalog:
   - Use the **CRTIMGCLG** command, as shown in Example 2-21.

     Example 2-21 Create Image Catalog and add volumes

     ```
     CRTIMGCLG IMGCLG(catalog-name) DIR(catalog-path) ADDVRTVOL(*DIR)
     IMGTYPE(*ALL) TEXT(text-description)
     ```

   - Use the QVOIFIMG API when the image catalog already exists as shown in Example 2-22.

     Example 2-22 Add images to an existing image catalog with QVOIFIMG

     ```
     CALL PGM(QVOIFIMG) PARM('catalog-name' '*ALL' 0)
     ```

5. Add an image catalog entry for each physical media or optical image file. You must repeat this step for each volume of media. Add the physical media or optical image files in the same order as though you were going to install from them. Start with the first media in the list and continue until all of the media are loaded.

   You can add the entries from an optical device or from an optical image file. Select one of the following methods:

   - From an image file

     This method is the fastest way. To add an image entry to an image catalog from an Integrated File System file that is in the image catalog directory, enter the command as shown in Example 2-23.

Example 2-23 Add an image entry to an image catalog

```
ADDIMGCLGE IMGCLG(catalog-name) FROMFILE(file-name) TOFILE(*fromfile)
TEXT(text-description)
```

Note: If you need to add multiple images, see the **CRTIMGCLG** command and the QVOIFIMG API to add all of the images at the same time.

– To add an image catalog entry to an image catalog from an integrated file system optical image file that is from a directory other than the image catalog directory, enter the commands as shown in Example 2-24.

Example 2-24 Add image catalog entry from optical image

```
ADDIMGCLGE IMGCLG(catalog-name) FROMFILE(/directory-name/directory-name/file-name)
TOFILE(file-name or *FROMFILE) TEXT(text-description)
```

– From a physical device

Add an image catalog entry to an image catalog from a physical optical media by using the optical device that is named `OPT01`, as shown in Example 2-25.

Example 2-25 Add Image Catalog Entry from a physical device

```
ADDIMGCLGE IMGCLG(catalog-name) FROMDEV(OPT01) TOFILE(file-name or *GEN)
TEXT(text-description)
```

Note: To generate a name for the TOFILE parameter and a text description from the media, specify `*GEN`.

6. Load the image catalog.

This step associates the virtual optical device to the image catalog. Only one image catalog at a time can be associated with a specific virtual optical device. To load the image catalog, run the command as shown in Example 2-26.

Example 2-26 Load the Virtual Optical Device to the Image Catalog

```
LODIMGCLG IMGCLG(catalog-name) DEV(virtual-device-name) OPTION(*LOAD)
```

7. Verify that the images are in the correct order.

Note: If you are preselecting the licensed programs to install, do *not* perform this step now. You are directed to perform this step later.

If you are preparing for an upgrade, you must verify that the required media for an upgrade exists and is sorted in the correct sequence. Also verify that your software agreements were accepted, and that reserved storage is available for the Licensed Internal Code.

Enter the command as shown in Example 2-27.

Example 2-27 Verify image catalog

```
VFYIMGCLG IMGCLG(catalog-name) TYPE(*UPGRADE) SORT(*YES)
```

To verify that images are added, another method is shown in Example 2-28.

Example 2-28 Work with Image Catalog

```
WRKIMGCLGE IMGCLG(catalog-name)
```

Then press **F7** to prompt for the **`VFYIMGCLG`** command. Enter `*UPGRADE` for the type and `*YES` for the sort field.

The system puts the images in the correct order. (If you are not successful, refer to Image catalog recovery.) By default, the volume with the lowest index is mounted. Then, all other volumes are loaded.

To see the order of the images, use the Work with Image Catalog Entries (**WRKIMGCLGE**) command:

```
WRKIMGCLGE IMGCLG(catalog-name)
```

After completing these steps, your image catalog is ready for use.

Required components in the IBM Cloud

The following components must be configured in the IBM Cloud UI:

- An IBM Cloud account (beyond the scope of this document).
- Create a Power Systems Virtual Server service and a private subnet.
- Provision IBM i.
- Virtual Server Instances (VSIs).
- Order Direct Link Connect Classic to connect each Power Systems Virtual Server location to IBM Cloud Classic, which is required for IBM Cloud Object Storage access for saves and restores.
- Create and configure a Reverse-proxy CentOS VSI to connect each Power Systems Virtual Server location to IBM Cloud Classic. For more information, see Using IBM Cloud Direct Link to connect to IBM Cloud Object Storage.

Full-system backups from IBM i source by using BRMS and IBM Cloud Storage

The following save data must be restored from physical media before BRMS can begin restoring save data directly from the cloud:

- SAVSYS to install the operating system.
- IBM Backup, Recovery, and Media Services for i and BRMS save information.
- IBM TCP/IP Connectivity for i and configuration information to allow communications with cloud storage providers.
- IBM Cloud Storage Solutions for i and configuration information to establish connections with cloud storage providers.
- BRMS provides specific control groups that can be used to automatically save this data to media in the cloud and the cloud media can be used to create physical media. The control groups create cloud media that is formatted, so it can be downloaded and burned directly to physical optical media. All remaining data on the system can be backed up to media in the cloud and restored directly from the cloud without a need to create physical media.
- Control group QCLDBIPLnn can be used to perform full backups of all data that must be recovered from physical media. Likewise, QCLDBGRPnn can be used to perform cumulative incremental saves of the data that must be recovered from physical media.
- Control group QCLDBSYSnn can be used to perform a full backup of the data that can be restored directly from the cloud. Likewise, control group QCLDBUSRnn can be used to perform cumulative incremental backups of the data that can be restored directly from the cloud.

Note: The Journaled objects control group field must be changed to `*YES` for a QCLDBGRPnn control group before the control group is used to perform an incremental backup.

Run the **WRKCTLGBRM** command and change the Journaled objects field by specifying `option 8=Change attributes for QCLDBGRPnn`.

It is critical to run the cloud control groups in the correct order. Otherwise, all necessary media information is not available to perform a recovery.

Running a full-system backup from the console

In this example, you run a restricted state backup by using your HMC or LAN Console connection.

Complete the following steps to perform a full save of your entire system:

Note: Processing time for each backup depends on the size of your system processor, device capability, and the amount of data that you want to save.

You cannot perform other activities during these backups because your system is in a restricted state.

1. Run the BRMS Control Group `QCLDBSYS01`.
2. Put the system in a restricted state as shown in Example 2-29.

Example 2-29 Putting the system in a restricted state

```
ENDSBS SBS(*ALL) DELAY(120)
```

3. Display `QSYSOPR MSGQ`, run the DSPMSG QSECOFR command, and look for the following messages:
   - `System ended to restricted condition.`
   - `A request to end TCP/IP has completed.`
4. Change subsystems to process for Control Group `QCLDBSYS01`:

1. Run the WRKCTLGBRM command.
2. Find `QCLDBSYS01` (see Figure 2-34).
3. Select **Option 9=Subsystems** to process.
4. Change Restart to `*NO for Seq 10 Subsystem *ALL`.

Note: If you see a break message during the backup, press Enter to return to the window in which you entered the STRBKUBRM command so that you can review the backup progress.

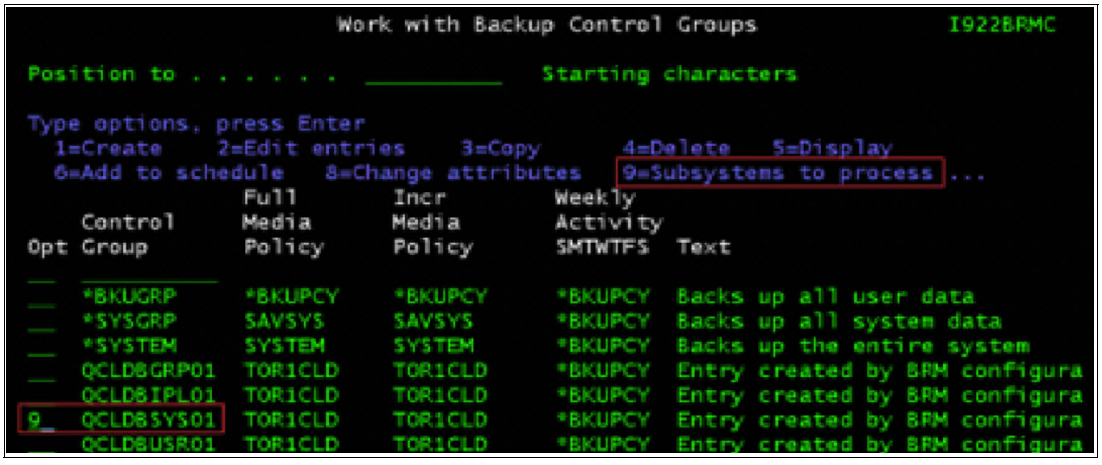

Figure 2-34 WRKCTLGBRM - Find QCLDBSYS01

5. Run the First backup from the console, as shown in Figure 2-34 on page 58.

Example 2-30 Run first backup from console

```
STRBKUBRM CTLGRP(QCLDBSYS01) SBMJOB(*NO)
```

6. Check the backup for errors. It is normal to see some errors, such as the following examples:
   - `Objects not saved` (some objects are not required for the recovery).
   - `Media not transferred` (complete this step manually after the second backup).
7. Check the subsystems after the backup has finished. You see only subsystem QCTL in a status of RSTD. If not, end all subsystems again, as shown in Example 2-31.

Example 2-31 End all Subsystems again when QCTL is not the only one running

```
ENDSBS SBS(*ALL) DELAY(120)
```

8. Change BRMS Control Group QCLDBIPL01:
   a. Use the WRKCTLGBRM command.
   b. Select **Option 8=Change attributes**.
   c. Page down, change Automatically backup media information to *LIB, also change Append to media to *NO.
   d. Select Option 9=Subsystems to process.
   e. Change Restart to *YES for Seq 10 Subsystem *ALL.

Note: Control groups that have a QCLD prefix enable BRMS to automatically create and transfer media to the cloud. Control groups `QCLDBIPLxx` or `QCLDBGRPxx` must be used to burn DVDs for recovery. If the backup uses media class `QCLDVRTOPT`, the BRMS default is to create 10 virtual volumes to back up to optical.

   - Because optical devices do not have an exit program interface to handle media switching while a backup is running, the backup command must provide enough volumes to successfully hold the backup data. A symptom that indicates the backup does not fit on the initial number of volumes that are provided is that the backup fails with message `BRM4301 Volume list exhausted`.
   - When control groups `QCLDBIPLxx` or `QCLDGRPxx` are used, the required volume size is 4.7 GB to burn DVDs for manual recovery. The control group data size cannot exceed 350 GB because of the 75 volume restriction.
9. Set the number of optical volumes for automatic transfers:
   - The number of virtual optical volumes to create can be specified by running the command that is shown in Example 2-32.

Example 2-32 Set number of Optical Volumes for Automatic Transfer

```
CALL PGM(QBRM/Q1AOLD) PARM('NUMOPTVOLS' '*SET ' 'nn')
```

   - *nn* is the number of volumes to auto create. This value must be 1 - 75.
   - The number of volumes to create can be displayed by running the command that is shown in Example 2-33 on page 60.

Example 2-33 Display the number of volumes to create

```
CALL PGM(QBRM/Q1AOLD) PARM('NUMOPTVOLS' '*DISPLAY')
```

– The number of volumes to create can be reset to the default value of 10 by running the command that is shown in Example 2-34.

Example 2-34 Reset the number of volumes

```
CALL PGM(QBRM/Q1AOLD) PARM('NUMOPTVOLS' '*REMOVE')
```

– Run the second backup from the console, as shown in Example 2-35.

Example 2-35 Run the second backup

```
STRBKUBRM CTLGRP(QCLDBIPL01) SBMJOB(*NO)
```

– Check the backup for errors. It is normal to see some errors, such as the following examples:

- `Objects not saved` (some objects are not required for the recovery).
- `Media not transferred` (complete manually after the second backup).

– Identify the volumes that are used for backups `QCLDBSYS01` and `QCLDBIPL001` and transfer to IBM Cloud Object Storage.

– Check the status of the transfer by running the **`WRKSTSICC STATUS(*ALL)`** command.

– If the status is `Failed`, this status is normal. The volumes are transferred in the next step.

– Identify which volumes were used for the backups, as shown in Example 2-36.

Example 2-36 Identify volumes transferred

```
WRKMEDBRM TYPE(*TRF)
```

10. Transfer the volumes to IBM Cloud Object Storage.

    Use the commands that are shown in Example 2-37 to transfer the volumes to IBM Cloud Object Storage.

Example 2-37 Transfer volumes to IBM Cloud Object Storage

```
STRMNTBRM
WRKSTSICC STATUS(*ALL)
```

    The volume name, status, and complete percentage for each file transfer is shown. Wait until all volumes are successfully completed to proceed to the next step.

11. Verify that all of the volumes that were used for the full-system backup no longer have a status of `*TRF`, as shown in Example 2-38.

Example 2-38 Identify volumes transferred

```
WRKMEDBRM TYPE(*TRF)
```

    You do not see any Volumes in the list.

12. As with other recoveries that are performed by using BRMS, a recovery report is used to assist with successful recoveries from save media that was transferred to the cloud. To generate a report for recovery from the cloud, run the command that is shown in Example 2-39 on page 61.

Example 2-39 Generate BRMS recovery report from the cloud

```
STRRCYBRM OPTION(*CTLGRP) ACTION(*REPORT) CTLGRP((QCLDBSYS01 1) (QCLDBIPL01 2))
```

Note: It is important to review the recovery report to ensure that it is complete. If any of the media that is produced during the backup process is not successfully transferred to the cloud, it is not included in the recovery report.

The **CTLGRP** and **PERIOD** parameters that are specified with the **STRRCYBRM** command help identify objects that are saved to volumes that were not transferred to the cloud.

If objects are on volumes that were not included in the recovery report, they are listed in a Missing Objects Attention section near the top of the report.

After the recovery report is verified, the report is stored in a safe location, so it can be referred to during a recovery.

13. Daily incremental backups can be run Monday–Saturday by using the control groups command, as shown in Example 2-40.

Example 2-40 BRMS - Run daily incremental backups

```
STRBKUBRM CTLGRP(QCLDBUSR01) SBMJOB(*NO)
STRBKUBRM CTLGRP(QCLDBGRP01) SBMJOB(*NO)
```

2.4 Migrating IBM i to Power Virtual Server using PowerVC

IBM PowerVC is an advanced virtualization and cloud management offering that is built on OpenStack. It provides simplified virtualization management and cloud deployments for IBM i VMs that are running on IBM Power Systems.

It also provides numerous operational benefits, such as one-click system evacuation for simplified server maintenance, dynamic resource optimization (DRO) to balance server usage during peak times, and automated VM restart to recover from failures.

Users can easily import and export VM images (in the standard Open Virtual Application (.OVA) format) from IBM PowerVC and upload them into IBM Cloud for easy back-and-forth image mobility.

For more information about migrating an *IBM i VM* from IBM Power Virtualization Center to IBM Power Systems Virtual Server, see *Red Hat OpenShift V4.X and IBM Cloud Pak on IBM Power Systems Volume 2*, SG24-8486.

IBM Cloud Pak for multicloud management

IBM Cloud Pak® for Multicloud Management provides consistent visibility, governance, and automation across many platforms. Whether you are considering managing VMs, containers, public clouds, or an on-premises infrastructure, IBM Cloud Pak for Multicloud Management features the tools to ensure that your enterprise runs smoothly. IBM Cloud Pak for Multicloud Management includes Infrastructure management and Service management components.

Infrastructure management delivers the insight, control, and automation enterprises that are needed to address the challenges of managing virtual environments, which can be more complex than physical environments.

Consider the following points:

- The Infrastructure management module was called *IBM Red Hat CloudForms*.
- The Managed services module was called *Terraform and Service Automation* or *IBM Cloud Automation Manager*.

The use of IBM Cloud Pak for Multicloud Management can be an end-to-end solution to manage IBM i workloads on-premises and off-premises. Therefore, you can migrate and deploy IBM i images, as shown in Figure 2-35.

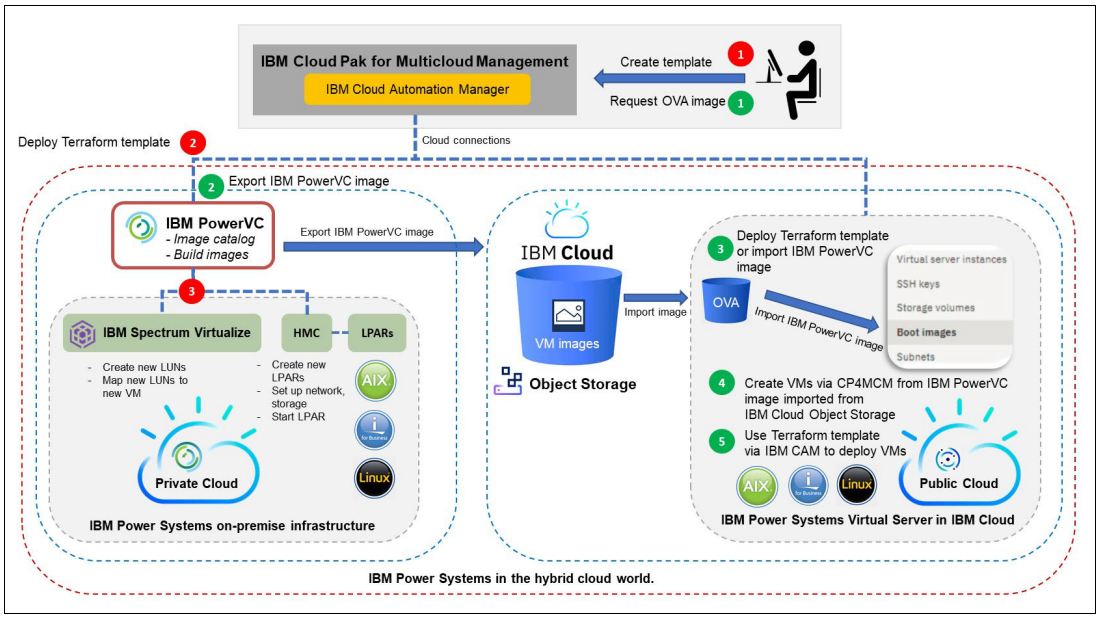

Figure 2-35 Export IBM i PowerVC images to IBM Power Virtual Server

IBM i images can be deployed on-premises or off-premises by using Infrastructure-as-Code (IaC) that uses HarshiCorp Terraform. You can also use Ansible for i for deployment and configuration management. This approach can be used for deploying an IBM i image by Terraform and configuration management that uses Ansible controller on IBM i as an in-house setup.

For an Enterprise solution, the use of IBM Cloud Pak for Multicloud Management that is integrated with Terraform, Red Hat Ansible Automation Platform, and Ansible Tower (compatible with IBM i systems), can be a solution for IBM i workloads management. It can be used as a DevOps strategy on IBM i, including application deployment, reduce delivery cycles, and so forth. A hybrid cloud approach also is committed for IBM i systems.

2.5 Other migration options in IBM Cloud catalog

The IBM Cloud Catalog contain several offerings that are related to products and services that can be used to migrate IBM i workloads to IBM Cloud. Figure 2-36 on page 63 shows multiple products for migration services that are offered by IBM Cloud that are specific for IBM i.

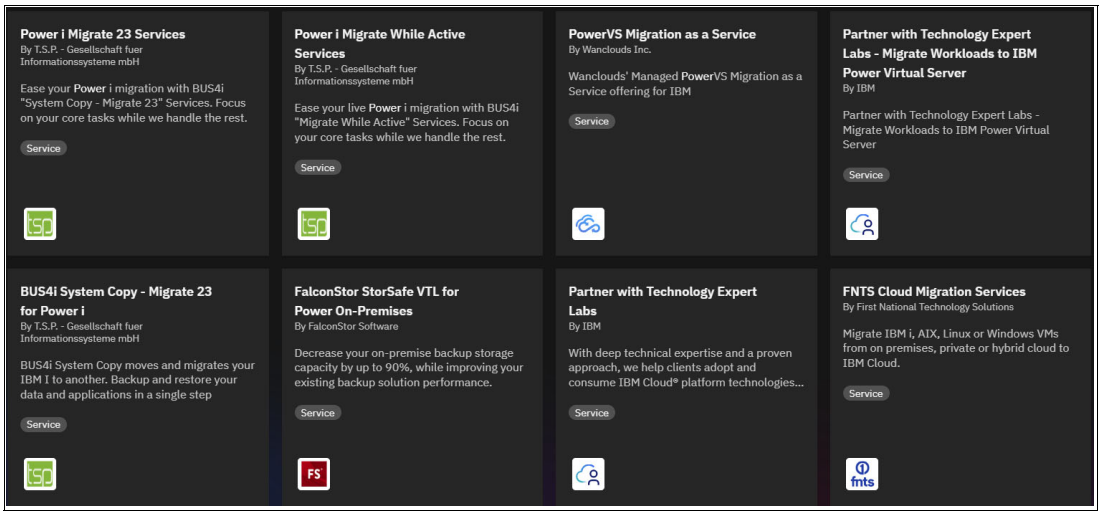

Figure 2-36 Product offer by IBM Cloud

T.S.P. Bus4i Migrate 23

BUS4i System Copy-Migrate 23 for IBM i can help simplify the process of moving and migrating your IBM i to another system, whether it's in the cloud or on-premises.

You can use this product to migrate an IBM i workload without the need for tapes or additional disk/storage space. Only a TCP/IP network connection is required, and the backup and restore processes begin simultaneously. The entire procedure is designed for the user's convenience and to ensure a smooth and efficient migration. Migration can be run unattended, and after it is finished, the user receives an email notification.

This product is useful for off-line migrations, that is, migrations in which the system to be migrated can be put in a restricted state while the migration lasts.

The basic offer includes only the license and does not include services. It is a one-time charge license for the migration tool in self -service modality (that is, without T.S.P. service hours). This license is valid for 30 days and for migrating one IBM i partition.

For more information, see BUS4i System Copy - Migrate 23 for Power i.

Also, when the Customer wants the offline migration with Bus4i System Copy-Migrate 23 to be run by the experts of T.S.P., it is available as an offering that complements the previous one, adding Professional Services of T.S.P. to provide the assistance required for migration.

This is a good option for companies that seek expert assistance with the migration of their IBM i systems, because the T.S.P. professional services. The services of T.S.P. offer a comprehensive and practice solution by experts with many years of experience, which will ensure perfect and efficient migration. For more information, see Power i Migrate 23 Services.

T.S.P. Bus4i Migrate While Active

BUS4i “Migrate While Active” Professional Services can ensure a smooth and efficient migration of your Power i system to an IBM PowerVM® in Power Virtual Server or any other Power i system anywhere in the world without disruption to your daily operations.

This professional services offering uses a proven methodology that combines advanced tools and expertise to migrate your IBM i system to an IBM i System in IBM Cloud (or in any other place) in a nondisruptive manner.

The team of experts from T.S.P. works in close collaboration with the customer to understand their specific requirements and develop a personalized migration plan. This plan incorporates the configuration, applications, and workload of your system to help provide a smooth transition.

Using specialized tools and techniques, data and applications are migrated from the source IBM i system to the new target system. This process is designed to minimize the interruption and to help guarantee the integrity of the data. Finally, after rigorous synchronized validation tests, a planned switchover is run for the enable of the productive system in its new environment, minimizing any potential impact.

This offering applies to online migrations in which users must continue working in the system while the migration lasts and that the only interruption is the planned switchover to move to the system in the IBM Cloud after both copies are synchronized.

This offer includes the license and services required in a single package. The license is valid for 30 days to migrate one IBM i partition.

For more information, see Power i Migrate While Active Services.

Rocket iCluster

Rocket Software's iCluster product is also available in the IBM Cloud catalog and can be used to migrate IBM i workloads to the IBM Cloud. Rocket iCluster high availability/disaster recovery (HA/DR) solutions ensure uninterrupted operation of IBM i applications, providing continuous access by monitoring, identifying, and self-correcting replication issues.

By using Rocket iCluster, a full replication from the on-premises system to a target system in IBM Cloud can be established so that both systems can be synchronized without affecting the availability of the source system. When both systems are synchronized, the migration to IBM Cloud can be run in minutes through a Planned Switchover more simply and more efficiently.

For more information, see Rocket iCluster HA/DR.

Power Virtual Server Migration as a Service by Wanclouds Inc

This managed service is led by Wanclouds cloud team by using their migrations and automation suite called VPC+ to help accelerate the entire migration and eliminate the long, complicated, and error-prone manual processes.

With this offering you can move from any of the following configurations:

- On-premises IBM Power System LPARs (AIX, IBM i, x86, Linux on Power)
- IBM Cloud - IBM Cloud Classic
- On Premise (including VMware)
- Amazon Web Services
- Google Cloud Platform
- Microsoft Azure

From the configurations in the previous list, you can move to any of the following configurations:

- IBM Cloud Power Virtual Server (AIX, IBM i)
- IBM Cloud - IBM Kubernetes Service
- IBM Cloud Satellite®
- VMware on IBM Cloud Class

Wanclouds presents IBM Power Virtual Server Migration as a Service and is designed to facilitate migration of your workloads to IBM Power Virtual Server cloud infrastructure.

With Power Virtual Server Migration as a Service, the Wanclouds team can craft a customized migration strategy based on each customer's unique requirements to help ensure minimal disruption. Also, the customized strategy can provide end-to-end migration support, from initial assessment and planning to execution and post-migration support.

This is done with a priority placed on mitigating risk by conducting thorough assessments of your existing infrastructure, and planning to ensure a smooth transition. Wanclouds also provides continuous monitoring and proactive support to address any issues promptly, ensuring the smooth operation of your workloads. For more information, see Power VS Migration as a Service - IBM Cloud.

FalconStor StorSafe for Power On-Premises by FalconStor Software

FalconStor StorSafe is a software solution that optimizes backup and restore to improve performance and significantly reduce backup storage costs without requiring changes to the existing backup solution currently being used. With its integrated de-duplication, the solution removes redundant copies of data, thereby reducing capacity requirements, decreasing storage costs, and minimizing replication and restore times. StorSafe can be used with all leading backup solutions, and enables both hybrid and native-cloud backup and both workload and tape migration to the cloud, including IBM Power Virtual Server.

For more detail on VTL, check out 4.4, “FalconStor overview” on page 165.

Partner with IBM Technology Expert Labs by IBM

Technology Expert Labs-Cloud (TEL-Cloud) is a professional services organization that is run by an experienced team of IBM Cloud experts with deep technical expertise, valuable tools, and proven methodologies to design and build IBM Cloud solutions.

TEL-Cloud provides professional services for all aspects of IBM Cloud adoption. The services include Cloud Native development, Hyper Protect Crypto Service (HPCS), IBM Kubernetes Services, Red Hat OpenShift Kubernetes Service (ROKS), IBM Satellite™, Security Compliance Center (SCC), VMware, and associated core capabilities in IBM Cloud for networking, firewalls, security, backup, and disaster recovery.

TEL-Cloud provides professional services for IBM Cloud in an offering framework of Architect and Build services. In addition, TEL-Cloud offers a monthly subscription called Expertise Connect for clients who want ongoing IBM Cloud subject matter expertise to plan and guide their IBM Cloud adoption.

FNTS Cloud Migration Services by First National Technology Solutions

First National Technology Solution (FNTS) is an IBM Business Partner and offers custom solutions based on business needs to migrate IBM i, AIX, Linux, and Windows workloads to IBM Cloud. FNTS works with you to build a migration plan and estimate the service credits required to deliver the project and ongoing services.

The use of FNTS provides the following features and capabilities:

- Create a Cloud Migration Plan. Build a custom migration plan based on current technology stack, Recovery Time Objective, Recovery Point Objective, and business requirements
- Prepare Target Cloud Environment. Provision and configure network, virtual servers, and related infrastructure that are required for the target environment
- Migrate Data to Cloud. Migrate data by using supported backup and replication tools
- Configure and Manage Cloud Backups. Configure a supported cloud backup solution

For information, contact FNTS.

2.6 Migrating IBM i by using IBM i Migrate While Active

Announced in October 2024, IBM i Migrate While Active provides the capability to safely migrate an IBM i partition. A new instance of the IBM i partition, otherwise known as the copy node, is established and maintained by Migrate While Active until the customer decides to finalize the migration.

IBM i Migrate While Active orchestrates the migration of an IBM i logical partition (LPAR) to a different location. That location can be within the same server, the same data center, or a remote location such as IBM Power Virtual Server on IBM Cloud.

The process starts with taking a source LPAR at IBM i 7.4, IBM i 7.5, or later, making a copy of that LPAR, and migrating data from the source to the copy. While the source node is still active, changed data is continuously migrated to the copy node. When the two nodes are in sync, a cutover to the new node can be performed to complete the migration.

Migrate While Active is optimized for migration whereas a physical tape cannot be used for a traditional D-Mode IPL and restore. Migrate While Active handles the save to optical media, so a D-mode network installation can be performed. Figure 2-37 provides a high-level view of the process.

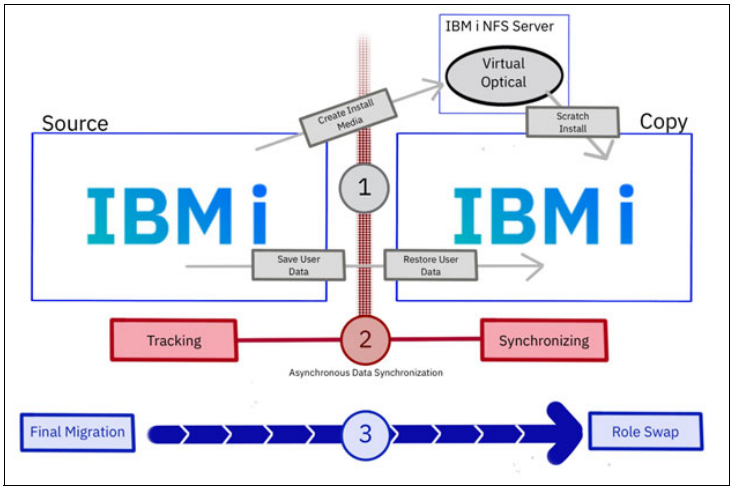

Figure 2-37 IBM i Migrate While Active

IBM i Migrate While Active builds upon the continuous availability technology of IBM Db2 Mirror for i. Migrate While Active (5700DBM - Option 2) is a separately licensed program option, and it is used independently from 5700DBM - Option 1 Db2 Mirror Enablement. If an IBM i system is already configured with IBM Db2 Mirror for i Option 1, Option 2 cannot be used for migration.

A traditional Db2 Mirror provides continuous availability with a two-node cluster, by using RDMA-capable hardware adapters to synchronize Db2 for i data, replication-eligible object types, and security attributes in real-time. The two IBM i nodes can operate at different maintenance levels with varying PTF levels and operating system levels. Migrate While Active creates a copy node that matches the source node's operating system level, PTF levels, and capabilities, as captured during the save process.

Instead of relying on RDMA capable hardware adapters, Migrate While Active uses TCP/IP as the communication protocol between the source and copy nodes. Another important difference is that traditional Db2 Mirror uses synchronous replication and Migrate While Active uses asynchronous replication. At any point during the migration process, the Migrate While

Active user can examine the remaining amount of migration needed before the copy node is in sync with the source node.

Lastly, Db2 Mirror can be used to achieve active/active, active/passive, or active/read-only availability solutions. Migrate While Active is similar to Db2 Mirror's active/read-only architecture. Changes to objects are synchronized only from the source node to the copy node. Migrate While Active goes beyond Db2 Mirror support by tracking and replicating changes to any object within *SYSBAS that is eligible to be saved. Figure 2-38 shows the basic overview of Migrate While Active.

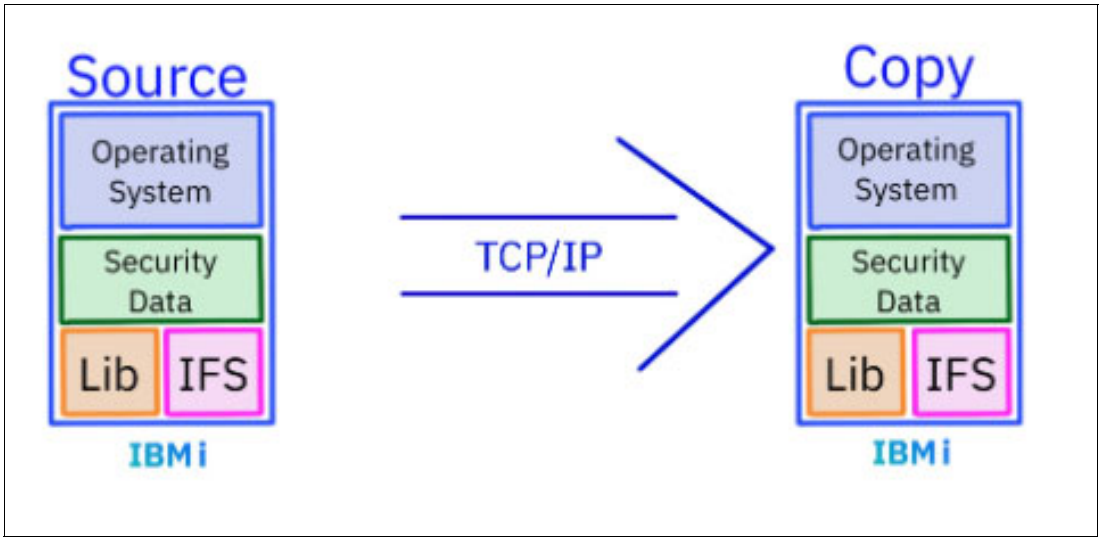

Figure 2-38 Basics of Migrate While Active

2.6.1 Migration use cases

Migrate While Active can be used to satisfy many business requirements and goals. The migration might be used to move an IBM i from one data center to another or be used to create a clone of an existing IBM i for some short-term project. This section explores the following cases that use Migrate While Active:

1. Migrate an on-premises IBM i from a local data center to IBM Power Virtual Server. The production node is moving to the IBM cloud.
2. Migrate an on-premises IBM i from a local data center to any IBM i cloud provider. The production node is moving to a cloud provider.
3. Migrate an on-premises IBM i to a newer and more advanced Power system. The production node is moving to new hardware within the same data center.
4. Migrate an IBM i to create a clone. The production node remains intact and is merely used as the master source of the copy node.

These and many other use cases are made possible with Migrate While Active.

2.6.2 Anatomy of a migration

Migrate While Active can make it easier than traditional approaches to migrate an IBM i. This section provides the steps that take place when migrating an IBM i.

A Migrate While Active migration consists of three distinct stages. (Figure 2-39)

1. Set up and migrate the system
2. Synchronize the data when the migration is nearly complete
3. Initiate a system cutover when the data is synchronized and validated

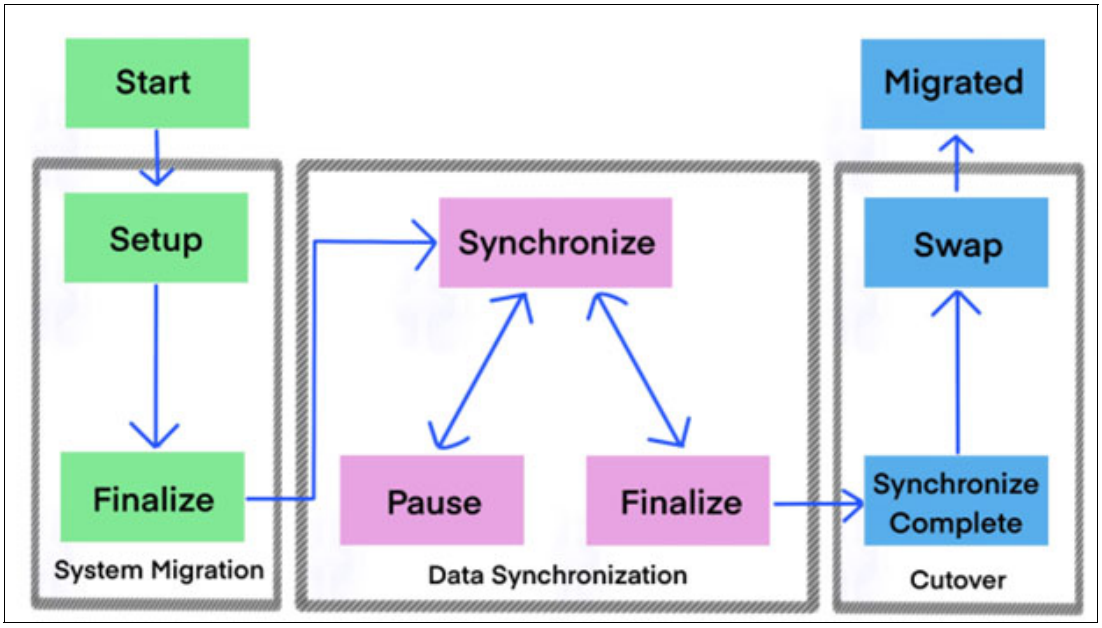

Figure 2-39 Migrate While Active Stages

1. System Migration. Creation of the copy node

   During the system migration stage, Migrate While Active is configured and a system save is performed on the source node. The save includes the operating system and some or all of the existing user data. The system save is done while in a restricted state. The Migrate While Active user indicates how much user data to save. All user data saves are focused on objects within *SYSBAS.

   A unique version of the install media is created during the system save and is used to install a temporary copy node. Then any saved user data is restored onto the copy node.

2. Data Synchronization. Keep the copy node synchronized with the source node

   After the system save of the source node is complete, users and applications can run production workloads on the source node. During this time, objects within *SYSBAS libraries and the integrated file system can be created, changed, or deleted. Any change to these objects is tracked by Migrate While Active.

   After the system migration is complete, Migrate While Active begins the continuous replication stage that is known as data synchronization. During this time, Migrate While Active is continuously but asynchronously keeping the copy node synchronized with the source node.

3. Cutover. Complete the migration

   The Migrate While Active user decides when to leave the data synchronization stage and finalize the migration to the copy node.

   During the cutover stage, the source node enters a restricted state and a final synchronization of any remaining tracked objects occurs. After the final synchronization is complete, the Migrate While Active user decides whether to complete the migration. Upon

successful completion of cutover, data synchronization between the source and copy nodes is ended and the copy node is ready to be used.

2.6.3 Data synchronization jobs

There are three batch jobs that drive the synchronization of changed objects:

1. QMRDBESYNC

   Handles tracked Db2 Mirror replication-eligible objects. The technology that is used is the Db2 Mirror resynchronization processing. For database files where rows changed, bundles of changed rows are synchronized to the copy node. For other changes to replication-eligible objects and for security attributes, Db2 Mirror uses its resynchronization processing. The RESYNC_STATUS view returns information about tracked changes to replicated objects on the object tracking list (OTL). Even though the OTL contains the complete record of tracked user actions, an object by object breakdown of tracking is also seen when querying the LIBRARY_MIGRATION_LIST view.

   For Db2 Mirror replication-eligible objects, Migrate While Active tracks every action that resulted in changes to replication-eligible objects.

   The resynchronization processing is focused on replicating the tracked user actions to the copy node.

2. QMRDBNSYNC

   Handles tracked non-Db2 Mirror replication-eligible objects within libraries. Multiple changes to the same object typically yield a single entry on the LIBRARY_MIGRATION_LIST view. The LIBRARY_MIGRATION_LIST view returns a list of all eligible library-based objects that require migration to the copy node by Migrate While Active.

For non-Db2 Mirror replication-eligible objects, tracking is focused on the library name, object name, and object type. Actions that cause a change to an object result in a single entry on the LIBRARY_MIGRATION_LIST.

The resynchronization processing saves and restores the object from the source node to the copy node.

3. QMRDBISYNC

   Handles tracked integrated file system objects. Multiple changes to the same object typically yield a single entry on the IFS_MIGRATION_LIST view. The returns a list of all the changed integrated file system objects that need to be migrated to the copy node.

For integrated file system objects, tracking is focused on the path name and object type. Actions that cause a change to a specific integrated file system object result in a single entry on the IFS_MIGRATION_LIST.

The resynchronization processing saves and restores the object from the source node to the copy node.

2.7 Migration example and use case

This section provides an overview of a migration that is based on an actual customer experience. The discussion includes an overview of the architecture of the solution and a discussion of the decision points in the planning for the migration. The discussion describes

how the final architecture meets the functional and non-functional requirements of the customer. An overview of an architecture diagram is presented.

2.7.1 Use case overview

In the use case, a customer wants IBM to move 10 IBM i virtual machines (VMs) that are in La Paz, Bolivia and named LPZ01 and LPZ02 to IBM Power Systems Virtual Server on IBM Cloud in SAO01 and SAO04, both in Sao Paulo, Brazil.

The customer requires that all IBM i VMs must be moved “as is”, with operating system release 7.1. No transformation is included in the project, except the operating system upgrade to 7.3.

The following VMs are duplicated to SAO04:

- Four production IBM i VMs
- One development VM
- One archive VM
- Four Disaster Recovery (DR) VMs

This environment includes storage, backups in IBM Cloud Object Storage, and an internal network, which features firewalls and a jump server for the IBM i VMs access management.

Also, this use case includes a third-party vendor logical replication solution between the source and target, which replicates across different zones in IBM Cloud (see Figure 2-40).

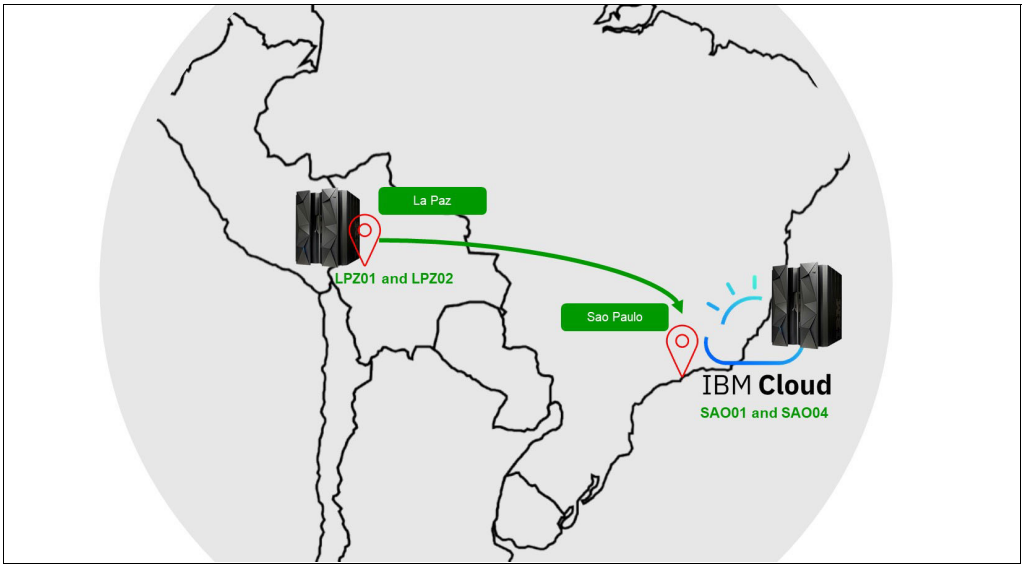

Figure 2-40 Migrating IBM i VMs from La Paz (on-premises) to Sao Paulo (off-premises)

Note: Bolivia and Brazil are in the same continent (South America). Therefore, this use case includes multiple zones. The cities of La Paz and Sao Paulo are separated by 3000 kilometers (1800 miles). For more information about multi-zone regions, see IBM Cloud region and data center locations for resource deployment.

2.7.2 Scope

In this use case, all 10 IBM i VMs are migrated from two data centers, LPZ01, which is the source, and LPZ02, which is the target and are managed by a customer in La Paz, Bolivia, to the IBM Cloud data center in SA001 (the source) and SA004 (the target).

Development and archive VMs that are in the La Paz-Bolivia data center are migrated to IBM Cloud by using IBM Cloud Object Storage backup. Then, they are restored to the target VM in IBM Power Virtual Server on IBM Cloud. The IBM i production VMs that are on LPZ01 are migrated by using the IBM Cloud Object Storage backup and logical replication method.

Note: A maintenance contract, an activation key of the logical replication, and installation and configuration done by the third-party vendor are often needed for this process.

After the IBM i VMs are active, the data is restored to the target location and the connectivity from La Paz data center is established by using IBM i Access Client Solution (ACS). After connectivity is validated, the customer's applications and their functions are configured in the environment.

Also, each VM and the overall function for each environment are tested.

IBM provides to customers the temporary technical information, such as IP addresses, DNS names, and a VM temporary name. This information is required so that the tests can be completed without any errors on IBM i production environments that are established in La Paz Bolivia.

The customer provides the firewall rules that are to be configured by IBM. User profiles are automatically transferred to the IBM Power Systems Virtual Server on IBM Cloud that is restoring user profiles (**RSTUSRPRF**). It is the customer's responsibility to administer their applications and user ID management.

After migration, IBM can run the suitable security scan for each IBM i environment that is based on the customer's security policy. IBM also provides the evaluation results to the customer.

Often, the following teams (among others) are involved in the transition and steady state support of the environment:

- IBM Cloud connect for the network
- IBM Cloud Object Storage
- IBM Power Virtual Server
- IBM i support team
- Customer's application team
- Third-party vendor software team support for logical replication

Image catalogs are created out of objects that are backed up by using optical devices. These catalogs must be restored on the IBM Power Systems Virtual Server instance by using some of the migration strategies that use IBM Cloud Object Storage and NFS server.

IBM Backup, Recovery, and Media Services (BRMS) is an IBM i product that can be used to automate tasks that help define and process your backup, recovery, and media management operations. IBM Content Collector can be integrated with IBM BRMS to move and retrieve objects from remote locations, including IBM Cloud Object Storage.

The following procedure shows how the customer's IBM i operating system and data is migrated from an on-premises system to the IBM Cloud environment. Most of these steps can be automated by using IBM BRMS and IBM Content Collector:

1. Install IBM BRMS and IBM Content Collector.
2. Install IBM i System Minimum PTF levels.
3. Install BRMS and IBM Cloud Object Storage PTFs.
4. Install Cloud-Init for IBM i and PTFs.
5. Save the system by using IBM BRMS and IBM Content Collector to IBM Cloud Object Storage.
6. Create IBM BRMS Recovery Reports.
7. Copy the data to IBM Cloud Object Storage.
8. Copy the data from IBM Cloud Object Storage.
9. Build a VM in Power Virtual Server by using NFS VM.
10.. Restore the user data by using the image catalog.

Infrastructure considerations

The technical infrastructure architecture and design consist of deploying several zones in IBM Cloud in Sao Paulo. In addition to the IBM Power Systems Virtual Servers zone that is named `Power Colo` (Power Colocation), a front-end zone is created in which jump servers are deployed for users access and other functions, such as a proxy for accessing the IBM Cloud Object Storage services. These services are hosted in a bare metal server.

A cluster of firewalls is deployed in the front-end zone within the `SAO01` site. A stand-alone firewall is deployed in `SAO04`.

Functional requirements

The solution must satisfy functional and nonfunctional requirements in a way that best balances the concerns of competing stakeholders and considers any relevant constraints (see Table 2-1).

Table 2-1 Functional requirements

Requirement	Description
Management services for IBM i	Support for IBM i operating systems.
Provide backup services	IBM Cloud Object Storage is used as backup-services.
Provide multi-site HA solution	Multi-site infrastructure to be provided in Sao Paulo (dual sites).
Dual sites high availability	The use of any third-party vendor solution on logical replication to establish DR between SAO01 and SAO04.
Provide fault tolerant LAN infrastructure in IBM Cloud	Provide network connectivity for applications and servers.
Data replication between IBM Cloud and customer data center	The customer uses a logical replication solution to replicate the data for the IBM i application.
Provide traffic isolation and segmentation	The use of jump servers and traffic filtering on IBM Cloud.

Requirement	Description
Provide WAN connectivity	Th customer provides WAN circuit and the POP network infrastructure. IBM provides the termination endpoint in Sao Paulo.

Nonfunctional requirements

The following nonfunctional requirements must be met:

- IBM Cloud portal access for IBM i VMs provisioning.
- Worldwide Tools Solutions for alert monitoring and reporting (IBM i).
- Traffic bandwidth in IBM Cloud infrastructure does not exceed 1 Gbps.
- Traffic bandwidth for replication is limited to 500 Mbps. Internet is used for preserving production traffic.
- Local network redundancy is provided in the primary IBM Cloud site (`SA001`). Firewall cluster in High Availability, dual ports connectivity.
- Manageability access for customer's users is provided by using jump servers.

Note: For more information, see WWTS (Worldwide Tool Solutions).

2.7.3 Architectural decisions

Architectural decisions help explain why the solution architecture is used across the scope of the solution. More than one conceivable arrangement choice is available for a specific architectural issue.

This choice incorporates components and their connections, innovation options, allotting usefulness to different components, making situation choices for components that are facilitated inside different infrastructure nodes, and so on.

The choices can have diverse costs that are related to them, the degree they fulfill different prerequisites and can show distinctive ways of adjusting competing stakeholders' concerns. Architects can perform the following tasks:

- Formally archive the basic choices that they make in creating the arrangement
- Agree as to why the arrangement looks the way it does.

In summary, architectural decisions focus on the following factors:

- Explaining the method that was chosen and why that method is being used.
- Ensure that the arrangement fulfills functional and nonfunctional prerequisites. If it does not fulfill those requirements, then offer assistance.
- Avoid superfluous adjustments through the arrangement conveyance lifecycle.
- The following tables list some example decisions about lifting and shifting IBM i VMs from the various architectural perspectives:
  - Table 2-2 on page 74: Infrastructure
  - Table 2-3 on page 75: Migration
  - Table 2-4 on page 75: Servers
  - Table 2-5 on page 76: Networking

Table 2-2 Infrastructure: Front-end accounts

Architectural decision	Front-end accounts that are used for accessibility and provisioning some services.
Problem statement	Providing a way for accessing the target Power Virtual Server, which is moved from the customer's data centers to IBM Cloud in Sao Paulo in a dual-site configuration.
Assumptions	The customer provides the WAN connectivity up to the network PoP Equinix[a] that is next to the data center.
Motivation	Standard design for this type of solution.
Alternatives	None.
Decision	Deploy front-end account and services.
Justifications	For accessibility of the Power Virtual Server, a front-end zone is needed. Some services are provided, such as the control of network traffic (firewalls), a relay environment for accessing the target IBM i images, and a Proxy for IBM Cloud Object Storage access.
Implications	Deploy WAN access and replication method for moving the data in the target environment.
Derived requirements	► Provide firewall services for VPN access and filtering of traffic. ► Provide IBM Cloud Object Storage services for Backup. ► Provide WAN network connectivity for customer's users and application connectivity. ► Provide bare metal servers to hosts relay applications and proxy.

a. For more information about Equinix, see Equinix America data centers.

Consider the following points:

- By using Equinix, you can get a direct link to reach IBM Power Systems Virtual Server over Direct Link Connect. From Equinix, you can also get a cross connection to Megaport and connect to IBM Power Systems Virtual Server directly.
- Before you begin, verify the location connection to IBM Cloud by verifying your colocation provider's capabilities or service provider's capabilities to reach the Meet-Me Room and cross-connect into IBM Cloud. For more information, see IBM Cloud Direct Link.

  For `SA001`, the location type is a data center, and the Meet-Me Room (MMR) Operator is Digital Realty (Ascenty). For `SA004`, the location is data center (AZ1) and the operator is ODATA.

  In this example, to transfer IBM i VMs from Bolivia to Brazil, the connection from Bolivia to `SA001` can be established. For example, you must contract directly with a carrier that can provide enough capacity and presence in any Ascenty data center. The solution might be LAN-to-LAN + Cross Connection Fiber Optic + IBM DirectLink 1 G or 10 G.

  For the LAN-to-LAN link, IBM must directly contract with the carriers for the private LAN-to-LAN circuits.

  For example, the following carriers and the location of the Ascenty data centers are used:

  - ALGAR TELECOM
  - ALOO TELECOM
  - AMERICANET
  - ANGOLA CABLES
  - BRDIGITAL, LUMEN

- COGENT
- SEABORN
- CLARO - PRIMESYS - EMBRATEL
- Durand/Tavola
- HOSTFIBER
- MEGATELECOM
- MUNDIVOX
- NETELL
- NEUTRONA NETWORKS
- NOVVACORE TELECOMUNICAÇÕES
- OI MOVEL
- Seaborn
- SILICA NETWORKS
- VIVO
- VOGEL
- WCS TELECOM

For more information about Digital Realty (Ascenty) and ODATA, see the following web pages:

- Ascenty
- ODATA

Table 2-3 Infrastructure: Dual site

Architectural decision	Dual site infrastructure is required for HA purposes.
Problem statement	If a major outage occurs, the customer's users can connect to the back up site. The use of DNS for servers translation is needed. The secondary site has a different TCP IP address.
Assumptions	Two sites are used for the solution: one in SAO01 and the other in SAO04 (in a different zone).
Motivation	Infrastructure recovery if a major outage occurs.
Alternatives	None.
Decision	Deploy a dual-site solution in an IBM Cloud Multi-Zone Region (Sao Paulo).
Justifications	If a major outage occurs at the primary site, the main goal is to restart part of the application and services in the secondary site.
Implications	Deploy a secondary site in addition to the Production environment.
Derived requirements	► Provide WAN network connectivity to the secondary site for customer's users and application connectivity. ► Duplicate part of the primary infrastructure in the backup site.

Table 2-4 Migration strategy and backup

Architectural decision	IBM Cloud Object Storage backup is used for migrating IBM i VMs to SAO01 and SAO04.
Problem statement	Back-up and data replication between client data center and IBM Cloud target infrastructure. No Automatic Tape Library (ATL) or VTS is available to perform a save and restore, which is a traditional migration method for IBM i operating system.

Architectural decision	IBM Cloud Object Storage backup is used for migrating IBM i VMs to SAO01 and SAO04.
Assumptions	The use of IBM Cloud Object Storage for the migration is one of the available methods to move the workload to IBM Power Systems Virtual Server in IBM Cloud.
Motivation	The use of IBM Cloud Object Storage to move IBM i workloads to SAO01 and SAO04.
Alternatives	The use of IBM Cloud Object Storage for migration. The use of a Master Data Management device for the migration. Transferring IBM i image OVA file to IBM Cloud Storage by using IBM Power Virtualization Center.
Decision	IBM Cloud Object Storage is used.
Justifications	► Master Data Management is excluded because of the delay for taking backups data on a 1 Gbps interface. ► Customer does not have a virtualization by PowerVC.
Implications	Network connectivity includes VPN WAN connectivity and Proxy in front-end account.
Derived requirements	► Deploy Proxy in Front-End zones and VPN access from client on IBM Cloud. ► Buckets are needed to create on IBM Cloud Object Storage for the data move. Needs more storage for the IBM Cloud Object Storage backup in the source IBM i VM.

Table 2-5 Delta data synchronization

Architectural decision	A third-party vendor replication solution is used.
Problem statement	The customer has IBM i VMs with storage of 10 TB–70 TB. Some of the IBM i VMs include a journal of 1 TB that is generated daily. The use of a logical replication tool is the best solution to remove the delta data after the IBM i VM restoration on the target is complete.
Assumptions	The logical replication tool syncs up the data between source system and target system.
Motivation	Customer uses a logical replication.
Alternatives	PowerHA for IBM i geographic mirroring.
Decision	The use of the current third-party vendor tool.
Justifications	In this case, this option is the most suggested because the customer has a third-party logical replication tool in support.
Implications	Third-party vendor provides license key for the logical replication tool to be deployed on IBM Power Virtual Server. The third-party vendor installs the tool in the cloud.
Derived requirements	The third-party vendor provides a temporary license key to migrate data to IBM Cloud.

Note: PowerHA Geo mirroring in the IBM Power Systems on IBM Cloud does not include any unique features. It is host-based and done over a TCP/IP connection. The only possible unique part is that you must design the network between the two servers, but that requirement is not unique to PowerHA. Geo Mirroring includes limitations, such as it works for IASP only, with which a customer cannot be configured.

Table 2-6 Networking: IBM Cloud Direct Link Dedicated on Classic

Architectural decision	WAN direct-link connectivity to be redundant: one primary and one secondary link.
Problem statement	WAN access connectivity to be recovered in case of a primary link outage.
Assumptions	WAN part is the customer's responsibility, IBM Cloud provides dual circuit connectivity on diverse physical devices.
Motivation	Maintain connectivity with customer's corporate network.
Alternatives	Doubling the WAN connectivity: a redundant connectivity in SA001 and a redundant connectivity in SA004.
Decision	Provide redundant connectivity in SA001 and use the IBM Cloud backbone for Inter-site communications.
Justifications	The provided service level is consistent and the option is available to connect the IBM Cloud site by using VPN.
Implications	Deploy Sao Paulo site-to-site connectivity.
Derived requirements	Deploy GRE and Direct link connectivity for Front-End zones communications.

IBM Cloud Direct Link Dedicated

IBM Cloud Direct Link Dedicated is a single-tenant product (Table 2-6). It offers a dedicated port that is suited for banks, insurance companies, or anyone with strict compliance policies.

A fiber cross-connection is created through a network service provider (NSP) in an IBM Cloud network Point of Presence (PoP). IBM engineers facilitate end-to-end connectivity with your selected NSP, and you can access your cloud infrastructure in the local IBM Cloud data center.

The NSP runs last-mile links directly between a router on your network and an IBM Cloud router. As with all of the Direct Link products, you can add global routing that enables private network traffic to all IBM Cloud locations.

For more information, see this IBM Technology YouTube video.

IBM Cloud Direct Link is available in the following offerings:

- IBM Cloud Direct Link on Classic:
  - Direct Link Connect on Classic
  - Direct Link Dedicated on Classic
  - Direct Link Exchange on Classic
  - Direct Link Dedicate Hosting on Classic
- IBM Direct Link 2.0:
  - Direct Link Connect
  - Direct Link Dedicated

For more information about which Direct Link solution to order, see the following IBM Cloud Docs web pages:

- Getting started with IBM Cloud Direct Link on Classic
- Getting started with IBM Cloud Direct Link (2.0)

Important: Tables Table 2-2 on page 74 through Table 2-5 on page 76 list decisions only as an example. Real world decisions can vary according to the customer, the scenario, third-party vendor applications, in-house applications, region, networking, and so on.

A certified IBM i architect can help you decide about your scenario.

2.7.4 Architecture overview

Figure 2-41 shows an overview for IBM Cloud in Sao Paulo. The division of the different zones is delineated.

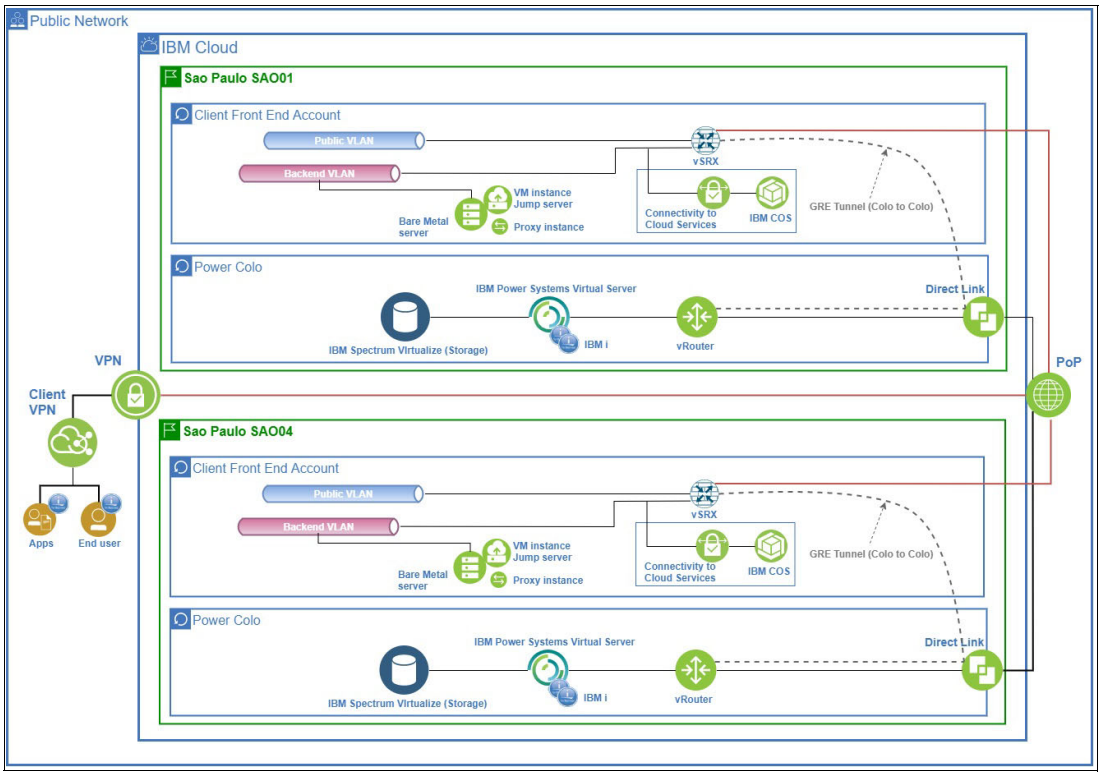

Figure 2-41 IBM i on IBM Power Virtual Server architecture

These zones are separated from each other by physical switches, and access to the zones is controlled by firewalls. A dual-site infrastructure is used. Consider the following points:

- `SAO01` is where the IBM i production runs, and it comprises two subzones:
  - The first subzone is named Client Front End Account. Here, you can see Jump servers and services, such as IBM Cloud firewalls and proxy. The two zones communicate by using the internal IBM Cloud network backbone.
  - The second subzone is named `Power Colo`, and it is here where IBM Power Virtual Server is installed.
  - `SAO04` is used for DR purposes if an outage occurs at `SAO01`.

Consider the following points:

- The System Context Diagram (SCD) provides the networking environment and the IBM i VMs and jump servers. It also identifies the interfaces between entities.
- An Operational Model provides a high-level description of an abstraction that specifies how the components of the application are organized and connected, where they are to be located and hosted, and what business requirements the application resolves.
- The backup and recovery strategy. The BRMS and IBM Cloud Storage Solutions for i are intended for small environments that save less than 2 TB of data per day. For more information, see BRMS and IBM Cloud Storage Solutions for i.
- PowerHA for IBM i Geographic Mirroring or a third-party vendor for logical replication solutions can be strategies for IBM i data to be replicated across zones in the IBM Cloud.
- For more information about Architectural decisions, see Architectural Thinking.

3

Networking Solutions for Power Virtual Server

This chapter includes the following topics:

- “Power Virtual Server virtual private network connectivity introduction”
- “VPN in cloud environments”
- “IBM Power Virtual Server Network overview”
- “Integration with IBM Cloud and other cloud services”
- “Security measures incorporated in Power Virtual Server networks”
- “Comparing IBM Power Virtual Server network features to similar cloud offerings”
- “IBM Power Virtual Server Network Scenarios”

© Copyright IBM Corp. 2025.

3.1 Power Virtual Server virtual private network connectivity introduction

IBM Power Systems Virtual Server offers a robust, scalable cloud infrastructure that is designed for enterprises that require high performance, reliability, and security. One of the foundational elements of ensuring secure communication within cloud environments is virtual private network (VPN) connectivity. This introduction covers the essential role VPNs play in IBM Power Virtual Server and their contribution to secure data transmission.

A VPN creates an encrypted tunnel between the on-premises environment and the cloud infrastructure, ensuring that data transferred across this connection is protected from unauthorized access and potential breaches. This tunnel enables enterprises to extend their private network into the cloud securely, maintaining data integrity and confidentiality. VPN connectivity within IBM Power Virtual Server includes the following key benefits:

- Data security and encryption. VPNs provide end-to-end encryption, which is crucial for protecting sensitive information as it travels between on-premises systems and the Power Virtual Server cloud.
- Private connectivity. Unlike public internet connections, VPNs allow for secure and private access, reducing exposure to potential attacks.
- Seamless integration. VPNs in IBM Power Virtual Server are designed to integrate smoothly with existing on-premises networks and support hybrid cloud deployments.
- Cost-effectiveness. Establishing a VPN connection is often more economical than other private networking solutions, such as dedicated lines, but still provides a reliable, secure link.
- Scalability. All VPN connections can be adjusted to meet the evolving needs of businesses, which makes them suitable for both small-scale projects and large, complex workloads.

IBM Power Virtual Server supports various VPN configurations to cater to different business requirements, such as site-to-site VPNs, which connect entire networks, and point-to-site VPNs, which allow individual devices to connect securely. These options ensure that enterprises can tailor their connectivity solutions to their specific use cases, whether they involve development, testing, or full-scale production.

The next sections delve deeper into the network overview and specific scenarios for deploying and managing VPN connectivity within the IBM Power Virtual Server ecosystem. This comprehensive understanding helps enterprises optimize their cloud strategies while maintaining security and operational efficiency.

3.2 VPN in cloud environments

A Virtual Private Network (VPN) is a technology that creates a secure and encrypted connection over a less secure network, such as the internet. In the context of cloud environments, VPNs are crucial for enabling protected communication between on-premises infrastructure, remote users, and cloud-based services. This technology ensures that data sent between these endpoints remains private and secure, which is essential for businesses handling sensitive or proprietary information.

Definition of VPN

A VPN acts as an encrypted tunnel through which data travels between two or more network points. This tunnel masks the data packets, making it difficult for unauthorized parties to intercept or read the transmitted information. VPNs use protocols such as IPsec (Internet Protocol Security) and SSL/TLS (Secure Sockets Layer/Transport Layer Security) to authenticate and encrypt communications, maintaining confidentiality and data integrity.

3.2.1 Significance of VPN in cloud environments

The increasing reliance on cloud computing has transformed the way organizations manage and store data. While cloud services offer scalability, flexibility, and cost-effectiveness, they also pose security challenges. VPNs play a significant role in addressing these challenges by providing a secure method of accessing cloud resources.

1. Secure data transmission. In cloud environments, data transmission often takes place over the public internet, which exposes it to potential security threats such as data breaches and cyberattacks. A VPN ensures that data transferred between on-premises systems and cloud services is encrypted, protecting it from interception. This encryption makes it virtually impossible for unauthorized parties to decipher the data, ensuring that sensitive information remains confidential.
2. Privacy and anonymity. VPNs not only encrypt data but also anonymize the user's IP address, making it difficult for third parties to trace or track online activities. This feature is essential for businesses that need to maintain privacy when accessing or managing cloud applications.
3. Regulatory compliance. Many industries, such as healthcare, finance, and government, are subject to strict data protection regulations like HIPAA, GDPR, and others. Using a VPN helps organizations comply with these regulations by ensuring secure data handling practices when transmitting information to and from cloud environments.
4. Enhanced security for remote access. The growing trend of remote work has highlighted the importance of securing access to cloud resources. A VPN provides a secure way for remote employees to connect to their company's internal cloud-based systems. By doing so, it mitigates the risks associated with unsecured public wifi networks or non-secure connections that might potentially expose corporate data.
5. Hybrid cloud integration. For organizations adopting a hybrid cloud approach, VPNs offer seamless integration between on-premises data centers and cloud services. This connectivity ensures that data and applications can move securely between private and public clouds, facilitating hybrid workloads and operational continuity.
6. Cost-effective security solution. Implementing private connections like leased lines can be costly and complex. VPNs offer a more economical alternative without compromising on security. They provide a flexible solution for scaling network security as business needs evolve, making them ideal for both small businesses and large enterprises.
7. Resilience and reliability. VPNs can enhance the reliability of cloud environments by providing consistent and secure connections. This reliability is critical for maintaining business operations, especially when multiple sites need to access centralized cloud applications. Advanced VPN configurations can support load balancing and failover mechanisms to prevent downtime and maintain continuous access.

Thus, VPNs are an integral part of cloud security architecture, enabling organizations to confidently use cloud technologies while safeguarding their data. The protection, privacy, compliance, and seamless integration that VPNs provide make them a valuable tool in today's cloud-first business landscape. This ensures that organizations can focus on their core operations without worrying about the safety and privacy of their data and applications in the cloud.

3.2.2 Core components of VPN connectivity for IBM Power Systems Virtual Servers

To fully harness the capabilities of IBM Power Systems Virtual Servers while maintaining robust security, VPN connectivity serves as a vital component. Establishing and managing VPN connections involves a set of core components that ensure secure, efficient, and reliable data transmission between the on-premises infrastructure and IBM Power Virtual Server cloud. Understanding these components is essential for deploying VPNs effectively in a cloud environment.

1. VPN Gateway

   The VPN gateway is a network device or service responsible for connecting two or more networks over a secure, encrypted connection. In IBM Power Virtual Server, the VPN gateway acts as the central point that facilitates the link between the on-premises network and the cloud infrastructure. The gateway handles the following tasks:

   – Encryption and decryption of data packets
   – Routing traffic securely between networks
   – Establishing and maintaining secure tunnels using VPN protocols

   The VPN gateway is pivotal for ensuring that data transferred over the connection remains encrypted and secure from interception.

2. Encryption protocols

   Encryption protocols are fundamental to the secure operation of VPNs, as they define the way data is encrypted and decrypted during transmission. IBM Power Virtual Server supports industry-standard VPN encryption protocols to ensure data confidentiality and integrity:

   – IPsec (Internet Protocol Security). A widely-used protocol that offers strong encryption and authentication, suitable for site-to-site VPNs connecting entire networks.
   – SSL/TLS (Secure Sockets Layer/Transport Layer Security). Often used for point-to-site VPNs, SSL/TLS provides a secure tunnel for individual users to access cloud resources securely.

   These protocols help protect data from potential eavesdropping and tampering, ensuring compliance with stringent security standards.

3. Tunneling mechanisms

   Tunneling mechanisms are techniques used to encapsulate and transmit data packets securely through a VPN. In the context of IBM Power Virtual Server, these mechanisms help facilitate seamless data transmission over shared or public networks. Common tunneling protocols are described in the following list:

   – L2TP (Layer 2 Tunneling Protocol). Often combined with IPsec to provide encryption and secure tunneling for site-to-site and remote access VPNs.
   – GRE (Generic Routing Encapsulation). Useful for scenarios where multiprotocol traffic is transmitted across the VPN, supporting advanced routing features.

   The choice of tunneling mechanism depends on the specific needs of the enterprise, such as the type of traffic and the level of encryption required.

4. Authentication methods

   Authentication is crucial for verifying the identity of users and devices connecting to the VPN. IBM Power Virtual Server supports multiple authentication methods to ensure that only authorized users and devices gain access:

   – Pre-shared keys (PSKs). A simpler authentication method where both ends of the VPN connection use a shared secret key for verification.

- Digital certificates. Provides a more secure and scalable authentication method by using public key infrastructure (PKI) to validate identities.
- User authentication. In point-to-site VPNs, user credentials (such as username and password) are used to authenticate individuals connecting to the Power Virtual Server environment.

5. Network address translation (NAT)

   NAT is a key component that manages IP address mapping to allow communication between different networks while maintaining the privacy of internal network addresses. In IBM Power Virtual Server, NAT provides the following capabilities:

   - Preserve IP addresses by allowing multiple devices to share a single public IP address
   - Provide an additional layer of security by masking the internal network structure from external users

   NAT ensures that traffic is routed correctly between the on-premises network and the Power Virtual Server cloud, enhancing connectivity and protecting internal network configurations.

6. Routing and IP configuration

   Proper routing and IP configuration are essential for seamless data flow across VPN connections. IBM Power Virtual Server supports dynamic and static routing configurations to manage how data packets are forwarded between on-premises infrastructure and the cloud. Routing and IP configuration includes the following key elements:

   - Static Routing. Involves manually configuring routes for traffic to follow, offering more control over the path of data packets
   - Dynamic Routing. Uses protocols such as BGP (Border Gateway Protocol) to automatically adjust routing paths based on network conditions, improving efficiency and resilience

   Correct IP configuration ensures that resources in the cloud and on-premises networks can communicate without conflicts or connectivity issues.

7. Redundancy and high availability (HA)

   High availability is critical to maintaining uninterrupted VPN connectivity. IBM Power Virtual Server allows configurations that support redundant VPN gateways to ensure that if one connection fails, another can take over without disrupting the service. This redundancy is vital for the following issues:

   - Minimizing downtime and maintaining service reliability
   - Supporting failover mechanisms to quickly switch to backup connections during outages or maintenance

8. Monitoring and logging tools

Monitoring and logging tools help maintain the performance, reliability, and security of VPN connections. IBM Power Virtual Server provides tools and dashboards to support the following activities:

- Real-time monitoring of connection status, latency, and data transfer rates.
- Logging and auditing VPN activity to detect anomalies, troubleshoot issues, and ensure compliance with security policies.

These tools assist administrators in maintaining optimal network performance and securing data traffic over the VPN.

The core components of VPN connectivity for IBM Power Systems Virtual Servers include the VPN gateway, encryption protocols, tunneling mechanisms, authentication methods, NAT, routing and IP configuration, redundancy, and monitoring tools. These elements work together to create a secure, scalable, and reliable VPN solution that enables seamless connectivity between on-premises networks and the Power Virtual Server cloud. By understanding these components, organizations can effectively deploy and manage VPN connections, ensuring data security and operational efficiency in their cloud environments.

3.2.3 Role of security and data encryption in VPN connections

Security and data encryption are fundamental aspects of VPN connections, ensuring the confidentiality, integrity, and authenticity of data transmitted between different networks, such as an on-premises infrastructure and a cloud environment like IBM Power Systems Virtual Server. In today's digital landscape, where cyberthreats and data breaches are prevalent, strong security measures and encryption protocols are indispensable for protecting sensitive information and maintaining trust.

1. Ensuring data confidentiality

   The primary role of encryption in a VPN is to maintain data confidentiality. Encryption algorithms transform readable data (plaintext) into an unreadable format (ciphertext) that can be deciphered only by authorized entities possessing the correct decryption keys. This process ensures that even if the data is intercepted during transmission, it cannot be understood or misused by unauthorized individuals.

   VPNs use strong encryption protocols such as IPsec and SSL/TLS to safeguard data as it moves through the public internet or shared networks. These protocols provide end-to-end encryption, ensuring that data remains secure from its source to its destination.

2. Data integrity

   Beyond confidentiality, maintaining the integrity of data is critical. Data integrity ensures that information is not altered during transmission. VPNs achieve this through mechanisms such as cryptographic hash functions, which create unique *fingerprints* for data packets. If any changes occur to the data during transmission, the receiving end can detect the discrepancy through these hash values and reject any altered data. This helps protect organizations from man-in-the-middle attacks, where malicious actors attempt to intercept and modify data in transit.

3. Authentication and access control

   VPN connections use authentication protocols to verify the identities of devices and users accessing the network. This process ensures that only authorized parties can establish a VPN connection and access data. VPNs include the following methods of authentication:

   – Pre-shared keys (PSK). A shared secret key is configured on both sides of the VPN to authenticate the connection.

   – Digital certificates. Certificates issued by a trusted Certificate Authority (CA) validate the identity of the VPN peers and ensure a higher level of security.

   – User credentials. Point-to-site VPNs often require individual user authentication, where users provide a username and password to establish a secure connection.

   Proper authentication methods prevent unauthorized access and help maintain the security posture of the network.

4. Preventing eavesdropping and data breaches

   Data transmitted over the internet or public networks is vulnerable to eavesdropping and interception by cybercriminals. By employing strong encryption, VPNs create a secure tunnel through which data passes. This tunnel is protected by complex encryption

algorithms, making it extremely difficult for attackers to decipher the contents even if they manage to capture the data.

IPsec and SSL/TLS protocols use symmetric and asymmetric cryptography to secure data transmission:

- Symmetric Cryptography. The same key is used for both encryption and decryption, providing efficient and fast data encryption.
- Asymmetric Cryptography. Uses a pair of keys (public and private) to encrypt and decrypt data, enhancing security by ensuring that only the holder of the private key can decrypt the data.

This robust encryption reduces the risk of data breaches and ensures that sensitive information, such as business transactions, client data, and confidential files, remains protected.

5. Protection against cyberthreats

   VPNs play a vital role in safeguarding networks against a variety of cyberthreats:

   - Man-in-the-middle (MitM) attacks. Where an attacker intercepts and possibly alters communication between two parties. VPNs prevent these attacks by encrypting data, which makes it unreadable to unauthorized interceptors.
   - Packet sniffing. VPNs protect against packet sniffers that attempt to capture unencrypted network traffic. Encrypted VPN traffic ensures that captured data cannot be interpreted.
   - Denial of service (DoS) attacks. Some VPNs offer additional security features that can help mitigate DoS attacks by ensuring that malicious traffic is identified and blocked before it can impact network performance.

6. Compliance with regulations

   Organizations in industries such as healthcare, finance, and government are often required to comply with stringent data protection regulations, such as HIPAA, GDPR, and PCI-DSS. VPNs play an essential role in meeting these requirements by providing secure data encryption and transmission standards. This helps organizations avoid penalties, maintain customer trust, and demonstrate their commitment to safeguarding sensitive data.

7. Use of advanced encryption standards

   VPN connections use advanced encryption standards to ensure robust security, which can include the following commonly used standards:

   - Advanced Encryption Standard (AES). A widely trusted encryption algorithm used by VPNs due to its high security and efficiency. AES with 256-bit keys is considered military-grade encryption, providing a significant level of protection.
   - Triple Data Encryption Standard (3DES). Although less common today, 3DES is an older encryption standard that adds an extra layer of security by encrypting data three times. It has largely been replaced by AES for better performance and security.

8. Secure remote access

   VPNs are particularly crucial for enabling secure remote access to cloud resources. This has become increasingly important with the rise of remote work, where employees connect to corporate networks from various locations. By using VPNs, organizations ensure that remote connections are secure and encrypted, protecting data from potential vulnerabilities associated with public or unsecured networks.

The role of security and data encryption in VPN connections cannot be overstated. These elements form the backbone of a secure network, protecting data from interception,

modification, and unauthorized access. VPNs enable organizations to extend their private networks into the cloud with confidence, knowing that data integrity, confidentiality, and compliance with regulatory standards are maintained. In the context of IBM Power Systems Virtual Servers, VPN connectivity is a critical feature for ensuring secure, reliable, and encrypted communication, which helps to foster trust and operational excellence in cloud environments.

3.2.4 Scenarios that use VPN for initial setups and scaling environments

VPN connectivity plays a crucial role in the various stages of cloud adoption, from initial setup to scaling and optimizing the environment. Organizations use VPNs to ensure secure, encrypted communication between their on-premises infrastructure and the cloud, facilitating seamless integration and expansion as business needs grow. This section includes different scenarios that illustrate the use of VPNs in both initial setups and scaling of IBM Power Virtual Server environments.

Initial setup of hybrid cloud environments

When an organization first deploys its IBM Power Virtual Server infrastructure, a VPN is essential for securely connecting its on-premises data center to the cloud. This secure connection supports various network activities:

- Data migration. VPNs enable safe and encrypted data transfer during the migration phase, ensuring that sensitive information is protected against potential cyberthreats.
- Testing and validation. Before starting production workloads, enterprises can use VPNs to test and validate applications in the cloud environment. This helps identify and resolve connectivity, security, or performance issues before the applications are used in production.
- Network configuration and synchronization. VPNs facilitate the configuration of cloud resources to mirror on-premises setups, making it easier to integrate existing systems, services, and applications into the new hybrid environment.

Example 1

A financial services company expanding its operations to the cloud uses VPN to securely migrate sensitive customer data from its on-premises servers to IBM Power Virtual Server. During the initial setup, VPN encryption ensures data confidentiality, while secure tunnels maintain consistent communication between local and cloud networks.

Development and testing environments

Organizations often create isolated development and testing environments in the cloud to avoid disruptions to production systems. VPNs help securely connect these environments to on-premises infrastructure, enabling developers and testers to perform the following actions:

- Access on-premises resources. Teams working in the cloud can securely access legacy systems and databases hosted on-premises, facilitating comprehensive application testing and debugging.
- Collaborate remotely. VPNs allow distributed development teams to work collaboratively on cloud-based projects with secure access, promoting remote work and global collaboration.

Example 2

A software development company initiates a VPN connection between its local office and IBM Power Virtual Server to enable developers to test new applications against real-world data

from on-premises databases. This setup ensures data protection and allows for efficient testing cycles.

Scaling to meet increased workloads

As business requirements evolve, organizations must scale their cloud resources to handle larger workloads and increased data traffic. Without compromising data security, VPNs provide a secure and flexible solution for scaling up and integrating additional cloud services:

- Dynamic workload management. VPNs enable the secure distribution of workloads between on-premises and cloud environments, balancing computational loads and optimizing resource usage.
- Temporary capacity expansion. During peak periods or specific projects, enterprises can use the cloud to temporarily expand capacity. VPNs ensure secure, seamless data flow between on-premises resources and cloud-based extensions.

Example 3

An e-commerce company experiencing seasonal spikes in traffic can scale its IBM Power Virtual Server resources to accommodate the increase. A VPN maintains secure, consistent communication between the cloud environment and the main data center, ensuring smooth operation during high-demand periods.

Secure remote work and collaboration

VPNs play a vital role in enabling remote work by providing secure access to cloud environments from various locations. Organizations can use VPNs to establish connections so that employees can do the following tasks:

- Access internal systems securely: Employees working from home or other remote locations can access cloud-based systems and data securely through VPNs, protecting company assets from public network threats.
- Collaborate on projects: VPNs facilitate secure collaboration on projects hosted on IBM Power Virtual Server, allowing remote teams to share resources, access files, and collaborate in real time.

Example 4

A global consulting firm uses VPNs to enable its remote workforce to securely access IBM Power Virtual Server-hosted applications and data. This ensures that consultants can collaborate on client projects without exposing sensitive information to potential security risks

Proof-of-concept (PoC) deployments

Before committing to full-scale cloud adoption, organizations often start with PoC deployments to assess feasibility and performance. VPNs enable secure testing and evaluation of PoC setups by connecting them to the existing on-premises infrastructure in the following situations:

- Simulating production environments. VPNs make it possible to simulate real-world scenarios by integrating on-premises and cloud resources, helping organizations evaluate cloud performance under actual workload conditions.
- Data security during evaluation. Even during temporary PoC deployments, VPNs ensure that data remains secure, preventing potential breaches or data leaks.

Example 5

A healthcare organization planning to move its patient management system to the cloud sets up a PoC on IBM Power Virtual Server. A VPN is used to connect the PoC environment to its

on-premises database, allowing the IT team to evaluate data transfer performance and security without risking patient data exposure.

Disaster recovery and business continuity

VPNs are essential for setting up disaster recovery (DR) solutions that ensure business continuity during outages or failures. By connecting on-premises infrastructure with IBM Power Virtual Server, organizations can maintain real-time data synchronization and backup through secure VPN connections in the following situations:

- Data replication. VPNs facilitate the replication of data between primary data centers and DR sites in the cloud, ensuring that data remains up-to-date and ready for failover.
- Secure access to backup resources. In the event of an outage, VPNs provide secure access to backup systems hosted on IBM Power Virtual Server, minimizing downtime and maintaining productivity.

Example 5

A manufacturing company sets up a VPN to link its on-premises data centre with IBM Power Virtual Server for a disaster recovery plan. If there is a disruption, employees can securely access backup resources and continue operations without significant impact.

VPNs are indispensable in enabling secure communication between on-premises and cloud environments for initial setups and scaling operations. From data migration and testing to scaling workloads and remote access, VPNs ensure the security, reliability, and efficiency of cloud adoption processes. For organizations that use IBM Power Systems Virtual Servers, the use of VPNs helps maintain data protection and optimize performance while providing the flexibility to expand and adapt to business needs.

3.3 IBM Power Virtual Server Network overview

The IBM Power Systems Virtual Server platform offers a flexible and secure cloud infrastructure that is highly optimized for running mission-critical workloads. A robust networking framework is essential to ensure seamless connectivity, data protection, and efficient resource management within this cloud environment. The IBM Power Virtual Server network is designed to meet the diverse needs of modern enterprises, providing features that support secure and scalable operations.

The Power Virtual Server network framework integrates smoothly with IBM Cloud, offering both virtual and physical connectivity options to support complex enterprise deployments. This ensures that organizations can connect their on-premises infrastructure to the cloud with minimal disruption while maintaining high security and performance standards. The architecture is built to enable seamless communication between on-premises systems, cloud resources, and third-party services, facilitating hybrid cloud and multi-cloud strategies.

3.3.1 Key aspects of the IBM Power Virtual Server network

The IBM Power Virtual Server network has several characteristics to consider:

- Network Topology. The architecture includes private subnets, virtual private clouds (VPCs), and public and private IP addressing, providing flexibility in network design to match specific business needs.
- Security features. Built-in security measures, such as network isolation and firewalls, ensure that data is protected against unauthorized access and cyberthreats. The network

incorporates best practices for secure data transfer, including VPN and encryption protocols.

- Connectivity options. IBM Power Virtual Server supports various connectivity options, such as VPNs for secure encrypted communication, Direct Link for private dedicated connections, and Transit Gateway for interconnecting different network environments.
- High availability and redundancy. The network infrastructure supports high availability configurations to ensure uptime and reliability. Redundant connections and failover mechanisms are built-in to provide continuous service even during network disruptions.
- Integration with IBM Cloud Services. Power Virtual Server is designed to work in tandem with other IBM Cloud services, allowing enterprises to build and scale solutions that span different environments seamlessly.
- Performance optimization. The network includes features to optimize bandwidth and reduce latency, which is crucial for running high-performance workloads and real-time applications.

The IBM Power Virtual Server network overview provides a comprehensive understanding of the underlying architecture and capabilities that make it suitable for handling diverse and demanding workloads. Whether an organization is seeking to migrate legacy applications to the cloud, build new cloud-native applications, or implement hybrid cloud solutions, Power Virtual Server offers the tools and network capabilities needed for success. In the following sections, we delve into specific network scenarios, from proof-of-concept to production setups, and explore the technologies and configurations that support them.

3.3.2 Detailed description of the IBM Power Virtual Server network structure

The IBM Power Systems Virtual Server network structure is carefully designed to support enterprise workloads, providing a secure, scalable, and high-performance environment within the IBM Cloud. The network structure comprises several key components, including private and public networks, virtual private clouds (VPCs), and connectivity options that allow seamless integration with on-premises infrastructure and other cloud services. This structure enables organizations to deploy and manage their applications in a cloud environment without compromising on security or performance.

The IBM Power Virtual Server network structure has several significant elements.

Virtual Private Cloud (VPC)

The Virtual Private Cloud (VPC) is a foundational element in the IBM Power Virtual Server network structure, providing a logically isolated section of the IBM Cloud where organizations can initiate and manage resources in a virtualized environment. VPCs offer several benefits:

- Network isolation. Each VPC is isolated from other VPCs, ensuring that resources within one VPC are not accessible by others unless explicitly configured. This isolation enhances security, making it ideal for sensitive workloads.
- Customization and control. Organizations have full control over their VPC's network configurations, including IP address ranges, subnets, and routing tables, which allows them to tailor the network to their specific requirements.
- Scalability. VPCs are scalable, enabling businesses to expand their network resources as their workloads grow without disruption.

VPCs in IBM Power Virtual Server provide the flexibility of private networking with the scalability and convenience of a public cloud, making them suitable for hybrid cloud and multi-cloud deployments.

Subnets and IP addressing

Within a VPC, networks are further segmented into subnets, each representing a distinct range of IP addresses. Subnets in IBM Power Virtual Server can be private or public:

- Private subnets. Resources in private subnets do not have direct access to the internet and are typically used for sensitive workloads that need restricted access, such as database servers or application backends.
- Public subnets. Resources in public subnets can access the internet directly and are generally used for workloads that require internet connectivity, such as web servers or API gateways.

This division allows organizations to control traffic flow between public and private resources, enhancing security and optimizing network performance. IP addressing within these subnets can be configured as needed, with support for both IPv4 and IPv6 addresses, providing flexibility for enterprises to design their network layout.

Network security and access control

Security is a core component of the IBM Power Virtual Server network structure. Multiple layers of security features are built into the network to safeguard resources and data:

- Security groups. Security groups act as virtual firewalls for resources within a VPC, controlling inbound and outbound traffic. They are configured with specific rules to allow or deny traffic based on criteria such as IP address, port, and protocol to ensure that only authorized traffic can access the network resources.
- Network access control lists (ACLs). In addition to security groups, network ACLs provide another layer of security by controlling traffic at the subnet level. ACLs can be used to set up stateless filtering rules that allow or deny specific types of traffic, enhancing protection against unwanted network traffic.
- Encryption. IBM Power Virtual Server supports data encryption at various layers, including encrypted VPN connections, to ensure that data remains confidential and secure during transmission. Encryption protocols such as IPsec and SSL/TLS are commonly used to protect sensitive data.

These security features enable organizations to implement a zero-trust model, where access is restricted based on specific rules and continuously monitored, ensuring a secure network environment for their workloads.

VPN and Direct Link Connections

IBM Power Virtual Server offers multiple connectivity options to link on-premises infrastructure with the IBM Cloud:

- VPN. VPNs provide a secure, encrypted tunnel between on-premises networks and the Power Virtual Server environment, enabling private access to cloud resources over the internet. VPNs are ideal for organizations that need secure but cost-effective connectivity for development, testing, and production workloads.
- Direct Link. For organizations requiring dedicated, high-performance connections, IBM offers Direct Link, a private, high-speed connection between the on-premises data center and the IBM Cloud. Direct Link bypasses the public internet, reducing latency and improving data transfer speeds, making it suitable for latency-sensitive applications and large data transfers.

These connectivity options allow organizations to build hybrid cloud environments, extending their existing networks into the cloud while maintaining control over security, performance, and data flow.

Transit Gateway

The IBM Cloud Transit Gateway is a service that facilitates the connection of multiple VPCs within the IBM Cloud, enabling organizations to create complex network topologies across different regions. The IBM Cloud Transit Gateway includes the following key features:

- Centralized routing. Transit Gateway acts as a central hub that simplifies routing between VPCs, reducing the complexity of managing separate connections and enabling efficient data flow across multiple environments.
- Inter-region connectivity. Organizations with resources in multiple regions can use Transit Gateway to ensure secure and low-latency connections between those regions, supporting disaster recovery, load balancing, and multi-region deployments.
- Support for hybrid and multi-cloud architectures. Transit Gateway can integrate with Direct Link, enabling connections between IBM Cloud and on-premises environments, as well as other cloud providers, supporting complex hybrid and multi-cloud strategies.

The Transit Gateway adds flexibility to the IBM Power Virtual Server network structure, allowing businesses to build large-scale architectures that can span regions and integrate multiple cloud resources.

Power Edge Router (PER)

The Power Edge Router (PER) is an IBM Power Virtual Server-specific networking component that is designed to handle routing and connectivity tasks for Power Virtual Server instances. It provides essential functions:

- Efficient traffic routing. PER routes traffic between subnets, VPCs, and external networks, ensuring efficient communication within and outside the Power Virtual Server environment.
- Enhanced security. PER integrates with security features such as firewalls and security groups, adding an extra layer of protection against unauthorized access and cyberthreats.
- Connectivity to external resources. PER can be used to connect Power Virtual Server instances to external resources, such as on-premises infrastructure or other cloud services, facilitating hybrid and multi-cloud setups.

PER is optimized for the unique requirements of Power Virtual Server, enabling smooth and secure communication between different network components.

Load balancing

IBM Power Virtual Server supports load balancing to distribute incoming network traffic across multiple resources, optimizing resource usage, enhancing reliability, and ensuring high availability. Load balancers are configured for various reasons:

- Balance traffic across instances. Distribute traffic across virtual server instances in different subnets, ensuring no single instance is overwhelmed.
- Improve availability. Load balancers can detect instance health and redirect traffic away from unresponsive or overloaded instances, ensuring uninterrupted service.
- Support for different protocols. IBM Power Virtual Server load balancers support multiple protocols, such as HTTP, HTTPS, and TCP, making them suitable for various types of applications and services.

Load balancing is essential for scaling applications in the cloud, ensuring a smooth user experience even during traffic spikes.

Bandwidth and Performance Optimization

IBM Power Virtual Server network structure is designed to offer high-performance networking capabilities that can be optimized based on workload requirements and includes the following key aspects:

- Configurable bandwidth. Organizations can choose bandwidth options that align with their application needs, ranging from standard speeds to high-throughput options for data-intensive applications.
- Low latency. By using options like Direct Link and Transit Gateway, Power Virtual Server provides low-latency connections suitable for real-time applications, such as financial transactions or video streaming.
- Quality of service (QoS). IBM Power Virtual Server enables traffic prioritization to ensure that critical applications receive the necessary bandwidth and performance, minimizing delays and ensuring reliable service.

Performance optimization options in IBM Power Virtual Server make it suitable for workloads with high throughput and low-latency requirements, providing a seamless experience for end-users.

The IBM Power Virtual Server network structure is built with flexibility, security, and performance in mind, making it suitable for a range of enterprise applications. It includes VPCs for isolated environments, subnets for structured IP addressing, robust security features, flexible connectivity options (VPN, Direct Link, and Transit Gateway), Power Edge Router (PER) for optimized routing, load balancing for high availability, and performance optimizations for bandwidth and latency control. This comprehensive network structure empowers organizations to build, deploy, and scale their applications in the IBM Power Virtual Server cloud with confidence, meeting the needs of both hybrid cloud and multi-cloud strategies.

3.4 Integration with IBM Cloud and other cloud services

The IBM Power Systems Virtual Server platform is designed to provide seamless integration with IBM Cloud services and other cloud environments, enabling enterprises to use a hybrid or multi-cloud strategy effectively. This integration capability allows organizations to take advantage of IBM's comprehensive suite of cloud services, while also connecting with other cloud providers to build versatile, scalable, and resilient infrastructure solutions. By integrating Power Virtual Server with IBM Cloud and other platforms, businesses can achieve greater flexibility, optimize costs, and ensure robust performance for various workloads.

3.4.1 Seamless hybrid cloud connectivity

IBM Power Virtual Server integrates natively with IBM Cloud, which can help make it easier for organizations to extend their on-premises infrastructure to the cloud without the need for complex reconfigurations. This integration allows companies to create hybrid cloud environments that combine the benefits of on-premises Power Systems with the scalability and flexibility of IBM Cloud. This integration includes the following key benefits:

- Consistent network policies. IBM Cloud provides a unified framework that enables consistent security and network policies across on-premises and cloud environments.
- Unified management. Power Virtual Server can be managed from the IBM Cloud Console, allowing administrators to monitor, provision, and manage both cloud and on-premises resources from a single interface.

- Secure data flow. Through services like IBM Cloud Direct Link and VPN, secure and encrypted connections can be established between on-premises infrastructure and the IBM Cloud, ensuring that data remains private and protected.

This seamless connectivity makes IBM Power Virtual Server an ideal solution for enterprises that want to gradually migrate workloads to the cloud while maintaining critical applications on-premises.

3.4.2 Direct Link for high-performance, low-latency connections

IBM Power Virtual Server integrates with IBM Cloud Direct Link, a service that provides dedicated, high-speed, low-latency connections between on-premises data centers and the IBM Cloud. Direct Link offers a private and secure pathway for data transfer, bypassing the public internet and reducing the risk of data breaches or latency issues. This service is beneficial for data-intensive applications and latency-sensitive workloads:

- Data-intensive applications. Direct Link supports large-scale data transfers between Power Virtual Server and on-premises systems, making it ideal for applications requiring high throughput, such as big data analytics and database replication.
- Latency-sensitive workloads. By reducing latency, Direct Link ensures that latency-sensitive applications, such as financial transactions and real-time communications, run smoothly and efficiently.

Direct Link enhances the hybrid cloud experience, allowing enterprises to combine the reliability and performance of their on-premises infrastructure with the scalability of IBM Cloud.

3.4.3 Integration with IBM Cloud Services

One of the key advantages of Power Virtual Server is its integration with the wide range of IBM cloud services, which allows organizations to enhance their applications with advanced cloud-native functionalities. The following examples are IBM Cloud services that can be integrated with Power Virtual Server:

- IBM Cloud Object Storage. Power Virtual Server workloads can store and access large datasets in IBM Cloud Object Storage, which provides scalable and cost-effective data storage solutions. This service is useful for applications requiring data archiving, backup, or storage for big data workloads.
- IBM Watson® AI Services. By integrating Power Virtual Server with IBM Watson, organizations can enrich their applications with AI capabilities such as natural language processing, machine learning, and data analysis. This is beneficial for applications in industries such as healthcare, finance, and retail that can use AI for predictive analytics and customer insights.
- IBM Cloud databases. Power Virtual Server can connect to IBM Cloud-managed databases, such as Db2, PosPower Virtual Server SQL, and Mongo DB, to facilitate data storage and management without the need to maintain databases on-premises.
- IBM Security® services. Power Virtual Server workloads benefit from the security tools of IBM Cloud, such as IBM Key Protect for encryption key management, IBM Cloud Security Advisor for monitoring, and IBM Cloud Identity and Access Management (IAM) for user authentication and access control.

Integrating these IBM Cloud services with Power Virtual Server enhances the functionality and security of applications, enabling enterprises to build advanced, cloud-native applications within a secure and managed environment.

3.4.4 Multi-Cloud Connectivity Using IBM Cloud Transit Gateway

IBM Cloud Transit Gateway is a service that facilitates the connection between multiple VPCs (Virtual Private Clouds) within IBM Cloud and enables organizations to extend connectivity to other cloud providers. Through Transit Gateway, Power Virtual Server users can integrate their Power Virtual Server instances with services hosted on other cloud providers, creating a truly multi-cloud architecture. Transit Gateway offers connection and management benefits:

- Centralized management of multi-cloud environments. Transit Gateway acts as a central hub for routing traffic between multiple VPCs, simplifying network management and allowing consistent policies across different cloud platforms.
- Cross-region connectivity. Power Virtual Server resources can use Transit Gateway to communicate with other IBM Cloud services across multiple regions, supporting applications with global reach and disaster recovery setups.
- Hybrid and multi-cloud integration. Organizations that use other cloud providers, such as AWS, Google Cloud, or Microsoft Azure, can use Transit Gateway to create secure, low-latency connections between IBM Power Virtual Server and those providers. This allows businesses to adopt a multi-cloud strategy where they can run different parts of their application stack in various clouds based on performance, cost, or regulatory requirements.

Multi-cloud connectivity enables organizations to choose the best platform for each workload and avoid vendor lock-in, providing greater flexibility and resilience.

3.4.5 Interoperability with Open-Source and Containerized Environments

IBM Power Virtual Server supports open source technologies and containerized workloads, enabling integration with container management platforms like Kubernetes and Red Hat OpenShift. This interoperability is particularly beneficial for organizations adopting cloud-native architectures and DevOps practices and includes the following key features:

- Kubernetes and OpenShift integration. Power Virtual Server can be integrated with IBM Cloud Kubernetes Service or Red Hat OpenShift on IBM Cloud, enabling seamless deployment and management of containerized applications. This allows developers to run workloads in containers, improve portability, and streamline the development lifecycle.
- Hybrid Cloud with OpenShift. IBM Power Virtual Server supports Red Hat OpenShift for creating and managing hybrid cloud environments, allowing organizations to deploy consistent Kubernetes environments across on-premises, Power Virtual Server, and other cloud platforms. This consistency simplifies application deployment, management, and scaling across different environments.
- Support for open source software. Power Virtual Server is compatible with a wide array of open source software, such as databases, development tools, and middleware, making it easier to integrate existing applications with IBM Cloud services.

By supporting containerization and open source standards, Power Virtual Server allows organizations to use modern application architectures, which can increase agility and operational efficiency.

3.4.6 API-driven automation and integration

IBM Power Virtual Server offers a robust set of APIs that enable automated management, provisioning, and integration of cloud resources. Organizations can use these APIs to programmatically interact with Power Virtual Server, making it easier to integrate with other

cloud services and automate cloud operations. API-driven integration includes several benefits:

- Automation of routine tasks. APIs allow teams to automate routine tasks, such as scaling resources, setting up VPN connections, or deploying virtual machines, reducing the need for manual intervention and increasing operational efficiency.
- Integration with third-party tools. Organizations can integrate Power Virtual Server with third-party tools for monitoring, logging, and continuous integration and continuous deployment (CI/CD) pipelines. This is particularly useful for DevOps teams that rely on tools like Jenkins, Terraform, or Ansible for managing cloud infrastructure.
- Customizable workflows. API-driven integration enables businesses to build customized workflows that suit their unique needs, such as automatically scaling applications based on demand, managing network configurations, and monitoring performance.

APIs enhance the flexibility of Power Virtual Server so that businesses can build highly customized and automated cloud environments that can scale and adapt to changing requirements.

3.4.7 Data and application mobility across environments

IBM Power Virtual Server enables data and application mobility between on-premises infrastructure, IBM Cloud, and other cloud environments. This flexibility is essential for enterprises that need to move workloads between environments for reasons such as compliance, cost optimization, or performance. Data and application mobility includes the following key aspects:

- Backup and disaster recovery. Power Virtual Server enables backup and recovery solutions that ensure data can be replicated and recovered across multiple environments, supporting business continuity and disaster recovery strategies.
- Flexible workload placement. Organizations can easily move workloads from on-premises systems to IBM Power Virtual Server or other cloud platforms, optimizing resource allocation based on cost or performance needs.
- Data synchronization and integration. Power Virtual Server supports data synchronization tools that facilitate the movement and integration of data across different environments, ensuring data consistency and availability.

Data and application mobility helps organizations avoid cloud lock-in so that they can maintain flexibility in choosing the most cost-effective and suitable environment for each workload.

The integration capabilities of IBM Power Virtual Server with IBM Cloud and other cloud services provide enterprises with a powerful, versatile platform for building hybrid and multi-cloud environments. From high-performance Direct Link connections and multi-cloud routing with Transit Gateway to container support with Kubernetes and OpenShift, Power Virtual Server offers a wide range of options for connecting, securing, and scaling applications across different environments. API-driven automation and data mobility further enhance Power Virtual Server's adaptability, making it a valuable solution for organizations seeking a flexible and secure cloud platform that integrates seamlessly with other cloud providers and services.

3.4.8 Network topology: Subnets, routing, and IP addressing

The network topology of IBM Power Systems Virtual Server is designed to provide a flexible, secure, and scalable environment for cloud workloads. This topology includes key

components such as subnets, routing, and IP addressing, which collectively enable efficient communication, network segmentation, and access control across the cloud infrastructure. A well-structured network topology is essential for ensuring that resources are properly isolated, accessible, and manageable, and that traffic flows efficiently within and outside the IBM Cloud environment.

Core elements of IBM Power Virtual Server network topology include subnets, routing, IP addressing, NAT, integration with VPC, and security and access control within the topology.

Subnets

Subnets are subdivisions within a network, each representing a distinct range of IP addresses. In IBM Power Virtual Server, subnets are used to segment network resources, allowing organizations to control traffic flow and isolate workloads based on their specific requirements.

Private and public subnets

IBM Power Virtual Server supports both private and public subnets:

- Private subnets. These are used for resources that do not require direct internet access. Private subnets are ideal for sensitive workloads, such as databases, application backends, or internal services, as they are not exposed to the public internet. Traffic to and from private subnets can be controlled through routing tables and security groups.
- Public subnets. Resources in public subnets have direct internet access, making them suitable for applications that need to interact with the internet, such as web servers or API gateways. Public subnets are often used for front-end services or other workloads that require external connectivity.

Network segmentation and security

By dividing networks into subnets, organizations can create distinct zones within their environment, each with its own security policies and access controls. This segmentation helps contain and limit the potential impact of security incidents and allows for a granular approach to traffic control.

Scalability

Subnets in IBM Power Virtual Server are highly scalable, meaning that they can be expanded as needed to accommodate growing workloads. Organizations can add more subnets within a Virtual Private Cloud (VPC) to support different types of applications, environments (development, testing, production), or geographic locations.

Routing

Routing is a critical component of the IBM Power Virtual Server network topology, as it determines how data packets travel between different subnets, VPCs, and external networks. Routing tables in Power Virtual Server define the rules for directing traffic, ensuring that data reaches its intended destination efficiently and securely.

Routing tables

In Power Virtual Server, routing tables are used to control the flow of traffic within a VPC. Each routing table consists of a set of rules that specify the paths data packets should take based on their destination IP addresses. Routing tables are associated with specific subnets, allowing administrators to customize traffic paths for different parts of the network.

Static and dynamic routing

IBM Power Virtual Server supports both static and dynamic routing options:

- Static routing. With static routing, administrators manually configure routes within the routing table. This method offers control over traffic paths but requires manual updates if the network changes. Static routing is suitable for smaller or simpler networks where traffic paths remain relatively consistent.
- Dynamic routing. Dynamic routing uses protocols like BGP (Border Gateway Protocol) to automatically adjust routes based on network conditions. This is useful for larger, more complex networks, where dynamic routing can optimize traffic flow, improve resilience, and reduce the administrative burden.

Route propagation with Transit Gateway

For organizations with multiple VPCs or hybrid cloud environments, IBM Cloud Transit Gateway can facilitate route propagation, allowing VPCs to share routing information and connect seamlessly. This simplifies network management in multi-VPC or multi-region setups by centralizing and automating route management, that can make it easier to build complex network topologies.

IP addressing

IP addressing is a fundamental aspect of network topology, as it assigns unique identifiers to each device or resource within the network. Proper IP addressing in IBM Power Virtual Server ensures that resources can communicate with each other and with external networks without conflicts or accessibility issues.

IPv4 and IPv6 support

IBM Power Virtual Server supports both IPv4 and IPv6 addressing, allowing organizations to choose the IP version that best fits their needs. IPv4 addresses are more commonly used and understood, but IPv6 provides a larger address space, which can be beneficial for organizations that require extensive IP addresses or plan for future scalability.

Private IP addresses

Resources in private subnets are typically assigned private IP addresses, which are only accessible within the VPC or over secure VPN connections. Private IP addresses enhance security by keeping sensitive resources hidden from the public internet, reducing exposure to external threats.

Public IP addresses

Resources in public subnets can be assigned public IP addresses, making them directly accessible from the internet. Public IPs are necessary for services that need to be reachable from outside the network, such as web servers. IBM Power Virtual Server allows administrators to manage and allocate public IP addresses efficiently, ensuring that only authorized resources have internet connectivity.

IP address management (IPAM)

IBM Power Virtual Server offers tools for IP address management (IPAM) to help organizations track their IP allocations, avoid conflicts, and maintain a well-organized network. IPAM helps administrators reserve specific address ranges, manage subnet CIDR (Classless Inter-Domain Routing) blocks, and reassign IPs as workloads evolve.

Network address translation (NAT)

Network address translation (NAT) is an essential component in Power Virtual Server network topology, as it enables resources in private subnets to access the internet without exposing their private IP addresses.

NAT gateway

IBM Power Virtual Server allows the configuration of NAT gateways for outbound traffic from private subnets. With a NAT gateway, resources in a private subnet can initiate internet connections (for software updates, API calls, etc.) without needing a public IP. The NAT gateway translates the private IP address to a public IP, enabling communication with the internet while keeping the resource's private IP address hidden from external access.

Security benefits

NAT provides an additional layer of security by preventing inbound connections from the internet to resources in private subnets. Only outbound traffic initiated by these resources is permitted, helping to protect internal resources from external threats.

Integration with Virtual Private Cloud

IBM Power Virtual Server network topology is built within the context of a Virtual Private Cloud (VPC), which is a logically isolated section of the IBM Cloud where organizations can deploy and manage resources securely. The VPC structure provides a private network environment with control over IP ranges, routing tables, subnets, and security settings.

VPC peering and Transit Gateway

IBM Power Virtual Server allows VPC peering, enabling direct network connectivity between VPCs within the same IBM Cloud region. For cross-region or multi-cloud connectivity, IBM Cloud Transit Gateway serves as a central hub to route traffic between multiple VPCs, making it easier to build and manage complex network topologies.

Flexible network design

Organizations can use VPCs to design flexible and scalable network architectures tailored to their applications and workloads. Power Virtual Server users can create multiple VPCs to isolate different environments, such as development, testing, and production or different departments, each with its own network policies and controls.

Security and access control within the topology

Network topology in IBM Power Virtual Server is further strengthened by security measures and access controls designed to protect resources and manage traffic within and outside the VPC.

Security groups

Security groups act as virtual firewalls at the instance level, allowing administrators to specify inbound and outbound rules for each instance. Security groups enable fine-grained control over traffic, ensuring that only authorized traffic can reach the resources.

Access control lists

Access control lists (ACLs) operate at the subnet level, providing an additional layer of security by filtering traffic entering and exiting the subnet. ACLs are useful for controlling traffic between different segments of the network, such as restricting access from one subnet to another.

Firewall rules

Administrators can use Power Virtual Server to define firewall rules to block or allow specific types of traffic. Firewall configurations can be applied to both public and private subnets to protect against unauthorized access and reduce the risk of attacks.

The network topology in IBM Power Virtual Server, comprising subnets, routing, and IP addressing, provides a flexible and secure foundation for deploying cloud resources. Subnets allow for effective network segmentation, and routing tables and protocols manage traffic flow across different network segments. IP addressing ensures unique and organized resource identification, and NAT enhances security for private subnets. Together, these elements enable organizations to create isolated, highly controlled, and scalable network environments within IBM Cloud. The result is a robust network infrastructure that supports a wide variety of applications and workloads, ranging from secure, private deployments to public-facing applications.

3.5 Security measures incorporated in Power Virtual Server networks

IBM Power Systems Virtual Server networks are built with multiple layers of security measures designed to protect sensitive data, ensure compliance with industry regulations, and provide organizations with the tools they need to safeguard their cloud environments. Given the complex and dynamic nature of cloud computing, these security features cover a wide range of areas, including network isolation, access control, data encryption, and monitoring. These measures enable organizations to confidently run mission-critical workloads on Power Virtual Server without compromising on security or control.

Network isolation with Virtual Private Cloud

IBM Power Virtual Server uses Virtual Private Cloud (VPC) environments to provide network isolation for each tenant, ensuring that resources in one VPC are securely separated from those in another. VPCs are logically isolated environments within IBM Cloud, enabling organizations to manage and control their network resources independently:

- Dedicated IP ranges. Each VPC is assigned a unique range of IP addresses, preventing IP conflicts and ensuring that network resources are isolated.
- Segmentation with subnets. Within a VPC, organizations can further segment their network into private and public subnets, isolating sensitive resources from those exposed to the internet. This segmentation minimizes the risk of unauthorized access to critical workloads.

Network isolation in Power Virtual Server offers robust protection against threats by preventing cross-tenant access and enforcing secure boundaries between different environments within IBM Cloud.

Security groups

Security groups in Power Virtual Server act as virtual firewalls at the instance level, allowing organizations to control inbound and outbound traffic to each virtual machine or resource within the VPC.

Granular access control

Administrators can set rules in security groups that define which types of traffic are allowed or denied based on IP address, port, and protocol. This enables fine-grained control over access, limiting exposure to only authorized users and applications.

Customizable rules

Security groups can be customized based on the needs of each application or workload. For example, an application server might allow HTTP and HTTPS traffic from the internet, while a database server might permit only internal traffic from specific subnets or IP addresses.

Stateful filtering

Security groups in Power Virtual Server are stateful, meaning that if a request is allowed outbound, the corresponding response is automatically allowed inbound. This simplifies configuration and enhances security by tracking connections.

Security groups help limit the network attack surface by ensuring that only specified traffic can reach Power Virtual Server resources, reducing the risk of unauthorized access.

Network access control lists

Network ACLs in Power Virtual Server provide an additional layer of security at the subnet level, complementing security groups by filtering traffic that is entering and exiting specific subnets.

Stateless filtering

Unlike security groups, ACLs are stateless, meaning that both inbound and outbound rules must be explicitly configured. This allows administrators to have more precise control over traffic for specific subnet zones.

Layered security

ACLs work alongside security groups, creating a layered security model that enhances overall protection. For example, ACLs can be used to allow or deny traffic at the subnet level, while security groups further control access to individual resources within the subnet.

Custom rule management

ACLs can be configured with rules that permit or deny traffic based on criteria such as IP range, port, and protocol, allowing for highly specific control of network traffic.

By providing traffic control at the subnet level, ACLs help protect resources from malicious traffic, restrict access, and add redundancy to security configurations.

Data encryption

Data encryption in Power Virtual Server ensures that sensitive data remains protected both in transit and at rest, preventing unauthorized access to information even if it is intercepted.

Encryption in transit

IBM Power Virtual Server supports encrypted VPN connections by using protocols like IPsec and SSL/TLS to secure data traveling between on-premises environments and the Power Virtual Server cloud. This ensures that data remains private and protected from interception as it moves through public networks.

Encryption at rest

Power Virtual Server can use encryption mechanisms such as AES-256 to secure stored data, preventing unauthorized access to data stored on virtual servers and cloud storage. This is particularly important for protecting sensitive information, like customer data or financial records, that might be required to comply with data protection regulations.

Key management with IBM Key Protect

Power Virtual Server integrates with IBM Key Protect, a key management service that allows organizations to manage encryption keys securely. Key Protect offers options for automatic or user-controlled key management, allowing organizations to control access to their encrypted data more effectively.

Encryption in Power Virtual Server provides comprehensive data protection, ensuring that sensitive information is secure at all stages-whether in transit across networks or stored on the cloud.

Network address translation

Network address translation (NAT) enhances security by allowing resources in private subnets to access the internet securely without exposing their private IP addresses.

Outbound Internet Access

Resources in private subnets can use NAT gateways to initiate connections to the internet, such as for software updates or API access, without needing public IP addresses, which keeps their private IPs hidden from external users.

Inbound traffic blocking

NAT provides a layer of protection by blocking all inbound traffic from the internet, ensuring that resources in private subnets cannot be directly accessed from outside the network.

NAT enhances security for private resources in Power Virtual Server by limiting their exposure to the internet and preventing unauthorized inbound access.

Firewalls

Power Virtual Server supports firewalls to provide an additional layer of protection for applications and resources. Firewalls help regulate and monitor traffic, allowing or blocking data packets based on predefined security rules.

Customizable firewall rules

Firewalls in Power Virtual Server allow organizations to define rules for inbound and outbound traffic at various levels, such as IP address, port, and protocol, based on application requirements.

Preventing unauthorized access

By setting strict rules, firewalls prevent unauthorized access to applications and resources, providing a defense against common attacks such as DDoS (Distributed Denial of Service) and SQL injection.

Enhanced monitoring

Firewalls offer logging and monitoring features that provide visibility into network traffic and potential security events, enabling administrators to quickly respond to suspicious activity.

Firewalls add another security layer by controlling data flow at the network perimeter, helping protect applications and data from external threats.

Identity and Access Management

IBM Power Virtual Server uses IBM Cloud Identity and Access Management (IAM) to enforce access control and manage permissions, ensuring that only authorized users can access specific resources.

Role-based access control (RBAC)

IAM allows organizations to assign roles to users, granting only the necessary permissions for their job functions. This principle of least privilege reduces the risk of unauthorized access and accidental data exposure.

Granular permissions

IAM provides fine-grained control over access to specific resources, enabling organizations to restrict access to sensitive resources like VMs, networks, or storage based on user roles.

Multi-factor authentication

For added security, IAM can require multi-factor authentication (MFA) for accessing sensitive resources, ensuring that only verified users gain access.

IAM enables comprehensive access control for Power Virtual Server, allowing administrators to manage and secure user permissions effectively.

Monitoring and logging

Power Virtual Server includes monitoring and logging tools that enable continuous tracking of network activity, providing insights into potential security incidents and performance issues.

Activity logs

Power Virtual Server logs network and access activity, allowing administrators to track actions within the network, including configuration changes, login attempts, and traffic patterns. Logs can be used to audit activity, identify unusual behavior, and ensure compliance with regulatory standards.

Real-time monitoring

IBM Cloud offers monitoring tools that provide real-time insights into network performance and security events, helping administrators detect and respond to security threats quickly.

Alerts and notifications

Power Virtual Server monitoring tools allow administrators to set up alerts for suspicious activity, such as unexpected traffic spikes or failed login attempts. These notifications enable rapid response to potential security breaches.

By providing monitoring and logging capabilities, Power Virtual Server enhances visibility into the cloud environment, allowing organizations to maintain security oversight and respond promptly to incidents.

Compliance and regulatory security

IBM Power Virtual Server is designed to help organizations meet regulatory and compliance requirements across various industries, such as healthcare, finance, and government.

Compliance certifications

IBM Cloud, including Power Virtual Server, holds certifications for various industry standards, such as HIPAA, GDPR, and SOC 2. This ensures that Power Virtual Server meets stringent regulatory requirements for data protection, security, and privacy.

Auditing and reporting tools

Power Virtual Server provides auditing tools to track compliance with industry regulations, offering detailed reports and logs for regulatory inspections.

Data sovereignty controls

IBM Power Virtual Server enables organizations to specify data storage locations, ensuring that data is stored and processed in compliance with local data sovereignty laws.

Power Virtual Server's compliance capabilities can help make it easier for organizations to operate securely within regulated environments and meet specific legal obligations for data protection.

IBM Power Virtual Server incorporates a comprehensive set of security measures designed to protect cloud environments, data, and applications. These include network isolation with VPCs, security groups, ACLs, encryption, NAT, firewalls, IAM, and monitoring tools that provide layered security across the network. By combining these features, Power Virtual Server helps to ensure that organizations have the tools to maintain a secure, compliant, and resilient network environment. This multi-layered security approach enables businesses to focus on their core operations while trusting that their data and resources are safeguarded against unauthorized access and cyberthreats.

3.6 Comparing IBM Power Virtual Server network features to similar cloud offerings

IBM Power Systems Virtual Server provides a unique cloud environment designed for enterprise-grade workloads, particularly those relying on IBM's Power architecture, such as AIX, IBM i, and Linux applications. Power Virtual Server stands out due to its specialized support for legacy systems and mission-critical workloads, with networking features optimized to meet high performance, security, and integration requirements. IBM Power Virtual Server has several distinguishing features and capabilities when compared with other similar cloud offerings such as AWS, Microsoft Azure, and Google Cloud Platform (GCP).

IBM Power Virtual Server has key network features that can be compared with these other cloud providers, with a focus on VPC design, connectivity options, network security, and hybrid and multi-cloud support.

VPC structure and customization

IBM Power Virtual Server

IBM Power Virtual Server offers VPCs that are logically isolated environments within the IBM Cloud, which is designed to provide flexible control over network configurations, including IP address ranges, subnets, and routing. Power Virtual Server allows organizations to segment their workloads effectively with private and public subnets, which is especially valuable for isolating mission-critical applications from public internet access.

AWS

AWS VPCs offer similarly robust controls, allowing users to create isolated networks, assign CIDR blocks, and configure subnets and routing tables. AWS VPCs are highly customizable and integrate seamlessly with a wide range of AWS services but are generally optimized for x86 workloads.

Azure

Azure VNet provides comparable features to IBM Power Virtual Server VPCs, including network segmentation, IP address management, and routing. Azure VNets are also highly integrated with Azure services but are predominantly geared toward Windows and Linux workloads.

Google Cloud Platform

Google Cloud Platform (GCP) VPCs offer flexible, global networks that allow for multi-region deployments within a single VPC. GCP also supports custom routing and IP ranges similar to Power Virtual Server, but is generally optimized for Google's cloud-native services instead of Power workloads.

IBM Power Virtual Server offers VPC capabilities that are tailored for IBM Power workloads, providing essential controls for legacy and high-compute applications. Although AWS, Azure, and GCP also offer customizable VPCs, Power Virtual Server stands out by natively supporting AIX and IBM i workloads, which is not available in other cloud providers.

Connectivity options

Each provider offers a connectivity option to compare with the options of Power Virtual Servers, which include Direct Link, VPN and hybrid cloud support.

IBM Power Virtual Server

Power Virtual Server offers IBM Cloud Direct Link, which provides high-speed, dedicated connectivity between on-premises data centers and IBM Cloud, bypassing the public internet for secure, low-latency connections. Power Virtual Server also supports VPN connections for secure data transmission, making it well suited for hybrid cloud deployments. Also, Power Virtual Server integrates with IBM Cloud Transit Gateway for connecting multiple VPCs or regions.

AWS

AWS provides Direct Connect, a similar service to IBM Direct Link, which enables private connectivity to AWS cloud resources. AWS also offers VPN solutions for secure connections and AWS Transit Gateway for inter-VPC routing. AWS hybrid cloud solutions integrate well with on-premises systems through AWS Outposts, but it lacks native support for IBM Power workloads.

Azure

Azure ExpressRoute offers private, dedicated connections similar to IBM Direct Link and AWS Direct Connect. Azure also supports VPN Gateway for encrypted connections and Virtual WAN for large-scale network architectures across multiple regions. Azure Arc provides hybrid and multi-cloud capabilities, but its integration with IBM Power environments is limited.

GCP

GCP offers Dedicated Interconnect for private connections and Cloud VPN for secure network links. GCP's Network Connectivity Center functions similarly to IBM's Transit Gateway, helping connect VPCs across regions. While GCP supports multi-cloud and hybrid options, it primarily targets x86 and containerized workloads rather than Power workloads.

IBM Power Virtual Server provides comparable connectivity options with a strong emphasis on integrating IBM Power workloads into hybrid cloud environments, particularly through IBM Direct Link and VPN solutions. Although AWS, Azure, and GCP offer robust private connectivity options, Power Virtual Server is uniquely suited to companies with existing IBM Power infrastructure.

Network security features

Each provider includes security features and options to prevent unauthorized access to support industry standards, and to support regulatory compliance.

IBM Power Virtual Server

Power Virtual Server incorporates comprehensive security features, including security groups, network ACLs, encryption for data at rest and in transit, and IAM for role-based access control. Additionally, IBM Power Virtual Server supports compliance with industry standards, making it suitable for highly regulated industries. Integration with IBM Cloud Security Advisor provides centralized monitoring and security insights, enhancing security management.

AWS

AWS offers a broad range of security features, including security groups, network ACLs, and extensive encryption options. AWS Identity and Access Management (IAM) provides fine-grained access controls. Services like AWS Shield and GuardDuty offer advanced security and threat detection. AWS is well established in meeting regulatory requirements but lacks native IBM Power support.

Azure

Azure provides security features like Network Security Groups (NSGs), encryption, and Azure Active Directory for access control. Azure Security Center offers monitoring and threat detection similar to IBM Cloud Security Advisor. Azure is highly compliant with regulatory standards and is widely used in regulated sectors, but it is not optimized for IBM Power workloads.

GCP

GCP's security offerings include firewall rules, encryption, and Identity and Access Management (IAM). Google Cloud Armor provides protection against DDoS attacks, and Cloud Security Command Center offers monitoring and security insights. GCP focuses on x86 workloads and cloud-native services with no specific optimizations for IBM Power.

IBM Power Virtual Server's security features are comparable to those of AWS, Azure, and GCP, with the added advantage of IBM's experience in supporting mission-critical IBM Power workloads. Power Virtual Server's security and compliance capabilities make it especially suitable for industries with stringent regulatory requirements, while AWS, Azure, and GCP provide more generalized security options for x86 and cloud-native workloads.

Hybrid and multi-cloud integration

Each provider offers methods to connect to hybrid cloud and multi-cloud environments.

IBM Power Virtual Server

IBM Power Virtual Server supports hybrid cloud deployments by integrating seamlessly with IBM Cloud and on-premises IBM Power systems. IBM Cloud Satellite extends IBM Cloud services to any location, which helps to make it easier to deploy and manage Power Virtual Server in a multi-cloud environment. The partnership of IBM with cloud providers like AWS enables multi-cloud strategies, though Power Virtual Server itself is hosted within IBM Cloud.

AWS

AWS has extensive multi-cloud capabilities through its support for third-party tools and integrations. AWS Outposts provides on-premises cloud solutions similar to IBM Cloud Satellite, and AWS's partnerships with other cloud providers allow for inter operate multi-cloud architectures. However, AWS lacks native support for IBM Power workloads.

Azure

Azure Arc enables hybrid and multi-cloud management across on-premises, Azure, and third-party cloud providers, allowing for consistent management across different

environments. Azure Stack brings Azure services to on-premises environments. While Azure provides robust multi-cloud capabilities, it does not natively support IBM Power workloads.

GCP

Google Anthos supports multi-cloud management across GCP, on-premises, and third-party cloud environments, enabling organizations to deploy and manage containerized applications in a consistent manner. However, GCP's focus is primarily on Kubernetes and containerized workloads, with limited options for IBM Power systems.

IBM Power Virtual Server is uniquely positioned for IBM Power hybrid and multi-cloud scenario because of its native support for IBM Power workloads and integration with IBM Cloud Satellite. Although AWS, Azure, and GCP provide strong multi-cloud capabilities, Power Virtual Server is the only option optimized for IBM Power applications.

Performance and scalability

Each provider has features to increase performance and to expand existing configurations.

IBM Power Virtual Server

IBM Power Virtual Server is optimized for the high-performance, compute-intensive workloads that are typical of IBM Power systems with specific optimizations for AIX, IBM i, and Linux. Power Virtual Server offers scalable options within IBM Cloud and with Direct Link and high-throughput networking capabilities to support demanding enterprise applications.

AWS

AWS offers scalable networking solutions and performance features across multiple instance types optimized for high-performance computing. However, AWS's compute options are generally optimized for x86 architectures, which limits direct compatibility with IBM Power workloads.

Azure

Azure provides scalable solutions across a range of instance types for Windows and Linux workloads. Azure's high-performance networking and ExpressRoute integration can support demanding applications but lack the specific optimizations for IBM Power environments.

GCP

GCP is well suited for scalable, high-performance cloud-native applications, with support for global load balancing and dedicated instances for compute-intensive workloads. However, GCP's offerings focus primarily on x86 and containerized workloads, not IBM Power.

IBM Power Virtual Server is the preferred choice for organizations seeking high performance for IBM Power-specific workloads. Although AWS, Azure, and GCP offer scalable and high-performance solutions, Power Virtual Server's optimizations for AIX, IBM i, and Linux on Power distinguish it as the best option for companies heavily invested in IBM Power architecture.

IBM Power Virtual Server provides a robust cloud platform with networking features tailored to IBM Power workloads, delivering secure, high-performance, and flexible options for hybrid and multi-cloud deployments. Compared to AWS, Azure, and GCP, Power Virtual Server excels in its specialized support for IBM Power systems, making it the top choice for organizations running mission-critical applications on AIX, IBM i, or Linux on Power. Although other providers offer competitive features, Power Virtual Server remains unmatched in its ability to integrate IBM Power workloads seamlessly into a modern cloud environment, ensuring enterprise-grade security, scalability, and compliance tailored to the needs of IBM Power customers.

3.7 IBM Power Virtual Server Network Scenarios

IBM Power Systems Virtual Server provides a versatile and robust networking framework designed to support a wide range of deployment scenarios, from development and testing to full-scale production. To meet the unique needs of various business environments, Power Virtual Server offers flexible network configurations that allow organizations to create secure, isolated, and high-performance network setups. These scenarios are essential in helping organizations achieve their operational goals, whether they are validating new applications, running mission-critical workloads, or scaling environments to support increasing demands.

Power Virtual Server network scenarios can be broadly categorized into non-production, proof-of-concept environments, and production deployments, each with specific requirements and considerations. In non-production configurations, the focus is often on rapid deployment and cost-effective solutions to validate applications. In contrast, production scenarios demand high availability, security, and optimized performance for stable, uninterrupted operations. Also, specialized components such as Power Edge Router (PER), Transit Gateway, and VPN on VPC provide advanced capabilities to meet the needs of enterprise environments, including local and global routing, hybrid connectivity, and options for seamless migration from legacy systems.

This section describes different Power Virtual Server network scenarios in detail, examining configurations for proof-of-concept (PoC) environments, production scenarios, and specific elements like PER, Transit Gateway, and enhanced connectivity options. By understanding these scenarios, organizations can make informed decisions about how best to use Power Virtual Server's network features to optimize performance, maintain security, and scale with confidence in a cloud-first world.

3.7.1 Non-production proof-of-concept scenarios

Proof-of-concept (PoC) scenarios in IBM Power Systems Virtual Server are designed to allow organizations to explore, validate, and test their applications and infrastructure configurations in a non-production environment before committing to full-scale deployments. PoC setups provide a valuable sandbox for evaluating new technologies, assessing performance, and ensuring compatibility within IBM Power Virtual Server. These non-production environments allow teams to experiment and make adjustments without risking disruption to production workloads, making them ideal for innovation and strategic planning.

Key aspects of running PoC in non-production environments

There are several key aspects to consider when you run PoC in non-production environments. Factors to consider include cost effectiveness, isolation of the PoC, resource flexibility, speed of provisioning, and scalability.

Cost-effective evaluation

Non-production PoC environments enable companies to evaluate IBM Power Virtual Server capabilities at a lower cost than production environments, as they are often smaller-scale, fewer resources and simplified network setups.

Isolation and flexibility

PoC setups in Power Virtual Server are isolated from production environments, which allows teams to test changes or upgrades without affecting business operations. This isolation also enhances security, as sensitive data is typically excluded, reducing the risk of exposure in experimental settings.

Resource flexibility

PoC environments allow teams to configure resources-such as virtual machines, storage, and network settings-to mimic production configurations. This flexibility enables teams to test different performance settings, troubleshoot issues, and fine-tune their infrastructure before transitioning to a live environment.

Rapid provisioning and scalability

Power Virtual Server PoC environments can be provisioned quickly, enabling rapid testing of new applications or architectures. These environments are also scalable, allowing teams to increase or decrease resources as needed to assess application performance under different loads.

Benefits of PoC setups in IBM Power Virtual Server

Among the benefits of using Power Virtual Server in PoC configurations are safe spaces for testing, for application evaluation, for scaling and cost estimations, and for testing for compatibility.

Low-risk testing

PoC environments in Power Virtual Server offer a safe space to test new applications, configurations, or infrastructure changes without the risk of affecting production workloads. This encourages innovation and experimentation.

Performance evaluation

Organizations can test the performance of their applications on IBM Power architecture and evaluate potential bottlenecks or latency issues. This is important for workloads with unique performance requirements, such as those running on AIX or IBM i.

Predictive scaling and cost planning

By using PoC environments to simulate expected usage patterns, organizations can predict the resources required for production deployments and optimize their budget.

Compatibility testing

PoCs enable companies to assess compatibility with other applications, cloud services, and security tools within the IBM Power Virtual Server ecosystem, ensuring smooth integration with existing infrastructure.

Limitations of PoC setups in IBM Power Virtual Server

PoC configurations can also have limits and restrictions to consider.

Limited resources

PoC setups often have limited resources compared to full production environments, which might not reflect the true performance of an application under full-scale operations.

Temporary setup

As PoCs are generally temporary, data persistence and long-term configuration management are usually not prioritized, making it necessary to reconfigure aspects when moving to production.

Simplified network configurations

Network security and configurations in PoC environments might be simplified to reduce complexity and costs, which can create discrepancies when transitioning to a production environment with stricter security policies.

Use case scenario

A financial services company wants to migrate a risk analysis application running on IBM i to IBM Power Virtual Server to take advantage of cloud scalability and high availability. The organization is uncertain about the application's performance on Power Virtual Server and wants to test the migration process and cloud performance under simulated conditions.

Setting up the PoC

The company creates a PoC environment on Power Virtual Server, configuring it to mirror their on-premises setup for IBM i, with comparable virtual machines, storage, and network configurations.

Testing and validation

The development team deploys the risk analysis application to the PoC environment, running tests to assess performance, latency, and security within Power Virtual Server. They also simulate traffic spikes to evaluate how the application scales under different loads.

Troubleshooting and optimization

During testing, the team identifies minor latency issues. To optimize performance, they adjust network configurations and resource allocations, such as increasing CPU and memory. They can use the PoC environment to make adjustments quickly without impacting production.

Migration process simulation

The company practices the migration process in the PoC setup, ensuring all data transfer, application configuration, and network integration steps are clear and efficient. They document the process and identify improvements to streamline the actual migration.

Results and analysis

After completing the PoC, the team concludes that the application performs optimally on IBM Power Virtual Server with the adjusted settings. They calculate expected resource needs and refine their budget and timeline for a smooth transition to a production environment.

This PoC deployment enabled the financial services company to confidently migrate its application to IBM Power Virtual Server, having verified performance, optimized configurations, and tested the migration process-all while maintaining production uptime.

3.7.2 Production Scenarios

In production scenarios, IBM Power Systems Virtual Server provides the infrastructure needed to support large-scale, mission-critical applications with high levels of performance, reliability, and security. When deploying Power Virtual Server for production, organizations must consider scalability, network configuration, security, and compliance to ensure uninterrupted service and optimal performance. This section explores key considerations, best practices, and in-depth features to help organizations deploy and manage Power Virtual Server in production environments.

Resource considerations for full-scale production

To scale production environments effectively in IBM Power Virtual Server, you must consider availability of resources, scalability that uses both local and remote resources, and how to balance loads:

- Resource planning

  Assessing the anticipated workload and provisioning adequate CPU, memory, and storage resources to handle peak traffic and demand.

- Elastic scalability

  Power Virtual Server enables the dynamic scaling of resources to adapt to changing demands. This flexibility is crucial for businesses with seasonal or unpredictable traffic patterns, ensuring that resources scale up during high-demand periods and scale down when demand decreases.

- Load balancing

  Load balancing helps distribute incoming traffic across multiple instances to prevent any single server from becoming overused. This strategy improves both performance and resilience by redirecting traffic in case of failures.

- Multi-region scaling

  For organizations with global operations, deploying instances across multiple regions can improve response times and provide a fallback in case of regional outages.

Network configuration best practices for reliability and uptime

To ensure high reliability and uptime in Power Virtual Server production environments, consider the following network configuration best practices:

- Redundant connections

  Configuring redundant connections between instances and data centers, such as multiple VPN connections or using IBM Cloud Direct Link, can help ensure continuity even if one connection fails.

- Use of high-availability zones

  Deploying resources across multiple availability zones minimizes the impact of potential failures in a single zone, which can enhance resilience.

- Regular testing of failover mechanisms

  Regularly testing failover configurations ensures that they function correctly during a real failure. These tests help verify that the routing can fail over to backup connections smoothly without interrupting service.

- Performance optimization

  Configuring network parameters to optimize latency and bandwidth based on application needs can enhance user experience, especially for latency-sensitive applications.

Security and compliance for production deployments

production deployments in power virtual server require robust security measures to protect sensitive data and maintain compliance with regulatory standards:

- Encryption

  Encrypting data both in transit, by using a VPN for example, and at rest is essential for ensuring data security.

- Strict access control

  Implementing role-based access control (RBAC) through IBM Cloud Identity and Access Management (IAM) restricts access to sensitive resources, ensuring that only authorized users can make changes.

- Network security policies

  Configuring security groups and network ACLs helps protect production environments by controlling traffic flow and preventing unauthorized access.

- Compliance with industry standards

  Power Virtual Server meets various industry compliance standards, such as HIPAA and GDPR, making it suitable for highly regulated industries. Organizations must ensure that they adhere to relevant regulations and document security policies to maintain compliance.

Example scenario for high-availability setup

For a financial services application that requires 24x7 uptime, Power Virtual Server can be configured with multiple instances spread across availability zones. Load balancers distribute traffic, and automated failover mechanisms reroute traffic during an instance failure, ensuring continuous operation.

Example scenario for multi-region architecture

An e-commerce business with a global user base deploys Power Virtual Server instances in North America, Europe, and Asia-Pacific regions. Each region has its own redundant network setup, including VPN and Direct Link connections. This configuration reduces latency for regional users and maintains service even if one region experiences an outage.

Power Edge Router

The Power Edge Router (PER) serves as the gateway for routing and managing network traffic within IBM Power Virtual Server. PER enhances connectivity by facilitating secure and efficient data transmission between subnets, VPCs, and external networks, making it an essential component for production environments. PER includes the following advantages:

- Optimized traffic routing. PER routes traffic efficiently across subnets and VPCs, supporting high-performance networking.
- Enhanced security. PER integrates with security policies, adding a layer of protection for sensitive data.

Configuring PER involves defining routing rules for different subnets and connecting it to external networks like on-premises systems or third-party cloud services. Routing tables can be customized to manage the flow of data and prioritize traffic based on business needs.

IBM Cloud Transit Gateway

The local and global routing capabilities of IBM Cloud Transit Gateway enables centralized routing across VPCs within IBM Cloud. This service simplifies network management by acting as a central hub for routing, allowing organizations to connect multiple VPCs efficiently:

- Local routing. Within the same region, Transit Gateway connects VPCs seamlessly, supporting high-speed, low-latency connections.
- Global routing. Organizations can use the global routing capabilities of Transit Gateway to connect resources across regions to create a unified, interconnected network architecture.
- Integration benefits with other cloud services. Transit Gateway integrates with IBM Direct Link and VPN, enabling hybrid connectivity to on-premises environments and multi-cloud setups. This integration allows businesses to build complex, flexible network topologies that support applications spanning multiple clouds.

VPN on VPC as a replacement to Power-VPNaaS (withdrawal)

IBM is phasing out Power-VPNaaS in favor of VPN on VPC, which provides a more modern, scalable, and flexible solution for secure network connectivity. This transition aligns with IBM's focus on enhancing VPC capabilities and simplifying network management.

The migration to VPN on VPC provides enhanced security, better performance, and easier integration with other IBM Cloud services, making it ideal for production workloads that require robust security.

Review the migration steps and best practices:

1. Plan the transition. Assess existing Power-VPNaaS configurations and prepare a migration plan to minimize downtime.
2. Set up VPN on VPC. Configure VPN on VPC with new routing and security settings and mirror previous configurations if possible.
3. Test and validate. Test the new VPN setup to ensure stable connectivity and troubleshoot any issues before decommissioning Power-VPNaaS.

Cloud Connection (Direct Link + Transit Gateway) withdrawal

IBM is gradually retiring the legacy Cloud Connection (Direct Link + Transit Gateway) in favor of more advanced connectivity options like enhanced Direct Link and VPC Transit Gateway. The updated options offer better performance, security, and compatibility with modern network requirements.

There are alternative solutions;

- Direct Link with Transit Gateway. For organizations that need private connections, Direct Link provides high-speed, dedicated access. Integrating with Transit Gateway further extends connectivity to multi-region and multi-cloud setups.
- Dedicated VPN Solutions. VPNs on VPC offer secure, scalable options for organizations that do not require dedicated private lines but still want secure access to IBM Power Virtual Server.

More bandwidth options

IBM Power Virtual Server offers a range of bandwidth options to meet the needs of different workloads, from standard configurations for typical applications to high-throughput options for data-intensive workloads:

- Standard bandwidth. Suitable for applications with moderate network requirements, such as web applications and internal tools.
- High-throughput bandwidth. Ideal for applications needing large data transfers, such as analytics, data warehousing, and real-time processing.

Use cases that benefit from higher bandwidths

Several situations can benefit by using configurations that provide higher bandwidths:

- Big data and analytics. Applications that process large datasets require high bandwidth to transfer data quickly and efficiently, enabling real-time insights.
- Media and content delivery. Streaming applications and content delivery networks benefit from higher bandwidth to deliver smooth, high-quality experiences to users.
- Disaster recovery and data backup. Higher bandwidths expedite data replication between production and backup environments, reducing recovery times after a failure.

Strategies for choosing the right bandwidth configuration

jConsider the following factors to determine your best bandwidth configuration:

- Assess application needs. Identify the bandwidth requirements of each application to ensure that they align with available options.

- Consider Redundancy. For critical applications, consider dual bandwidth configurations to ensure service continuity during a failure.
- Optimize Costs. Balance performance with budget by choosing bandwidth configurations that match peak requirements without excessive over-provisioning.

Summary

Production scenarios in IBM Power Virtual Server benefit from thoughtful network architecture, robust security, and high availability. By taking advantage of features like Power Edge Router, Transit Gateway, VPN on VPC, and flexible bandwidth options, organizations can build scalable, secure, and resilient cloud environments that are tailored to meet the demands of their mission-critical applications.

4

Backup for IBM i on IBM Power Virtual Server

An effective backup process is a vital part of any data center operation. It ensures business continuity and allows specific data recovery for historical and auditing purposes.

IBM Power Systems Virtual Server provides different tools to fit operational needs and includes some constraints that are related to its cloud nature.

This chapter describes different backup options that can be used for normal operations, makes suggestions for backup strategies, and discusses the differences between methods.

This chapter includes the following topics:

- “Backup and restore for IBM Power Virtual Server instances”
- “IBM BRMS”
- “Creating object-level backups”
- “FalconStor overview”
- “Power Virtual Instance”

© Copyright IBM Corp. 2025.

4.1 Backup and restore for IBM Power Virtual Server instances

Backup and restore operations on Power Virtual Server instances are different from the regular on-premises environment. Most IBM i operations on-premises rely on tape backups due to their capacity, transfer rate, and reliability, but physical tape is not an option on Power Virtual Server instances. In the IBM Power Systems Virtual Server environment, backups can be transported only over the network to either IBM Cloud Object Storage (COS) or to a Virtual Tape Library. IBM i natively offers two solutions:

- Native commands
- IBM Backup, Recovery, and Media Services (IBM BRMS) for complete recovery

Third-party software solutions for IBM i backup are also available, but not all of them can back up IBM i by using network facilities.

Backing up over a network connection is typically not as fast as backing up to physical tape, so consider that backups might take longer in the cloud compared to your on-premises environment. Fast connections increase the transfer rate, but backup windows must be sized and considered in this environment.

IBM COS can support storing large amounts of data from multiple backups. This parallel storage system provides concurrent access from anywhere, reducing bottlenecks and offering almost endless scalability. IBM COS stores data in primary or auxiliary storage layers, allowing different service levels for backups and archives and reducing costs when needed.

IBM Power Systems Virtual Server provides snapshots and clones, which are used to take a complete or partial copy of disks to improve recovery. Disks can also be copied by using a gold image and stored on IBM COS for DR activities, which can simplify and improve the full system backup procedure.

When sizing COS containers for your backup environment, the following issues must be considered:

- Size and frequency of ingress and egress requests
- Amount of data that is transferred
- Type of communications between the server instance and IBM COS
- Needed redundancy

Server-side data compression can use IBM Power Systems Virtual Server hardware compression and uncapped capacity when available. Using server-side data compression can reduce the amount of data that is transferred and stored on IBM COS. This reduction in data also reduces transfer times and storage costs.

The use of incremental and differential backups, which is supported by IBM i and other backup solutions can also reduce the amount of data that must be saved and stored. This is one of the most used mechanisms when data is backed up to the cloud. By using differential backups, the backup window and size of backup media can be reduced dramatically. However, the time that it takes to restore increases because the process starts with the full backup and continues in partial increments. It is recommended that differential backups are used with full backups once a week or once a month and are combined with full disk clones or Disaster Recovery (DR) software.

Today, some open source tools can be used to create backup scripts. Compression and encryption capabilities can be used that are included with IBM i Cloud Storage Solutions software to secure and improve transfers.

The limited capability to start an instance from physical or virtual media, along with the pre and post-processing requirements, makes it difficult to use a backup and restore process for your disaster recovery strategy.

4.2 IBM BRMS

This section describes backup and restore procedures using IBM BRMS.

4.2.1 IBM BRMS overview

IBM BRMS is the IBM strategic solution for planning and managing save and restore operations on your IBM i instance. The IBM BRMS base product provides all of the functions that most IBM i users need to implement a fully automated, single system, backup, recovery, and media management strategy. IBM BRMS helps you implement a disciplined approach to managing your backups and provides you with an orderly way to retrieve lost or damaged data. By using IBM BRMS, you can manage your most critical and complex save operations, which includes online backups of Lotus servers. It also supports parallel save operations of a library or of a single object by using up to 32 tape devices, which shortens the save window.

You can also recover your system fully during a disaster or failure, or restore single objects or libraries from your save media. IBM BRMS can also perform some daily maintenance activities that are related to your backup routine.

In addition to these backup and recovery functions, IBM BRMS can support and manage an unlimited number of media, shared tape devices, automated tape libraries, virtual tape devices, optical devices, and IBM Storage Protect (previously IBM Tivoli® Storage Manager) servers. You can use IBM BRMS to track all of your media from creation to expiration. You no longer must track which items are on which volumes, nor be concerned that you might accidentally write over active data.

As your business needs change and grow, you can add functions to the IBM BRMS base product by purchasing and installing other options.

Note: For more information about IBM BRMS, see *Systems management Backup, Recovery, and Media Services for i*.

4.2.2 IBM Cloud Storage Solutions

IBM Cloud Storage Solutions can be used to store IBM i data securely in the cloud, which eliminates the need for tape drives, tape cartridges, and courier services for off-site storage. To store data, you must obtain server space from a cloud service provider. Then you must create a Cloud Storage Solutions resource that defines the cloud server location and authorization information that is needed to access it. You can then copy files between the integrated file system (IFS) on the IBM i system and the cloud server.

You can work with resources and files directly by using Cloud Storage Solutions commands, or your applications can work with resources and files by using the Cloud Storage Solutions API.

Cloud Storage Solutions passes information about file transfers to and from the cloud to a registered IBM i exit point. To have your applications receive that information, you can register

the applications as exit programs and associate them with the Cloud Storage Solutions exit point.

The following licensed programs must be installed:

- 5733ICC *BASE IBM Cloud Storage Solutions for i
- 5733ICC 1 Cloud Storage

Important: You can add an IBM Cloud Storage Solutions license to your IBM i instance by selecting the checkbox on the instance creation wizard.

IBM BRMS can be used to transfer virtual save media from tape or optical image catalogs to and from the cloud by using product IBM Cloud Storage Solutions for i (5733ICC). IBM Cloud Storage Solutions for i allows cloud connector resources to be defined for cloud storage providers, such as IBM COS with S3 protocol, and for private interfaces, such as file transfer protocol (FTP).

IBM BRMS creates storage locations for each cloud resource that is defined on a system. When virtual media is moved to a cloud storage location, the media is transferred to the cloud by using the cloud resource. When that media is moved from a cloud location, the media is transferred back to the IBM i system. Media is also automatically transferred back to the system during a restore when no local save media is available to the restore.

IBM i components for backup in IBM Power Virtual Server

The following components are required to create IBM i backups on IBM Power Systems Virtual Server on IBM Cloud:

- 5733ICC *BASE IBM Cloud Storage Solutions for i

  Unless you have twice the storage capacity on the IBM i virtual machine (VM) to hold the backup images
- Communications between IBM Power Virtual Server workspace and IBM VPC
- VPC at the same region
- Transit Gateway and Power Edge Router (PER) to VPC
- Virtual Private Endpoint gateway for VPC
  - 5770-BR1 *BASE
  - 5770-BR1 Option 1: Network feature
  - 5770-BR1 Option 2: Advanced Functions feature

Important: BRMS is installed in IBM i stock images by default.

- If Falconstor Storsafe is used, a VM instance that uses Red Hat Enterprise Linux is provisioned by using the “FalconStor StorSafe VTL for Power Virtual Server Cloud” option from the catalog:
  - Order the required resources based on the Falconstor Sizing Calculator
  - BRMS can maintain IBM COS virtual media based on your backup and retention policies

Note: Consider the following points:

- Media and Storage Extensions provide an API to enable application monitoring and to control media usage, which includes volumes to be selected and volume expiration dates. This feature can be useful for software developers who want to customize their own storage management applications.
- An API is provided to manage the interruption that occurs when an application attempts to open a database file that was migrated to offline media. The API enables an on-demand recall of a database file from offline media to a direct access storage device (DASD) and resumes the application (although application changes are *not* required).

These APIs provide support to use or build applications to manage tape use and data recall from offline media to DASD.

For more information about IBM Cloud Storage Solutions for i, see *IBM Cloud Storage Solutions for i User's Guide*.

Terminology

The following terminologies are used:

- *Cloud location* is an IBM BRMS storage location that is associated with a Cloud Storage Solutions cloud resource.
- *Cloud resource* is a Cloud Storage Solutions cloud resource.
- *Location* is an IBM BRMS storage location.
- *Media* is used for Cloud Storage Solutions topics, or virtual media from an image catalog.
- *Move and movement* refers to media that change from one IBM BRMS storage location to another. This logical movement is reflected in IBM BRMS databases only. It does *not* imply that the media is physically transferred to other storage.
- *System* refers to an IBM i system that uses Cloud Storage Solutions.
- *Transfer* refers to media that is physically changing from IBM i storage to storage that is associated with a cloud connector. IBM BRMS media movement to or from a cloud location causes media to be transferred to or from a cloud resource.

4.2.3 Cloud solutions for i characteristic

Solutions in the cloud include the following IBM i characteristics:

- Creates and maintain IBM COS and FTP targets.
- Provides copy to cloud, copy from cloud, and delete from cloud functions for IBM COS and FTP servers.
- Supports IBM i and Linux FTP servers.
- Manages TCP/IP networking communications to IBM COS or FTP Servers on IBM i or Linux.
- Authenticates resource users with IBM COS or FTP servers, and updates expired IBM COS tokens.
- Calls registered exit programs after each copy operation.
- Maintains file codices.
- Creates local or remote directories when needed.
- Tracks transfer progress.

- Identifies server errors.
- Provides the **WRKSTSICC** tool to identify active, failed, and successful transfers, and to show the progress of active transfers:
  - Operations are run in jobs.
  - View the status of those jobs.
  - Work with those jobs (for example, end).
- Uses the **CRTS3RICC** command to create a Cloud Storage Solutions AWS S3 or IBM COS resource. A resource defines a cloud server location and the credentials that are needed to access that location. After a resource is created, the files can be copied between IFS directories and the cloud server location.
- Uses the **CHGS3RICC** command to change an AWS S3 or IBM COS resource. A resource defines an AWS S3 or IBM COS cloud server location and the credentials that are needed to access that location. It also changes a resource to use different credentials to access the same bucket, or to specify a different bucket.
- Uses the **DSPS3RICC** command to display an AWS S3 or IBM COS resource. A resource defines an AWS S3 or IBM COS cloud server location and the credentials that are needed to access that location.
- Overrides database files to the scope of the activation group so that the user does not need to be concerned about library lists.
- APIs implemented with `Qlg_Path_Name_T` path name formats so that native IBM i programs can work seamlessly with the APIs.
- Provides synchronous or asynchronous copy and delete operations.
- Seamless handles CCSID string conversions (IBM i interfaces works in EBCDIC, IBM COS, and FTP Linux in UTF-8).
- Works with the native IBM i messaging system to return server errors as diagnostic or escape messages.

For more information about the IBM i Cloud storage solution, see the following resources:

- Cloud Storage Solutions for i 5733ICC Support Documentation summary
- IBM BRMS Cloud Education videos:
  - Automatic Transfers of Media to Cloud Storage
  - User Initiated Transfers of Media to Cloud Storage

4.2.4 IBM BRMS turn-key

IBM BRMS turn-key is an automated, pre-scripted solution within IBM BRMS that is designed to make it easier to set up and run.

Overview of IBM BRMS turn-key

IBM Cloud Storage Solutions for i must be present with at least one Cloud Resource.

All required PTFs must be present. See link references below for additional information.

When you provide the cloud resource name (that is, the cloud storage resource name), IBM BRMS performs the following tasks:

- Creates the virtual tape
- Shows the image catalog
- Mounts the image catalog

- Saves the image catalog
- Backs up to the cloud

When the policy is set, IBM BRMS turn-key automatically runs Cloud Storage Solutions commands.

IBM BRMS turn-key setup

One of the following commands can create the required IBM BRMS objects for new cloud resources:

- Run STRMNTBRM RUNCLNUP(*YES)
- Run INZBRM (*DATA)

IBM BRMS creates the following objects when it detects the IBM COS or FTP Resource Name:

- Media class. One class for Virtual Tape (if MSE is on the system) and one class for Virtual Optical.
- Storage location. Based on the resource name.
- Move policy. Based on resource name.
- Media policy. Based on resource name.
- Four backup control groups:
  - `QCLDBIPL`*nn*. Backs up what is minimally needed for a system D-IPL (must be burned to a DVD).
  - `QCLDBSYS`*nn*. Backs up all system data, except `*SAVSYS`. Paired with `QCLDBIPL`*nn*.
  - `QCLDBUSR`*nn*. Backs up all user data.
  - `QCLDBGRP`*nn*. Backs up what is minimally needed to and from the cloud, except `*SAVSYS`. Paired with `QCLDBUSR`*nn* (must be burned to a DVD).

The backup requires the following separate recovery reports:

- Cloud Connection Recovery (DVD).
- Recovering user data from cloud

Important: You can use any of these Control Groups as templates to create your own, but must preserve the prefix.

You can get updated information from the following links:

- BRMS Cloud Storage Solutions Requirements and Considerations
- Quick Start Guide for Cloud Storage Solutions
- Using BRMS with Cloud Storage Solutions for i

4.2.5 Backing up the IBM i system to the cloud

IBM BRMS stores system backup media in cloud locations the same way that it stores media in physical media storage devices.

This backup solution is an example of how to use IBM BRMS with IBM Cloud Storage Solutions so that it can save your entire system to virtual media in the cloud. Consider the following points:

- If IBM Cloud Storage Solutions are used for i V1.2.0 with compression or encryption, it is not possible to recover the system from the cloud if a disaster occurs.

- The *nn* in the control group name is a number that IBM BRMS assigns to the connector for which the control group was created.
- Backup strategies often require full system backups at specific intervals and daily, incremental backups to capture changes. It is assumed in this example that a full backup of the system was run with the QCLDBSYS*nn* and QCLDBIPL*nn* control groups to obtain a full backup of the user data.

When IBM BRMS is used to perform a system backup to media at a conventional location, control groups run in succession to save specific objects and data. The control groups obtain the data that is required to restore the entire system.

The use of IBM BRMS to back up a system and store the data to media in the cloud requires a similar process. Special control groups that are created by IBM BRMS store the media in the cloud.

System backups to the cloud follow the same procedure as transferring data to the cloud. The only difference is that the entire system is backed up at a specific point by using the cloud resource control groups. As with backups to a conventional location, the cloud backup must specifically save the system and user group objects and other production data can be saved on other media.

Beginning with the full backup control groups, QCLDBSYS*nn* and QCLDBIPL*nn*, complete the following steps:

1. Sign onto the console.
2. Verify that the cloud location is available by using the **WRKLOCBRM** command.
3. Begin with the QCLDBSYS*nn* control group by running the **STRBKUBRM CTLGRP(QCLDBSYS01) SBMJOB(*NO)** command to start the system backup.
4. Run the QCLDBIPL*nn* control group by using the **STRBKUBRM CTLGRP(QCLDBIPL01)SBMJOB(*NO)** command.
5. Begin the backup by running the **STRBKUBRM CTLGRP(QCLDBUSR01) SBMJOB(*NO)** command.
6. Run the **STRBKUBRM CTLGRP(QCLDBGRP01) SBMJOB(*NO)** command.
7. After these backups are done, review the job logs to ensure that the full backup was successful.

After obtaining a full backup of a system, move to a backup plan that is similar to the following example:

- On Sunday, run the following commands to obtain a full backup of the system:
  - **STRBKUBRM CTLGRP(QCLDBSYS01) SBMJOB(*NO)**
  - **STRBKUBRM CTLGRP(QCLDBIPL01) SBMJOB(*NO)**
- Run the following commands Monday - Saturday to obtain incremental backups of user data:
  - **STRBKUBRM CTLGRP(QCLDBUSR01) SBMJOB(*NO)**
  - **STRBKUBRM CTLGRP(QCLDBGRP01) SBMJOB(*NO)**

Important: It is critical to run these cloud control groups in the order that is indicated. Otherwise, all necessary media information is not available to perform a recovery.

As a part of your backup process, it can be advantageous to have the backup media quickly accessible on the system for a period after the transfer. If cloud resources are the only location where your system backup is stored, consider copying your QCLDBGRP*nn* and QCLDBIPL*nn* control group data to optical media and storing the discs with other physical

media. By keeping this backup data accessible, you can perform restores from the media without the time and expense of transferring the media that is in the cloud back to the system.

IBM BRMS permits media that is associated with a move policy to be retained on the system for a period after it was transferred to the cloud. This retention is done by changing the Retain media field of the move policy to keep the media for a specified number of days after the move. Use the Work with IBM BRMS policies **WRKPCYBRM *MOV** command.

Customizing the backup of an entire IBM i system to the cloud

The example that is presented in this section shows how to override default settings that are used by IBM BRMS with IBM Cloud Storage Solutions for IBM i to save your entire system to virtual media in the cloud.

The same restrictions and precautions that must be observed when backing up your entire system to the cloud also apply when customized control groups are used to perform the system backup to the cloud.

Consider the following points:

- If IBM Cloud Storage Solutions for i V1.2.0 is used with compression or encryption, the system cannot be recovered from the cloud if a disaster occurs.
- The *nn* in the control group name is a number that IBM BRMS assigns to the connector for which the control group was created.
- Backup strategies often require full system backups at specific intervals and daily, incremental backups to capture changes. It is assumed in this example that a full backup of the system was run with the QCLDCSYS*nn* and QCLDCIPL*nn* control groups to obtain a full backup of the user data.

In addition, when creating custom control groups for a system backup to the cloud, some restrictions exist on the naming conventions. Consider the following points when control groups are copied and modified:

- The new control group name must begin with a QCLD prefix to enable automatic transfers to the cloud.
- The new control group names cannot begin with QCLDB, QCLDUIPL, or QCLDUGRP.

When IBM BRMS is used to perform a system backup to media in a cloud location, the control group or multiple control groups run in succession and save specific objects and data.

In system backup that uses the turnkey settings, IBM BRMS stores the data in the cloud in volumes of a default size that is enforced in the media class.

The IBM BRMS implementation enforces the use of the default media class QCLDVRTOPT with a value of **IMGSIZ(*DVD4700)** for virtual optical media volumes. Copying a control group for use as a modified user control group named, for example, QCLDUGRP*nn*, also results in the same implementation that enforces the default media class value.

If this default size is too small for the data users to store, then modified, custom control groups must be created with larger media sizes. To use a user-defined media size that is controlled by a user-defined media class, a control group must be created and used that does not follow the IBM BRMS turnkey automated control group naming conventions.

Note: Control groups cannot use the QCLD*U* prefix but must instead adopt a new name with the QCLD*C* prefix.

To customize the system backup and use modified control groups that can use nondefault media sizes, complete the following steps:

1. Sign onto the console.
2. Verify that the cloud location is available by using the **WRKLOCBRM** command.
3. Copy control group QCLDBSYS*nn* to a new custom control group and update any entries as needed. Also, rename the control group to, for example, QCLDCSYS01.
4. Copy the default cloud-named media policy to a new custom name and set the wanted cloud virtual media class that is to be used.
5. Run the QCLDCSYS01 control group by using the **STRBKUBRM CTLGRP(QCLDCSYS01) SBMJOB(*NO)** command to start the system backup.
6. Copy the control group QCLDBIPL01 to a new custom control group and update any entries as needed. Also, rename the control group to, for example, QCLDCIPL01.
7. Update the custom control group attributes to use the new, modified media policy.
8. Run the QCLDCIPL01 control group with **STRBKUBRM CTLGRP(QCLDCIPL01) SBMJOB(*NO)**.
9. Copy the control group QCLDBUSR01 to a new custom control group and update any entries as needed. Also, rename the control group to, for example, QCLDCUSR01.
10. Update the custom control group attributes to use the new, modified media policy. Begin a backup that uses this control group by running the **STRBKUBRM CTLGRP(QCLDCUSR01) SBMJOB(*NO)** command.
11. Copy the control group QCLDBGRP01 to a new custom control group and update any entries as needed. Also, rename the control group to, for example, QCLDCGRP01.
12. Update the custom control group attributes to use the new, modified media policy.
13. Run a backup that uses this control group by running the **STRBKUBRM CTLGRP(QCLDCGRP01) SBMJOB(*NO)** command.
14. After these backups complete, review the job logs to ensure that the backup was successful.
15. Run the cloud control groups regarding the previous sample.

As a part of your backup process, it can be advantageous to have the backup media quickly accessible on the system for a period after the transfer.

If cloud resources are the only location where your system backup is stored, consider copying your QCLDCGRP01 and QCLDCIPL01 control group data to physical media. Keeping this backup data accessible allows restores to be performed from the media without the time and expense of transferring the media that is in the cloud back to the system.

4.2.6 Full-system backups from the cloud

Complete the following steps to run the full-system backups from the console by using IBM BRMS:

1. Log in to the IBM Cloud by using your IBM user ID and password.
2. Go to your Service and select your virtual server instance. Open the console, as shown in Figure 4-1 on page 127.

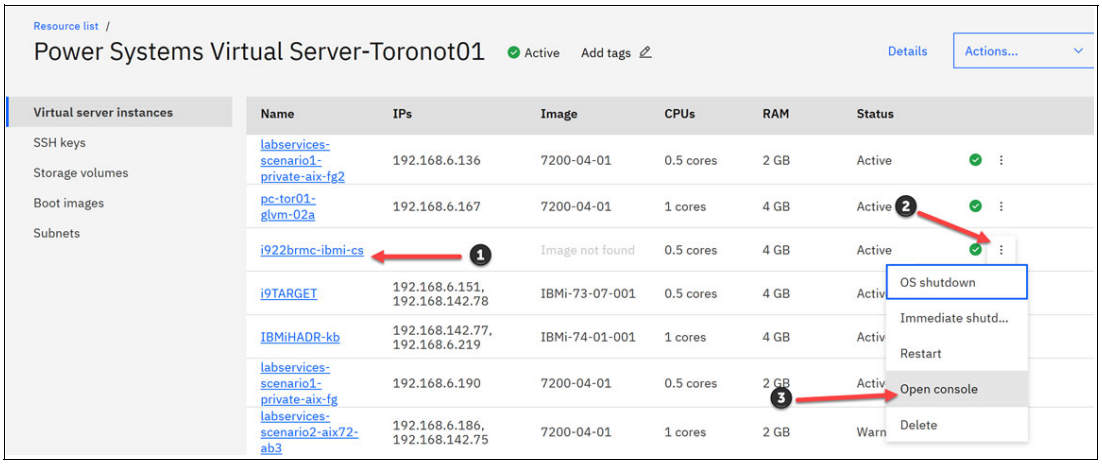

Figure 4-1 Open 5250 console for IBM i

3. Sign on by using your Dedicated Service Tools (DST) username and password, or select **PF18** to bypass the use of the DST tool. Select **Next** and then press **PF18**.

> **Note:** The console times out if it is inactive after 5 minutes. Close your console browser and start a new console connection.
>
> If you see a break message during the backup process, press **Enter** to return to the window in which you entered the **STRBKUBRM** command so that you can see the progress of the backup.

4. Run the IBM BRMS control group QCLDBSYS01. Put the system in a restricted state and then run the **ENDSBS SBS(*ALL) DELAY(120)** command.
5. Display QSYSOPR MSGQ on the command line. Run the **DSPMSG QSECOFR** command and look for the following messages:
   - System ended to restricted condition.
   - A request to end TCP/IP has completed.
6. Change the subsystems to process for control group QCLDBSYS0:
   a. Use the **WRKCTLGBRM** command, and find QCLDBSYS01.
   b. Select **Option 9=Subsystems to process**.
   c. Change the restart to ***NO** for Seq 10 Subsystem *ALL, as shown in Figure 4-2.

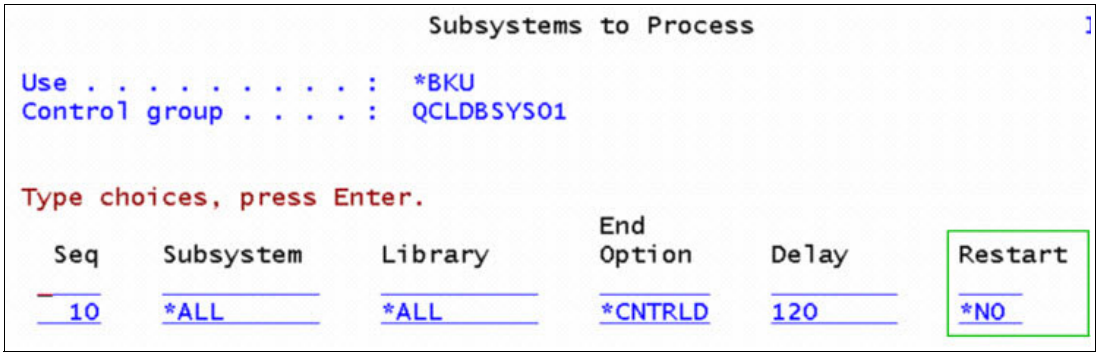

*Figure 4-2 Change restart to *No on subsystem to process*

7. Run the First backup from the console: **STRBKUBRM CTLGRP(QCLDBSYS01) SBMJOB(*NO)**.
8. Check the backup for errors. It is normal to have some errors (see Figure 4-3), including the following examples:
   - Objects not saved (Some objects are not required for the recovery).
   - Media not transferred (You will complete this step manually after the Second backup).

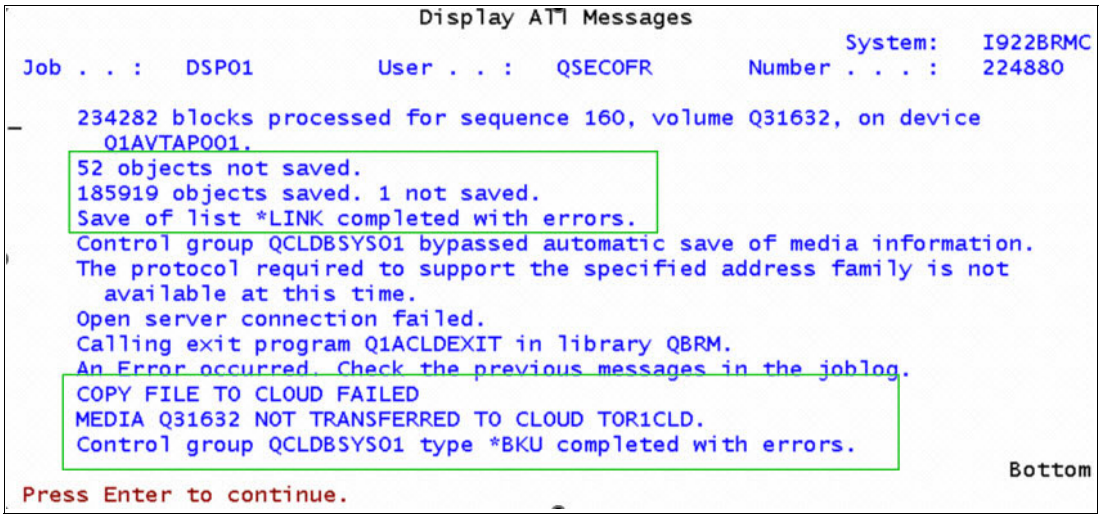

Figure 4-3 Display possible errors

9. Check the subsystems after the backup is completed. You see only subsystem QCTL in an RSTD status. If it is not in this status, end all subsystems again and run the **ENDSBS SBS(*ALL) DELAY(120)** command.
10. Change IBM BRMS control group QCLDBIPL01:
    a. Run the **WRKCTLGBRM** command.
    b. Select **Option 8=Change attributes**.
    c. Page down, change the Automatically backup media information to ***LIB** and the Append to media to ***NO**.
    d. Select **Option 9=Subsystems to process**.
    e. Change the restart to ***YES** for **Seq** 10 Subsystem *ALL (see Figure 4-4).

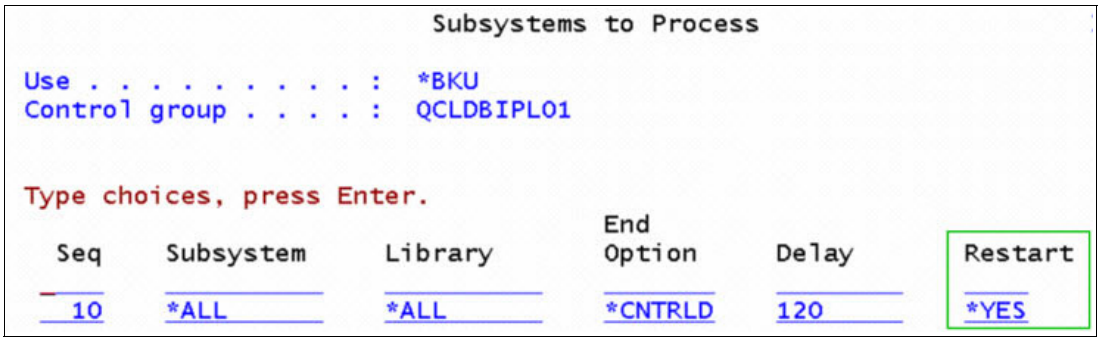

*Figure 4-4 Change restart to *Yes on subsystem to process*

11. Issue the second backup from the console **STRBKUBRM CTLGRP(QCLDBIPL01) SBMJOB(*NO)**.

12. Check the backup for errors. It is normal to have some errors, including the following examples:
    - `Objects not saved (Some objects are not required for the recovery).`
    - `Media not transferred (You will complete this step manually after the Second backup).`
13. Identify the volumes that used *both* backups `QCLDBSYS01` and `QCLDBIPL001` and transfer to IBM COS.
14. Check the status of the transfer by running the **`WRKSTSICC STATUS(*ALL)`** command (a status of `Failed` is normal). The volumes are transferred in the next step.
15. Identify which volumes were used for the backups: **`WRKMEDBRM TYPE(*TRF)`**, as shown in Figure 4-5.

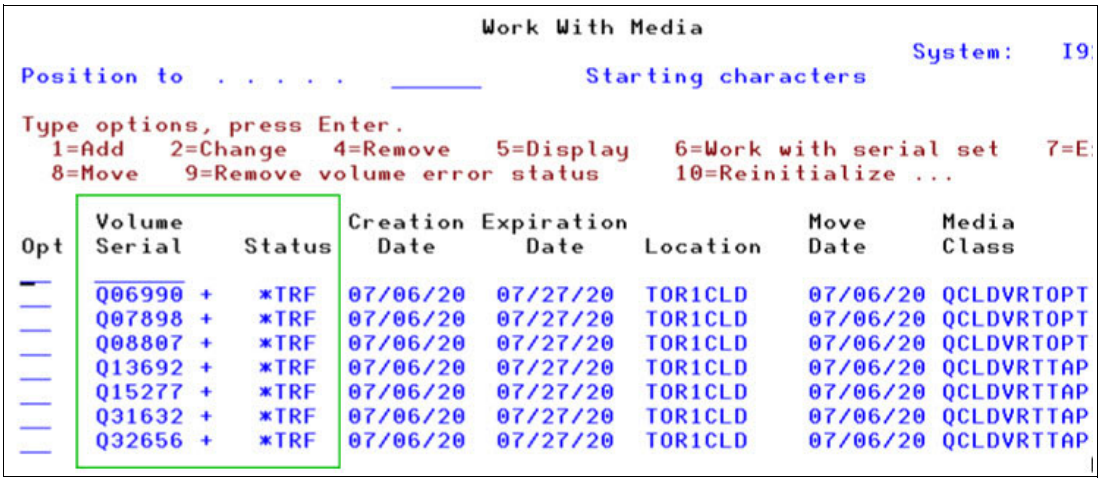

Figure 4-5 Identifying the volumes that were used

16. Transfer the volumes to IBM COS. Run the **`STRMNTBRM`** command and the **`WRKSTSICC STATUS(*ALL)`** command. You see the volume name, status, and completion percentage for each file transfer. Wait until all volumes are successfully completed before proceeding to the next step.
17. Verify that all of the volumes that were used for the full-system backup no longer feature a status of **`*TRF`**. Run the **`WRKMEDBRM TYPE(*TRF)`** command to verify that no volumes are listed.
18. As with other recoveries that are performed by using IBM BRMS, a recovery report is used to assist with successful recoveries from save media that was transferred to the cloud. To generate a report for recovery from the cloud, run the following command:

    `STRRCYBRM OPTION(*CTLGRP) ACTION(*REPORT) CTLGRP((QCLDBSYS01 1) (QCLDBIPL01 2)`

Important: It is important to review the recovery report to ensure that it is complete. If any of the media that was produced during the backup process was successfully transferred to the cloud, it is not included in the recovery report.

The **CTLGRP** and **PERIOD** parameters that were specified in the **STRRCYBRM** command help identify objects that were saved to volumes that were not transferred to the cloud. If objects are on volumes that were not included in the recovery report, they are listed in a missing objects Attention section that is near the top of the report.

After the recovery report is verified, the report is stored in a safe location so that it can be referred to during a recovery.

Daily incremental backups can be run Monday - Saturday by using the following control groups:

```
STRBKUBRM CTLGRP(QCLDBUSR01) SBMJOB(*NO)
STRBKUBRM CTLGRP(QCLDBGRP01) SBMJOB(*NO)
```

Figure 4-6 shows how to perform full system saves by using IBM BRMS, IBM Cloud Storage (the IBM Cloud Storage GUI), and IBM COS.

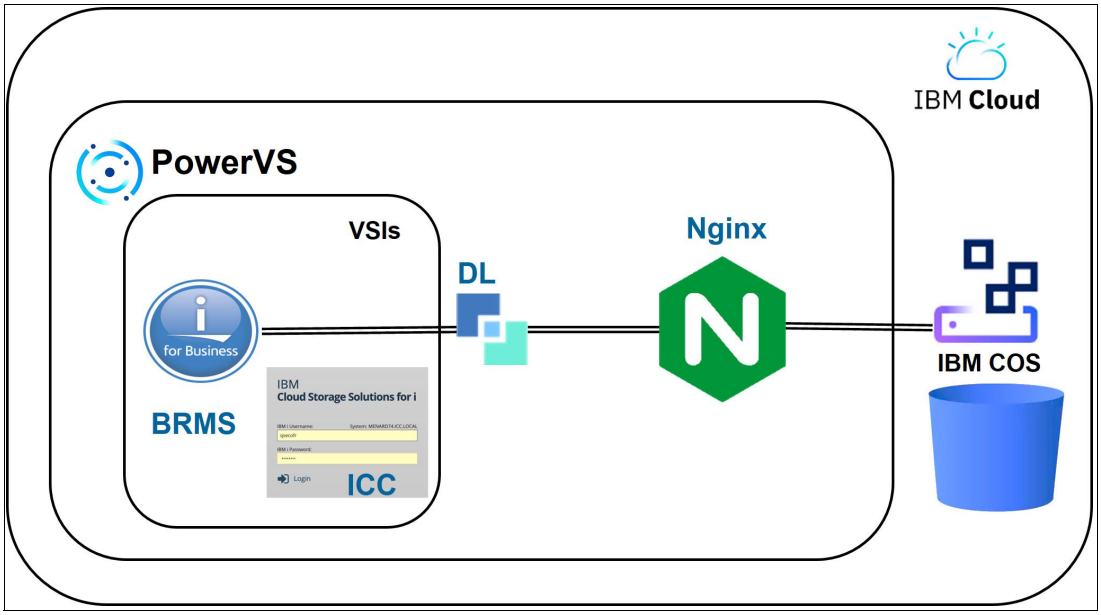

Figure 4-6 IBM i source that uses IBM BRMS and IBM Cloud Storage

Note: Consider the following points about the use of IBM COS over IBM Cloud Direct link:

- IBM Cloud customers that purchase IBM COS and Direct Link can make remote connections to IBM COS “private” endpoints.
- HTTPS IBM COS requests are started from a remote site.
- The requests are transmitted through IBM Cloud Direct Link and target a Proxy server.
- The requests are passed to an IBM COS private endpoint.
- The request is processed.
- The request results are returned to the user.

4.2.7 Recovering the IBM i system from the cloud

IBM BRMS restores backup media from cloud locations by using a unique process:

- As with other recoveries that are performed by using IBM BRMS, a recovery report is used to assist with successful recoveries from save media that was transferred to the cloud. To generate a report for recovery from the cloud, the following command is used:

  STRRCYBRM OPTION(*CTLGRP) ACTION(*REPORT) CTLGRP((QCLDBSYS*xx* 1) (QCLDBIPL*xx* 2) (QCLDBUSR*xx* 3) (QCLDBGRP*xx* 4)) PERIOD((<time> <date>))

  where ***xx*** identifies the cloud where the media is stored and ***<timestamp*** before to the last full backup.

- The following save data must be restored from physical media before IBM BRMS can begin restoring save data directly from the cloud:
  - SAVSYS, which is used to install the operating system.
  - Product IBM BRMS for i and IBM BRMS save information before automatic recovery can be performed.
  - Product IBM TCP/IP Connectivity Utilities for i and configuration information to allow communications with cloud storage providers.
  - Product IBM Cloud Storage Solutions for i and configuration information to establish connections with cloud storage providers.

System recoveries cannot be performed directly from IBM BRMS media that was transferred to the cloud. To perform a system recovery from cloud media, special procedures must be followed to create physical optical installation media.

The physical optical media contains SAVSYS data and objects from other libraries, such as **QUSRSYS**, **QBRM**, **QUSRBRM**. After the optical media is used to restore Licensed Internal Code (LIC), the operating system, and other required objects, subsequent restores can be performed directly from the cloud media by using IBM BRMS.

Some media that is listed in your recovery report must be recovered from physical media. Complete the following steps:

1. Locate the media that requires conversion to a physical copy. Use the volume identifiers to locate the following items:
   - LIC
   - Operating system objects
   - IBM BRMS product and associated libraries
   - User profiles and configuration data
   - Required system libraries

   Normally, these media are stored in the `QIBM BRMS_XXXXXXXX` directory. The files in this directory use the same name as the media volume identifiers that are listed in your recovery report.
2. By using your connection to the cloud, transfer this media to your system to write the files to optical discs. A .iso extension must be added to the file name if it is required by the image burning software that is used.

   Note: Before restoring from media in the cloud, the IBM BRMS media database must be updated. Cloud volumes must be registered with IBM BRMS.

3. Run the program **ADDLIBLE LIB(QICC) QICC/REGEXTPTS ACTION(R)** to register the cloud exit programs.

4. Ensure that the media library name is correct for the Device prompt and enter `*YES` in the Create parent directories prompt.

   After this step, the physical media is not required. The system is now in a restricted state. TCP/IP must be started to allow IBM BRMS to download the media that is required for cloud recovery.

5. To start TCP/IP, enter the following commands:

   ```
   STRTCP STRSVR(*NO) STRIFC(*NO) STRPTPPRF(*NO) STRIP6(*YES)
   STRTCPIFC INTNETADR('nnn.nnn.nnn.nnn')
   ```

   where *'nnn.nnn.nnn.nnn'* is the IP address of the recovery system.

6. Press **Enter**.

IBM BRMS downloads the media from the cloud and begins recovering all remaining system data. The restored system performs verification as a final stage of the recovery. To allow the system to verify the system information, end TCP/IP by using the **ENDTCP** command.

After system verification is complete, restart TCP/IP by using the **STRTCP** command and then IPL the system.

4.2.8 Full-system recovery from the cloud using IBM i as an NFS server

This section provides the steps to perform a full-system recovery from the cloud.

Setting up IBM i Network Install Server with NFS Server and NFS Client

To set up IBM i Network Install Server with NFS server and NFS client, provision an IBM i VSI in the target Power Systems Virtual Server location to be an NFS Server.

Consider the following points:

- IBM i NFS Server must be at a minimum at Version 7.2 with current PTFs.
- To use virtual optical images through an NFS server, the IBM i NFS client must meet the following requirements:
  - The IBM i has a Version 4 IP address.
  - During set up, the shared NFS server directory is mounted over a directory on the IBM i client.
  - An IBM i service tools server or a LAN console connection is configured by using a Version 4 IP address.
  - A 632B-003 virtual optical device is created by using the IP address of the NFS server.

Notes:

The IBM i IP address and the IBM i service tools server (LAN console connection) IP addresses **cannot** be the same.

NFS V3 with UDP is required to perform remote Install procedure.

IBM i NFS Server V7.5 and later uses TCP instead of UDP ports, and needs additional settings to work as an Install server.

For more information, see Setting up an IBM i network install server.

Requirements to create a Power Virtual Server instance for NFS Server

The following requirements must be met to create a Power Systems Virtual Server Instance NFS Server:

- An IBM i NFS Server is created in the IBM Power Systems Virtual Server location.

  Ensure that enough disk storage is available to support the size of the SAVSYS or Control Group `QCLDBIPLxx`.

- An IBM i Client Server is created in the same IBM Power Systems Virtual Server location.

  Ensure that *twice* the disk storage is available to support the size of the IMAGE Catalog volumes (User Data), and the Full-System Restore.

- Verify that the IBM i Client Server can ping itself and the IP address of the IBM i NFS Server.

- Verify that the IBM i NFS Server can ping itself and the IP address of the IBM i Client Server.

Creating an IBM i instance for NFS Server

Complete the following steps to create an IBM i IBM Power Systems Virtual Server instance NFS Server:

1. Create an IBM i instance in your Power Systems Virtual Server location.
2. Provision and add your storage.
3. Create a mount directory for the IBM i NFS server:
   a. Run the **`MKDIR DIR('/install')`** command.
   b. Run the **`MKDIR DIR('/install/sysipl')`** command.

      A directory is created in the IFS that is called `/install/sysipl`.
4. From the Cloud, copy the volumes that were created by using the `QCLDBIPL01` backup. In this example, three volumes are available (`Q06990`, `Q07898`, and `Q08807`).

   To copy the volumes from the Cloud, you need the following information:

   – Resource name: The cloud resource that was created.
   – Submit to batch: `*YES` (submit all three simultaneously).
   – Cloud file name: The full name of the volume that is on the cloud.
   – Local file name: The full directory path plus the volume name.
5. Display the volumes on the IBM Cloud Object Storage by clicking **Cloud Object Storage** → **Bucket** → **Object name** → **Details**.

Note: IBM BRMS stores media in the cloud as files in the `QIBM BRMS_`*`XXXXXXXX`* directory, where *`XXXXXXXX`* is the name of the system that performed the backups (see to Figure 4-7).

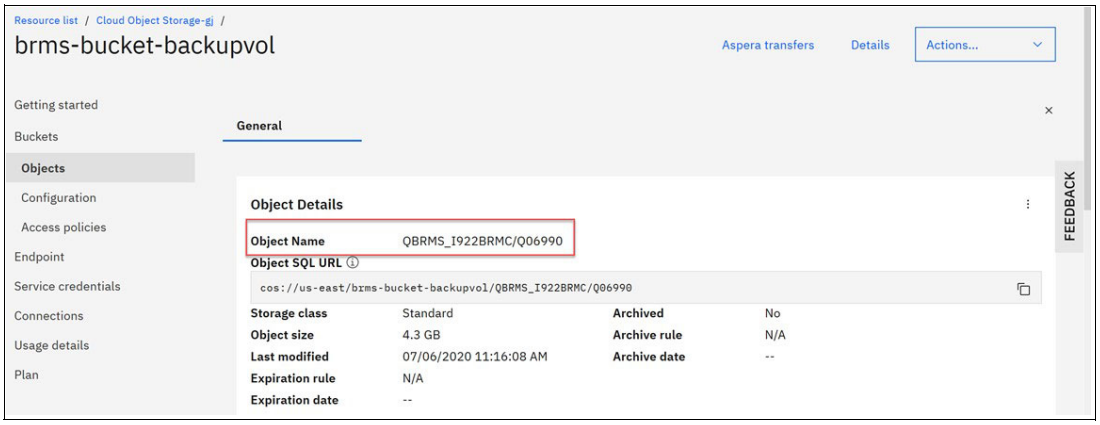

Figure 4-7 IBM BRMS storing media on IBM Cloud Object Storage

6. Run the **CPYFRMCLD** command, as shown in Figure 4-8.

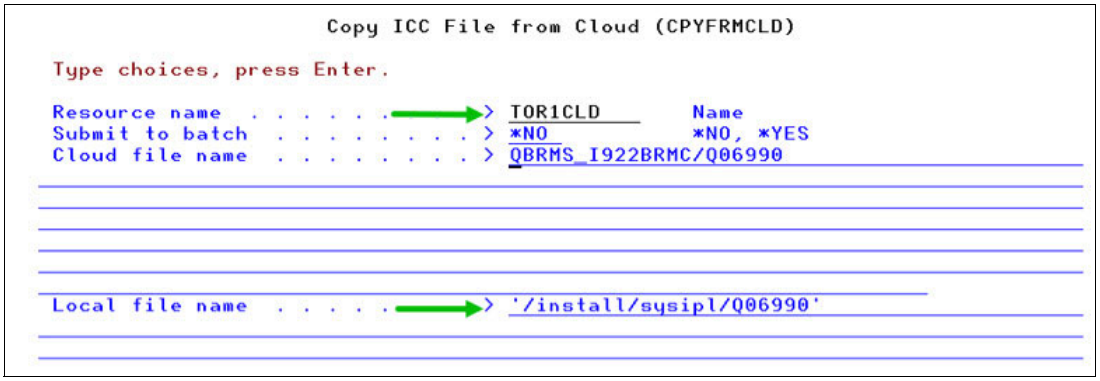

Figure 4-8 Copying IBM Cloud Storage file from cloud

7. Check the status of the transfer jobs by using the **WRKSTSICC STATUS(*ALL)** command. Then after all volumes complete the transfer and show a status of `Success`, see the next section, Creating a virtual optical device on an IBM i NFS Server.

Creating a virtual optical device on an IBM i NFS Server

Complete the following steps to create an optical device on IBM i NFS Server:

1. Create a Virtual Optical Drive that is named **"**`INSTALL`**"** `CRTDEVOPT DEVD(INSTALL) RSRCNAME(*VRT) LCLINTNETA(*N)`.
2. Vary on the Virtual Optical Device by using the **WRKCFGSTS *DEV INSTALL** command.
3. Create an image catalog: **CRTIMGCLG IMGCLG(SYSIPL) DIR('/install/sysipl')**.
4. Add an image catalog entry, as shown in Figure 4-9 on page 135 (or example, volume `Q06990`).

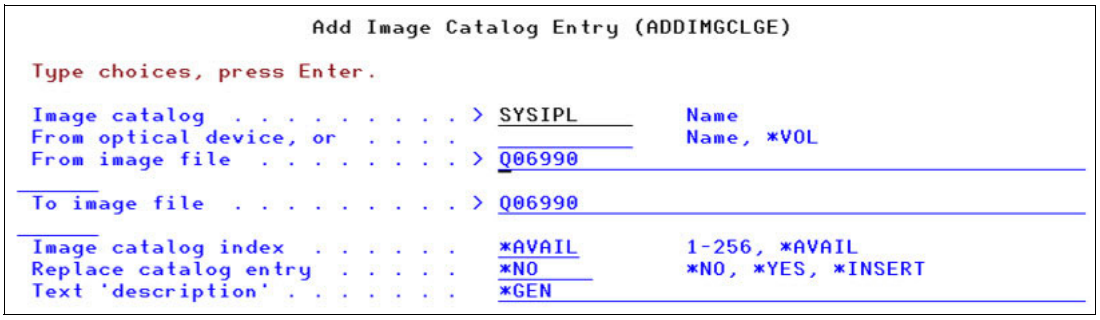

Figure 4-9 Add image catalog entry

5. Add the next two volumes by using the same name for the TOFILE, which is listed in the field name To image file.
6. Load the image catalog, as shown in Figure 4-10.

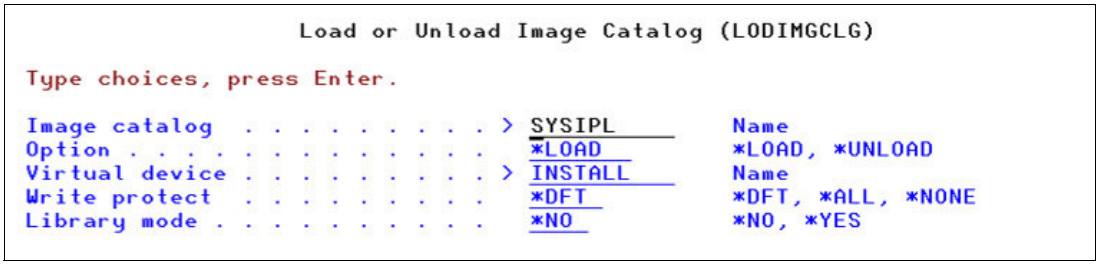

Figure 4-10 Load to image catalog

7. Check or verify the Image catalog:
   a. On the command line, run the **WRKIMGCLG** command.
   b. Select **Option 10**=Verify for image catalog “SYSIPL”.
   c. From the command output, verify that the type = `*LIC`.
   d. Verify from the command output that the sort image catalog = `*YES`.
   e. Verify that the network file server share = `*YES`.
8. Work with Entries:
   a. Select **Option 12=Work with entries for image catalog “SYSIPL”**.
   b. Verify the directory.
   c. Verify the Index order of the image file name The image file that contains the SAVSYS is the first image in the index order.
   d. Verify the status (`Mounted` or `Loaded`).
9. Work with Object Links by using the following command:

   `WRKLNK OBJ('/install/sysipl/*')`

   – Verify that a BOOTP DIR was created.
   – Verify that a VOLUME_LIST was created.
10. Run the **`STRNFSSVR SERVER(*ALL)`** command to start NFS Servers.
11. Run the **`CHGNFSEXP OPTIONS('-i -o ro') DIR('/install/sysipl')`** command to change NFS export options.
12. Run the **`CHGAUT OBJ('/install/sysipl') USER(*PUBLIC) DTAAUT(*RWX) SUBTREE(*ALL)`** command to change Object Authority.

13. Run the **CHGTFTPA AUTOSTART(*YES) ALTSRCDIR('/install/sysipl')** command to change the TFTP Attributes.
14. Specify the alternative source directory where the volumes are stored.
15. Run the **CHGAUT OBJ('/install/sysipl') USER(QTFTP) DTAAUT(*RX) SUBTREE(*ALL)** command to change Object Authority.
16. End the TCP Server TFTP by running the **ENDTCPSVR SERVER(*TFTP)** command.
17. Start the TCP Server TFTP by running the **STRTCPSVR SERVER(*TFTP)** command.

Creating and configuring an IBM i client server instance

Complete the following steps to create an IBM i Power Systems Virtual Server Instance Client Server:

1. Create an IBM i instance in your Power Systems Virtual Server location.
2. Provision and add your storage.
3. Log in to the Start System Service Tools (SST) by using the **STRSST** command.
4. Select **Option 8=Work with service tools user IDs and Devices**.
5. Select **F13=Select STS LAN adapter**. Click **Next**.
6. Click **F13** at the bottom of the console window.
7. Choose **Option 1=Select**, as shown in Figure 4-11. Use the same resource that was used to communicate with the IBM i NFS Server.

```
                         Select LAN Adapter Port
                                                                  System:
Type option, press Enter.
  1=Select   5=Display details

          Resource                        Serial
Option    Name          Type    Model     Number            LOC    Selected
  1       CMN03         268C    002       00-00000
```

Figure 4-11 Selecting a resource name for LAN Adapter

> **Note:** The CMNxx Resource *must* be on the same VLAN as your IBM i NSF Server. You must also end TCP/IP and Vary Off the Line Description using that CMNxx.

8. Configure Service Tools LAN Adapter (see Figure 4-12 on page 137):
   - IP version allowed. IPV4.
   - Internet address. Use a free IP address on the same subnet of the client (TARGET).
   - Gateway router address. Use the same gateway of the client (TARGET).
   - Subnet mask. Use the same as on the client.

   Complete the following steps:

   a. Select **F7=Store**.
   b. Select **F13= Deactivate**.
   c. Select **F14=Activate**.

   > **Note:** After you select **F14=Activate**, the adapter restarts and might not be ready immediately.

d. Select **F3=Exit** to exit the Service Tools window (see Figure 4-12).

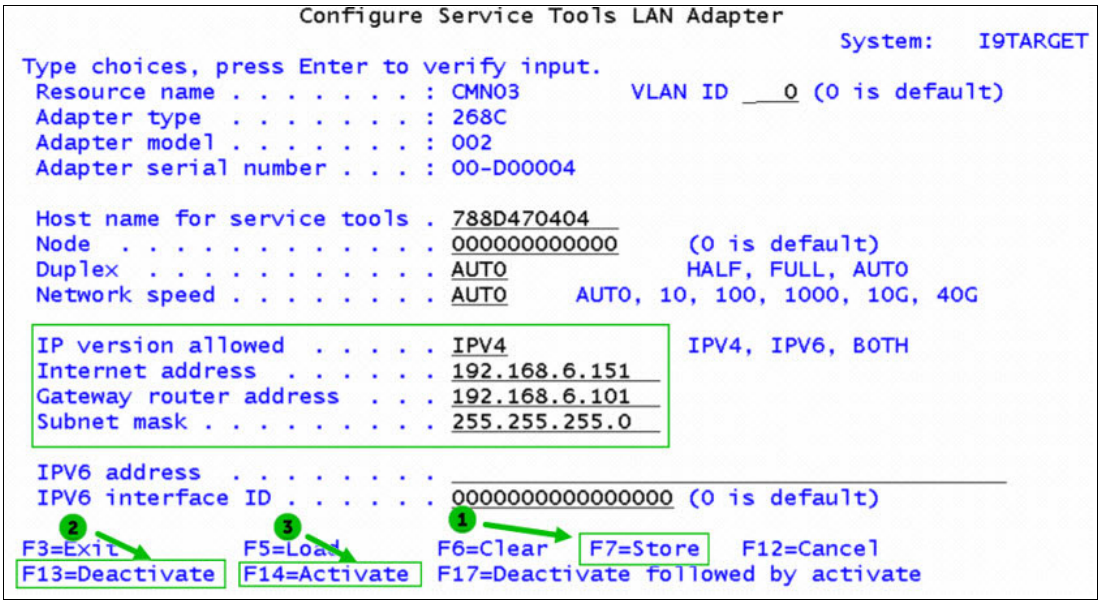

Figure 4-12 Set up TCP IP on LAN Adapter

9. Work with IP IPv4 Connection Status and verify the status of `port 3000`:
   a. Run the **`NETSTAT`** command.
   b. Select **Option 3=Work with IPv4 connection status**.
   c. Select **F14=Display port numbers**.
   d. Verify that port number 3000 (`as-sts`) is running, which is the Service Tools Server.
10. On the client server, create the optical device by running the **`CRTDEVOPT`** command:
    a. Select **F4** to prompt.
    b. Set the local internet address as `*SRVLAN`.
    c. Set the remote internet address as the IP address of the IBM i VSI NFS Server.
    d. Set the network image directory as '`/install/sysipl`'.
11. Run the **`WRKCFGSTS *DEV INSTALL`** command. Select **Option 1= Vary On for device INSTALL**.
12. Run the **`WRKIMGCLGE IMGCLG(*DEV) DEV(INSTALL)`** command to verify that you can access the remote image catalog.
13. Run the **`/install/sysipl`** command to verify the Catalog, Type, and Directory.
14. Verify the status of the volumes (`Mounted` or `Loaded`).

Installing LIC and operating system on IBM i client server from NFS server

Complete the following steps to install the LIC and operating system on IBM i Power Virtual Server on client server by using NFS server:

Important: Before you begin the scratch installation on the target IBM i, document all your network information. Re-create the network information after the installation is complete. For example, document the following information:

- **CFGTCP**:
  - Work with TCP/IP interfaces
  - Work with TCP/IP routes
- **DSPLIND**:
  - CLOUDINIT0
  - CLOUDINIT1
  - CLOUDINIT2
- **WRKHDWRSC *CMN**:

  Display all of the available CMNxx resource information and document the location and resource name:
  - Resource name . . . . . : CMN03
  - Location: U9009.22A.788D380-V5-C4-T1

1. Start the network installation:
   a. Run the **STRNETINS** command and then select **F4** to prompt.
   b. Use the following settings:
   - Network optical device: INSTALL.
   - Installation option: *LIC.
   - Keylock mode: MANUAL.

   c. Press **Enter**.
   d. Select **F16** to confirm power down.

COMM657 is displayed after 15 minutes (see Figure 4-13).

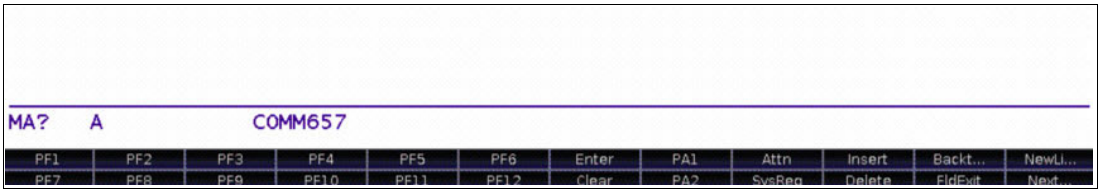

Figure 4-13 COMM657 after STRNETINS

2. Install the licensed internal code by completing steps 4–6. For a description of the steps, see Recovering your system to a different system.
3. In the menu, a prompt to IPL or install the System is displayed. Select **Option 2=Install the operating system**.
4. Select **Option 5=Network device**.
5. Configure the network device (see Figure 4-14 on page 139) by using the following settings:
   - Server IP: IP address of the SOURCE NFS IBM i Server.
   - Path Name: Name of the Directory where the image volumes are located.

Complete the following steps:

a. Select **F10 =Continue**.
b. Press **Enter** to confirm.

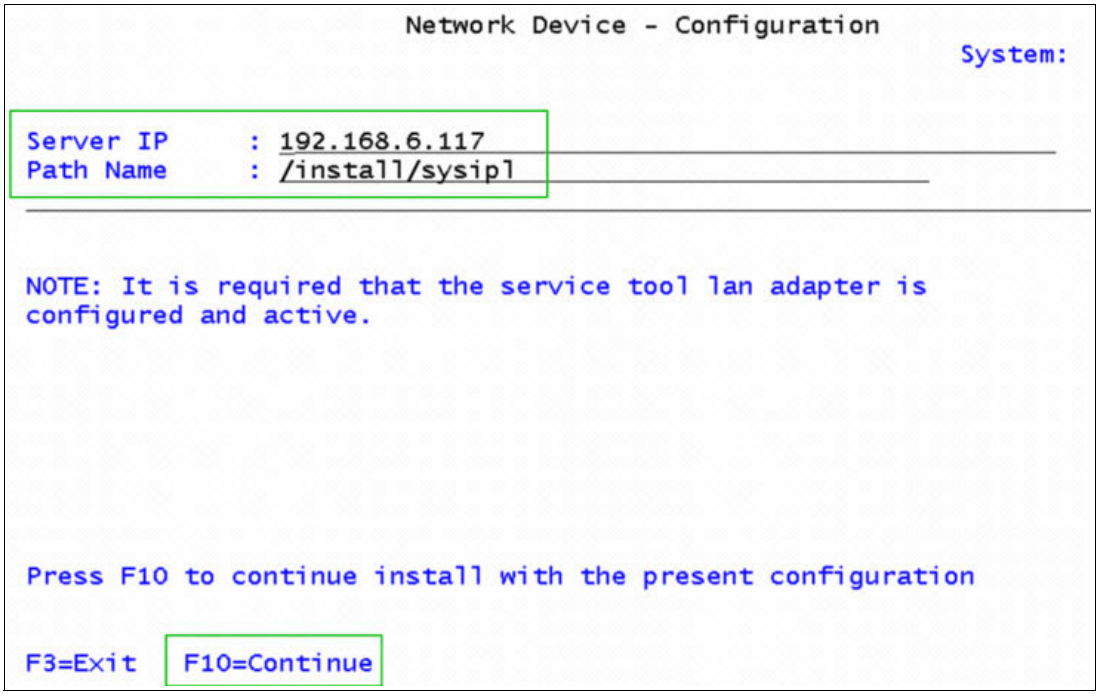

Figure 4-14 Configure the network device

6. Check whether the language group is correct. Then press **Enter** twice, and confirm.
7. Complete steps 19 - 28 at this IBM Documentation web page.

   After the first login at the IBM i main menu, the System name is changed.
8. At the command line, enter `GO LICPGM`.
9. Select **Option 10=Display Installed Licensed Programs**.

   The Base IBM i that is installed and the Library `QGPL` and `QUSRSYS` are `*BACKLEVEL`.
10. Recover the IBM BRMS product and associated libraries on the IBM i Power Virtual Server instance Client server by using NFS Server. On the Client Server Create Optical Device, run the **`CRTDEVOPT`** command. Then select **F4** to prompt and use the following settings:
    - Local internet address: *SRVLAN
    - Remote internet address: IP address of the IBM i VSI NFS Server
    - Network image directory: **‘**`/install/sysipl`**’**
11. Run the **`WRKCFGSTS *DEV INSTALL`** command. Then select **Option 1= Vary On for device INSTALL**.
12. Verify that you can access the remote image catalog by running the following command:

    `WRKIMGCLGE IMGCLG(*DEV) DEV(INSTALL)`
13. Verify the Catalog, Type, and Directory by running the **`/install/sysipl`** command.

    Verify the status of the volumes (`Mounted` or `Loaded`).

> **Tip:** Your IBM BRMS Recovery Report is available in the `QP1ARCY` file.

14. Starting with STEP004: Recover the IBM BRMS Product and Associated Libraries in your IBM BRMS Recovery Report, complete the following steps:

    a. Run the following command to change the QSYSOPR message queue to prevent messages that are not related to the recovery from interrupting the recovery process:

       CHGMSGQ MSGQ(QSYSOPR) DLVRY(*NOTIFY) SEV(99)

    b. Press **Enter**.

    c. Recover the libraries by specifying the saved item, the name of the stand-alone device, or media library that is used, and the volume identifiers and sequence numbers that are listed. For type *FULL, run the following command:

       RSTLIB SAVLIB(saved-item) DEV(device-name) VOL(volume-identifier) OPTFILE('')

Example 4-1 shows the QBRM, QMSE, and QUSRBRM saved item.

Example 4-1 QBRM, QMSE, and QUSRBRM saved items

```
Saved      Save    ----- ASP ------  Save     Save             Not    Sequence
Item       Type    Name      Number  Date     Time    Saved    Saved  Number
---------- ------- ---------- ----- -------- -------- -------- ------ ---------
QBRM       *FULL   *SYSBAS    00001  7/06/20 10:37:08     1184      0        48
    OPTFILE('/IBM BRMSOPTSAV22488012007061037061922BRMC/QBRM ')
QMSE       *FULL   *SYSBAS    00001  7/06/20 10:37:09        8      0        49
    OPTFILE('/IBM BRMSOPTSAV22488012007061037061922BRMC/QMSE ')
QUSRBRM    *FULL   *SYSBAS    00001  7/06/20 10:37:41      264      0        58
    OPTFILE('/IBM BRMSOPTSAV22488012007061037081922BRMC/QUSRBRM ')
```

In our example, run the commands that are shown in Example 4-2 to recover the IBM BRMS libraries.

Example 4-2 Restoring library to recover IBM BRMS libraries

```
RSTLIB SAVLIB(QBRM) DEV(INSTALL) VOL(QO6990) ENDOPT(*REWIND) OPTFILE('/IBM
BRMSOPTSAV22488012007061037061922BRMC/QBRM')
RSTLIB SAVLIB(QMSE) DEV(INSTALL) VOL(QO7898) ENDOPT(*REWIND) OPTFILE('/IBM
BRMSOPTSAV22488012007061037061922BRMC/QMSE')
RSTLIB SAVLIB(QUSRBRM) DEV(INSTALL) VOL(QO8807) ENDOPT(*REWIND) OPTFILE('/IBM
BRMSOPTSAV22488012007061037081922BRMC/QUSRBRM')
```

15. Complete Step 005 in your IBM BRMS Recovery Report.

    Recover this information for the IBM BRMS product to accurately guide you through the remaining recovery steps.

    Run the following command to recover the libraries that specify the saved item, name of the stand-alone device or media library that is used, and the volume identifiers and sequence numbers that are listed as shown in Example 4-3:

    RSTOBJ OBJ(*ALL) SAVLIB(saved-item) DEV(device-name) VOL(volume-identifier) OPTFILE('') MBROPT(*ALL) ALWOBJDIF(*COMPATIBLE)

Example 4-3 QUSRBRM saved item

```
Saved      Save    ----- ASP ------  Save     Save             Not    Sequence
Item       Type    Name      Number  Date     Time    Saved    Saved  Number
---------- ------- ---------- ----- -------- -------- -------- ------ ---------
QUSRBRM    *QBRM   *SYSBAS    00001  7/06/20 10:37:48       16      0        59
```

16. Press **Enter**.

    The information that is shown in Example 4-4 is used to restore the object.

Example 4-4 Sample to restore object

```
RSTOBJ OBJ(*ALL) SAVLIB(QUSRBRM) DEV(INSTALL) VOL(Q08807) MBROPT(*ALL) AL
OBJDIF(*COMPATIBLE) OPTFILE('/IBM BRMSOPTSAV22488012007061037081922BRMC/QUSRBRM')
```

17. Complete Step 006 in your IBM BRMS Recovery Report by running:

    ```
    INZBRM OPTION(*SETAUT) SETUSRBRM USER(QSECOFR) USAGE(*ADMIN)
    ```

18. Follow Step 009 in your IBM BRMS Recovery Report by running:

    ```
    INZBRM OPTION(*DEVICE)
    WRKDEVBRM
    ```

    Verify that you INSTALL device is included in the list.

19. Complete Step 010 in your IBM BRMS Recovery Report (see Example 4-5).

Example 4-5 Sample of a report that contains item from a cloud backup

```
-----  Attention  ------------------------------------------------
THIS REPORT CONTAINS ITEMS FROM A CLOUD BACKUP.
PLEASE RUN THE FOLLOWING PROGRAM CALL TO SET UP THE CLOUD VOLUMES:
CALL QBRM/Q1AOLD PARM('CLOUD     ' 'FIXDRVOL  ' 'Q06990' 'Q07898' 'Q08807'
'Q13692' 'Q32656')
```

> **Note:** You find your VOLUMES in the IBM BRMS Report `QP1A2RCY` file.

20. Run the following command to restore a current version of your user profiles:

    ```
    STRRCYBRM OPTION(*SYSTEM) ACTION(*RESTORE)
    ```

    Then press **Enter**.

> **Attention:** Press **F9** in the Select Recovery Items display to return to the Restore.
>
> Ensure that the tape device name or media library name is correct for the Device prompt and that the following prompts are specified:
>
> - ***SAVLIB** for Restore to library
> - ***SAVASP** for Auxiliary storage pool
> - ***YES** is specified for Create Parent Directories
>
> If you are recovering to a different system or logical partition, specify the following prompts:
>
> - ***ALL** for the Data base member option
> - ***COMPATIBLE** for the Allow object differences
> - ***NONE** for the System resource management
>
> Press **Enter** to return to the Select Recovery Items display.

Figure 4-15 shows a sample of **STRRCYBRM**.

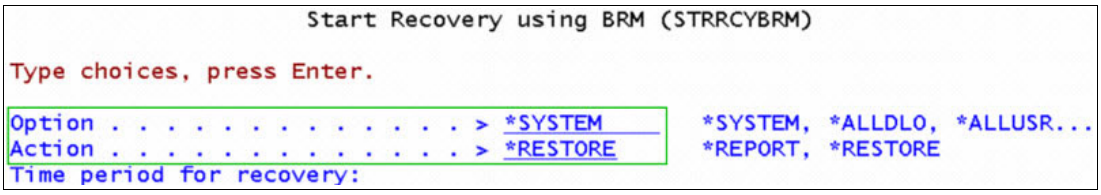

Figure 4-15 Sample of STRRCYBRM

21. Select **F9= Restore Command Defaults** and make the changes that are shown in Figure 4-16.

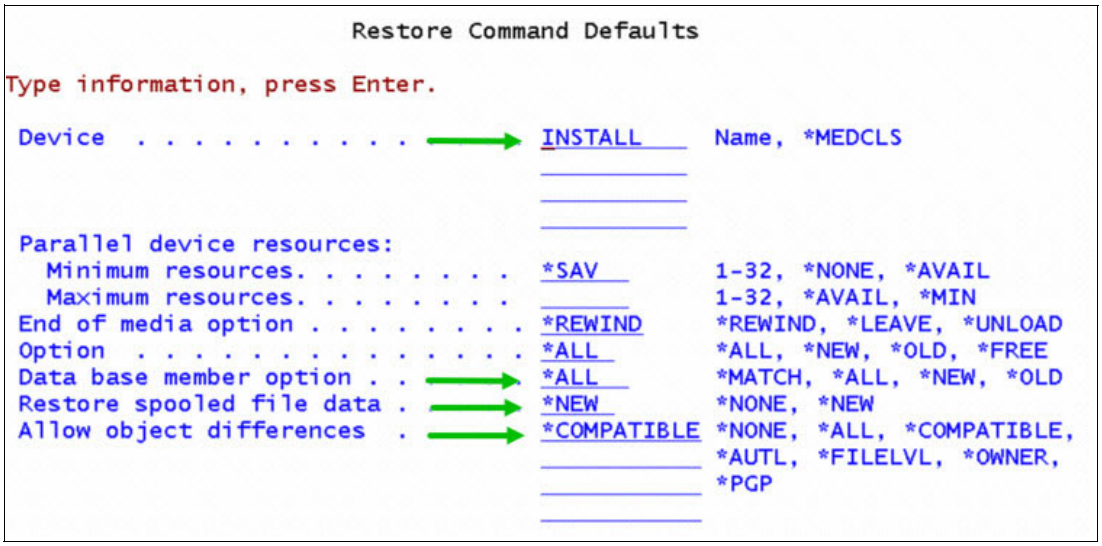

Figure 4-16 Sample to change parameters on restore commands

22. Page down and make the following changes:
   - System resource management: ***NONE** (when you restore to a different system)
   - Create parent directories: ***Yes**

   Press **Enter**.

23. Select **Option 1=Select "*SAVSECDTA"** and press **Enter**.

> **Note:** If you receive a message that reads: `Waiting for reply to message on message queue QSYSOPR`, select **SysReq** and **type 6** to display `QSYSOPR` system messages.
>
> If you are prompted to load volume `Qxxxxx` on INSTALL (`C G`), enter `C` to cancel and continue with the restore.

24. Follow Step 011 in your IBM BRMS Recovery Report and run the following command:

```
CHGUSRPRF USRPRF(QSECOFR) PASSWORD(new-password)
```

25. Follow Step 012 in your IBM BRMS Recovery Report. Select **Option 1=Select "*SAVCFG"** and press **Enter** (see Figure 4-17).

```
                              Select Recovery Items                    I922BRMC

                                                Select action  . . . . . : *ALL
                                                Select volume  . . . . . : ______
 Type options, press Enter.
   1=Select   4=Remove   5=Display   7=Specify object

    Saved        Save     Save    Save   Parallel Volume    File        Expire
Opt Item         Date     Time    Type   Devices  Serial    Sequence    Date
 1  *SAVCFG     7/06/20 10:36:12 *FULL           Q06990           46    7/27/20
 _  QSYS2       7/06/20 10:37:12 *FULL           Q07898           52    7/27/20
 _  QGPL        7/06/20 10:37:13 *FULL           Q07898 +         53    7/27/20
    QUSRSYS     7/06/20 10:37:08 *FULL           Q06990           47    7/27/20
```

Figure 4-17 Sample of SAVCFG

26. Follow Step 013 in your IBM BRMS Recovery Report. Select **Option 1=Select for all of the "Saved Item"** and press **Enter**.

27. Display the remaining objects during the restore (see Figure 4-18).

```
                              Select Recovery Items                    I922BRMC
.......................................................................................
: _                          Display Recovery Items                    I922BRMC  :
:                                                                      17:57:35  :
:    Remaining items . . . . :              8                                   :
:    Remaining objects . . . :          2,511                                   :
:    Remaining size  . . . . :     4,967.2540 M    98.7 %                       :
:                                                                                :
: Saved                              Save  Volume                 Exp    Objects :
: Item          Date      Time       Type  Serial   File Seq      Date    Saved  :
: QGPL          7/06/20 10:37:13     *FULL Q07898 +         53  7/27/20    165  :
: QUSRSYS       7/06/20 10:37:08     *FULL Q06990           47  7/27/20   1680  :
: QICC          7/06/20 10:37:12     *FULL Q07898           50  7/27/20    108  :
: QUSRICC       7/06/20 10:37:12     *FULL Q07898           51  7/27/20     36  :
: QIWS          7/06/20 10:37:13     *FULL Q08807           55  7/27/20    138  :
: QSYSDIR       7/06/20 10:37:13     *FULL Q08807           56  7/27/20     73  :
: QTCP          7/06/20 10:37:13     *FULL Q08807           54  7/27/20    281  :
: QCLDIPL       7/06/20 10:37:39     *FULL Q08807           57  7/27/20     30  :
:                                                                                :
:                                                                                :
:  Press ATTN key to cancel recovery after current item completes.               :
:                                                                                :
:.......................................................................................
```

Figure 4-18 Sample of display recovery items.

28. Work with TCP/IP interfaces and add internet addresses from target IBM i.

> **Important:** Before you began the scratch installation on the (target) IBM i VSI, document all your network information.

Re-create the following network information to complete the restore:

– *Work with TCP/IP Interfaces*.
– Command **CFGTCP**.

Select **Option 1= Work with TCP/IP interfaces**.

Select **Option 1=Add (**enter the IP address you used to connect to IBM Cloud Object Storage.

Note: Use the IP address information that you documented from the target IBM i VM that was created in the previous step. The following information is needed:

- Internet address: `IP address x.x.x.x`
- Subnet mask: `255.255.255.x`
- Line description: Use one of the three `CLOUDINITx` line descriptions that were restored or use the same description that was documented.

29. Work with TCP/IP Interfaces to start the Interface by running the **`CFGTCP`** command:

    Select **Option 1= Work with TCP/IP interfaces**.

    Select **Option 9= Start**.

30. Find the Resource URI that was used for IBM COS (it is where the volumes are stored), as shown in Figure 4-19:

    a. Work with IBM Cloud Storage Resources.
    b. Run the **`WRKCFGICC`** command and press **Enter**.
    c. Select **Option 5=Display**.

```
                        Display ICC AWS S3 Resource

Resource name  . . . . . :  TOR1CLD
Resource type  . . . . . :  AWSS3
Resource description . . . . . :  Cloud Object Storage Toronto 1

Use compression  . . . . . :  *NO
Use encryption . . . . . :  *NO

Bucket . . . . . :  brms-bucket-backupvol

Resource URI . . . . . :  s3.us-east.cloud-object-storage.appdomain.clo
ud
```

Figure 4-19 Sample of URI

31. Ping the Resource URI to verify the connection.

    If the ping was successful, a working network to the COS exists.

    If the ping was unsuccessful, END the IP Interface and Vary off the Ethernet Line. Change the Ethernet line description to another CMNxx Resource that is not yet used and attempt to ping again.

32. Create a virtual tape device by running the **`CRTDEVTAP DEVD(TOR1CLDTAP) RSRCNAME(*VRT)`** command.

33. Vary on the virtual tape device by running the following command:

    ```
    VRYCFG CFGOBJ(TOR1CLDTAP) CFGTYPE(*DEV) STATUS(*ON)
    ```

    At this point in the recovery, physical media is no longer required because media is downloaded from the cloud.

    Because the system is in restricted state, TCP/IP must be started so that IBM BRMS can transfer the media.

Attention: Consider the following points:

- If done in a previous step, this command is not required to be run again because TCP/IP must be started so that IBM BRMS can transfer media that is required by a recovery from the cloud. To continue the recovery in restricted state, run the following commands:

  ```
  STRTCP STRSVR(*NO) STRIFC(*NO) STRPTPPRF(*NO) STRIP6(*YES)
  STRTCPIFC INTNETADR('nnn.nnn.nnn.nnn')
  ```

  where *'nnn.nnn.nnn.nnn'* is the Internet address of the recovery system. Also, ***nnn*** is a decimal number 0 - 255.

- The use of restricted state TCP/IP requires that 5770SS1 option 3 was installed when the backups were run. Otherwise, start all subsystems by running the following commands:

  ```
  STRSBS SBSD(QCTL)
  STRTCP
  ```

34. Move volumes from IBM COS to the TARGET IBM i VSI Client:

    a. Run the **WRKMEDBRM** command and press **Enter**.

       Find your volumes that are listed in your IBM BRMS Recovery Report.

       If your volume has a plus (+) to the right, it is part of a serial set.

    b. Select **Option 6=Work with serial set**, as shown in Figure 4-20.

    Notice that the volumes are in the IBM COS location.

```
                                Work With Media
                                                                System:   I922BRM
Position to  . . . . .  ______        Starting characters

Type options, press Enter.
  1=Add   2=Change   4=Remove   5=Display   6=Work with serial set   7=Expire
  8=Move   9=Remove volume error status     10=Reinitialize ...

     Volume            Creation Expiration            Move     Media      Dup
Opt  Serial    Status   Date      Date     Location   Date     Class      Sts
__   ______
__   Q09148     *EXP  07/06/20  *NONE     *HOME      *NONE    QCLDVRTTAP
__   Q10030 +   *ACT  06/25/20  07/16/20  TOR1CLD    06/25/20 QCLDVRTOPT
__   Q10057     *EXP  07/06/20  *NONE     *HOME      *NONE    QCLDVRTTAP
__   Q10966     *EXP  07/06/20  *NONE     *HOME      *NONE    QCLDVRTTAP
__   Q11874     *EXP  07/06/20  *NONE     *HOME      *NONE    QCLDVRTTAP
__   Q12783     *EXP  07/06/20  *NONE     *HOME      *NONE    QCLDVRTTAP
6_   Q13692 +   *ACT  07/06/20  07/27/20  TOR1CLD    07/09/20 QCLDVRTTAP
                                                                        More..
```

Figure 4-20 Sample to work with serial set of volumes

35. Select **Option 8=Move**, as shown in Figure 4-21 (all the volumes at the same time) and press **Enter**.

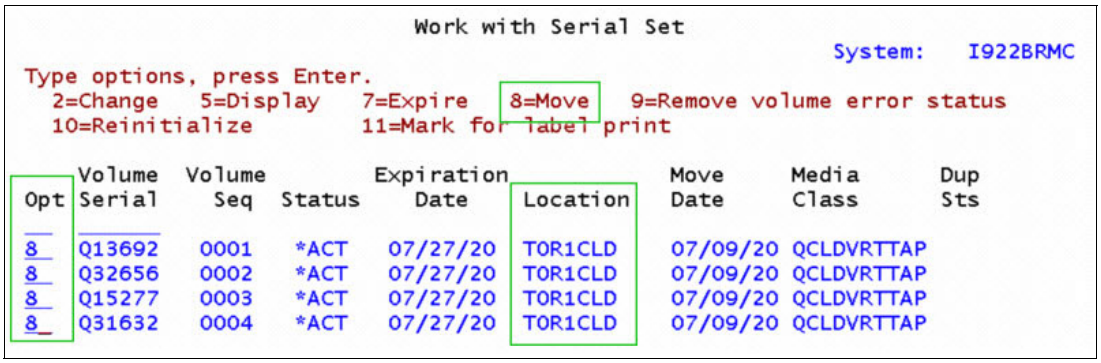

Figure 4-21 Sample of volumes to move from cloud

36. Change the Storage location to ***HOME** and then press **Enter**, as shown in Figure 4-22.

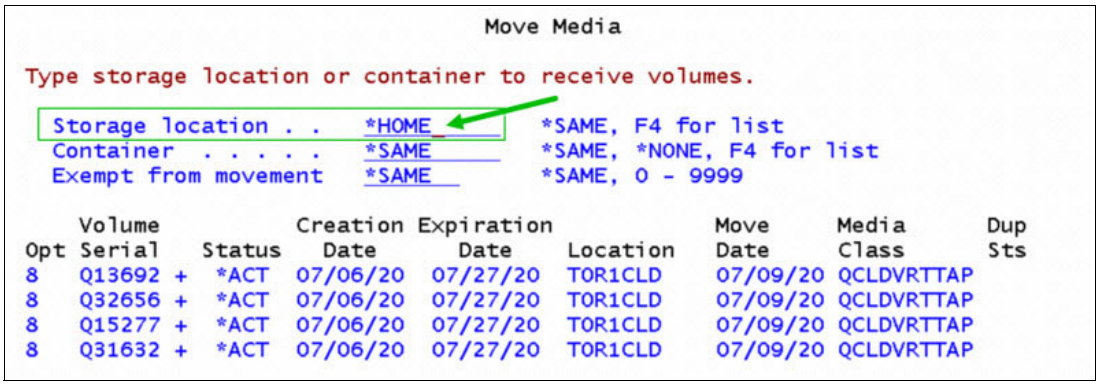

Figure 4-22 Sample to change storage location to HOME

37. Check the IBM Cloud Storage Status transfer by running the **WRKSTSICC STATUS(*ALL)** command:

    Check that the status is “Active”.

    The operation being run; for example, FRMCLD is a copy from the cloud operation. Oper = FRMCLD (From Cloud).

    After all of the jobs have a status of `Success`, you can continue with the IBM BRMS restore by running the **WRKSTSICC STATUS(*ALL)** command.

38. Follow Step 014 in your IBM BRMS Recovery Report and the following command:

    ```
    INZBRM OPTION(*DEVICE)
    WRKDEVBRM
    ```

    You see the Virtual Tape Device *“TOR1CLDTAP *VRTTAP”*.

39. Recover IBM Product Libraries on the IBM i Power Systems Virtual Server Instance Client Server (target) by using IBM COS:

    Follow Step 017 in your IBM BRMS Recovery Report and run the following command:

    ```
    STRRCYBRM OPTION(*IBM) ACTION(*RESTORE)
    ```

Note: The Restore Command default settings are used to specify the correct Device parameter and change the Create parent directories prompt back to ***NO**.

Attention: Press **F9** on the Select Recovery Items display to go to the Restore. The following Command Defaults display:

- Ensure that *MEDCLS is specified for the Device prompt.
- Ensure that *NO is specified for the Create parent directories prompt.

Press **Enter** to return to the Select Recovery Items display.

40. Review the list of Recovery Items, as shown in Figure 4-23, and Remove any that were restored. Select **Option 4=Remove** and press **Enter**.

    You can also see that the Volume Serial is from the optical media.

    After you remove the items, they are removed from the list.

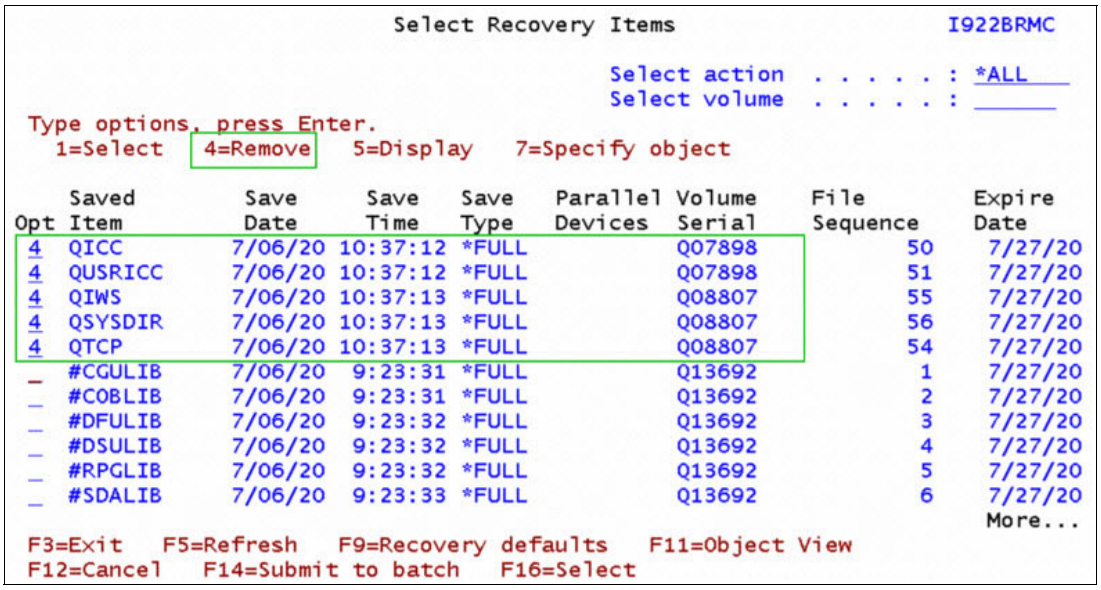

Figure 4-23 Sample to remove items already used

41. Change the Recovery Defaults by selecting **F9=Recovery defaults**:

    - Device: TOR1CLDTAP.
    - Option: **`*ALL`**.
    - Data base member option: **`*ALL`**.
    - Restore spooled file data: **`*NEW`**.
    - Allow object differences: **`*COMPATIBLE`**.

      Page down.

    - System resource management: **`*NONE`**.
    - Create parent directories: **`*NO`**.

    Press **Enter** to return to Select Recovery Items.

42. Complete the following steps:

    a. Select the saved items.
    b. Review the list again.
    c. Select **Option 1=Select for each item** or **F16= Select (this will select all the items)**.
    d. Press **Enter** to recover the saved items.

43. Recover User Libraries on the IBM i Power Systems Virtual Server Instance Client Server (target) by using IBM COS:

    Follow Step 018 in your IBM BRMS Recovery Report by running the **STRRCYBRM OPTION(*ALLUSR) ACTION(*RESTORE)** command.

44. Review the list of Recovery Items and remove any that were restored.

    Select **Option** 4=Remove and press **Enter**.

    You can also see that the Volume Serial is from the optical media.

    Press **F11=Object View** (shows you which Control Group created the saved item).

    Remove any items that were created by **QCLDBIPL01** (see Figure 4-24).

    After you remove the items, they are removed from the list.

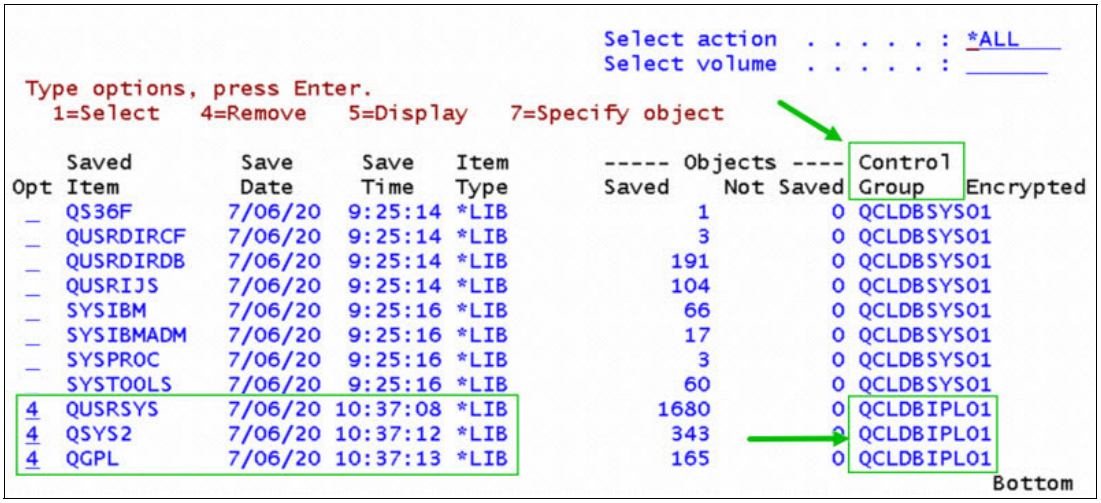

Figure 4-24 Sample of item to be removed created by QCLDBIPL01

45. Select the saved items:

    a. Review the list again.

    b. Select **Option 1=Select for each item** or **F16= Select (this will select all the items for you)**.

    c. Press **Enter** to recover the saved items.

46. To recover Document Library Objects on the IBM i Power Systems Virtual Server Instance Client Server (TARGET) by using IBM COS, follow Step 019 in your IBM BRMS Recovery Report:

    ```
    STRRCYBRM OPTION(*ALLDLO) ACTION(*RESTORE)
    ```

47. Select the saved item, as shown in Figure 4-25.

    Select **Option 1=Select**.

    Press **Enter** to recover the saved items.



*Figure 4-25 Sample of *ALLDLO saved item*

48. Recover Directories and Files on the IBM i Power Systems Virtual Server Instance Client Server (TARGET) by using IBM COS.

    Follow *Step 020* in your IBM BRMS Recovery Report:

    ```
    STRRCYBRM OPTION(*LNKLIST) ACTION(*RESTORE)
    ```

49. Review the list of Recovery Items and remove any that were restored:

    a. Select **Option** 4=Remove and press Enter.

       You can also see that the Volume Serial is from the optical media.

    b. Press **F11=Object View** (shows you which Control Group created the saved item).

    c. Remove any items that were created by **QCLDBIPL01**.

       After you remove the items, they are removed from the list.

50. Select ***LINK** the saved items:

    a. Select **Option 1=Select**.
    b. Press **Enter** to recover the saved items.

51. Follow Step 025 in your IBM BRMS Recovery Report:

    ```
    UPDPTFINF
    ```

52. Follow Step 026 in your IBM BRMS Recovery Report:

    ```
    RSTAUT USRPRF(*ALL)
    ```

53. Follow Step 027 in your IBM BRMS Recovery Report (restores your System Values):

    ```
    UPDSYSINF LIB(QUSRSYS) TYPE(*SYSVAL)
    ```

54. Follow `Step 030` in your IBM BRMS Recovery Report:

    ```
    DSPJOBLOG JOB(*) OUTPUT(*PRINT)
    ```

55. Change IPL Attributes:

    ```
    CHGIPLA STRRSTD(*YES)
    ```

    After the IPL, you can verify the system.

56. Change System Value for QIPLTYPE:

    ```
    WRKSYSVAL QIPLTYPE
    ```

    Select **Option 2=Change**.

    Select **0=Unattended IPL**.

57. Follow Step 031 in your IBM BRMS Recovery Report:

```
PWRDWNSYS OPTION(*IMMED) RESTART(*YES)
```

Press **F16=Confirm**.

58. Complete the following steps at the IPL or Install the System menu:

a. Select **Option 3=Use Dedicated Service Tools (DST)**.
b. Sign on to Dedicated Service Tools (DST).
c. Select **Option 7=Start a service tool**.
d. Select **Option 7=Operator panel functions**.

In IPL mode:

a. Select **Option** 2=Normal.
b. Press **F8** to set the IPL attributes and restart the system.
c. Press **Enter** to confirm.

> **Note:** Check that the IPL source is set to 2 and the IPL mode is set to 2.

After the IPL, the 5250 is ready to sign on.

Figure 4-26 shows a summary of your IBM i NFS client through an IBM i NFS Server by using a backup that was created by using IBM BRMS that is on IBM COS.

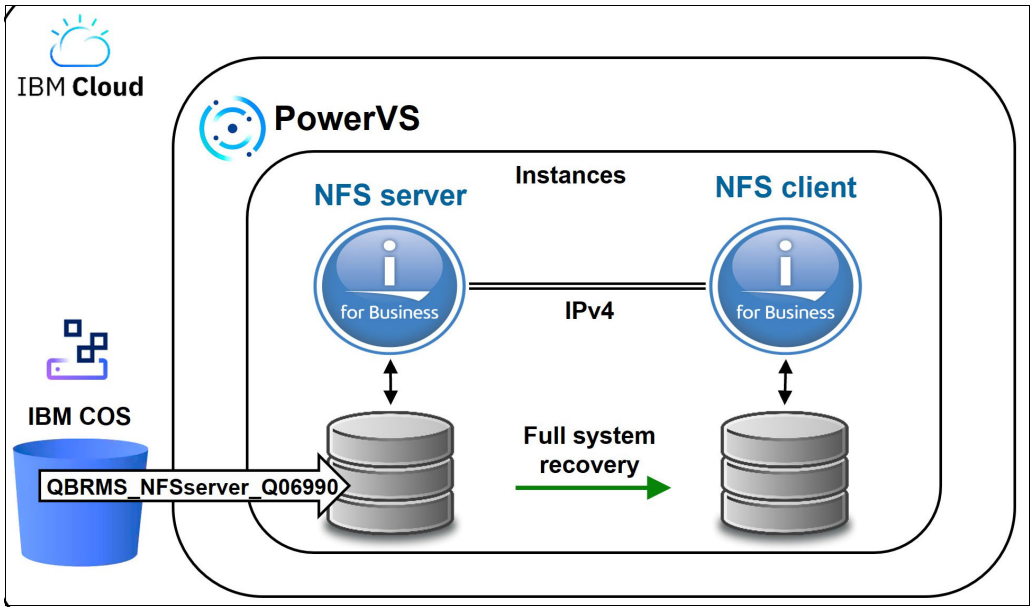

Figure 4-26 Full system recovery from the cloud by using IBM i

4.3 Creating object-level backups

This section describes how to create object-level backups.

4.3.1 Introduction

This method is the most popular backup and recovery method that is used by IBM i customers at the time of this writing.

Similar to the on-premises environment, day-to-day daily backups and full system backups can be done on IBM Power Systems Virtual Server by using native commands and CL scripts or **GO SAVE**, **GO BACKUP**, and **GO RESTORE** menus to run and schedule backups.

The following commands are used:

- SAVLIB
- SAVOBJ
- SAVCHGOBJ
- SAVDLO
- SAV
- SAVSECDTA
- SAVCFG

As a target device, use the Image Catalog facility with virtual optical or virtual tape media, and save files or the new offering from IBM (see 4.4, “FalconStor overview” on page 165).

This virtual media that is produced must be stored and archived. IBM COS is the native solution to save backups with almost endless capacity.

Communicate with IBM COS by using the IBM i Cloud Storage Solutions software offering, or the use of the open source capacity in the IBM i operating system.

4.3.2 Copying files to the cloud

Use the **CPYTOCLD** command to copy files from the IFS to a cloud resource. Consider the following points:

- A cloud resource must be created.
- An IBM i Read **(*R)** authority must be available on the file that you copy, and Execute **(*X)** authority on all directories in the path leading to the file. For example, if you copy the `/home/user/jdoe/file.txt` file, you must have *Execute* authority on the `home`, `user`, and `jdoe` directories, and then `Readauthority` on `file.txt`.

The copy command cannot use Cloud Storage Solutions to work with files in the `/QSYS.LIB` file system. The size of files that you can copy to a resource is determined by the cloud service provider.

When you copy files to or from the cloud, the operation is run asynchronously in its own batch job instead of in the same job as the command.

When you copy files asynchronously, you do not have to wait for large files to finish copying before other commands can be run. Also, you can use the IBM i facilities to work with asynchronous jobs to, for example, schedule when the jobs run.

In the Submit to batch field, you can specify instead that the copy operation be run in the same job as the command. This command is *not* thread safe.

Complete the following steps to copy a file to the cloud:

1. From the IBM i command line, enter **CPYTOCLD** and press **F4**.
2. Complete the required fields as listed in Table 4-1.

Table 4-1 CPYTOCLD required fields

Field	Action
Resource Name	Enter the name of a resource.

Field	Action
Submit to Batch	enter *NO to run the copy operation in the same job as the command. Leave the *YES value to run the copy operation in its own job.
Local file name	Enter the IFS path and the name of the file to copy; for example:/home/user/jdoe/file.txt. The path must begin with a forward slash (/) and is not case-sensitive.
Cloud file name	Enter a path and name for the cloud copy of the file; for example, dir1/dir2/file.txt. The path is created if it does not exist. When you specify this path, do not include the container, bucket, or root directory that is defined in the resource. Cloud Storage Solutions combines that directory with this path to create the full path in the cloud. The container, bucket, or root directory that is defined in the resource must exist before you copy files to the cloud. When overwriting a file, usually the directory and file name are case-sensitive. If you overwrite a file on an IBM i FTP computer, the directory and file name are not case-sensitive unless they are on the /QOpensys file system.

3. Press **Enter.**

4.3.3 Copying files from the cloud

Run the **CPYFRMCLD** command to copy files from a cloud resource to the IFS.

You must have execute **(*X)** authority on all directories in the path to which you copy the file, and write **(*W)** authority on the last directory in the path. I

f the file was copied before and exists in the path, the user must have Write access to it. For example, to copy **file.txt** to **/home/user/jdoe**, you must have *Execute* authority on the home, user, and jdoe directories, and Write authority on jdoe. If file.txt is available, you must have Write authority on it.

You cannot use Cloud Storage Solutions to work with files in the /QSYS.LIB file system. If you copy a file from an FTP cloud server to the IBM i computer, and that file was not originally copied to the FTP server by using Cloud Storage Solutions, Cloud Storage Solutions assigns the file a coded character set identifier (CCSID) of 65535. A CCSID of 65535 means that the operating system treats the file as binary data and it is unreadable in an editor.

If you copy a file from an Amazon S3 or IBM COS cloud server to the IBM i computer and that file was not originally copied to the cloud server by using Cloud Storage Solutions, Cloud Storage Solutions reads the data and from it assigns the file a coded character set identifier (CCSID) of 1208 (UTF-8) if it is text, or 65535 if it is binary.

When files are copied to or from the cloud, the operation is run asynchronously, so the copying is its own batch job instead of in the same job as the command. When files are copied asynchronously, you do not have to wait for large files to finish copying before running other commands.

Also, the IBM i facilities can be used to work with asynchronous jobs by, for example, scheduling when the jobs run. In the Submit to batch field, specify instead that the copy operation be run in the same job as the command. This command is *not* thread safe.

Complete the following steps:

1. From the IBM i command line, enter `CPYFRMCLD` and press **F4**.
2. Complete the required fields as listed in Table 4-2.

Table 4-2 CPYFRMCLD required fields

Field	Action
Resource Name	Enter the name of a resource.
Submit to Batch	Enter `*NO` to run the copy operation in the same job as the command. Leave the `*YES` value to run the copy operation in its own job.
Cloud file name	Enter the cloud path and the name of the file to be copied; for example: `dir1/dir2/file.txt`. Do not include the container or bucket that is defined in the resource. The cloud file name path is appended to those directories to construct the full path for the cloud copy of the file. Usually, the directory and file name are case-sensitive. The directory and file name are not case-sensitive with IBM i FTP resources unless they are on the `/QOpensys` file system.
Local file name	Enter the IFS path and file name of the file being copied, for example:`/home/user/jdoe/file.txt`. The path must begin with a forward slash (/). Path and file names are not case-sensitive. Path directories are created if they do not exist locally. You can enter a local file name that is different from the cloud file name.

3. Press **Enter.**

4.3.4 Backup and restore by using save files

When small to medium amounts of data must be moved between IBM i systems, the most common file transfer method is the use of save files with FTP, which has the following restrictions and features:

- Data from only one library can be saved on each save file.
- SAVSYS cannot be used on a save file.
- “D” IPL from a Save File cannot be performed to restore License Machine Code (LIC or SLIC).
- The maximum size for a save file on IBM i V7R1 and later is 2 TB.
- An empty save file must be created by using the target version operating system.
- Save files cannot be created on IFS or a mounted file system. However, they can be copied to IFS or mounted file system after the data is saved.
- Data can be compressed during or after the backup finishes by using the **JAR** command or Open Source tools.
- Data in save files can be uploaded to physical tapes or virtual tapes by using **SAVSAVFDTA** command.

- Parallel backups can be performed, which save to multiple save files simultaneously.
- When a large library must be saved to a single file, some PASE or QSHELL commands can be used to split and compress the data into several components.

Sample scenario

In this example, one library and the security information are saved in two different Save Files. The data is compressed and then uploaded to IBM COS, assuming communications exist.

Note: For more information about how to connect from IBM Power Systems Virtual Server and IBM Cloud Object Storage, see *IBM i Migration to Cloud with IBM Power Systems Virtual Server*.

Complete the following steps:

1. Create an empty Save Files for backups, libraries, and security data, as shown in Example 4-6.

Example 4-6 Steps to backup on save file

```
CRTLIB BACKUPSAV
CRTSAVF BACKUPSAV/MYLIB01
CRTSAVF BACKUPSAV/SAVSECDTA
```

2. Backup a library to the Save File, as shown in Example 4-7.

Example 4-7 Save library routine

```
SAVLIB LIB(MYLIB01) DEV(*SAVF) SAVF(BACKUPSAV/MYLIB01)
```

3. Backup security data as shown in Example 4-8.

Example 4-8 Save security data routine

```
SAVSECDTA DEV(*SAVF) SAVF(BACKUPSAV/SAVSECDTA)
```

4. Update the YUM environment and install some tools, as shown in Example 4-9.

Example 4-9 Update the YUM environment

```
CALL QP2TERM
TERM=aixterm
export TERM
PATH=$PATH:/QOpenSys/pkgs/bin
export PATH
yum -y update
yum -y install pigz zstd gzip gunzip python3.9 python39-pip readline curl
pip3.9 install awscli
```

5. Configure **awscli** to communicate with IBM COS by using **aws configure**.
6. Use your credentials from your IBM COS bucket to complete the configuration.

   For more information about how to create your IBM COS resource, bucket, and credentials and configure AWS CLI, see Using the AWS CLI.
7. Compress the files and upload to IBM COS, as shown in Example 4-10.

Example 4-10 Saving and compressing data to COS bucket

```
cd /QSYS.LIB/BACKUPSAV.LIB
```

```
cat MYLIB01.FILE | pigz -9 -p40 -c | aws
--endpoint-url=https://s3.eu-gb.cloud-object-storage.appdomain.cloud s3 cp -
s3://ibmi-backup/MYLIB01.gz

cat SAVSECDTA.FILE | pigz -9 -p40 -c | aws
--endpoint-url=https://s3.eu-gb.cloud-object-storage.appdomain.cloud s3 cp -
s3://ibmi-backup/SAVSECDTA.gz
```

8. List the backup on IBM Cloud Object Storage, as shown in Example 4-11.

Example 4-11 List the backup from IBM Cloud Object Storage

```
aws --endpoint-url=https://s3.eu-gb.cloud-object-storage.appdomain.cloud s3 ls s3://ibmi-backup/
```

4.3.5 Creating backups by using image catalog and IBM COS

The image catalog is a feature off IBM i as an image device starting with V5R1. Starting with V5R3, you can save to your disk by using optical images. By using V5R4 and later, you can create Virtual Tape images.

This useful feature can be used to install PTFs, perform system upgrades, and save the data to disk, which avoids save file limitations to only one library.

Media can be saved to a virtual tape image catalog and duplicated to a physical tape, or from a virtual optical image catalog to a CD or DVD (images must be created with the correct size and format).

For more information about maximum capacities, see Image Catalog and Virtual Media - Maximum Capacities and Limitations.

When you back up data on IBM Cloud Power Virtual Server instances, you can use both of these image catalog types. However, consider the following environment characteristics when restoring:

- A full system restore or recover cannot be made from SAVSYS by using tape images.
- Because media is transferred to IBM Cloud Object Storage or a gateway appliance, consider the number of image files and their size when transferring and restoring data.
- Large Image files can slow the restore process when only one file must be restored.
- Small Image files can be a bad idea when many terabytes must be backed up because of image catalog limits and the number of files that are created.
- Managing multiple backup sets can be complicated. Consider the use of IBM BRMS to create backups and to manage your image catalogs.
- Whenever possible, use the IBM Cloud internal network to avoid extra charges.
- The following commands can be used to manage the image catalog:
  - **WRKIMGCLG**. Allows the user to work with a list of image catalogs.
  - **CRTIMGCLG**. Creates an image catalog object (*IMGCLG) in the library QUSRSYS and associates the image catalog with a target directory.
  - **CHGIMGCLG**. Changes the attributes of an image catalog.
  - **DLTIMGCLG**. Deletes an image catalog.
  - **ADDIMGCLGE**. Creates a virtual image in the image catalog directory.

- **CHGIMGCLGE.** Changes the attributes of a virtual tape volume.
- **RMVIMGCLGE**. Removes a virtual volume from an image catalog and optionally deletes the virtual volume.
- **RTVIMGCLG.** Used in a CL procedure to retrieve the name of the image catalog that is loaded in a virtual device or used to retrieve the name of the virtual device where an image catalog is loaded.
- **STRNETINS**. Starts a network installation from the NFS server.
- **LODIMGCLGE**. Changes the status of a virtual tape or optical volume within an image catalog.
- **LODIMGCLG**. Loads an image catalog on a virtual tape or optical device to make the virtual volumes accessible by the device.
- **WRKIMGCLGE**. Works with the images in the specified image catalog.
- **VFYIMGCLG**. Used to work with the image catalogs that are on the system.

Using image catalogs

Complete the following steps:

1. Create an optical image catalog and virtual media, as shown in Example 4-12.

Example 4-12 Create image catalog and virtual media

```
CRTIMGCLG IMGCLG(OBACKUP) DIR('/IMGCLG/OBACKUP/') TEXT('Optical Image Catalog -
OBACKUP')

ADDIMGCLGE IMGCLG(OBACKUP) FROMFILE(*new) TOFILE(OBKP01) IMGSIZ(20480)
TEXT('Optical Backup Image 1 - 20GB') VOLNAM(OBKP01) NEWOWNID(IBMI01)
```

2. Create a virtual optical device, make it available, load it with the image catalog and initialize, as shown in Example 4-13.

Example 4-13 Create a virtual optical device

```
CRTDEVOPT DEVD(VOPTBK) RSRCNAME(*VRT) LCLINTNETA(*N) TEXT('Virtual Optical
Device for Backup')

VRYCFG CFGOBJ(VOPTBK) CFGTYPE(*DEV) STATUS(*ON)

LODIMGCLG IMGCLG(OBACKUP) DEV(VOPTBK)

INZOPT VOL(*MOUNTED) NEWVOL(OBKP01) DEV(VOPTBK) CHECK(*NO)
```

3. Start the use of the virtual device and backup to the new virtual media, as shown in Example 4-14.

Example 4-14 Save security data to the virtual optical device

```
SAVSECDTA DEV(VOPTBK)
```

When the virtual optical media must be transferred to IBM COS, use IBM Cloud Storage Solutions for i or the same procedure, as shown in Example 4-15.

Example 4-15 Upload file with IBM Cloud Storage Solutions for i

```
CPYTOCLD RESOURCE(ICOSO1) LOCALFILE('/IMGCLG/OBACKUP/OBKPO1') CLOUDFILE(OBACKUPO1)
```

The file is transferred to resource `ICOSO1`.

Using Example 4-15 with the configured AWS CLI, compress the file in gzip format, upload the file and list the contents to verify the correct file transfer, as shown in Example 4-16.

Example 4-16 Transferring file to IBM COS using PASE and AWSCLI

```
CALL QP2TERM

### Add environment variables ###
cd /IMGCLG/OBACKUP
TERM=aixterm
PATH=$PATH:/QOpenSys/pkgs/bin
export TERM
export PATH

### Compress and upload file ###
cat OBKPO1 | pigz -9 -p40 -c | aws
--endpoint-url=https://s3.eu-gb.cloud-object-storage.appdomain.cloud s3 cp -
s3://ibmi-backup/OBKPO1.gz

### List content on IBM COS bucket ###

aws --endpoint-url=https://s3.eu-gb.cloud-object-storage.appdomain.cloud s3 ls s3://ibmi-backup/
```

TIP: Use QSHELL and include **`awscli`** in the CLs to upload data to IBM COS. Run the following command from QCMD:

```
SBMJOB CMD(QSH CMD('/QOpenSys/usr/bin/sh -c  "cd /IMGCLG/OBACKUP ;
PATH=$PATH:/QOpenSys/pkgs/bin;export PATH;cat OBKPO1 | pigz -9 -p40 -c | aws
--endpoint-url=https://s3.eu-gb.cloud-object-storage.appdomain.cloud s3 cp -
s3://ibmi-backup/OBKPO1.gz" ')) JOB(S3CP)
```

4.3.6 Saving and restoring IBM i objects to IBM COS

This section includes an introduction to IBM COS and describes how to copy a save file directly from your IBM i partition to your own bucket in IBM COS.

This process can be useful not just as a backup solution, but also as a migration strategy to move a few applications and many small databases to IBM Power Virtual Server for your own PoC or testing.

Therefore, you can move a save file directly from IBM i by using the IBM Cloud Storage Solutions LPP. Your bucket includes public and private endpoints that can be used whether your source partition is within the IBM Cloud, or is external.

Note: The use of the private endpoints requires Direct Link, which is why this example uses the public interface to the IBM COS Bucket.

On IBM i, and in this example, IBM CONTENT COLLECTOR use the IBM Cloud (S3/COS) Connector (5733-ICC). This Copy From Cloud (**CPYFRMCLD**) command, and Copy to Cloud (**CPYTOCLD**) command are 5733-ICC CL commands that copy files between IBM COS and the IFS. These IFS files can be used with image catalogs to automate the process (as does IBM BRMS).

IBM BRMS is fully integrated with IBM Cloud Storage and can help in managing terabytes of data and thousands of objects in an HA/DR plan.

The following overall process is used in our example:

1. Save data to a save file.
2. Copy that save file from IBM i to IBM COS.
3. Verify the bucket contents.
4. Drop the file to the IBM COS bucket.

Backing up IBM i object to IBM COS

Complete the following steps:

1. Create a database to save:

```
RUNSQL
CALL QSYS.CREATE_SQL_SAMPLE('KMSAMPLE')
WRKLIB KMSAMPLE
```

2. Save that library to a save file:

```
CRTSAVF FILE(QGPL/KMSAMPLE)SAVLIB LIB(KMSAMPLE) DEV(*SAVF) SAVF(QGPL/KMSAMPLE)
```

3. Create your IBM COS Resource by using the **CRTS3RICC** command.

   Enter a Resource Name, the Access key ID, and the secret access key that was captured when you created your bucket. Press **Enter** and continue to complete other information (see Figure 4-27).

```
                    Create ICC AWS S3 Resource (CRTS3RICC)

 Type choices, press Enter.

 Resource name  . . . . . . . . .   COSRSC        Name
 Resource description . . . . . .   COSRSC
 Access key id  . . . . . . . . .   bea445a4d41f40d597ba859bbb3401d6

 Secret access key  . . . . . . .   ef786ab2c173f24f43e14cd75f0b4ab35526358788f0
0e49

 Use compression  . . . . . . . .   *NO           *NO, *YES
 Use encryption . . . . . . . . .   *NO           *NO, *YES
```

Figure 4-27 Create an IBM COS resource by way of CRTS3RICC

4. Enter your Bucket Name and your Resource URI, as shown in Figure 4-28. Press **Enter**.

```
 Bucket . . . . . . . . . . . . . > 'cos-bucket-ibmi01'

 Resource URI . . . . . . . . . . > 's3.us-east.cloud-object-storage.appdomain.c
loud'
```

Figure 4-28 Bucket name and resource URI

The resource is created successfully, and the following message received:

```
Resource COSRSC was created for resource type AWSS3.
```

5. From the console, run the **ENDTCP** command. The console closes.
6. Open a new console and run the following commands:

```
STRTCP
STRTCPSVR *TELNET
STRTCPSVR *SSHD
```

> **Note:** You must restart your SSH tunnel connections.

7. Start the COS Subsystem:

```
STRSBS SBSD(QICC/QICCSBS)
```

8. Copy the file to your IBM COS:

```
CPYTOCLD RESOURCE(COSRSC) ASYNC(*NO)
LOCALFILE('/QSYS.LIB/QGPL.LIB/KMSAMPLE.FILE') CLOUDFILE(KMSAMPLE)
```

The following message is displayed:

```
File copied.
```

9. Switch back to the bucket, and verify the file, as shown in Figure 4-29. Browse to **Buckets** and **Objects**.



Figure 4-29 IBM i object on IBM COS

Restoring IBM i object from IBM COS

This section describes how to restore from IBM COS. Combined with 4.3.6, “Saving and restoring IBM i objects to IBM COS” on page 157, the restore defines an application migration strategy.

You save your applications and data into a save file, move the file to IBM COS and restore the file from IBM COS. As a prerequisite, complete the steps that are described in “Backing up IBM i object to IBM COS” on page 158.

In this example, assume that a SAVLIB was performed of a business database and objects (ACMEAIR Library) to a save file `ACMEAIR.FILE` by using the **CPYTOCLD** command. Because this process is out of the scope of this example, assume that the file is in an IBM COS bucket.

The objective is to use an IBM COS service as a user with reader access and credentials to restore this data to your IBM i VM that is running in IBM Power Systems Virtual Server.

Note: Although the graphics that are used in the figures in this section show how the service works, you are encouraged to create your own IBM COS service from the IBM Cloud Catalog, create your first bucket and credentials, and start your first backup and restore.

Complete the following steps to transfer and restore a library from IBM COS to the Integrated File System (IFS):

1. Run the **CPYFRMCLD** command (see Figure 4-30) and enter the following information:
   - Resource Name. The resource name that was created.
   - Submit to Batch. ***NO** (for example, interactive; for large files, use ***YES**).
   - Cloud File Name. The file name that is in the IBM COS Bucket.

     In the example, `acmeair.file` is the name of the save file that is in the IBM COS service. It was sent to IBM COS by using a **CPYTOCLD** command or another upload process (HTTP, S3 client, Aspera, and so on).
   - Local file name. An IFS directory and file name as a destination such as a home directory `/home/QSECOFR/acmeair.file`. If the file exists, it is overwritten.

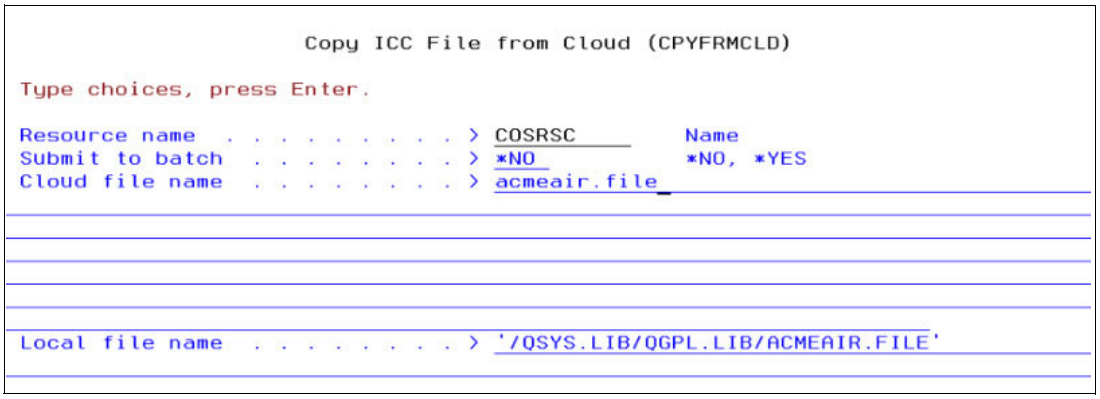

Figure 4-30 Copy IBM Cloud Storage file from cloud

   Issue CPYFRMCLD RESOURCE(COSRSC) ASYNC(*NO) CLOUDFILE(acmeair.file) LOCALFILE('/QSYS.LIB/QGPL.LIB/ACMEAIR.FILE').

   The following message is displayed:

   ```
   File copied.
   ```

2. Restore the file by running the following command:

   ```
   RSTLIB SAVLIB(ACMEAIR) DEV(*SAVF) SAVF(QGPL/ACMEAIR)
   ```

3. Display the restored content, as shown in Figure 4-31 on page 161:

   ```
   DSPSAVF FILE(QGPL/ACMEAIR)
   ```

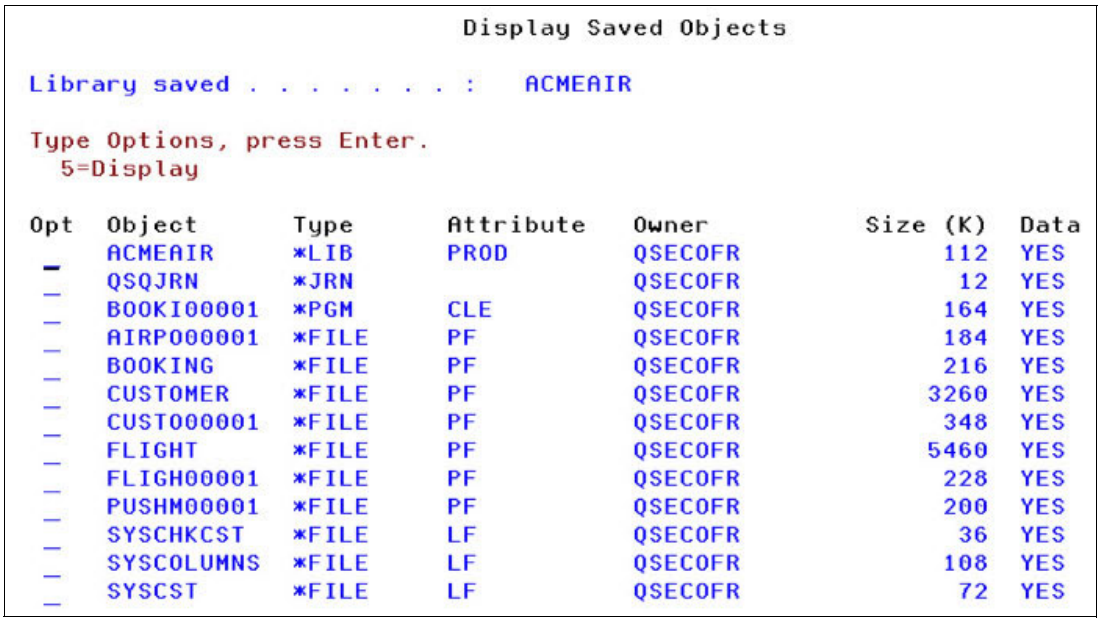

Figure 4-31 Display the objects restored from IBM COS

4.3.7 Transferring a save library to the cloud by using IBM BRMS

Complete the following steps to transfer a save library to the cloud by using IBM BRMS:

1. Create a cloud resource that is named `IBM BRMSCLD` by running *one* of the following commands:

   ```
   CRTFPRICC
   CRTSLRICC
   ```

   Note: The following commands can be used to verify that the cloud resource is working correctly:

   ```
   CPYTOCLD
   CPYFRMCLD
   ```

2. Configure IBM BRMS objects that are required to use the cloud resource:

   ```
   INZBRM OPTION(*DATA)
   ```

3. Create a virtual tape device:

   ```
   CRTDEVTAP DEVD(IBM BRMSCLDTAP) RSRCNAME(*VRT)
   ```

4. Vary on the virtual tape device:

   ```
   VRYCFG CFGOBJ(IBM BRMSCLDTAP) CFGTYPE(*DEV) STATUS(*ON)
   ```

5. Create an image catalog:

   ```
   CRTIMGCLG IMGCLG(IBM BRMSCLDTAP) TYPE(*TAP) DIR('/tmp/IBM BRMSCLDTAP')
   CRTDIR(*YES)
   ```

6. Add a volume to the image catalog:

   ```
   ADDIMGCLGE IMGCLG(IBM BRMSCLDTAP) FROMFILE(*NEW) TOFILE(*GEN) VOLNAM(BRMCLD)
   ```

7. Load the image catalog on the device:

   ```
   LODIMGCLG IMGCLG(IBM BRMSCLDTAP) DEV(IBM BRMSCLDTAP)
   ```

8. Configure the virtual tape device in IBM BRMS:

   ```
   WRKDEVBRM
   ```

   Specify the following information:

   – Opt field: 1= Create
   – Device field: `IBM BRMSCLDTAP`
   – Category field: *VRTTAP

9. Add a media class for the virtual tape media:

   ```
   WRKCLSBRM TYPE(*MED)
   ```

   Specify the following information:

   – Opt field: 1=Create
   – Class field: `IBM BRMSCLDTAP`
   – Density field: *VRT256K
   – Shared media field: *NO

10. Add volume BRMCLD to the IBM BRMS media inventory:

    ```
    ADDMEDBRM VOL(BRMCLD) MEDCLS(IBM BRMSCLDTAP) IMGCLG(IBM BRMSCLDTAP)
    ```

11. Add a media policy for the virtual tape media:

    ```
    WRKPCYBRM TYPE(*MED)
    ```

    Specify the following information:

    – Opt field: 1=Create
    – Policy field: `IBM BRMSCLDTAP`
    – Media class field: `IBM BRMSCLDTAP`

12. Create a library to save to the cloud:

    ```
    CRTLIB LIB(CLDLIB)
    ```

13. Save the library to the virtual tape volume:

    ```
    SAVLIBBRM LIB(CLDLIB) DEV(IBM BRMSCLDTAP) MEDPCY(IBM BRMSCLDTAP)
    ```

14. Move volume BRMCLD to the cloud location:

    ```
    WRKMEDBRM VOL(BRMCLD)
    ```

    Verify the following information:

    – Volume is in *ACT status
    – Location is *HOME
    – Volume is in *ACT status
    – Location is CLOUD

► Specify the following information:

  – Opt field: 8=Move
  – Storage location field: CLOUD

Manually restoring a library from cloud

To manually restore a library, use **`RSTLIBBRM`** command. IBM BRMS automatically moves the virtual media from the cloud and performs the restore. IBM BRMS then creates the image catalog and image stub to move the file or library into the cloud.

4.3.8 Automatically transferring media to IBM COS

IBM BRMS provides specific control groups that can be used to automatically save data to media in the cloud, so the cloud media can be used to create physical media.

The control groups create cloud media that is formatted, so it can be downloaded and burned directly to physical optical media. All remaining data on the system can be backed up to media in the cloud and restored directly from the cloud without a need to create physical media.

Automatic transfers of media to IBM COS are shown in Figure 4-32.

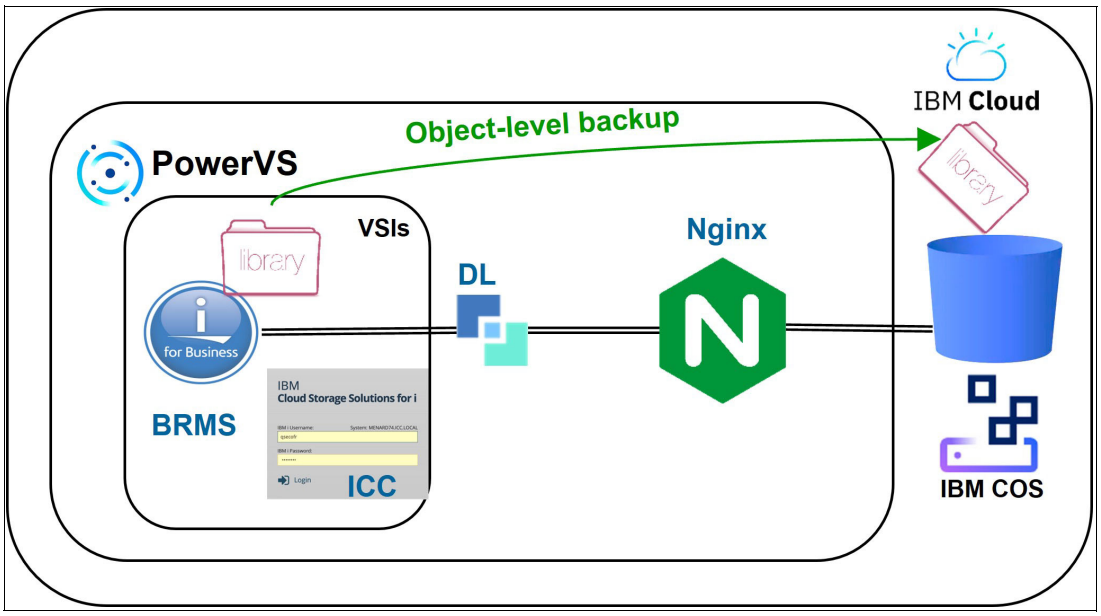

Figure 4-32 Automatic transfer of media to IBM COS

Control groups that feature a `QCLD` prefix cause IBM BRMS to automatically create and transfer media to the cloud during a backup.

For example, a library that is named `PAYROLL` can be used to automatically transfer the media to the cloud and transfer the media back to the system by restoring the library.

Sample automatic save and transfer

The predefined cloud control groups that IBM BRMS created for the cloud resource is set up to work as defined. To simplify this example, copy one of the predefined control groups:

```
WRKCTLGBRM TYPE(*BKU)
```

Complete the following steps:

1. Specify 3=Copy next to QCLDBUSR01.
2. Specify `CLOUDBKU` for the New Name field.
3. Change the entries in the control group to save the library **`WRKCTLGBRM TYPE(*BKU)`**.
4. Specify 2=Change entries next to `CLOUDBKU`.
5. Clear all entries from the control group by removing all of the numbers from the Seq column.
6. Add an entry for the `CLDLIB` library:
   - Seq: 20
   - Backup Items: `CLDLIB`

7. Use the default values in the remaining fields, as shown in Figure 4-33.

```
                    Edit Backup Control Group Entries                I750CL21

Group . . . . . . . . . . . : CLOUDBKU
Default activity  . . . . . *BKUPCY
Text  . . . . . . . . . . . Entry created by BRM configuration

Type information, press Enter.

                                    Weekly   Retain Save     SWA
     Backup      List ASP           Activity Object While    Message    Sync
Seq  Items       Type Device        SMTWTFS  Detail Active   Queue      ID
____ __________ ____ __________    _______
  10 *EXIT       ____               *DFTACT
  20 CLDLIB      ____ *SYSBAS       IIIIIII  *ERR   *LIB    *LIB       *NONE
  30 *EXIT       ____               *DFTACT
```

Figure 4-33 Using default values in the remaining fields

8. Issue the control group, as shown in Figure 4-34.

```
                    Start Backup using BRM (STRBKUBRM)

Type choices, press Enter.

Control group  . . . . . . . . . > CLOUDBKU     *BKUGRP, *SYSGRP, *SYSTEM...
Schedule time  . . . . . . . . .   *IMMED       hhmm, *IMMED
Submit to batch  . . . . . . . . > *NO          *YES, *CONSOLE, *CTLSBS, *NO
Starting sequence:
  Number . . . . . . . . . . . .   *FIRST       1-9999, *FIRST
  Library  . . . . . . . . . . .   *FIRST       Name, *FIRST
Append to media  . . . . . . . .   *CTLGRPATR   *CTLGRPATR, *BKUPCY, *NO...
Activity . . . . . . . . . . . .   *CTLGRPATR   *CTLGRPATR, *FULL, *INCR
Retention:
  Retention type . . . . . . . .   *CTLGRPATR   *CTLGRPATR, *DAYS, *PERM
  Retain media . . . . . . . . .   35           1-9999
Device . . . . . . . . . . . . .   *CTLGRPATR   Name, *CTLGRPATR, *MEDCLS
                + for more values   __________
Parallel device resources:
  Minimum resources  . . . . . .   *CTLGRPATR   1-32, *CTLGRPATR, *NONE...
  Maximum resources  . . . . . .   *MIN         1-32, *MIN, *AVAIL
```

Figure 4-34 Issuing the control group

As shown in Figure 4-35, the volume status is `*ACT` and the location point is a cloud resource.

```
                              Work With Media
                                                          System:   I750CL21
Position to  . . . . .  ______       Starting characters

Type options, press Enter.
  1=Add   2=Change   4=Remove   5=Display   6=Work with serial set   7=Expire
  8=Move   9=Remove volume error status     10=Reinitialize ...

     Volume            Creation Expiration            Move     Media      Dup
Opt  Serial     Status   Date      Date     Location   Date     Class      Sts
__   ______
__   BRMCL2     *ACT   12/09/16  12/30/16   CLOUD     12/09/16  BRMSCLDTAP
```

Figure 4-35 Volume status and location

4.4 FalconStor overview

IBM i customers require backups and many use tape or virtual tape, which requires a software-based Virtual Tape Library (VTL). IBM and FalconStor offer a VTL solution which can be use with Power Virtual Server.

FalconStor StorSafe is an optimized backup and deduplication solution that provides Virtual VTL emulation, high-speed backup and restore, data archival to supported S3 clouds for long-term storage, global data deduplication, and enterprise-wide replication, without requiring changes to the existing environment.

FalconStor StorSight is a single integrated platform that can simplify the management of data across legacy, modern, and virtual storage environments. StorSight gathers and consolidates information coming from different StorSafe servers into a scalable repository of services, tenants, users, predictive analytics, alert rules, reporting, and historical data. StorSight provides a web-based portal for centralized management and monitoring of multiple backup and deduplication servers.

Some of the StorSafe VTL Benefits are included in the following list:

- Up to 95% data reduction. FalconStor reduces the storage capacity required by backup data on-premises and in the cloud by up to 95%.
- Ransomware Protection. Recover data from any point in time with air-gapped, immutable backups.
- Offsite Protection. For offsite protection, StorSafe VTL exports virtual tapes as physical tapes, and tracks their location with a management dashboard. Virtual tapes can remain in the library for fast restores or can be stubbed when exported to local object storage.
- High Performance. Improve both backup and recovery.
- Modernize. FalconStor is 100% compatible with existing backup software, hardware, and operational procedures.

FalconStor StorSafe is offered through the IBM Cloud Catalog for both off premises and on premises with the following names:

- FalconStor StorSafe for Power On-Premises provides flexible deployment options for the user:
  - It can be installed on any X86 server that is in the FalconStor certification matrix.
  - It can be installed as a VM in Hypervisors such as VMware, HyperV, Xen, and KVM.
  - It can be installed inside Power Virtual Server as an LPAR.
- FalconStor StorSafe for Power Virtual Server Cloud is compatible with IBM i and can be deployed from the IBM Cloud catalog.

4.4.1 Installation and Configuration Prerequisites for FalconStor Storsafe

The purpose of supporting an iSCSI virtual tape library (VTL) is to provide an offering that improves the performance of save and restore operations over Ethernet. This support is only tested and supported with a VTL from FalconStor and with a library type of Falcon on the Power Virtual Server infrastructure.

IBM BRMS also requires the following PTFs:

- IBM i version 7.4
  - Must have technical refresh 9 installed (MF99309).

- MF69817

- To enable configuration using SQL, install the following updates:
  - SI85869
  - SI84630
  - To enable IPsec and VPN (which requires Navigator) install the latest cumulative PTF package.
- IBM i version 7.5

  Backup/Recovery group PTF SF99954 level 12, April 2 2024 update contains PTFS needed to run ISCSI and the latest cumulative PTFs enable IPsec and VPN (which requires Navigator).

4.4.2 Deployment Scenario with IBM Cloud

The following scenarios are examples of different deployment cases for StorSafe with IBM Cloud:

- Native backup in the cloud. StorSafe in IBM Cloud serves different IBM i host clients for backup.
- IBM cloud-to-cloud replication. StorSafe in IBM Cloud replicates data to another StorSafe.
- Hybrid Cloud. On-premises StorSafe replicates data to a StorSafe in IBM Cloud.
- Workload migration. On-premises StorSafe replicates data to StorSafe in IBM Cloud, migrates backup client workload, and then removes the on-premises StorSafe.

Native Backup Deployment

Several IBM i host clients are connected by using iSCSI to StorSafe in the cloud to back up data by using IBM BRMS to virtual tape libraries. StorSafe uses IBM COS as the data devices for the deduplication repository by using a Generic S3 object storage account.

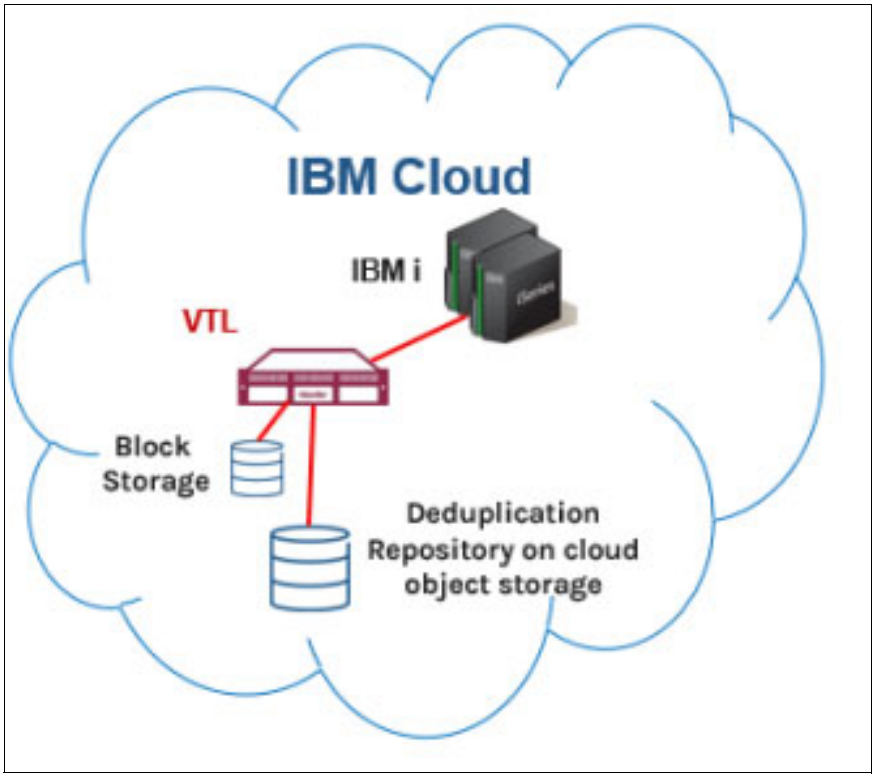

Figure 4-36 Native backup deployment

Connectivity

The following connections are required:

- Network connections between StorSafe and Storsight
- iSCSI connections between StorSafe and IBM i host clients
- IBM Cloud Direct Link connection between StorSafe and the IBM COS in the IBM Cloud Classic environment
- Internet connection with the FalconStor license server for online registration (optional)

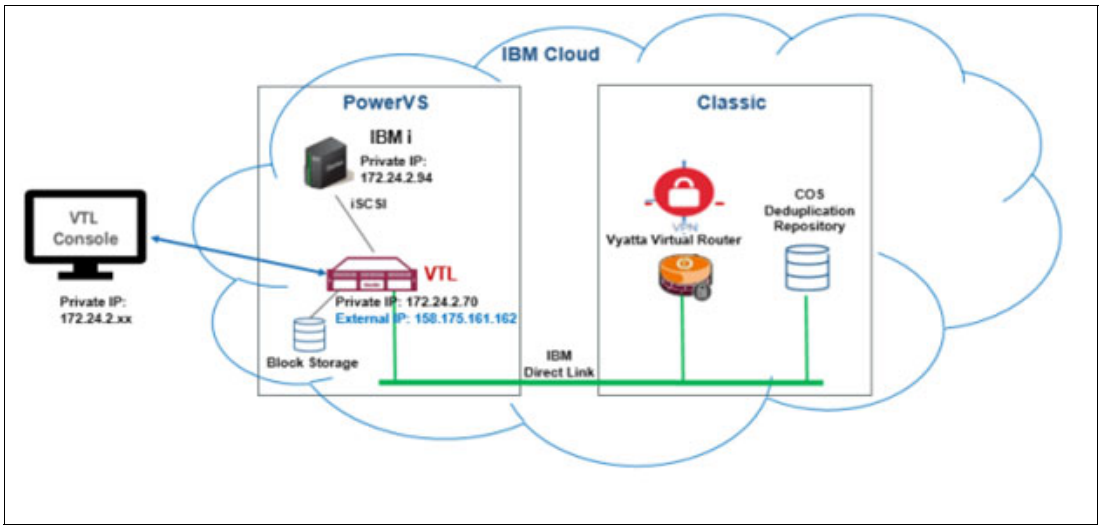

Figure 4-37 Connectivity for Native backup deployment

Scenario

Network default interface eth0 has a public IP address and is used for all external connections.

Network interface eth1 has a private IP address and is used to add routes to other networks.



Figure 4-38 public and private network

The system command displays the network configuration on StorSafe as follows:

```
# ifconfig
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
 inet 192.168.150.162 netmask 255.255.255.248 broadcast
192.168.150.167
eth1: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 9000
```

```
inet 172.24.2.70 netmask 255.255.255.224 broadcast 172.24.2.95
```

Create routes from the private IP to connect to the other networks. For example, ADDRESS1=10.0.0.0 is for routing to the Classic infrastructure and ADDRESS2=172.22.0.0 is added for routing to other private networks.

Any other customized networks to the route can be added using sequential numbers as a postfix.

```
# cat /etc/sysconfig/network-scripts/route-eth1
# Created by cloud-init on instance boot automatically, do not edit.
#
ADDRESS0=172.24.2.64
GATEWAY0=172.24.2.65
NETMASK0=255.255.255.224
ADDRESS1=10.0.0.0
GATEWAY1=172.24.2.65
NETMASK1=255.0.0.0
ADDRESS2=172.22.0.0
GATEWAY2=172.24.2.65
NETMASK2=255.255.0.0
```

Cloud-to-Cloud replication deployment

StorSafe in IBM Cloud replicates data over a WAN to another StorSafe in the cloud for data protection and disaster recovery. StorSafe source and replica servers use IBM COS as the data devices for the deduplication repository through Generic S3 object storage accounts.

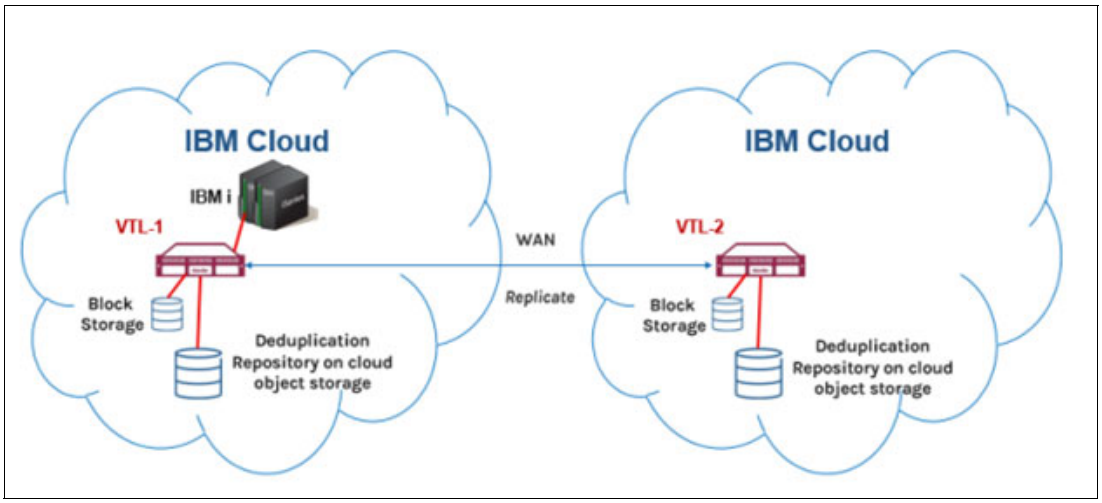

Figure 4-39 Cloud-to-Cloud Replication Deployment

The following connections are required for connectivity:

- Network connections between StorSafe and StorSight
- Network connections between source and the replica StorSafe servers by virtual routers or Megaport
- iSCSI connections between StorSafe and IBM i host clients
- IBM Cloud Direct Link connection between StorSafe and the IBM COS in the IBM Cloud Classic environment

- Internet connection with the FalconStor license server for online registration (optional)

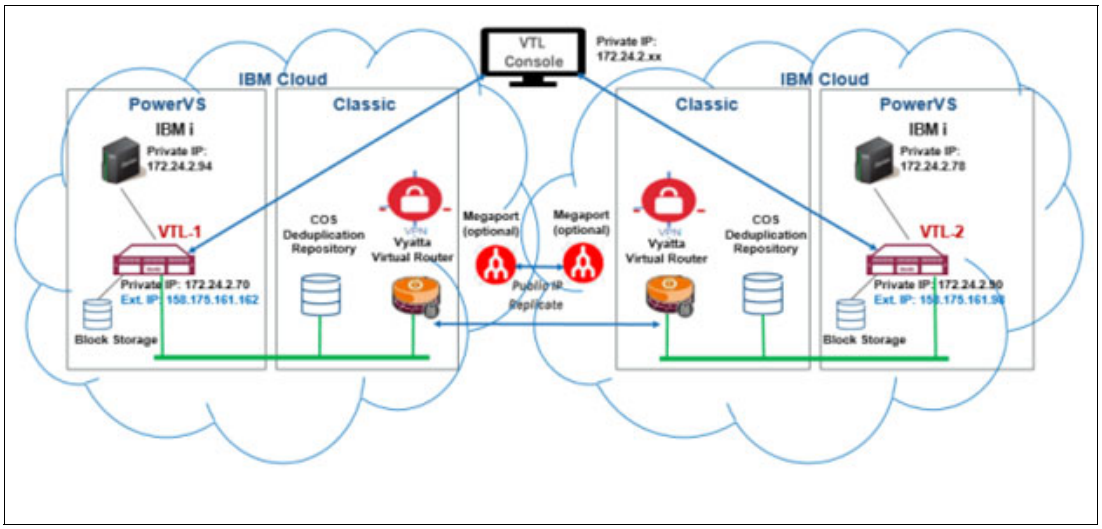

Figure 4-40 Connectivity for Cloud-to-Cloud Replication Deployment

Hybrid Cloud Deployment

An on-premises StorSafe server deduplicates data and replicates to a StorSafe server in IBM Cloud for data protection and disaster recovery purposes. The on-premises StorSafe server uses SCSI devices for the deduplication repository. The StorSafe target server in the cloud uses IBM COS as the data devices for the deduplication repository via a Generic S3 object storage account.

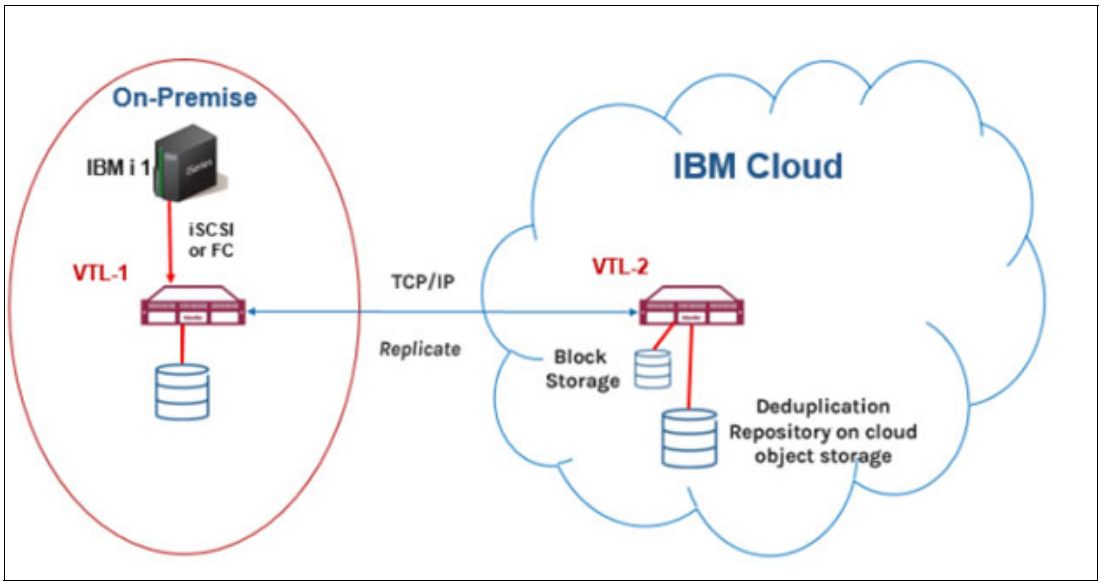

Figure 4-41 Hybrid Cloud Replication Deployment

The following connections are required for connectivity for hybrid cloud deployment:

- Network connection between the on-premises StorSafe and the Power Virtual Server StorSafe in IBM Cloud
- Network connection between StorSight in the x86 infrastructure and the on-premises StorSafe server

- Network connections between StorSight in the x86 infrastructure and the Power Virtual Server running StorSafe in IBM Cloud
- iSCSI or Fibre Channel connection between IBM i host clients and the on-premises StorSafe server
- iSCSI connections between IBM i host clients and the Power Virtual Server running StorSafe in IBM Cloud
- IBM Cloud Direct Link connection between the Power Virtual Server running StorSafe and the IBM Cloud Classic environment, if the IBM COS (COS) is used for the deduplication repository
- Internet connection with the FalconStor license server for online registration (optional)

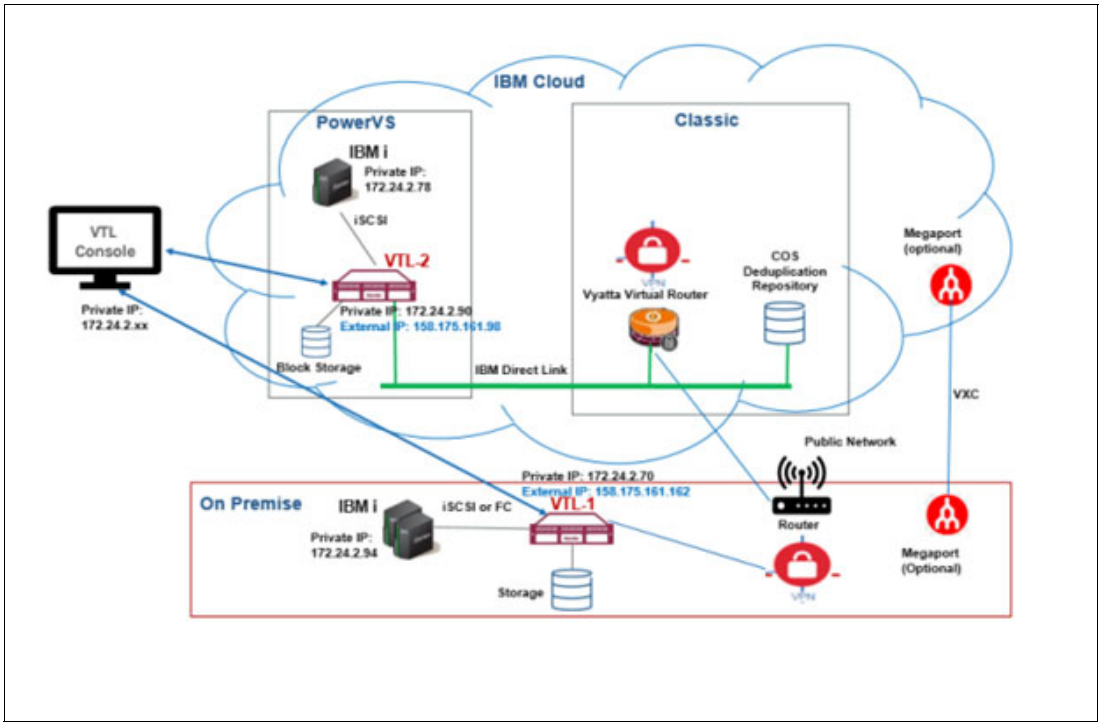

Figure 4-42 Connectivity for Hybrid Cloud Replication Deployment

Workload Migration Deployment

IBM i backup clients of on-premises StorSafe servers are to be moved to the cloud through StorSafe data replication. StorSafe in the cloud uses IBM COS as the data devices for the deduplication repository through a Generic S3 object storage account.

IBM Cloud can be used for workload migration in the following way:

1. IBM BRMS running at the primary site on an IBM i host client performs I/O to a FalconStor StorSafe server through iSCSI to back up the OS boot image and application data on virtual tapes.
2. The IBM i host client is prepared for migration. Two virtual tapes are needed. The first tape contains the base operating system for the initial remote boot. The second tape contains the remaining data of the host machine image.
3. The primary StorSafe server deduplicates data and replicates both tapes to another StorSafe server in IBM Cloud.
4. On the replica site, the replica virtual tapes are promoted. The first tape is assigned to an IBM i remote boot server via iSCSI.

5. The IBM i remote server in the cloud mounts the OS image on the first promoted tape on an NFS share, then it boots another IBM i virtual server from the NFS boot directory running the same OS as the IBM i of the primary site.
6. The second tape is assigned to the new IBM i via iSCSI. This new server reads data from this second tape to complete the workload migration and to become the clone of the IBM i server on the primary site.

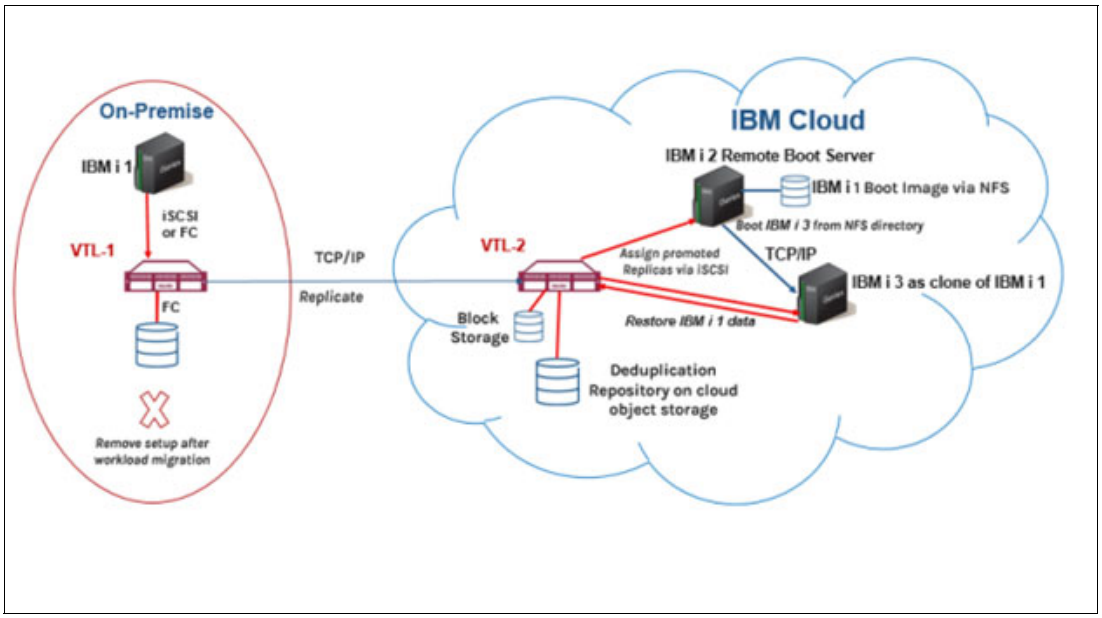

Figure 4-43 Workload Migration Deployment

The following connections are required for workload migration deployment connectivity:

- Network connection between the on-premises StorSafe and the Power Virtual Server running StorSafe in IBM Cloud
- Network connection between the StorSafe console in the classic infrastructure and the on-premises StorSafe server
- Network connections between the StorSafe console in the classic infrastructure and the Power Virtual Server running StorSafe in IBM Cloud
- iSCSI or Fibre Channel connection between IBM i host clients and the on-premises StorSafe server
- iSCSI connections between IBM i host clients and the Power Virtual Server running StorSafe in IBM Cloud
- IBM Cloud Direct Link connection between the Power Virtual Server running StorSafe and the IBM Cloud Classic environment if the IBM COS is used for the deduplication repository
- Internet connection with the FalconStor license server for online registration (optional)

When an IBM i host client is used for a workload migration scenario, it must be an NFS server to use the boot image for the cloned machine. Refer to IBM Documentation for more information.

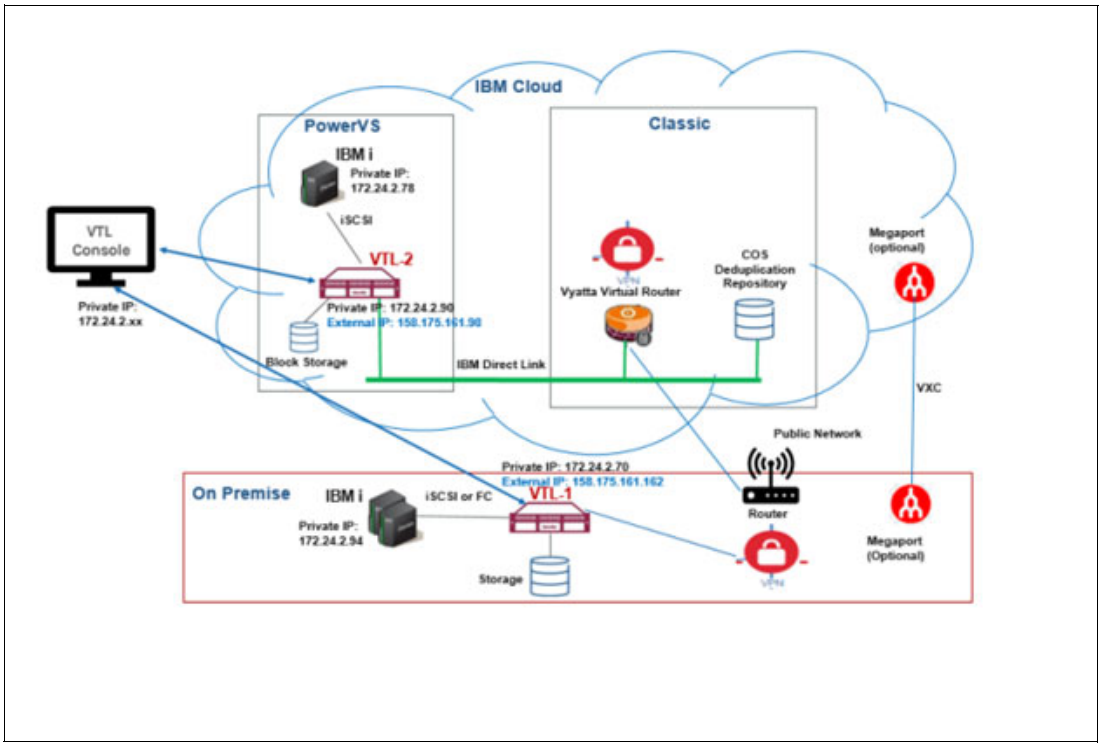

Figure 4-44 Connectivity for Workload Migration Deployment

4.4.3 IBM i migration example that uses FalconStor StorSafe

Figure 4-45 shows an overview for the 7 steps in this migration example.

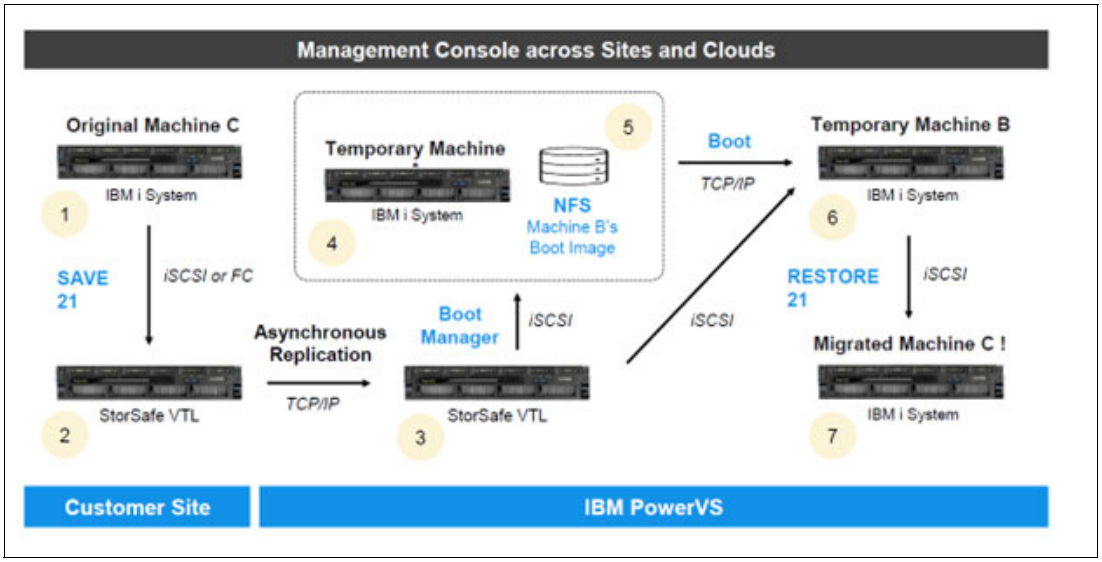

Figure 4-45 Seven key steps IBM i Migration Example

Run the seven key steps listed in Figure 4-45:

1. Prepare for migration using BLDMIGDSLO script on Machine C to make the first tape, which contains SAVSYS and the base IBM i operating system.

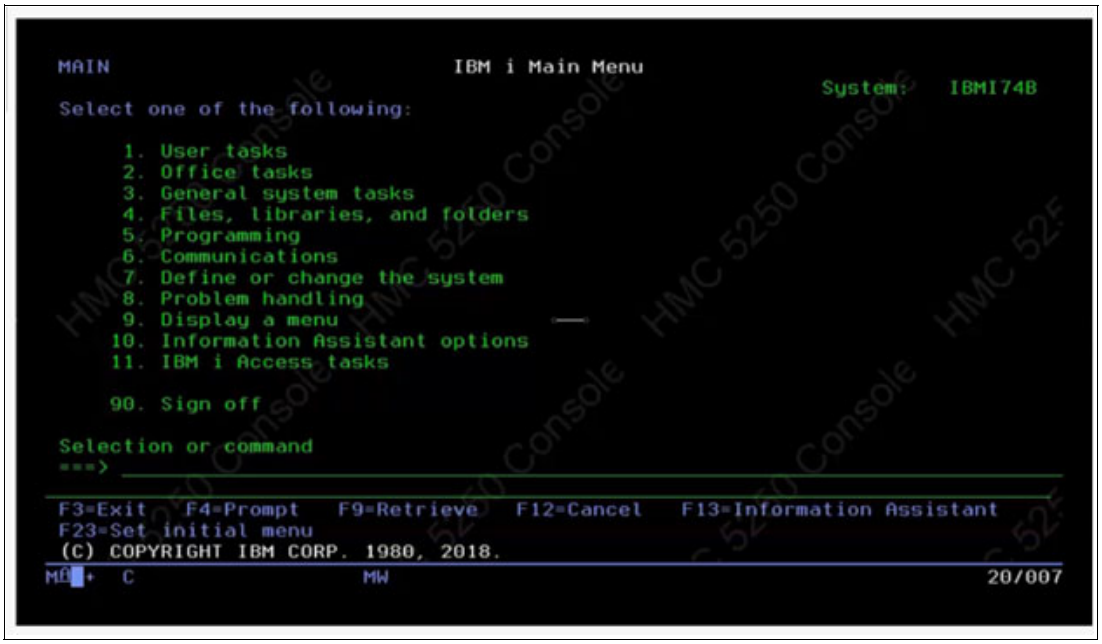

Figure 4-46 IBM i Main menu

a. Issue the following commands:

```
RSTLIB FALCONSTOR *SAVE SAVE(FALCONSTOR)
CHGCURLIB FALCONSTOR
```

Expected message showing in .



Figure 4-47 Message changed

b. Issue the following command:

```
ENDSYS *IMMED ENDSBSOPT(*NOJOBLOG)
```

Verify the following message:

```
ENDSYS command in progress.
```

Issue STRTCP and note the message as shown in the example in Figure 4-48.



Figure 4-48 STRTCP started

a. Issue the following command:

```
STRTCPIFC ‘10.7.2.54’
```

Verify the following message:

```
User QSECORFR started IP interface 10.7.2.54 on ETH8.
```

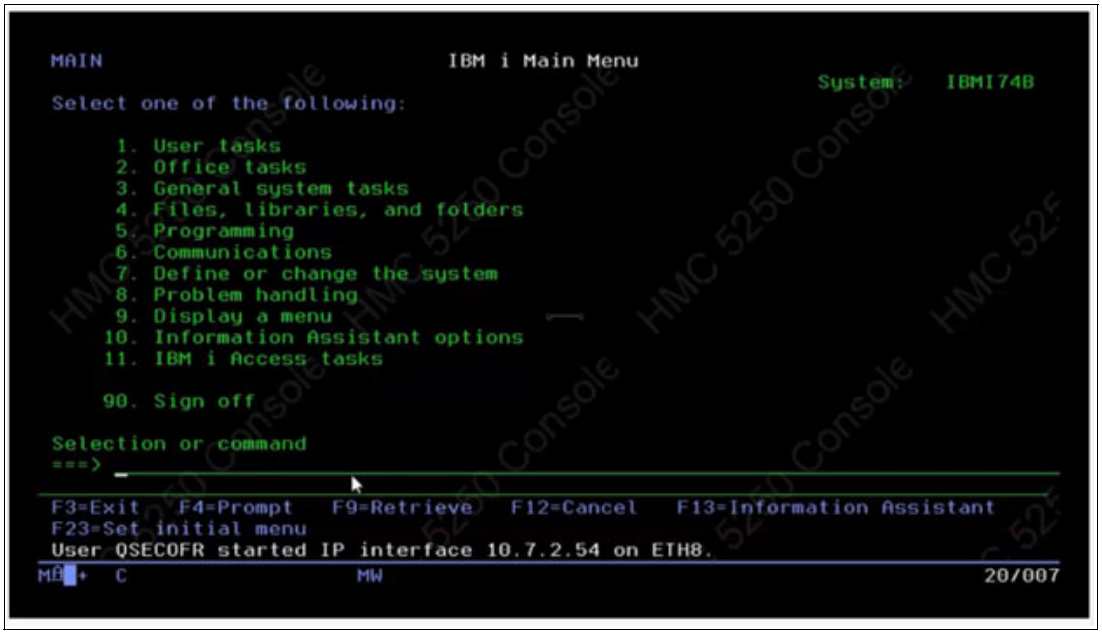

Figure 4-49 User connected

2. Perform a SAVE 21 on Machine C on-premises at the customer site. This makes a second tape.
   a. Issue the following command:

      ```
      BLDMIGDSLO
      ```

   b. Go to Configure TCP/IP (Figure 4-50).

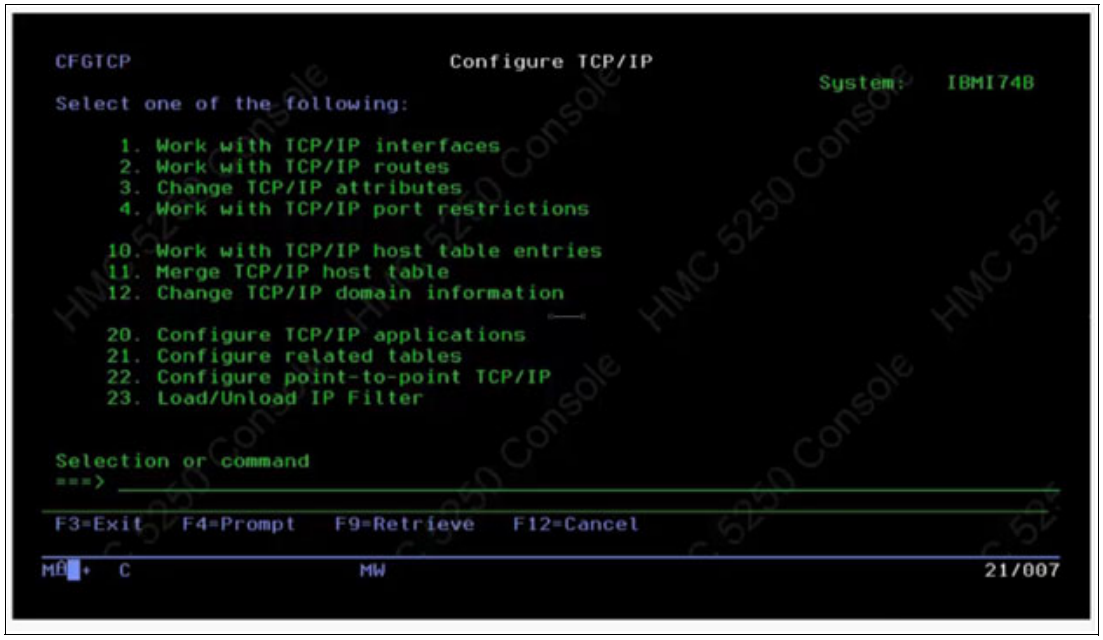

Figure 4-50 TCP/IP main menu

   c. Issue the command ENDSYS *IMMED
   d. Issue the command STRTCP
   e. Issue the command GO SAVE
   f. Enter **21** and verify the screen as shown in Figure 4-51 on page 176.

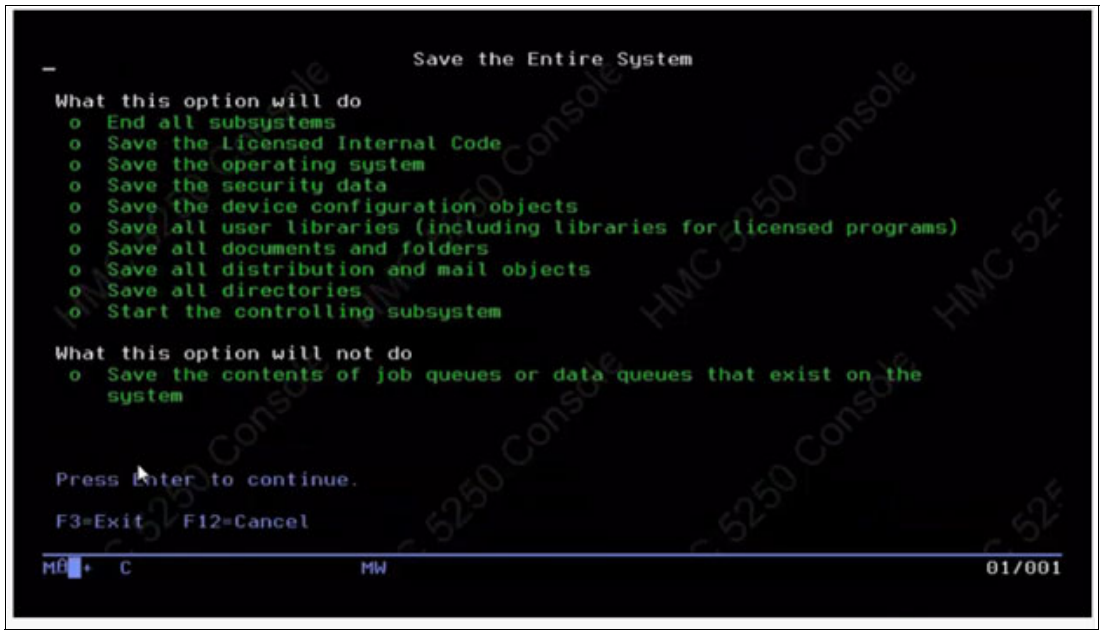

Figure 4-51 Save the entire system screen

g. Press Enter and verify the message as shown in Figure 4-52.

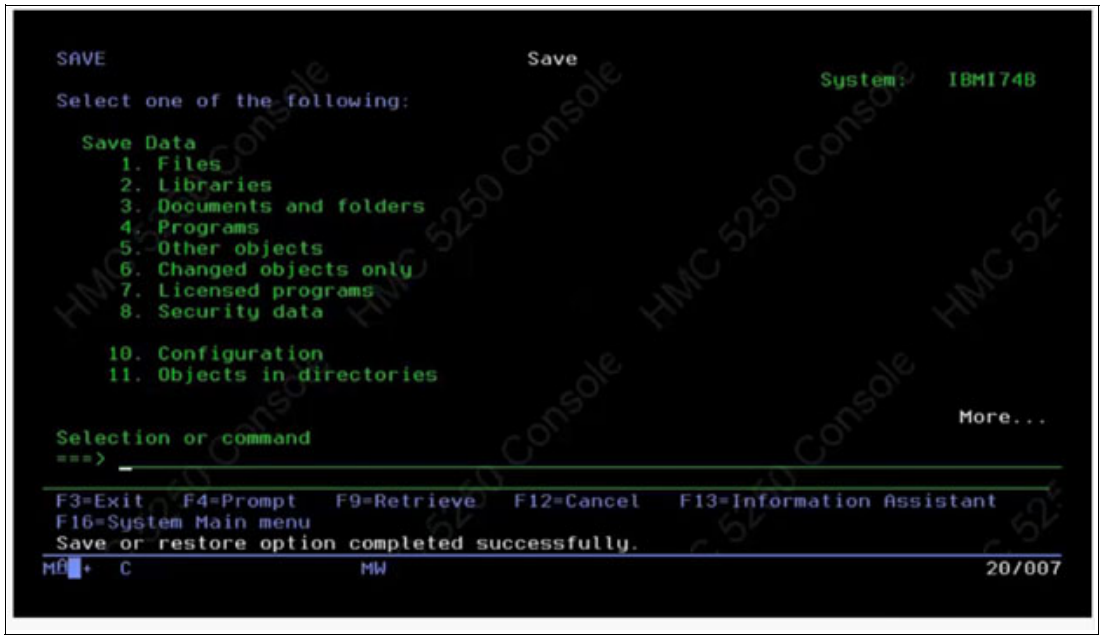

Figure 4-52 Save confirmed

3. Replicate both virtual tapes from the customer's site to StorSafe in Power Virtual Server. The Falconstor manage screen is shown in Figure 4-53 on page 177

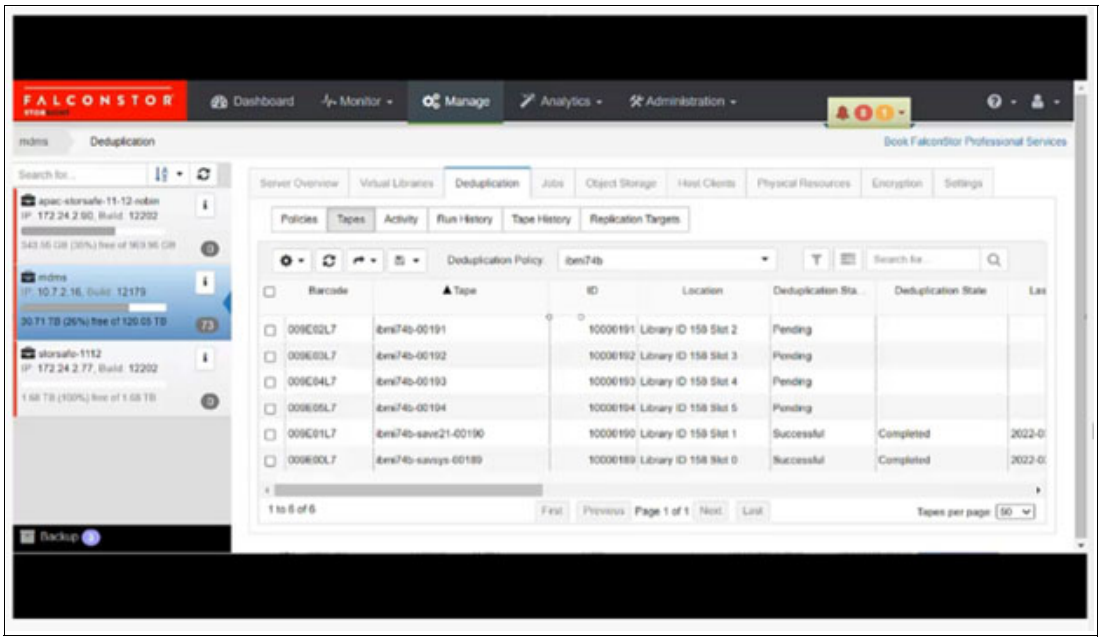

Figure 4-53 Duplication status in Falconstor

Operation success completed is shown in Figure 4-54.

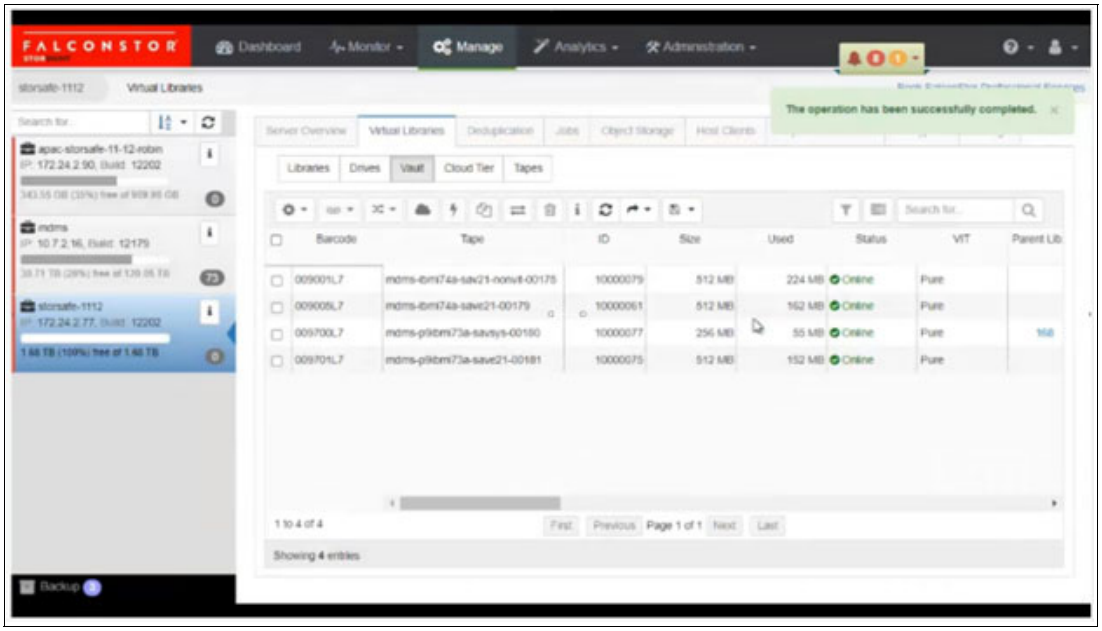

Figure 4-54 Duplication completed

4. Restore the boot manager from StorSafe virtual tape to Machine A in Power Virtual Server.
5. Run BLDBOOTENV on Machine A to read the data from the DSLO tape, which puts the boot image of Machine B onto an NFS share.
   a. From the IBM i Main Menu issue the command BLDBOOTENV.
   b. Verify that the status changes to `Current library changed to FALCONSTOR`.

c. Continue to press Enter until the screen shows the message `Restore and setup complete.` as shown in Figure 4-55.

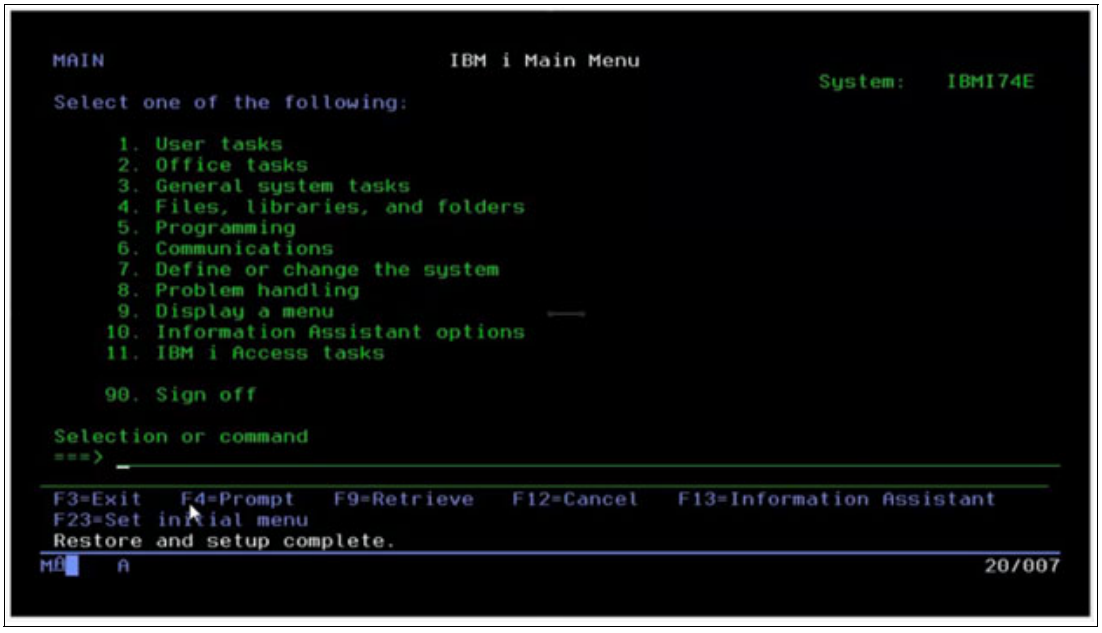

Figure 4-55 Restore and setup complete

6. Boot Machine B from the NFS boot directory on Machine A and run the CFGISCSITP script to apply the required PTFs onto Machine B for supporting iSCSI and configure StorSafe.

   a. Enter the information as shown in Figure 4-56.

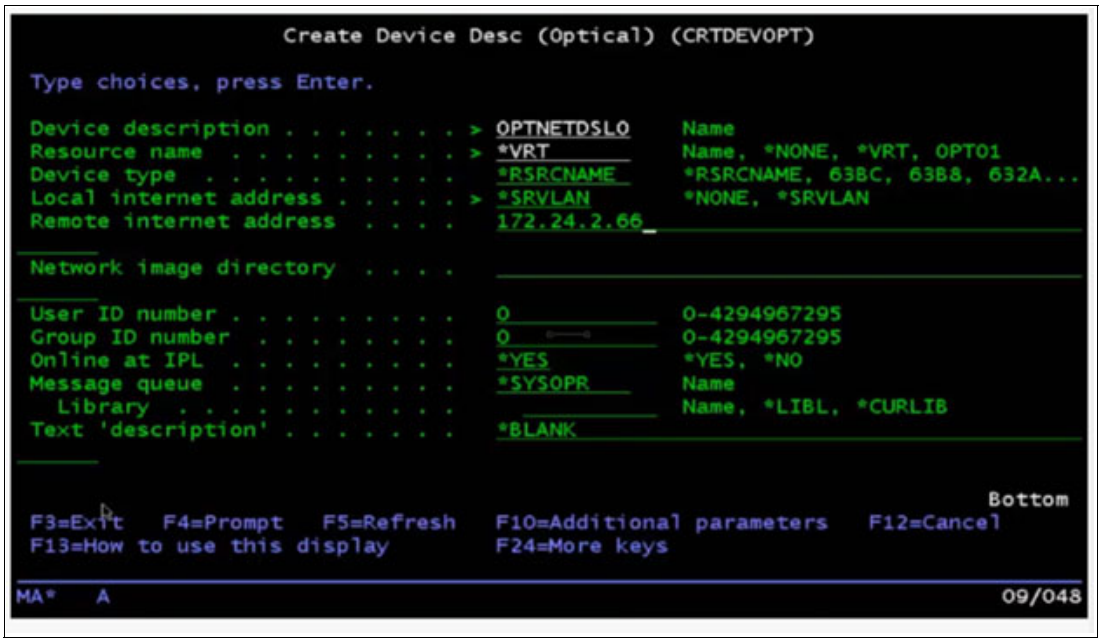

Figure 4-56 Create device Desc

   b. Press ‘F10’ on ‘Install LIC and Initialize System’ to continue.

   c. Select option ‘2’ to Install the operating system.

The Install Licensed Programs screen opens as shown in Figure 4-57.

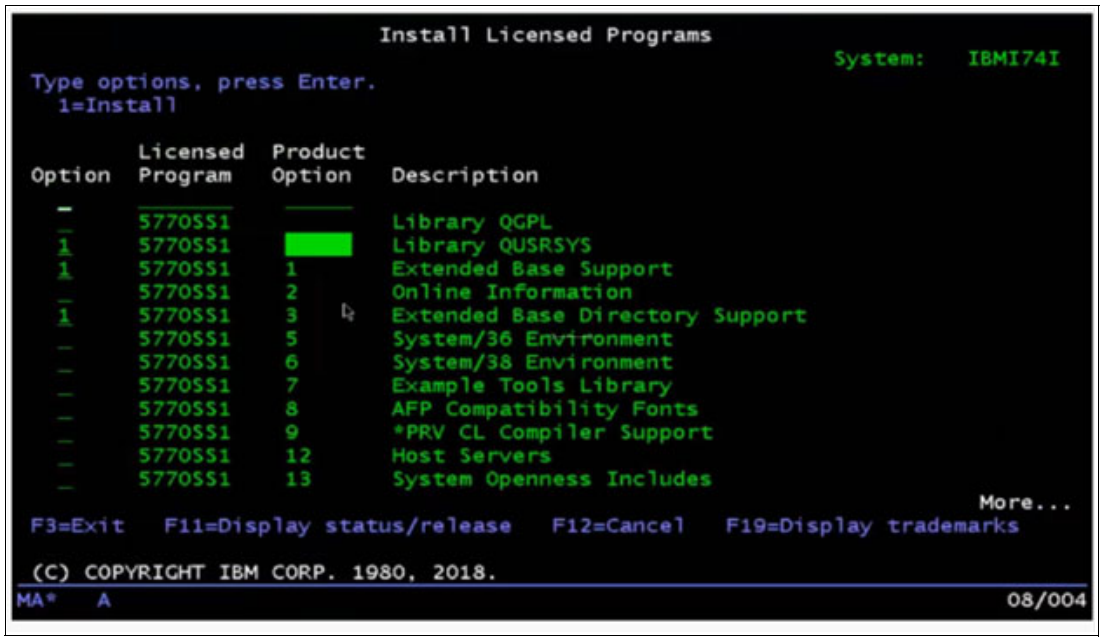

Figure 4-57 Install licensed programs

7. Perform a RESTORE 21 on machine B and reboot to complete the migration of Machine C to Power Virtual Server.

   a. Issue the CO RESTORE command from the IBM i main menu (Figure 4-58).

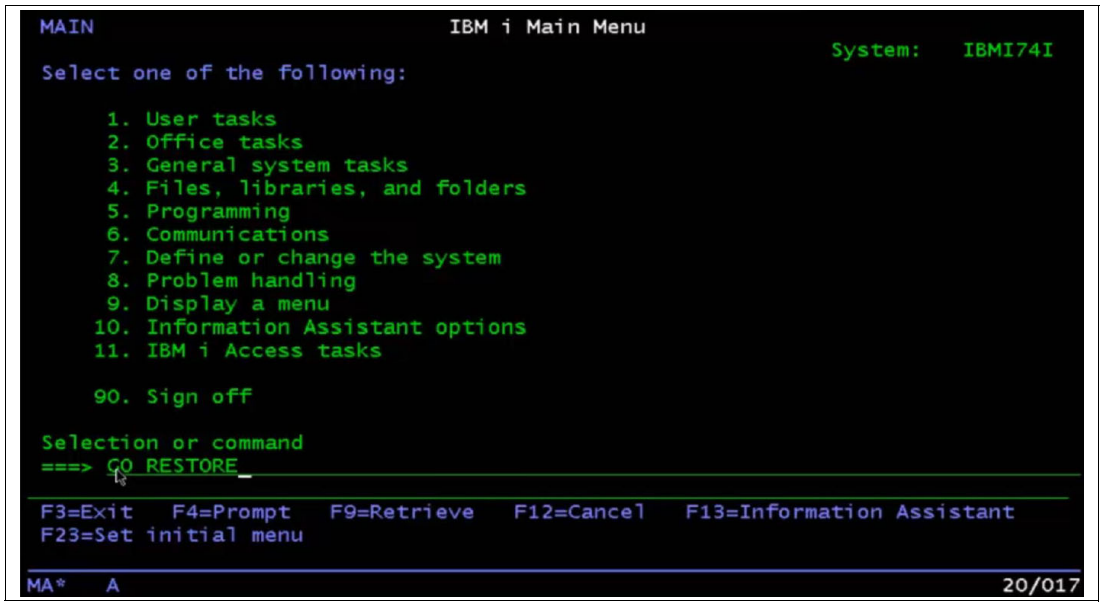

Figure 4-58 Issue CO RESTORE COMMAND

b. Select option ‘21’ - for Restore System and User Data (Figure 4-59).

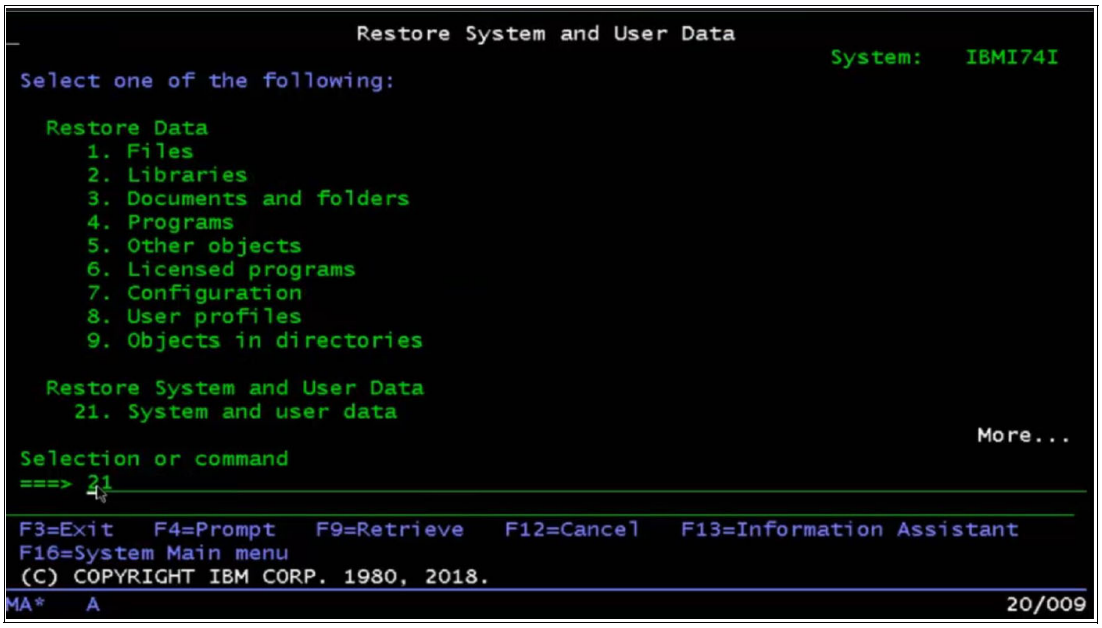

Figure 4-59 Select option 21

The Restore System and User Data screen opens (Figure 4-60).

c. Press Enter to continue.

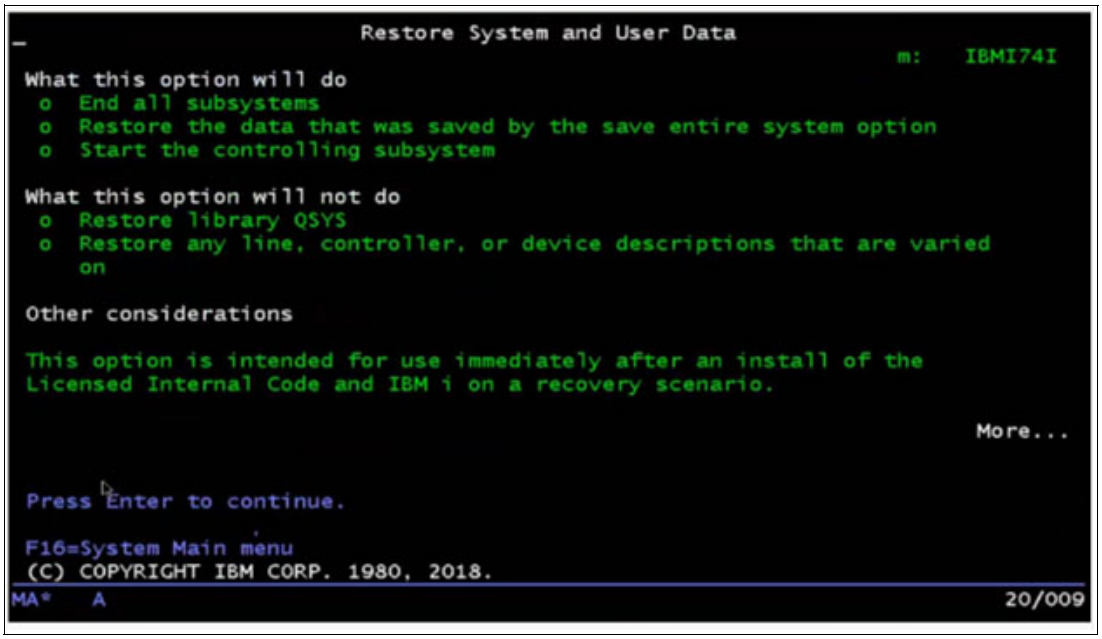

Figure 4-60 Restore System and User Data screen

4.4.4 Sizing and ordering FalconStor StorSafe

To identify the required system resources, such as the amount of memory, CPU cores, and storage capacity for your FalconStor servers, follow the sizing guidelines. VTL sizing depends on the amount of data to retain based on the size of ingest data, deduplication ratio, and backup parameters.

Refer to the FalconStor Solution Sizing page to view the sizing information about the deduplication repository, backup cache, memory, CPU, and the machine type.

Refer to the IBM Cost Estimation page:

1. Click Estimate costs on the upper right side panel.
2. Select Virtual Tape Library as OS.
3. Enter values for usage parameters based on the Solution and n Sizing tool results.
4. Click Calculate cost and Save to see the cost.
5. Click Review estimate to go the Cost estimator page.

Additional costs might apply based on extra capacity for COS, or additional network and infrastructure components. For the COS capacity use the following instructions:

1. Go to Catalog and type Object Storage in the Search box.
2. Select the Standard plan.
3. Click Estimate costs.
4. Enter a value for Monthly average capacity.
5. Calculate cost again.

The FalconStor StorSafe is activated when it is added to StorSight for management. StorSafe servers do not require individual licenses. A global unique license is provided for StorSight that manages all StorSafe servers. After you create a StorSafe Power Virtual Server, an email from FalconStor is sent. The email contains a license keycode that you can enter in Falcon StorSight.

4.4.5 Configuring FalconStor StoreSafe by using StorSight

Run FalconStor StorSight to configure StorSafe. The console is the graphical administration tool that you can use to manage your StorSafe servers.

Downloading and installing Falcon StorSafe

The following list describes the first steps to configure StorSafe:

1. From a supported browser, enter the IP address of your FalconStor StorSight server.

- Enter your username, domain (optional), and password to log in. The default login is:
  - Username: superadmin
  - Domain: <blank>
  - Password: freestor.

2. Start the console from your machine and confirm the End-User License Agreement (EULA) that opens the first time that the console is started.
3. Select **Administration** -> **StorSight**.
4. Click **License**, and the "**+**" icon to add your license keycode received from FalconStor. Register your license.

Configuring your StorSafe server

In the Manage tab select your StorSafe server:

1. Select the **Physical Resources** -> **Physical Devices** tab to virtualize storage devices.
2. Select unassigned physical devices.
3. Click the **Manage** icon, and select **Prepare**.
4. Set the category as Virtual and set the appropriate reservation type:
   - Configuration for the 20 GB device to be used as the Configuration Repository and tape database
   - Deduplication for index and folder devices of the Deduplication Repository
   - Tapes for backup cache devices. Click **Prepare**
5. Select the Object Storage tab and click the "+" icon to add the object storage account to be used for deduplication.
   a. Select the Generic S3 provide.
   b. Enter the IBM Cloud access account information noted above, and enable the "Use for Deduplication Repository" option.
   c. If you use a reverse proxy server to make remote connections to IBM COS private endpoints, enter the proxy IP address, user ID and password.
6. Select the Settings tab to configure the Configuration Repository, the Tape Database, and the Deduplication Repository:
   a. Click the **Configuration Repository** icon.
      i. Select **Create Configuration Repository**.
      ii. Select the physical device reserved for Configuration
      iii. Click **Create** to enable the Configuration Repository and use 10 GB of the reserved device.
   b. Click the **Tape Database** icon:
      i. Click **Create Tape Database.**
      ii. Select the physical device reserved for the configuration.
      iii. Click **Create** to enable the tape database using 10 GB of the reserved device.
   c. Click the **Deduplication Repository** icon:
      i. Click **Create Deduplication Repository**.
      ii. Select **Object Storage** for the Deduplication Repository data disks.
      iii. Select the object storage account that you had created with the "Use for Deduplication Repository" option.
      iv. Select the physical devices reserved for Deduplication that will be used for the deduplication index and folder.
      v. Click **Create** to enable the Deduplication Repository.
   d. Enable iSCSI target mode to enable host clients running backup software to communicate with StorSafe via the iSCSI protocol.

7. Select the Virtual Libraries tab and click the "+" icon to create virtual tape libraries with IBM drives:
   a. For IBM i host clients, select the virtual tape library type as FalconStor FALCON TS3500L32 (03584L32) or FALCON TS3500L32 (03584L32) and the media type as ULTRIUM3 (LT03) or newer. By using Falcon library types, you get the 3584-403 device types to configure on IBM i host clients.
   b. Set or define the library characteristics:
      - Library name
      - Drive name prefix
      - Number of drives
      - Barcode range for tapes
      - Set the number of Import/Export slots to **1**.
   c. Enable Tape Capacity On demand (COD) to create small resources for your tapes and then automatically allocate additional space when needed. The minimum value for the incremental size is determined by the formula:
      (Maximum Capacity - Initial Tape Size)/63.
   d. Leave default settings and do not enable any service options.

Configuring iSCSI from the client side

Before performing these steps, configure iSCSI on the client by following the steps in the section “Configuring an iSCSI client” on page 185. Add the clients before assigning the virtual tape libraries to them:

1. Select the Host Clients tab and click the "+" icon to add host clients.
2. Specify the client name and select **iSCSI** as the communication protocol.
3. Click the iSCSI tab to select the initiators for the client. If the initiator does not appear, you might need to rescan by clicking **Rescan Initiators**. The client iSCSI initiator name is the same as the one configured on the client side, for example, `iqn.1994-05.com.ibm:apacibmi74falcon`.

Important: To allow unauthenticated access, do *not* enable CHAP.

Creating an iSCSI target for the client

Select the newly created host client to create an iSCSI target for the client:

1. Click the **Manage** icon and select **Create iSCSI Target**.
2. Enter the iSCSI target name as the one configured on the client side, for example, `iqn.2000-03.com.falconstor:h21-47.ibmi94`.
3. Select the IP address of the adapter on the StorSafe server. Leave the starting LUN as **0**. Create one iSCSI target for each iSCSI client initiator. Refer to “Configure an iSCSI client” for more information.
4. Select the newly created host client so that you can assign virtual tape libraries to the new iSCSI target.
5. Click the **Assign** icon and select the libraries to assign them.

On the client side, to discover assigned devices, run the IPL I/O processor option to send a login request. You can run the IPL using an SQL command or the Start System Service Tools command (STRSST):

1. Run the SQL command with the IPL I/O processor option:

```
CALL QSYS2.CHANGE_IOP(IOP=>'ISCSI', OPTION=>'IPL')
```

2. Run STRSST on the client to check the system bus resource and initiate a bus reset to the iSCSI bus resource.
3. Select option 6 I/O debug to 298A-001 IOP resource (Figure 4-61)



Figure 4-61 Logical hardware resources under SST

4. Run the IPL I/O processor option that sends a login request from the client iSCSI initiator to the VTL server with a nonexistent target. Use option IPL I/O processor (Figure 4-62)



Figure 4-62 IPL I/O processor

If everything is successful, the StorSafe server receives the iSCSI login request. Example 4-17 is an example of messages in the system log.

Example 4-17 System log messages example

```
May 30 14:34:23 h21-47 fsiscsid[9033]: IPSTOR||1653892463||I||0x0000c351||Login
to the target %1 from the initiator %2||iqn.2000-03.com.falconstor:h21-
47.ibmi94||iqn.1994-05.com.ibm:apacibmi74falcon
May 30 14:34:23 h21-47 kernel: FSISCSI client 14 initiator iqn.1994-
05.com.ibm:apacibmi74falcon type 2 login request to target iqn.2000-
03.com.falconstor:h21-47.ibmi94 from 172.24.2.94, conn 4020293, new tcp session
May 30 14:34:23 h21-47 kernel: FSISCSI conn 4020293 create new session 62603.
May 30 14:34:24 h21-47 kernel: svdp_get_cpu: vdev 536 cpu 31.
May 30 14:34:24 h21-47 kernel: svdp_get_cpu: vdev 537 cpu 32.
May 30 14:34:25 h21-47 kernel: svdp_get_cpu: vdev 538 cpu 33.
May 30 14:34:25 h21-47 kernel: [vtl_tde_537|4018142] TLE_INFO: VDrive 537
bPowerOnReset is set to 1, CDB[0]=0h vtape=-1 [n/a]
May 30 14:34:25 h21-47 kernel: IOCORE1 [kworker/31:1|3523459] release_vdev,
releasing vdev 536 without a reservation
May 30 14:34:26 h21-47 kernel: [vtl_tde_538|4018147] TLE_INFO: VDrive 538
bPowerOnReset is set to 1, CDB[0]=0h vtape=-1 [n/a]
```

5. Confirm that the iSCSI device is available to the client. For example, the 3584-403 device type is displayed as a FalconStor vendor device configured for the client (Figure 4-63 on page 185).

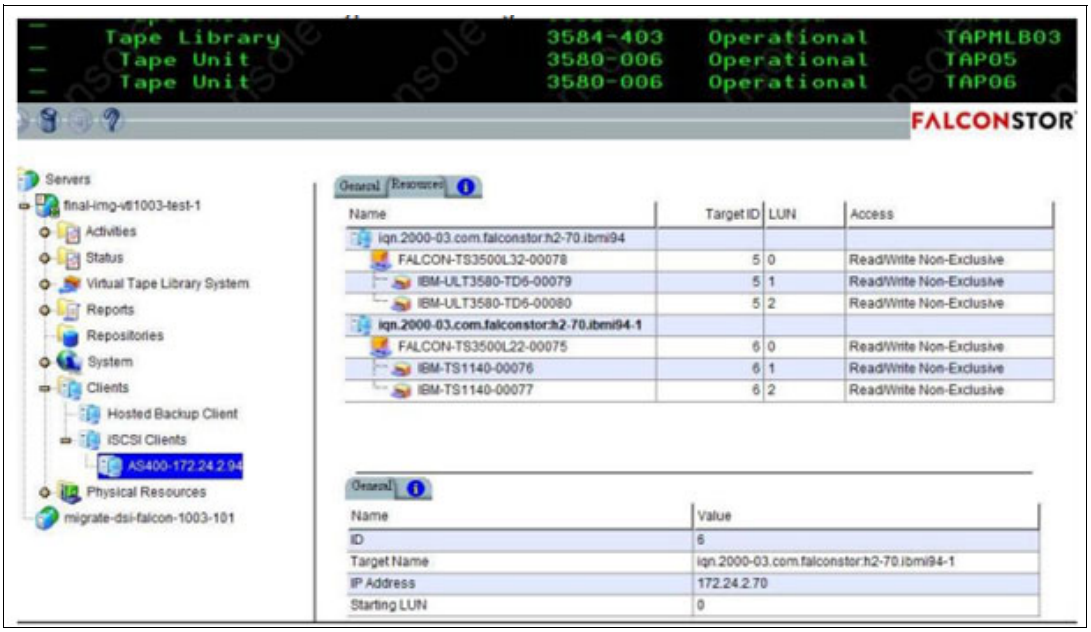

Figure 4-63 iSCSI device on StorSafe

Configuring deduplication

You can use the default deduplication policy or create deduplication policies by using the Deduplication tab.

You can define a new deduplication policy:

1. Enter a name for the policy.
2. Select the Inline deduplication trigger.
3. You can accept the default settings for priority and retry.
4. If applicable, on the primary StorSafe server, designate the replica server as the target server and create deduplication/replication policies.

On the IBM i host client, start your backup software to write data on virtual tapes.

Configuring an iSCSI client

This section provides an example of configuring an IBM i host client with StorSafe by using iSCSI. One iSCSI target is created in StorSafe for each iSCSI client initiator. The StorSafe Power Virtual Server instance is running Ethernet under an IBM Cloud private network that belongs to a VLAN subnet. No extra VLAN is set on the StorSafe server.

Verifying type and model

Confirm the IBM PTF is installed on the client by verifying that the listed Type-Model is 298A-001 (Figure 4-64). This indicates that the iSCSI bus IOP resource is operational on the client. If it does not exist, contact IBM to install the required PTF file.

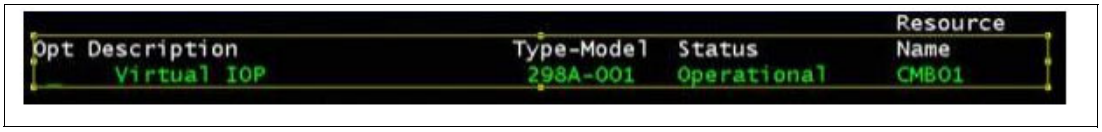

Figure 4-64 Type model for Virtual IOP

Creating an iSCSI target

For IBM i version 7.3 or later, use the IBM Navigator for i GUI to create an iSCSI target.

1. Select **System Services** →**iSCSI Tab** →**Action** →**Create iSCSI Target**
2. Enter the iSCSI Qualified name (IQN) and the IP address or hostname of the target.

For versions earlier than IBM i version 7.3, use the SQL service commands (Figure 4-65):

1. Confirm the SQL service is operational by running the STRSQL command. This service is used for configuring communication between the client iSCSI initiator and the StorSafe iSCSI target.



Figure 4-65 SQL statements

2. Run the SQL command to add the iSCSI target and the initiator information on the client (Example 4-18). You can set the target to any value matching IQN patterns.

Example 4-18 SQL command to add the iSCSI target

```
CALL QSYS2.ADD_ISCSI_TARGET(
TARGET_NAME=>'iqn.2000-03.com.falconstor:h21-47.ibmi94',
TARGET_HOST_NAME=>'172.22.21.47',
INITIATOR_NAME=>'iqn.1994-05.com.ibm:apacibmi74falcon');
```

The iSCSI target name does not exist on StorSafe yet. After the SQL command runs, use the same target name on StorSafe when configuring the iSCSI client later.

3. Run the SQL command with the IPL I/O processor option that sends a login request from the client iSCSI initiator to the StorSafe server with a nonexistent target:

   ```
   CALL QSYS2.CHANGE_IOP(IOP=>'ISCSI', OPTION=>'IPL')
   ```

4. Examine the StorSafe system log, `/var/log/messages` to confirm that the IPL I/O processor ran successfully.
5. Record the target name that is listed in the log. The target name is used when you configure the iSCSI target in the StorSight. Example 4-19 provides an example message.

Example 4-19

```
May 30 14:31:36 h21-47 fsiscsid[9033]:
IPSTOR||1653892296||E||0x0000c352||Login request to nonexistent target %1
from initiator %2||iqn.2000-03.com.falconstor:h21-47.ibmi94 (ip
172.22.21.47)||iqn.1994-05.com.ibm:apacibmi74falcon
```

4.5 Power Virtual Instance

This section introduces the capability to perform snapshots, restores, and clones of Power Virtual Server Instances (VSIs).

4.5.1 Snapshot

By using the snapshot interface, a relationship can be created between source disks and target disksat time T1. Target disks are created as part of the snapshot API. The snapshot API tracks the delta changes that are made to the source disk beyond time T1. This feature enables the user to restore the source disks to their T1 state later.

Several use cases are available for the snapshot feature. For example, an administrator plans to upgrade the middleware on their system, but they also want to revert the middleware to its original state before proceeding with an upgrade.

If the upgrade fails, the source disk can be restored to its previous state by completing the following steps:

1. Start the snapshot API with the source disks where the middleware information is stored.
2. Upgrade the middleware.
3. If the upgrade fails, restore the source disks by using the snapshot that was created before the upgrade.
4. If the upgrade succeeds, delete the snapshot.

Note: Multiple snapshot operations can be conducted simultaneously. Concurrent snapshot operations can occur on a different set of disks.

Best practices

Consider the following best practices:

- Before you take a snapshot, ensure that all of the data is flushed to the disk. If you take a snapshot on a running VM and did not flush the file system, you might lose some content that is in memory.
- It is recommended to quiesce all of the applications that are on the snapshot volume.

Restrictions

Consider the following restrictions:

- Parallel VM snapshot operations from different VM nodes for the same shared volume are not allowed.
- A VM cannot be restored if a snapshot is taken and if cloned (full-copy) IBM FlashCopy® operations are running in the background. The FlashCopy operations must finish first.
- Some of the attributes of source disks cannot be changed while the disks are in a snapshot relationship. For example, the source disks cannot be resized when snapshot relationships are in place for those disks.
- Volumes that are in a snapshot relationship cannot be detached from the VM.

4.5.2 Cloning a volume

The clone operation creates a full copy of the volume. You can select multiple volumes and start a group clone operation. When multiple volumes are selected, the clone operation ensures that a consistent data copy is created.

The clone operation continues to copy data from the source disks to target disks in the background. Depending on the size of the source disks and the amount of data to be copied, the clone operation can take a significant amount of time.

You cannot modify the source or target disk attributes, such as disk size, while the clone operation is in progress.

Best practice

Quiesce all of the applications on the volume that you want to clone.

Restriction and consideration

When the clone operation is performed on a volume that is in-use, the Power Systems Virtual Server service creates a consistent group snapshot and re-creates the copy of the cloned volume by using the group snapshot.

4.5.3 Restoring a snapshot

The restore operation restores all of the volumes that are part of a VM snapshot back to the source disks. When it restores the VM, the Power Systems Virtual Server service creates a backup snapshot, which can be used if the restore operation fails.

If the restore operation succeeds, the backup snapshots are deleted. If the restore operation fails, pass in the `restore_fail_action` query parameter with a value of `retry` to attempt the restore operation again.

To roll back a previous disk state, include the **restore_fail_action** query parameter with a value of **rollback**. When the restore operation fails, the VM enters an Error state.

Best practice

During the restore operation, it is critical that your source disks be quiesced. Your source disks cannot be in use.

Considerations

Consider the following points:

- If the restore operation fails, contact your storage support administrator. A failed restore operation can leave behind incomplete states, which might require a cleanup initiative from an IBM operations team.
- If you choose to restore a shared volume on one VM, then while the restore is running, you cannot perform the snapshot, restore, clone, or capture operations on the other VMs that use the shared volume.

4.5.4 Snapshot operations by using IBM Cloud CLI

Important: The following examples are as-is and no support is provided.

Example 4-20 shows how to use IBM Cloud CLI and power-iaas plug-in to take a Snapshot

This example assumes that you installed IBM Cloud CLI and the power-iaas plugin in a Linux server. For detailed information about IBM Cloud CLI setup, see: IBM Power Virtual Server CLI Reference.

Example 4-20 Using IBM Cloud CLI to take a Snapshot

```
#!/bin/bash
# Constants #
error=0;
errmsg="";
# Cloud Account data
apikey="<APIKEY>";
Resource="<CRN Resource ID>";
# SSH Server
SvrAddr="<SSH Server Address>";
UserID="<SSH UserID>";
# Instance ID
InstanceName="<InstanceName>";
# Object Storage
Region="<REGION>";
AccessKey="<ICOS AccessKey>";
SecretKey="<ICOS SecretKey>";
############################################
SnapshotName="Snap_$InstanceName_"$(date +"%Y-%m-%d_%I%M%S");
############################################




echo '================';
echo 'Suspend Db2 for i';
echo '================';
#  Quiesce the database - Suspend Database disk activity
ssh $UserID@$SvrAddr  "system 'CHGASPACT ASPDEV(*SYSBAS) OPTION(*FRCWRT)';system
'CHGASPACT ASPDEV(*SYSBAS) OPTION(*SUSPEND) SSPTIMO(120)'"

# Login to the account
ibmcloud login --apikey $apikey -r $Region
if (($error == 0))
    then
    echo '================'
    echo 'Starting Snapshot'
    echo '================'
    # List instance volumes to an auxiliary file
        ibmcloud pi ins vol list $InstanceName |cut  -c1-36|grep -v
ID>/tmp/volume_list.lst

    # Create variable with volume names
```

```
        vols='';
    while read line
        do
        vols=$vols$line','
    done < /tmp/volume_list.lst
    # Remove last ','
    vols=${vols:0:-1}

    # Start the snapshot process
    rm /tmp/snapshot_result.json
    ibmcloud pi snap create $InstanceName --volumes "$vols" --name $SnapshotName
--json> /tmp/snapshot_result.json || { error=1;errmsg="Error while taking the
snapshot";echo $errmsg; }
    SnapshotID=$(cat /tmp/snapshot_result.json|jq .snapshotID)
    # Remove the double quote on this variable
    SnapID=${SnapshotID//'"'/''}
fi
if (($error == 0))
    then
    echo '========================================='
    echo 'Wait 2 minutes to reach a consistent state'
    sleep 2m
    echo '========================================='
    echo 'Resuming database activity'
    # Resume database activity
    ssh $UserID@$SvrAddr "system 'CHGASPACT ASPDEV(*SYSBAS) OPTION(*RESUME)'" || {
error=1;errmsg="Error resuming database activity";echo $errmsg; }
fi
if (($error == 0))
    then
echo 'Snapshot process has finished'
else
echo 'Error during the snapshot process'
    echo ''
    echo $errmsg
fi
```

Obtaining snapshot details

By using values from Example 4-20 on page 189, you can use the commands in Example 4-21 to list the available snapshots in an instance.

Example 4-21 Listing available snapshots

```
ibmcloud login --apikey $apikey -r $Region
ibmcloud pi snap list
```

Restoring to snapshot

Example 4-22 shows an example of restoring to a snapshot.

Example 4-22 Restoring to snapshot

```
ibmcloud login --apikey $apikey -r $Region
ibmcloud pi snap restore $InstanceName --snapshot $SnapID
```

Important: The instance must be in a Power Off state before running a snapshot restore operation.

Exporting your instance

You can use the export procedure to save your instance data to the boot image repository or to create a compressed OVA in your IBM COS bucket.

Both can be used to create a new instance, or as a Full System Backup alternative.

Example 4-23 shows how to export your instance to your IBM COS by using the IBM Cloud CLI. Size restrictions apply.

Example 4-23 Exporting your instance with IBM Cloud CLI

```
#!/bin/bash
# Constants #
error=0;
errmsg="";
# Cloud Account data
apikey="<APIKEY>";
Resource="<CRN Resource ID>";
# SSH Server
SvrAddr="<SSH Server Address>";
UserID="<SSH UserID>";
# Instance ID
InstanceName="<InstanceName>";
# Object Storage
Region="<REGION>;
AccessKey="<ICOS AccessKey>";
SecretKey="<ICOS SecretKey>";
ImageName='IBMiPVS_'$InstanceName'_'$(date -I);
###########################################
echo '================';
echo 'Suspend Db2 for i';
echo '================';
#  Quiesce the database - Suspend Database disk activity
ssh $UserID@$SvrAddr  "system 'CHGASPACT ASPDEV(*SYSBAS) OPTION(*FRCWRT)';system
'CHGASPACT ASPDEV(*SYSBAS) OPTION(*SUSPEND) SSPTIMO(120)'"

# Login to the account
ibmcloud login --apikey $apikey -r $Region

if (($error == 0))
    then
    echo '============================='
    echo 'Start IBM i instance export job'
    echo '============================='
    # List instance volumes to an auxiliary file
        ibmcloud pi ins vol list $InstanceName |cut  -c1-36|grep -v
ID>/tmp/volume_list.lst

    # Create variable with volume names
        vols='';
    while read line
        do
```

```
        vols=$vols$line','
    done < /tmp/volume_list.lst
    # Remove last ','
    vols=${vols:0:-1}

    # Start the export process
    ibmcloud pi ins cap cr $InstanceName --destination $InCapDest --name
$ImageName --volumes "$vols" --image-path $BucketName --region $Region
--access-key $AccessKey --secret-key $SecretKey >
/tmp/$InstanceName_capture_job.txt || { error=1;errmsg="Error exporting
instance";echo $errmsg; }
fi
if (($error == 0))
    then
    echo '========================================'
    echo 'Wait 2 minutes to reach consistency point'
    sleep 2m
    echo '========================================'
    echo 'Resuming database activity'
    # Resume database activity
   ssh $UserID@$SvrAddr "system 'CHGASPACT ASPDEV(*SYSBAS) OPTION(*RESUME)'" || {
error=1;errmsg="Error resuming database activity";echo $errmsg; }
fi
if (($error == 0))
    then
    errmsg="Export: OK!"
    echo ""
    echo "OVA File >>"$ImageName"<< in process"
    echo "================================================================================"
    ibmcloud pi job ls   --operation-action vmCapture
    echo "================================================================================"
else
    if [[ "$errmsg" ==  "" ]]
    then
        errmsg="Export: With errors!"
        echo $errmsg
    fi
fi
```

4.5.5 Use cases

This section provides some use cases.

Snapshot and restore use case

This section describes the snapshot and restore use case, which is shown in Figure 4-66 on page 193:

- Capture the LUN's state before an upgrade, of, for example, an IBM i release:
  - A successful upgrade results in the snapshot relation being deleted.
  - A failed upgrade results in a rollback to the pre-upgrade state.

Capture and Export use case

A Capture and Export use case scenario is shown in Figure 4-67 on page 193:

- A production environment is exported to create a new QA instance.
- Save the exported OVA to IBM COS. Use IBM COS to regularly save as a DR strategy.

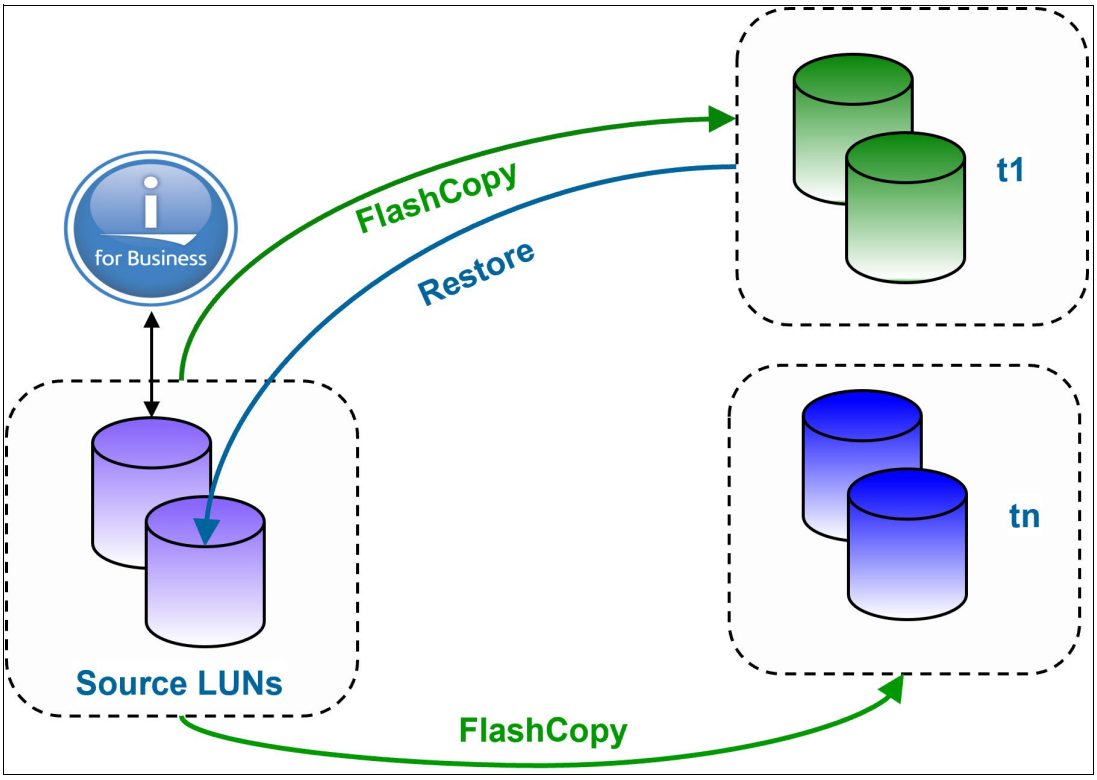

Figure 4-66 Use case - Snapshot and restore

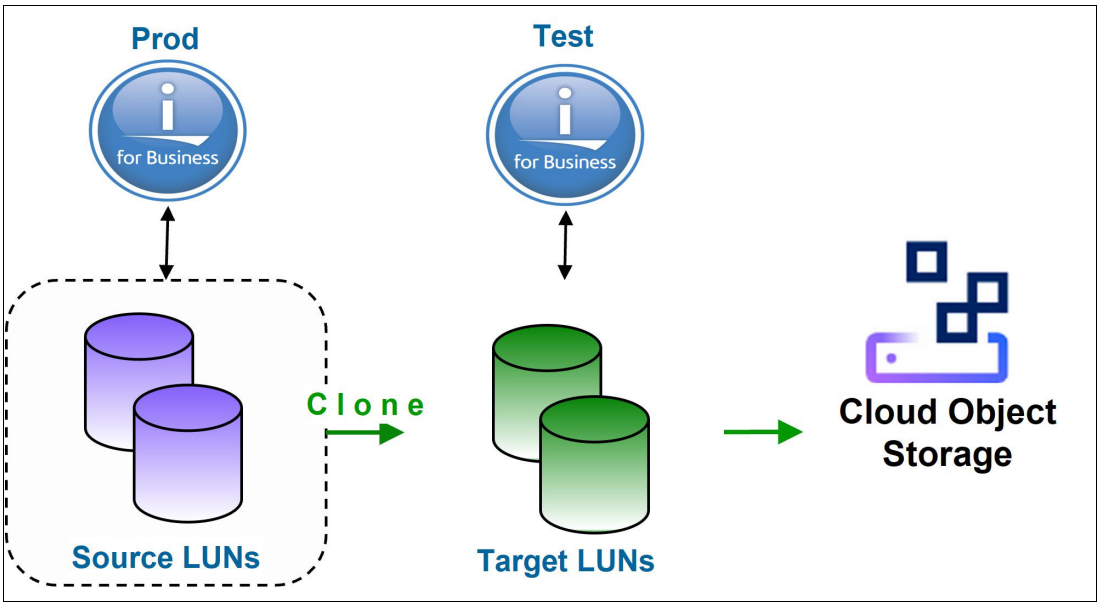

Figure 4-67 Use case - cloning

5

Hints and tips for IBM i deployments on IBM Power Virtual Server

This chapter is derived from real-life experiences and challenges while deploying and managing workloads on-premises and in the Cloud with the IBM Power Systems Virtual Server. The content in this chapter can help you become more familiar with the VNC console, SSH tunneling, accessing the console by way of the LAN adapter, and take snapshots on IBM i instances.

This chapter includes the following topics:

- “Connecting to an IBM i virtual machine”
- “Using snapshots on IBM i instances”
- “Power Virtual Server with VPC Landing Zone”

© Copyright IBM Corp. 2025.

5.1 Connecting to an IBM i virtual machine

For more information about configuring your IBM i VM before connecting to an IBM i VM, see Configuring your IBM i virtual machine (VM).

Complete the following steps to connect to an IBM i virtual machine (VM) after the system is deployed:

1. From the dashboard, click the button that features three stacked dashes that are at the end of the entry for `ibmiVS` (see Figure 5-1) and click **Open console**. Use the VNC to access the 5250 console in the IBM Power Systems Virtual Server for IBM i.

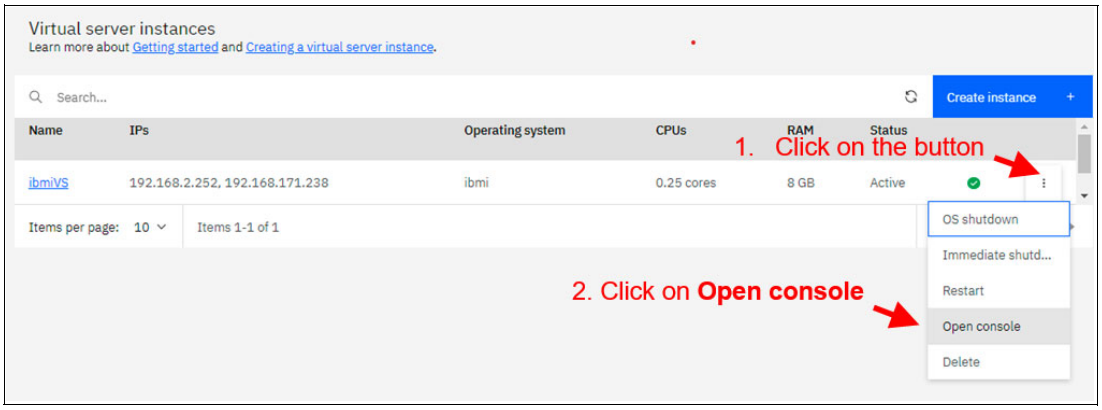

Figure 5-1 Opening console on Power Virtual Server for IBM i

Remember: When using IBM i stock images, the default values enforce case-sensitive passwords up to 15 characters.

2. After a new window opens to access the default credentials for the initial login by way of the web console, enter the user profile **`QSECOFR`** and password **`QSECOFR`** (as written, in uppercase), as shown in Figure 5-2.

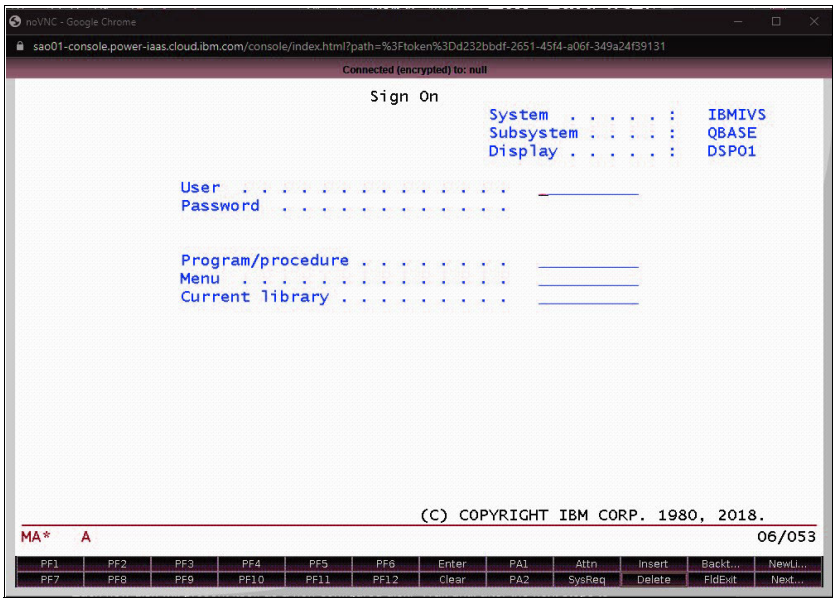

Figure 5-2 Opened console for IBM i on Power Virtual Server

3. Change the password for user profile QSECOFR, as shown in Figure 5-3.

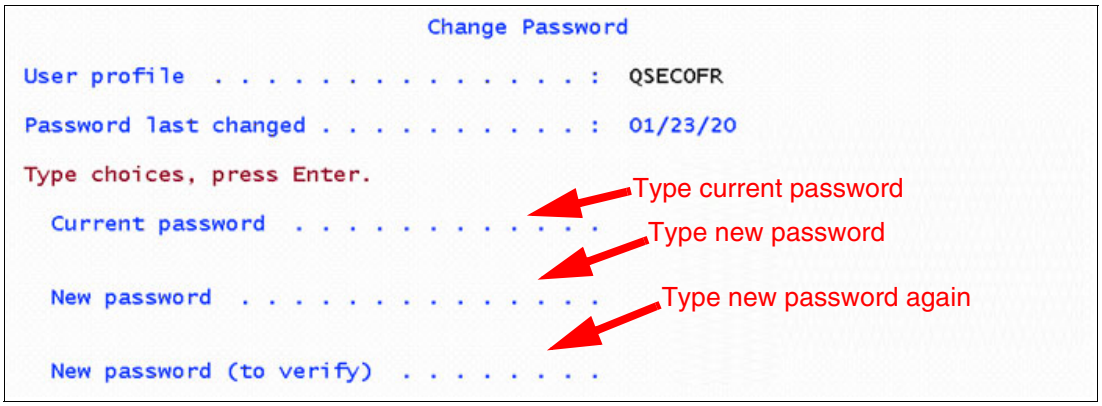

Figure 5-3 Change QSECOFR's password

4. Accept the software agreements by selecting Option 5 for each. Then, press Enter as shown in Figure 5-4.

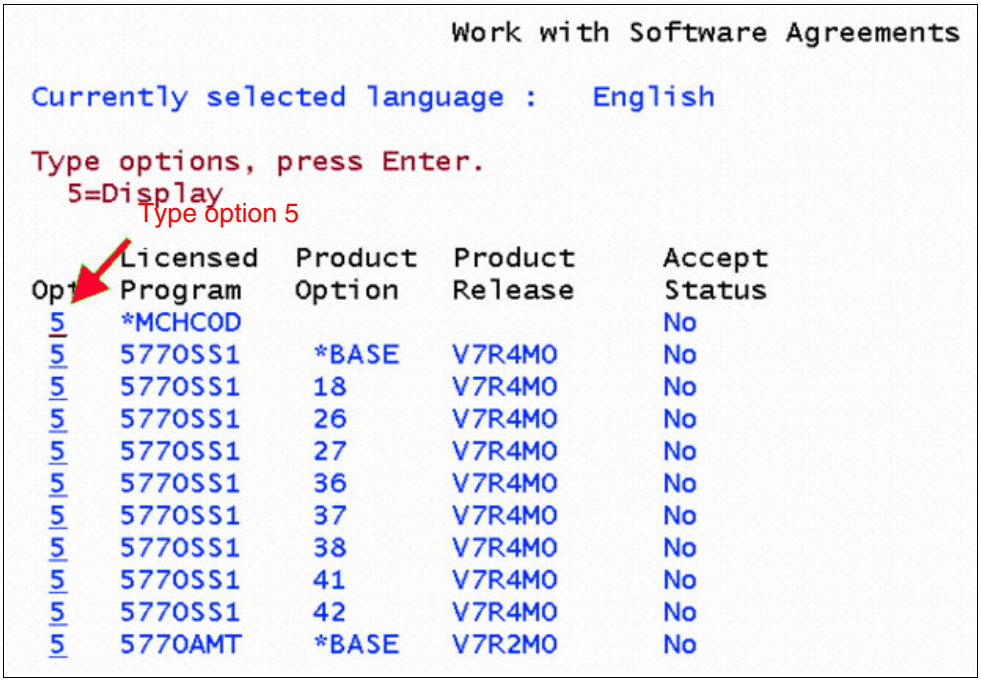

Figure 5-4 Accept software agreements

5. To use the keys that are available from VNC, click **Next**. Then, select **PF15= Accept All**, as shown in Figure 5-5 on page 198.

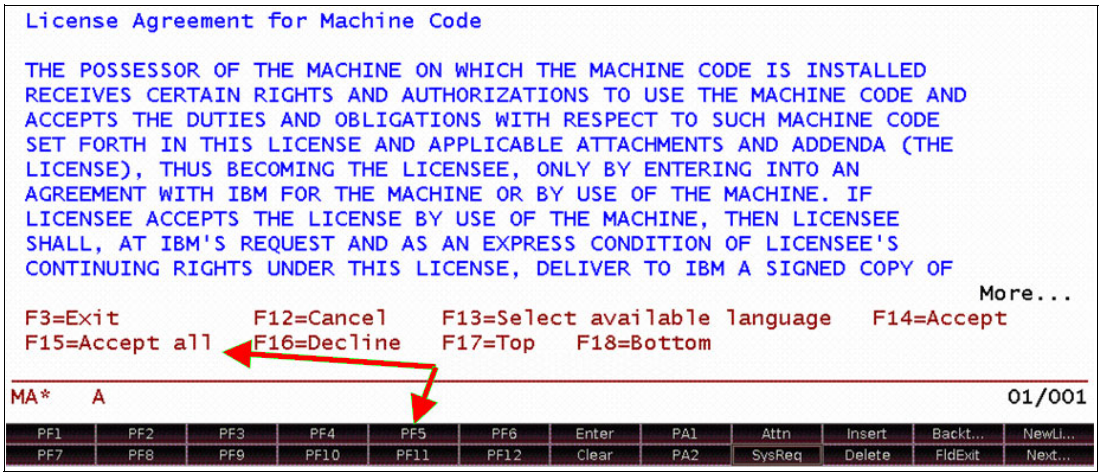

Figure 5-5 Accept the licenses agreements for Machine code

6. Repeat Step 5 until the display shows all of the software agreements are in the Accept status (see Figure 5-6) and Press **PF3**.

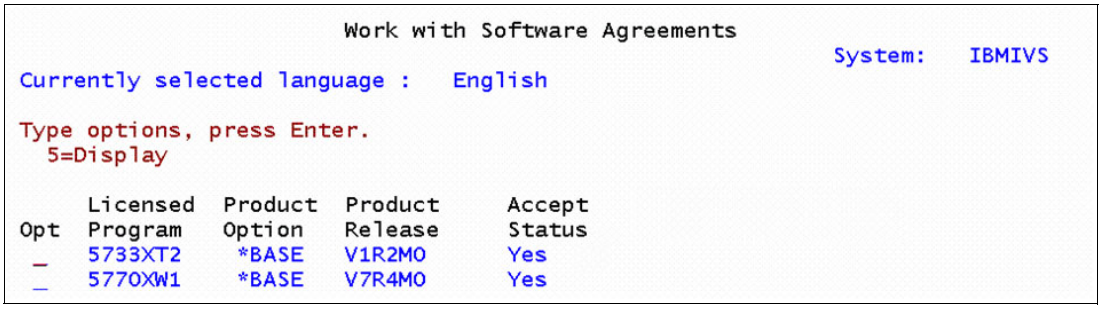

Figure 5-6 Work with Software agreements window

The Main Menu is displayed by the VNC.

7. Connect by SSH by using the `QSECOFR` profile. Then, enable the attribute by running the following command:

   ```
   EDTF '/QOpenSys/QIBM/UserData/SC1/OpenSSH/etc/sshd_config'
   ```

8. Find the **#PermitRootLogin yes** and **#UseDNS no** in the file and remove the # symbol for each.
9. Set the correct value when required.
10. At the end of the file, add **ibmpaseforienv PASE_USRGRP_LIMITED=N**, and press **PF3** to save and exit.

> **Attention:** This `QSECOFR` profile is for demonstration purposes. The `QSECOFR` profile is used for the SSH connection. However, the preferred practices suggest to *not* use `QSECOFR` profile with SSH on IBM i. Use a different user profile with similar authority. The new user profile also has a Home directory, which is usually `/home/<userprofile>`.

11. Run the **WRKSYSVAL QAUTOVRT** command and press Enter.
12. Select **option 2 = Change** and enter the number of auto-configured consoles to which you want to connect concurrently. In this example, four consoles are used. Press Enter (see Figure 5-7 on page 199).

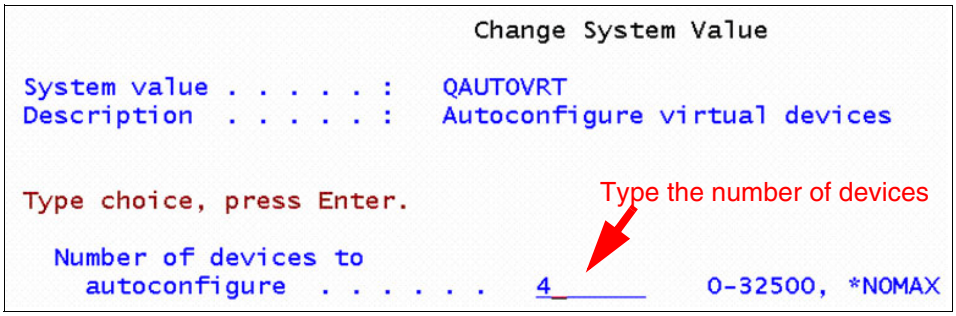

Figure 5-7 Change QAUTOVRT system value

Note: You can add as many devices that you require on `QAUTOVRT`.

13. Issue the **`WRKSYSVAL QLMTSECOFR`** command and press Enter
14. Select **option 2 = Change**. Then, change the value from `1` to `0` and press Enter.
15. Verify that `cloud-init` is configured as shown in Figure 5-8. Run the **`CFGTCP`** command, select **Option 1** and press Enter.

```
192.168.2.17      255.255.255.0     CLOUDINIT1    *ELAN
192.168.171.236   255.255.255.248   CLOUDINIT2    *ELAN
```

Figure 5-8 TCP/IP on IBM i configured by cloud-init

16. Issue the following commands to start the Telnet and SSH services:

```
STRTCPSVR *TELNET
STRTCPSVR *SSH
```

17. Verify that the ports for SSH and Telnet are listening, as shown in Figure 5-9. Issue the **`NETSTAT *CNN`** command and press Enter.

```
Remote          Remote     Local
Address         Port       Port       Idle Time   State
*               *          ssh        000:02:58   Listen
*               *          telnet     000:03:04   Listen
```

Figure 5-9 Display NETSTAT on IBM i

18. Press **Shift + PF2**.

5.1.1 Remote access to IBM i by using tunneling

The public IP address blocks most ports. Therefore, you must use SSH tunneling or configure your certificates and use SSL to allow IBM Access Client Solution to connect over public IP.

Before you use an SSH tunnel, you must create a user profile with `USRCLS(*SECOFR)` specified.

Complete the following steps to access the solution by using SSH tunneling:

1. Open a PuTTY terminal and create a session that uses the public IP, as shown in Figure 5-10 on page 200.
   a. Enter an external IP address

b. Select **SSH**.

c. Enter a name in the Saved Sessions field.

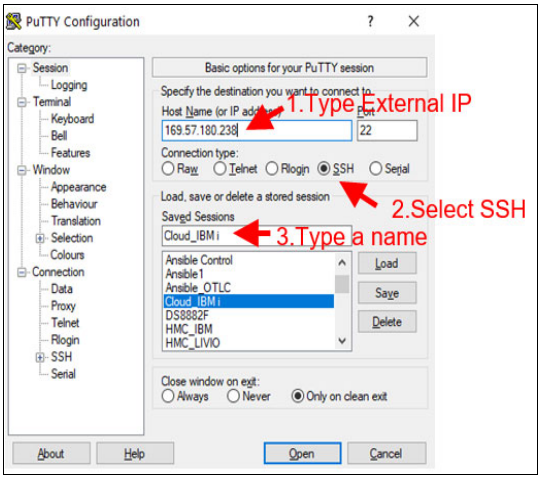

Figure 5-10 Set PuTTY terminal

The ports that are listed in Table 5-1 must be configured on PuTTY.

Table 5-1 Ports that must be configured on PuTTY

Source port	Destination
449	localhost:449
50000	localhost:23
8470	localhost:8470
8471	localhost:8471
8472	localhost:8472
8473	localhost:8473
8474	localhost:8474
8475	localhost:8475
8476	localhost:8476
9470	localhost:9470
9471	localhost:9471
9472	localhost:9472
9473	localhost:9473
9474	localhost:9474
9475	localhost:9475
9476	localhost:9476

2. Complete the following steps:

a. Set up the source port and destination as shown in Figure 5-11 on page 201. Select **Connection** → **SSH** → **Tunnels**.

b. Under Port forwarding, ensure both checkboxes are selected.

3. For each port listed in Table 5-1 on page 200 use the following steps:
   a. In the Source port field, enter a value from the Source port column.
   b. In the Destination field, enter the corresponding value from the Destination port column.
   c. Click **Add**.

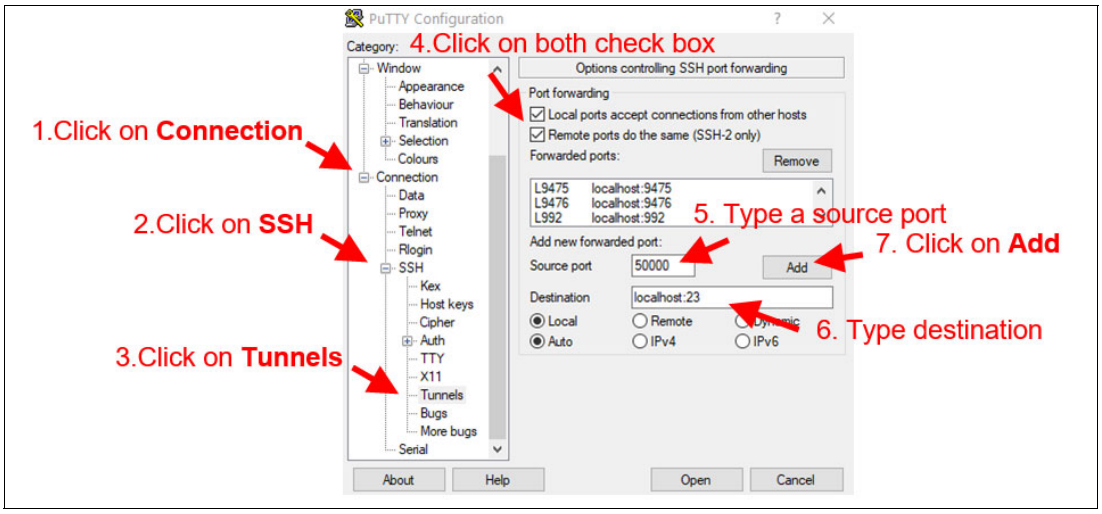

Figure 5-11 Set the source ports on PuTTY

> **Note:** Steps 5, 6, and 7 that are shown in Figure 5-11 are repeated for each port that is in Table 5-1 on page 200.

   d. Scroll up to the beginning of the Category section that is on the left side of PuTTY and select **Session**.
   e. Click **Save**.
   f. Click **Open** at the session in PuTTY and click **Yes** to trust this host and connection. A new window terminal opens.
   g. Enter `QSECOFR` and the password that was changed, and press **Enter**. The tunneling is now ready. Ensure the terminal window (Figure 5-12) stays open.

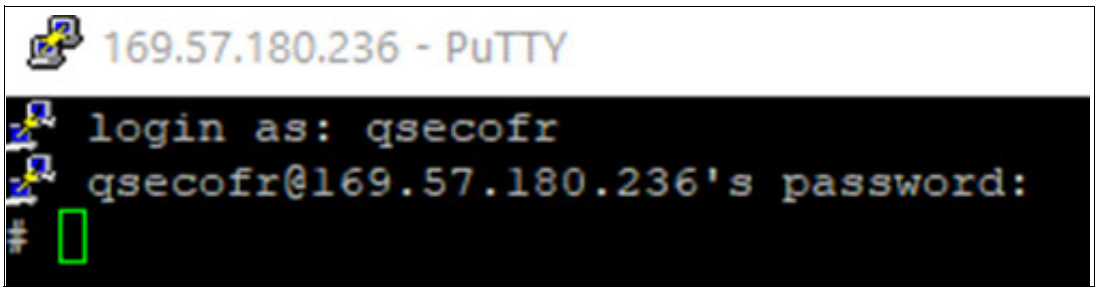

Figure 5-12 PuTTY log in for QSECOFR

Important: The tunneling that is done by using PuTTY is for a Windows system. If you use another operating system, such as Linux or Apple iOS, the SSH tunneling to allow ACS to connect over the External IP is different. For more information, see Using SSH tunneling to allow ACS to connect over the public IP.

4. Select **Access client solutions** → **Management** → **5250 Session Manager** → **New Display Session**. A new window opens, and the 5250 display setup is shown (see Figure 5-13). In the figure, the Destination address is LOCALHOST and the Destination port is 50000.
5. After entering the values, click **OK**.

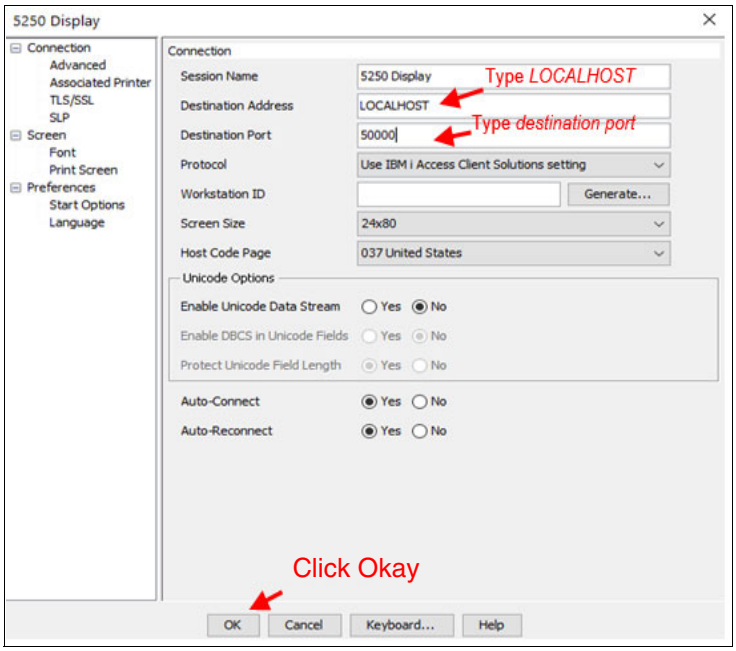

Figure 5-13 5250 display configurations

Important: For this example, 50000 was chosen as the source port number. This value was also configured in PuTTY. Do not change the source port numbers. When telnetting, avoid making the source port the same as the destination.

6. A window opens in which the connection with the PuTTY tunnel is established. Enter the credentials of the QSECOFR profile and their password.
7. Select **File** → **Save as** to save the 5250 display.

Note: A secure Telnet session is commonly used on-premises. However, because the use of SSH tunnel encryption is working in this example, port 992 is *not* necessary.

5.1.2 IBM i 5250 console through LAN adapter

When you work with IBM i instances on IBM Power Systems Virtual Server, a 5250 console is required to perform migration, backups, and maintenance procedures.

IBM Power Systems Virtual Server instances provide a web-based console from the portal, which is based on noVNC and HTML5. Any HTML5 browser can be used to open our console session.

Figure 5-3 on page 197 shows the web-based console that presents the following limitations:

- Session expires after approximately 5 minutes of inactivity.
- Options that are available at the bottom of the window can be confusing because they include function keys and other special keys.
- Cut and paste cannot be used in the window.

Tip: IBM i instance connects to a remote LAN console by using IBM i Access Client Solutions. With IBM i ACS, full access is available to the physical keyboard, sessions do not drop unless the connection fails, data can be inserted by using cut and paste.

Configuration procedure

IBM i Access Client Solutions (ACS) must be installed so that the IBM i ACS console can be configured.

Download the installation package from one of the following web pages:

- IBM Support: IBM i Access - Client Solutions
- IBM Entitled Systems Support

Note: The client software is available free of charge. The license is installed on the IBM i instance.

Configure a Private Network and add it to the instance. In this example, the `192.168.80.0/24` network was added.

Tip: A ticket in IBM Cloud is required to activate the VLAN and subnet. For more information, see IBM Cloud Support.

Configuring the Service Tools Server LAN adapter

Complete the following steps:

1. Enter the command **`STRSST`** to start the Service Tools by using `QSECOFR` or a user with *SERVICE authority.
2. In the console window, connect by using the `QSECOFR` user (the default password is `QSECOFR`).
3. In the DST menu, select **option 8 – Work with Service Tools Server Security and Devices**. Select **F13=Select STS LAN Adapter**, as shown in Figure 5-14 on page 204. (Use NEXT to see more Fn options.)

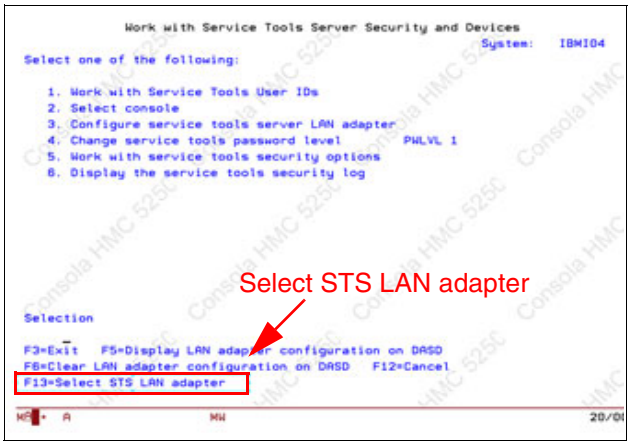

Figure 5-14 Select STS LAN adapter at SST on IBM i

4. Select the Resource Name that is connected to your Private Network. If nothing is shown, press **F21**, and the adapters that are in use are shown (see Figure 5-15).



Figure 5-15 Select Service Tools Server LAN adapter

5. Assign an IP address to the LAN Adapter, as shown in Figure 5-16. Press **F7** to save the changes. Press **F13** to deactivate the adapter, then press **F14** to reactivate the adapter.

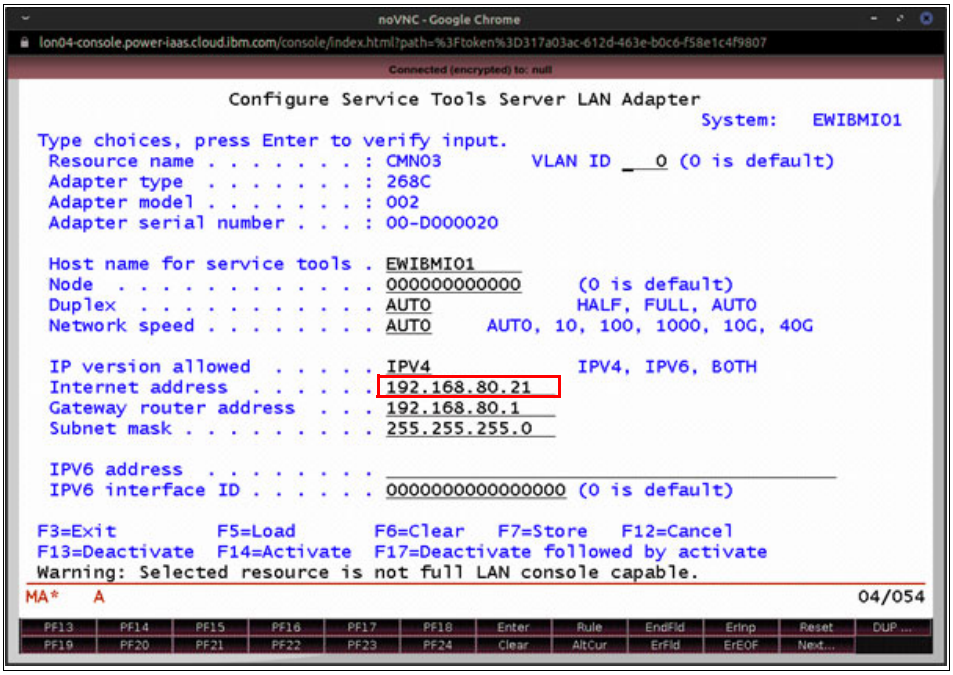

Figure 5-16 Assigning IP on LAN adapter

Note: When an adapter is in use, IPL your IBM i system before using the LAN Console.

6. Configure the console at IBM Access Client Solutions, as shown in Figure 5-17. Click **System Configuration** → **New**. Enter the System Name, which is often the IP address and click **OK**.

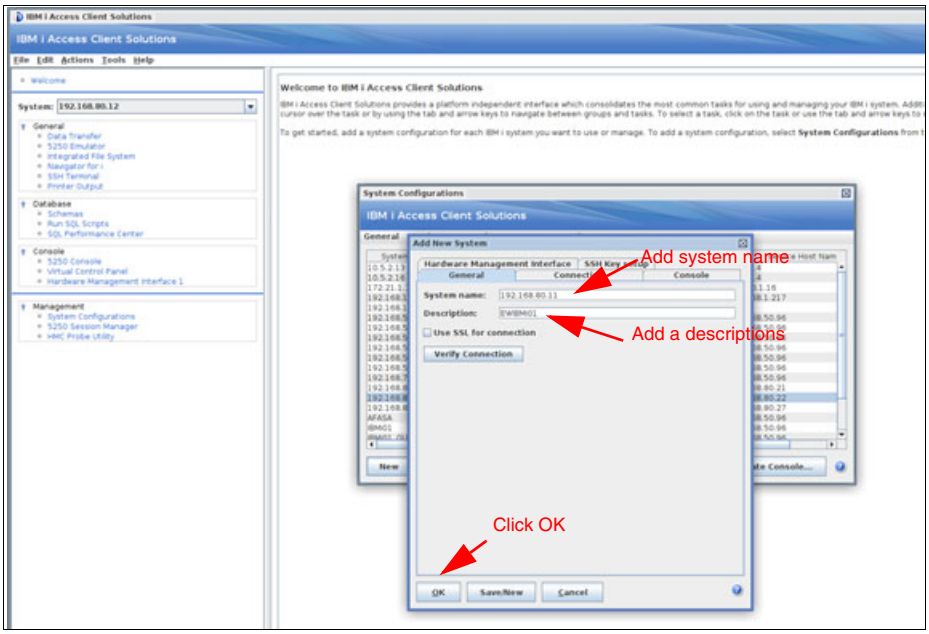

Figure 5-17 Adding a system on Access Client Solutions

7. Select the Console tab and select the Service hostname, as shown in Figure 5-18. This name is the IP address for the Service Tools Server LAN Adapter.

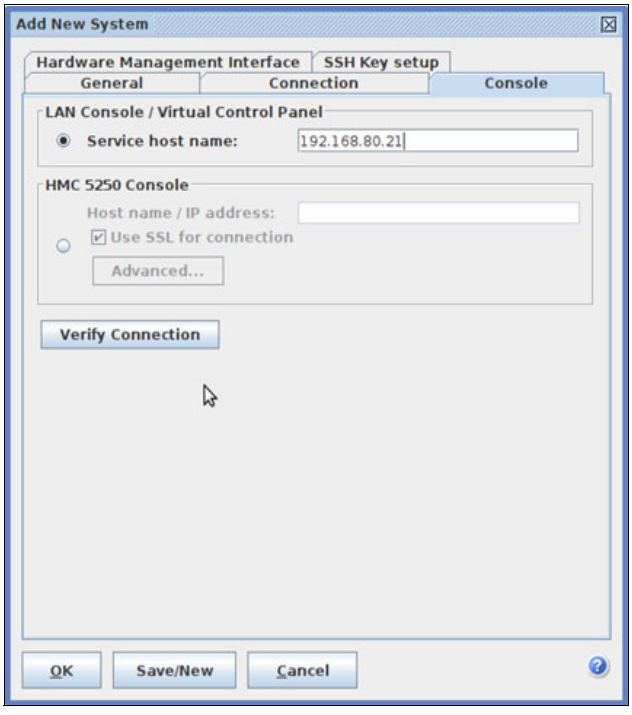

Figure 5-18 Insert the IP address

8. Verify the connection. If all of the information is correct, click **OK** and start the console connection.

5.2 Using snapshots on IBM i instances

Snapshot is a resource used to create a checkpoint in IBM i VMs for a possible future rollback. It uses copy-on-write procedures to minimize snapshot time to near zero, which allows the VM to be restored quickly.

5.2.1 Uses for snapshots

To ensure data integrity, a disk stage must be completed on IBM i to save data that is cached in memory.

Snapshots are useful *before* the following change management tasks are conducted:

- Performing an OS upgrade
- Installing PTFs
- Changing system values
- Updating application programs

Remember: Snapshot is *not* a backup mechanism. It cannot use the snapshot in a different system than the source instance. Also, snapshot data cannot be moved to other medium and snapshots cannot be mounted on a new system as is done with PowerHA SystemMirror for backup purposes.

5.2.2 Taking snapshots by installing and using IBM Cloud CLI

Snapshots can be taken from the command line only by using APIs or IBM Cloud CLI. Users must be familiar with IBM Cloud CLI to perform this task.

For more information about installing the IBM Cloud CLI on Windows, see Getting started with the IBM Cloud CLI.

Note: For Windows, some functions require the latest operating system release.

Complete the following steps:

1. Open a PowerShell window as Administrator and run the command that is shown in Example 5-1.

Example 5-1 Running the installation command

```
[Net.ServicePointManager]::SecurityProtocol = "Tls12, Tls11, Tls, Ssl3";
iex(New-Object
Net.WebClient).DownloadString('https://raw.githubusercontent.com/IBM-Cloud/ibm-clo
ud-developer-tools/master/windows-installer/idt-win-installer.ps1')
```

   This process can take some minutes. Wait until the end of the process and restart the system.

2. In a command window, run the command that is shown in Example 5-2 on page 207 to determine whether the setup processes succeeded.

Example 5-2 Checking ibmcloud cli setup process

```
ibmcloud dev help
```

If the help is shown, the process can continue.

3. Connect to your account as shown in Example 5-3 and continue installing the plug-ins.

Example 5-3 Logging in to your account

```
ibmcloud login
```

The system prompts you for your account's email and password.

If you use more than one account, select the account you use and the region by using the item number that is shown in the window.

4. To install the required plug-in to work with Power Systems Virtual Server, run the command that is shown in Example 5-4.

Example 5-4 Install Power IaaS plug-in

```
ibmcloud plugin install power-iaas
```

5. Click **Y** when you are prompted to continue with the setup process.

   The available workspaces in this account are then listed (see Example 5-5).

Example 5-5 List available workspaces

```
ibmcloud pi workspace list
```

6. Target the service by running the **ibmcloud pi workspace target <crn>** command, as shown in Example 5-6.

Example 5-6 Target the workspace to use

```
ibmcloud pi workspace target
crn:v1:bluemix:public:power-iaas:lon04:a/af339adbfd124f99a5cea8271bf030cc:5fa3d138-6b
c4-409f-ba6d-641c071b8e51::
```

7. To list the available instances, run the command that is shown in Example 5-7.

Example 5-7 List instances

```
ibmcloud pi instance list
```

8. By using this list, you can copy the instance ID that is frozen by using snapshot.
9. Return to the green window console or terminal session and perform a disk stage. Then, quiesce the database, as shown in Example 5-8. All data that is cached in memory is written to disk, and the transactions are held in memory until the snapshot is done.

Example 5-8 Perform a disk stage and quiesce database

```
CHGASPACT ASPDEV(*SYSBAS) OPTION(*FRCWRT)
CHGASPACT ASPDEV(*SYSBAS) OPTION(*SUSPEND) SSPTIMO(60)
```

This action must be performed on any available ASP before the snapshot command is run.

10. After the data is written to disk and transactions are held in memory, use the ibmcloud CLI to take the snapshot. Run the **ibmcloud pi instance snapshot create** command and target your VM instance. Then, choose a name to identify the snapshot (see Example 5-9 on page 208).

Example 5-9 Taking the snapshot by using ibmcloud cli

```
ibmcloud pi instance snapshot create EWO1-IBMiO1 --name SNPO1_EWO1-EWIBMiO1
```

Important: If you have more volumes in addition to the boot disk, create a comma separated list with the volumes id, using the --volumes parameter.

11. Run the **ibmcloud pi instance snapshot list <InstanceName>** command to list the snapshot. The snapshot state is next to the Instance ID and Snapshot ID. Wait for the available status.
12. When the status is available, resume database activity, as shown in Example 5-10.

Example 5-10 Resuming database activity

```
CHGASPACT ASPDEV(*SYSBAS) OPTION(*RESUME)
```

13. To restore the snapshot, power-off your VM instance.
14. Run the following command to restore the snapshot data, as shown in Example 5-11:

    i**bmcloud pi instance snapshot restore <instance_id> --snapshot <snapshot_id>**

Example 5-11 Restoring the snapshot

```
ibmcloud pi instance snapshot restore EWO1-IBMiO1 --snapshot
57122996-ac9b-4aec-a356-YYYYYYYYYYYY
```

15. Run the **ibmcloud pi instance snapshot list** command to see the restore operation status.
16. Start the VM instance after the snapshot is restored.

5.3 Power Virtual Server with VPC Landing Zone

The Power Virtual Server with VPC Landing Zone deployable architectures are a set of IaC assets based on the IBM Cloud for Financial Services® reference architecture. You can use the deployable architectures to create a secure and customizable IBM Cloud Power Virtual Server Infrastructure offering according to the best practices and requirements from the IBM Cloud.

Provisioning Power Virtual Server with VPC landing zone by using deployable architectures provides an automated deployment method to create an isolated Power Virtual Server workspace and connect it with IBM Cloud services and public internet. Network management components like DNS, NTP, proxy servers and NFS as a Service might be installed. Comparing the provisioning through the projects UI, user interaction is minimized and ready-to-go deployment time of a Power Virtual Server workspace is reduced from days to less than 1 hour.

Automated Power Virtual Server with VPC landing zone provisioning that is described in this guide is based on IBM Cloud catalog deployable architectures. This document describes specifics that are related only to Power Virtual Server with VPC landing zone deployable architecture.

Figure 5-19 on page 209 describes the deployable architecture variants.

Variation	Available on IBM Catalog	Requires IBM Schematics Workspace ID	Creates VPC Landing Zone	Performs VPC VSI OS Config	Creates PowerVS Infrastructure	Creates PowerVS Instance	Performs PowerVS OS Config
Standard	✓	N/A	✓	✓	✓	N/A	N/A
Standard Extend	✓	✓	N/A	N/A	✓	N/A	N/A
Quickstart (Standard plus VSI)	✓	N/A	✓	✓	✓	✓	N/A
Import	✓	N/A	N/A	N/A	N/A	N/A	N/A

Figure 5-19 Solution Overview

5.3.1 Standard variation

This deployable architecture variation supports these features:

- A VPC Infrastructure with the following components:
  - One RHEL VSI for management (jump/bastion).
  - One RHEL VSI for network-services configured as squid proxy, NTP, and DNS servers using Ansible Galaxy collection roles ibm.power_linux_sap collection. This VSI also acts as a central ansible execution node.
  - Optional Client to site VPN server.
  - Optional File storage share.
  - Optional Application load balancer.
  - IBM Cloud Object Storage (COS) Virtual Private endpoint gateway (VPE).
  - IBM COS Instance and buckets.
  - VPC flow logs.
  - KMS keys.
  - Activity tracker.
  - Optional Secrets Manager Instance with private certificate
- A local or global transit gateway
- A Power Virtual Server workspace with the following network topology:
  - Creates two private networks: a management network and a backup network
  - Attaches the Power Virtual Server workspace to transit gateway
  - Creates an SSH key
  - Imports catalog stock images

Figure 5-20 on page 210 shows the mapping of the ‘Standard’ variation.

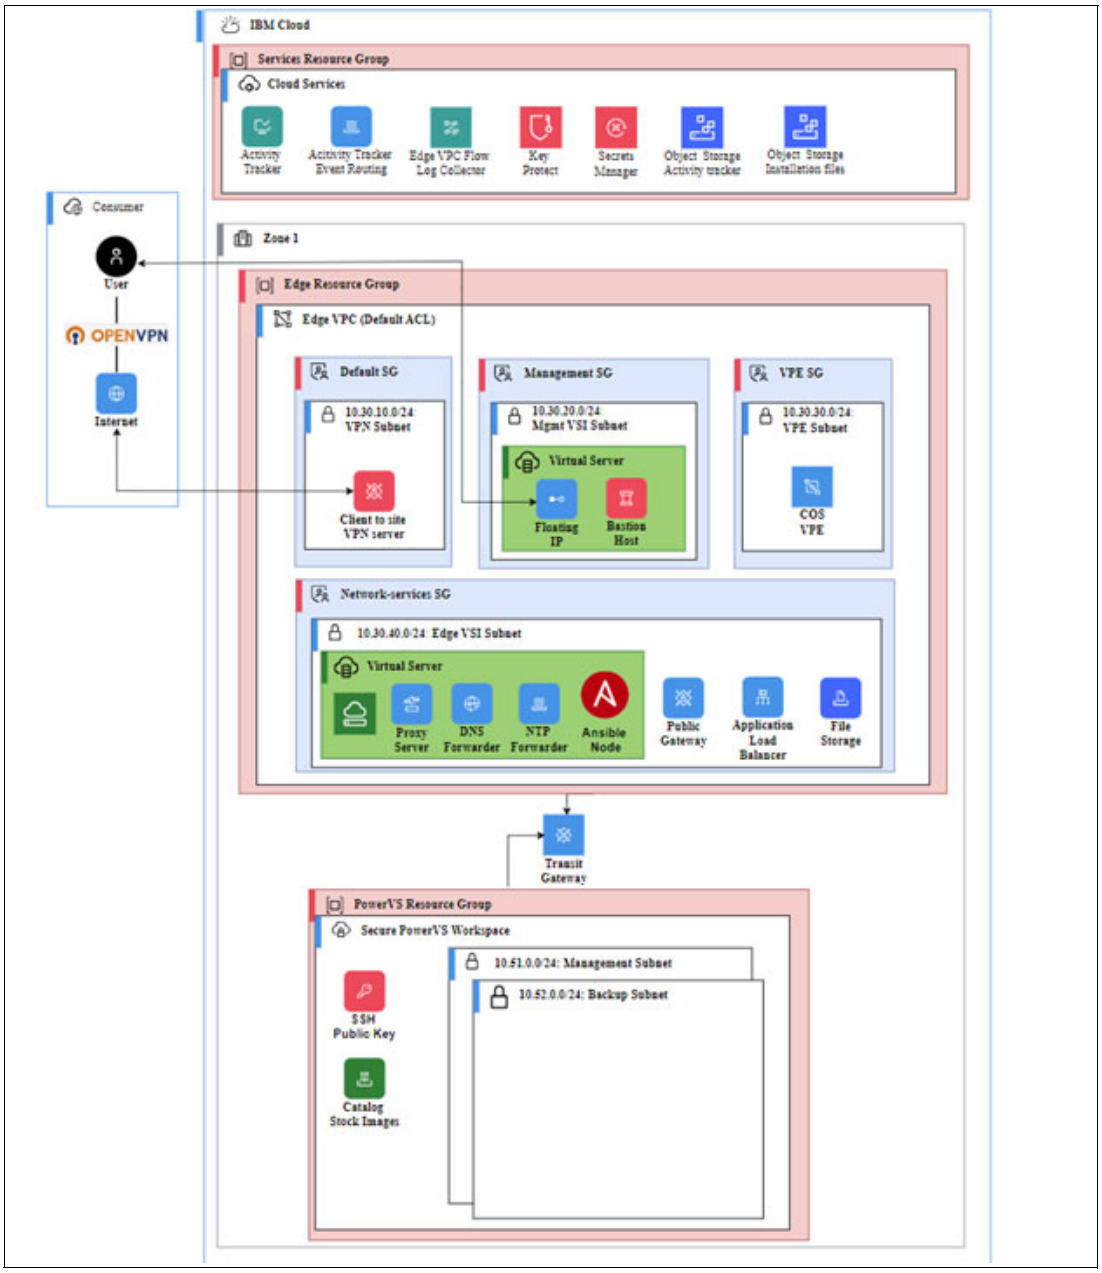

Figure 5-20 Power Virtual Server with VPC landing zone 'Standard' variation

5.3.2 Extend Power Virtual Server with VPC landing zone - standard variation

Important: This variation has a prerequisite. Deploy the 'Create a new architecture Standard' variant first.

The 'Extend Power Virtual Server with VPC landing zone' variation creates an additional Power Virtual Server workspace and connects it to the existing Power Virtual Server with VPC landing zone. It builds on existing Power Virtual Server with VPC landing zone deployed as a variation 'Create a new architecture'. This is typically used for High Availability scenarios in the same regions.

A Power Virtual Server workspace with the following network topology:

- Creates two private networks: a management network and a backup network
- Attaches the Power Virtual Server workspace to transit gateway
- Creates an SSH key
- Imports catalog stock images

5.3.3 Quickstart variation

This deployable architecture variation supports these features:

- A VPC Infrastructure with the following components:
  - One RHEL VSI for management (jump/bastion).
  - One RHEL VSI for network-services configured as squid proxy, NTP, and DNS servers using Ansible Galaxy collection roles (see ibm.power_linux_sap collection). This VSI also acts as central ansible execution node.
  - Optional Client to site VPN server.
  - Optional File storage share.
  - Optional Application load balancer.
  - IBM COS Virtual Private endpoint gateway(VPE).
  - IBM COS Instance and buckets.
  - VPC flow logs.
  - KMS keys.
  - Activity tracker.
  - Optional Secrets Manager Instance with private certificate
- A local or global transit gateway
- A Power Virtual Server workspace with the following network topology:
  - Creates two private networks: a management network and a backup network
  - Attaches the Power Virtual Server workspace to transit gateway
  - Creates an SSH key
  - Imports cloud catalog stock images
- A Power Virtual Server Instance with following options:
  - t-shirt profile (Aix/IBMi/SAP Image)
  - Custom profile (cores, memory, storage, and image)
  - 1 volume

Figure 5-21 on page 212 shows the structure of a ‘Quickstart’ variation.

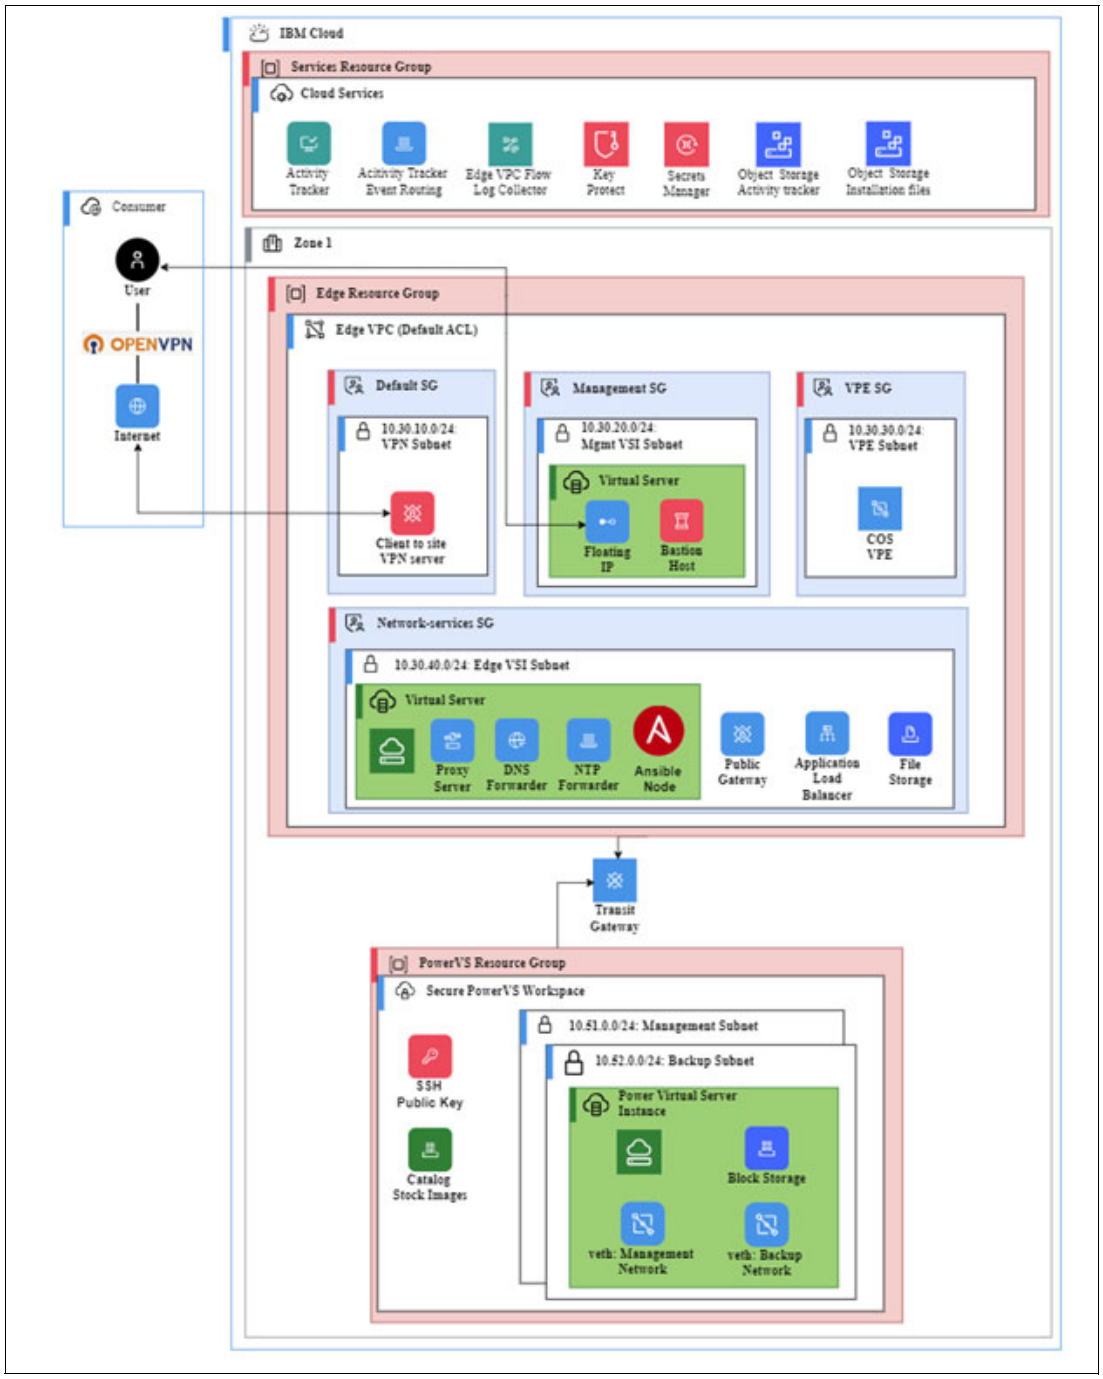

Figure 5-21 Power Virtual Server with VPC landing zone 'Quickstart' variation

You can run AIX, IBM i, and Linux images on your virtual server instances. Select the required T-shirt size and a virtual server instance with chosen T-shirt size or custom configuration is deployed. The T-shirt sizes and the configuration parameters mapping are shown in Figure 5-22.

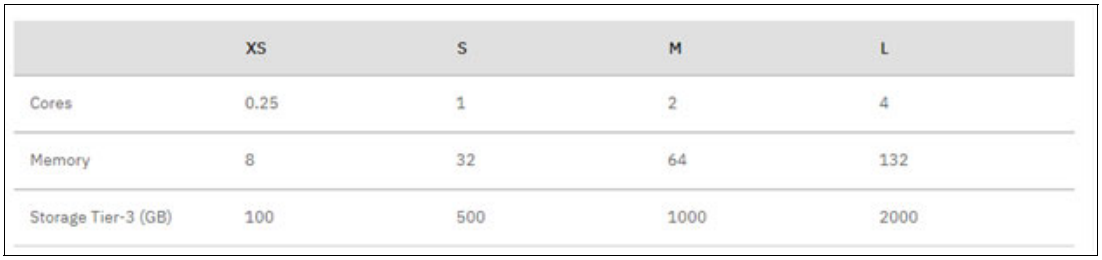

	XS	S	M	L
Cores	0.25	1	2	4
Memory	8	32	64	132
Storage Tier-3 (GB)	100	500	1000	2000

Figure 5-22 IBM i T-shirt size and the configuration mapping

5.3.4 Import variation

Create an IBM Cloud schematics workspace for your pre-existing VPC and Power Virtual Server infrastructure resources using the Power Virtual Server with VPC landing zone variation - 'Import Power Virtual Server Workspace'.

This variation helps to install the deployable architecture 'Power Virtual Server for SAP HANA' on top of a pre-existing Power Virtual Server landscape. 'Power Virtual Server for SAP HANA' automation requires a schematics workspace id for installation. The 'Import' solution creates a schematics workspace by taking pre-existing VPC and Power Virtual Server infrastructure resource details as inputs. The ID of this schematics workspace is the pre-requisite workspace id required by 'Power Virtual Server for SAP HANA' to create and configure the Power Virtual Server instances for SAP on top of the existing infrastructure.

For more information, see IBM Cloud solution for Power Virtual Server Pre-requisites.

6

Disaster Recovery with IBM Power Virtual Server

This chapter discusses the disaster recovery (DR) process that uses IBM Power Systems Virtual Servers. This chapter also describes the two phases of DR configuration for IBM i workloads on IBM Power Systems Virtual Servers, how to realize the most benefits from it, and how to improve the DR times by using this solution.

This chapter includes the following topics:

- “Disaster recovery by using system-level replication”
- “OS level replication with PowerHA SystemMirror for i geographic mirroring”
- “Logical replication use case with Bus4i”
- “Disaster Recovery with storage replication”

© Copyright IBM Corp. 2025.

6.1 Disaster recovery by using system-level replication

Uptime is a key client expectation for critical IBM i workloads. Whether the requirement is for high availability (HA) solutions within a local data center or for a multi-site DR solution between different geographical locations, these objectives are achieved through HA/DR solutions.

IBM Power Virtual Server meets this requirement so that clients can take advantage of DR solutions between two IBM i Virtual Server Instances (VSIs) in separate IBM Cloud data centers. DR solutions for IBM Power Virtual Server can be based on operating system level replication or logical replication.

Important: With the availability of Global Replication Service (GRS) within Power Virtual Server, IBM i can now use storage level replication between two Power Virtual Server data centers. This is discussed in section 6.4, “Disaster Recovery with storage replication” on page 295

Replication solutions between two data centers always require a secure network connection between the data centers so that the data can be transmitted securely. This requirement also applies to DR with IBM Power Virtual Server, which requires specific networking steps in IBM Cloud before implementing the replication solution.

This chapter addresses the two phases of DR configuration for IBM i workloads on IBM Power Systems Virtual Server:

1. Performing the required network configuration
2. Implementing the DR solution:
   - Operating system level replication use case with IBM PowerHA SystemMirror for i geographic mirroring
   - Logical Replication use case with TSP Bus4i

6.1.1 Solution components

DR solutions with IBM Power Systems Virtual Server use the following components:

- An IBM Cloud account
- Two Power Systems Virtual Server location services and a private subnet in each Power Systems Virtual Server location
- An IBM provisioned in each Power Systems Virtual Server location
- Orders for a Direct Link Connect Classic to connect each Power Systems Virtual Server location to IBM Cloud
- Orders for two Vyatta Gateways (one in each data center) to allow for Power Systems Virtual Server location-to-location communication
- Request for a Generic Routing Encapsulation (GRE) tunnel to be provisioned at each Power Systems Virtual Server location
- Three GRE tunnels in the Vyatta Gateways: two to connect Vyatta Gateway to the Power Systems Virtual Server location GRE tunnels and one across Vyatta Gateways to connect Vyatta-to-Vyatta
- A Reverse-proxy Centos VSI that is configured to allow access to Private Cloud Object Storage endpoint from Power Systems Virtual Server location

6.1.2 Solution requirements

For the solution requirements, follow these steps:

1. Open an IBM Cloud account.

   Log in to `https://cloud.ibm.com` and follow the procedure to open an internal or external account.

   For internal accounts, you can use your IBM intranet ID and password.

   For external accounts, you must provide a billing source, such as a credit card.

2. Create Power Systems Virtual Server location Service and Subnets.

   All Power Systems Virtual Server Instances are provisioned in what is called a Power Systems Virtual Server location. This separate data center is near IBM Cloud data centers.

   To set up your Power Systems Virtual Server location, you set up a Power Systems Virtual Server location service in the IBM Cloud UI. You can use the Power Systems Virtual Server location service within IBM Cloud to provision IBM i VSIs.

   Only one Power Systems Virtual Server location service can be set up per data center in IBM Cloud. This scenario has two Power Systems Virtual Server data center locations: one in Toronto and one in London.

   Before provisioning a Power Systems Virtual Server Instance in the Power Systems Virtual Server location, create at least one subnet. You can use as many subnets as you require in each Power Systems Virtual Server location service on which you can provision your Power Systems Virtual Server Instances. More subnets beyond the initial one can be added later after the VSIs are created.

   You also need an IBM Cloud account to start this process.

3. Go to the main IBM Cloud UI page and click **Catalog**, as shown in Figure 6-1.

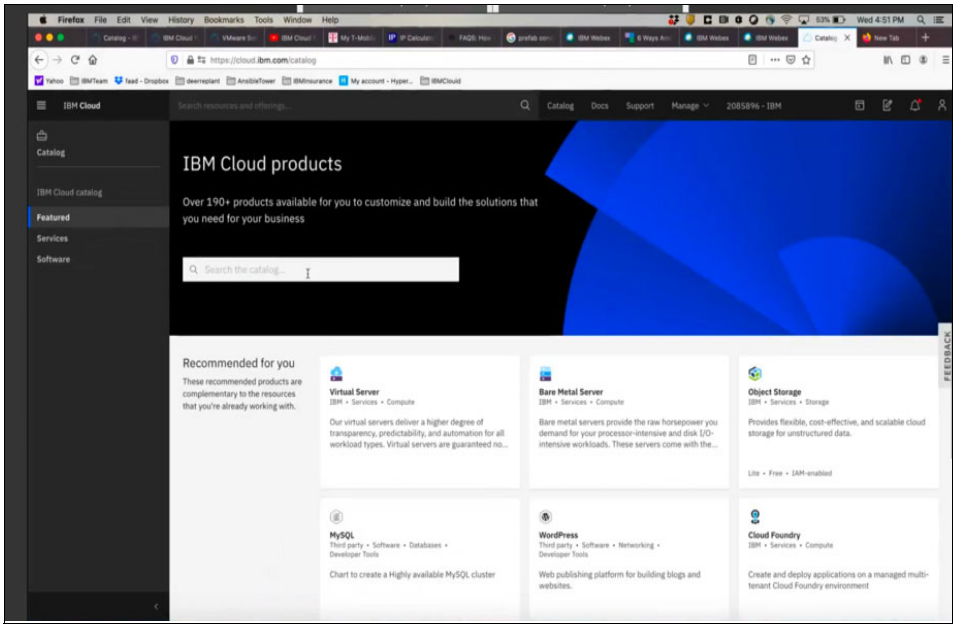

Figure 6-1 IBM Cloud products

4. Search the results for Power and select **Power Systems Virtual Servers**, as shown in Figure 6-2.

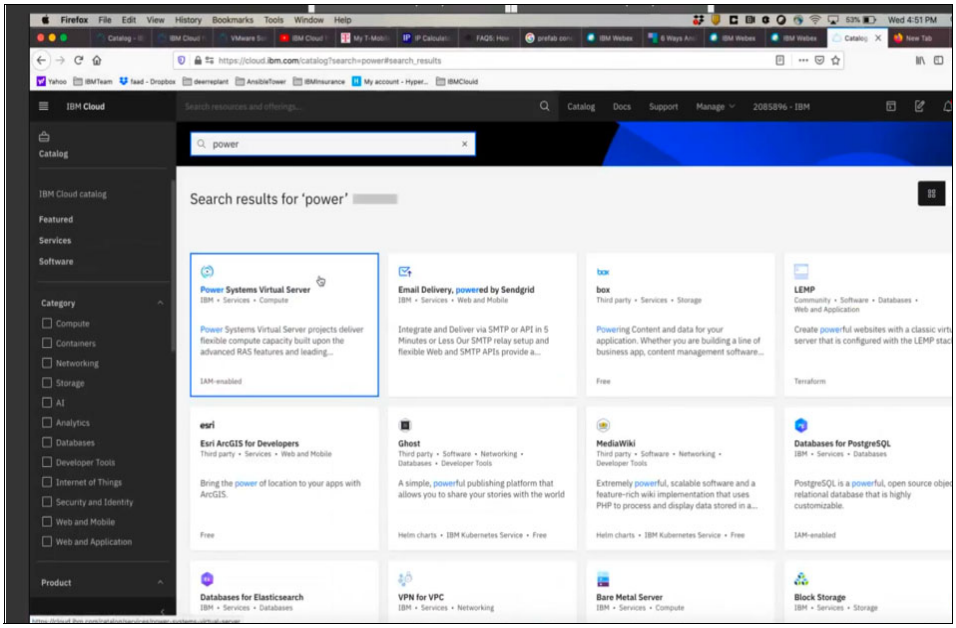

Figure 6-2 Results for Power

5. Under Select Region, choose your region, as shown in Figure 6-3. You are limited to only one service per region.

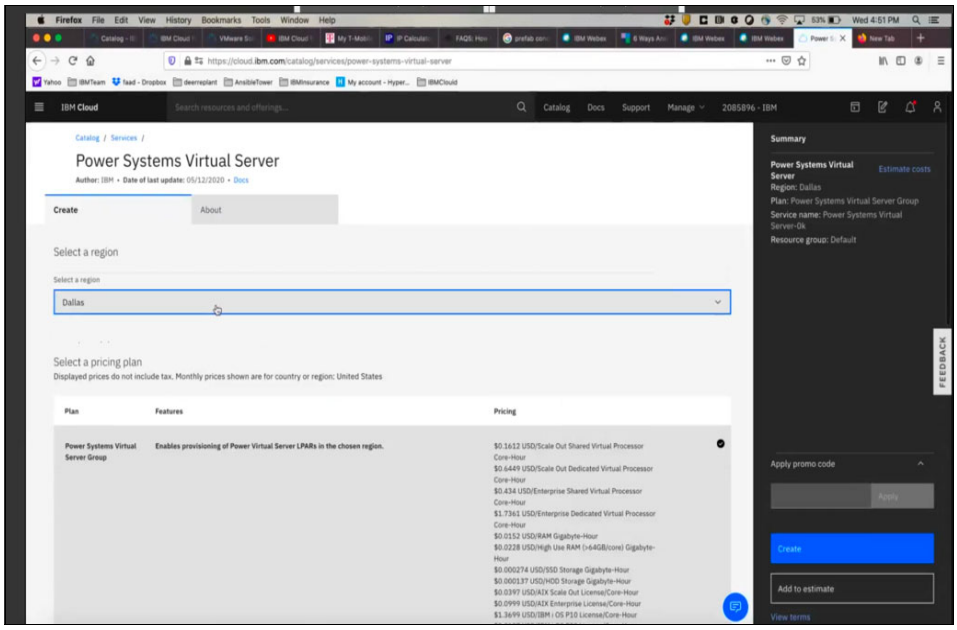

Figure 6-3 Selecting region

6. Select a Service Name or chose the default name.
7. Click **Create**, as shown in Figure 6-4.

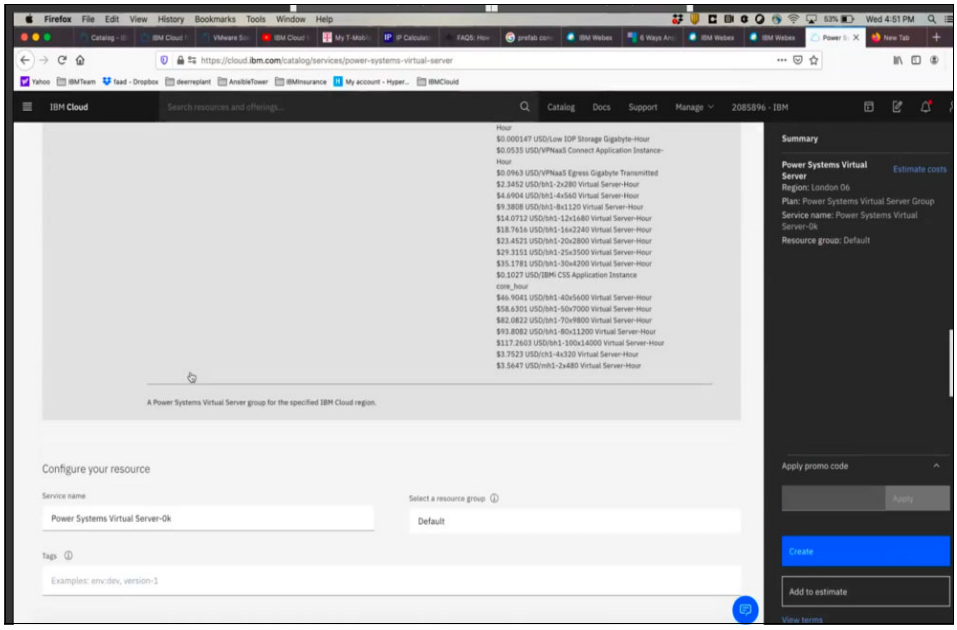

Figure 6-4 Selecting a service name

Your Power Systems Virtual Server location service now appears under the Services tab, as shown in Figure 6-5.

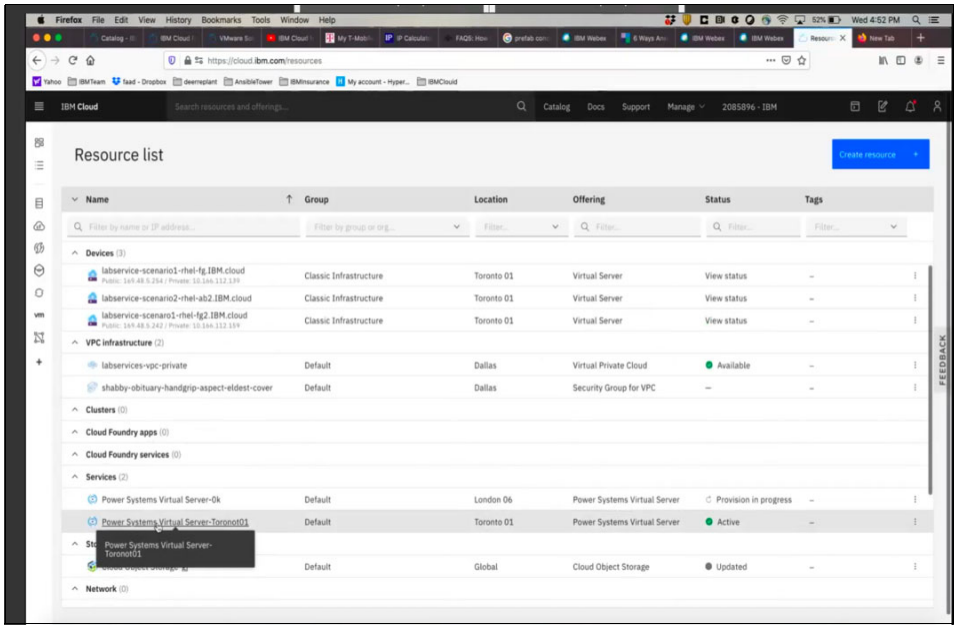

Figure 6-5 Power Systems Virtual Server instance

Repeat this process to create a second Power Systems Virtual Server location service.

8. Click the Power Systems Virtual Server location Service that you created and provision a subnet to be used by your Power Systems Virtual Server Instances (see Figure 6-6).

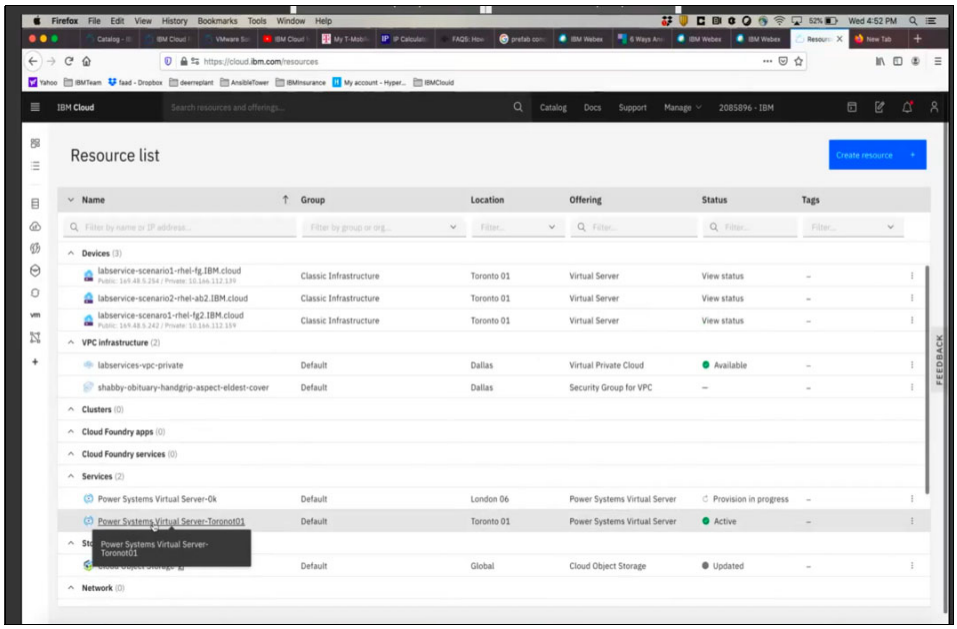

Figure 6-6 Resource list

9. Choose **Subnets** from the menu on the left, as shown in Figure 6-7.

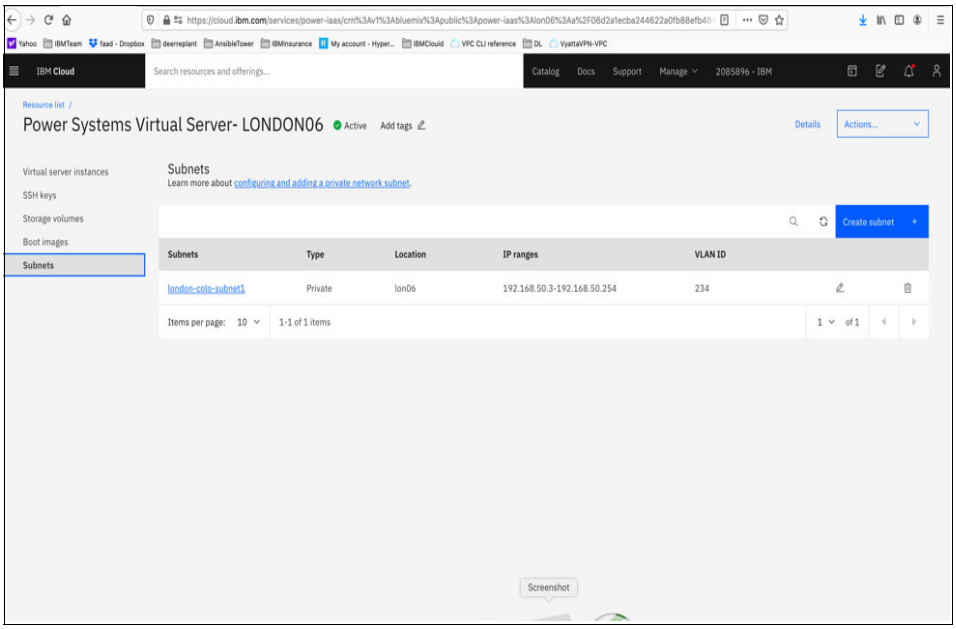

Figure 6-7 Subnets

10. Provide the following information:

- A name for your subnet, as shown in Figure 6-8 on page 221.
- CIDR range. This range can be any private IP subnet ranges. You choose /21 - /30 based on how many IPs you require. You can also use your own private CIDR if you want.

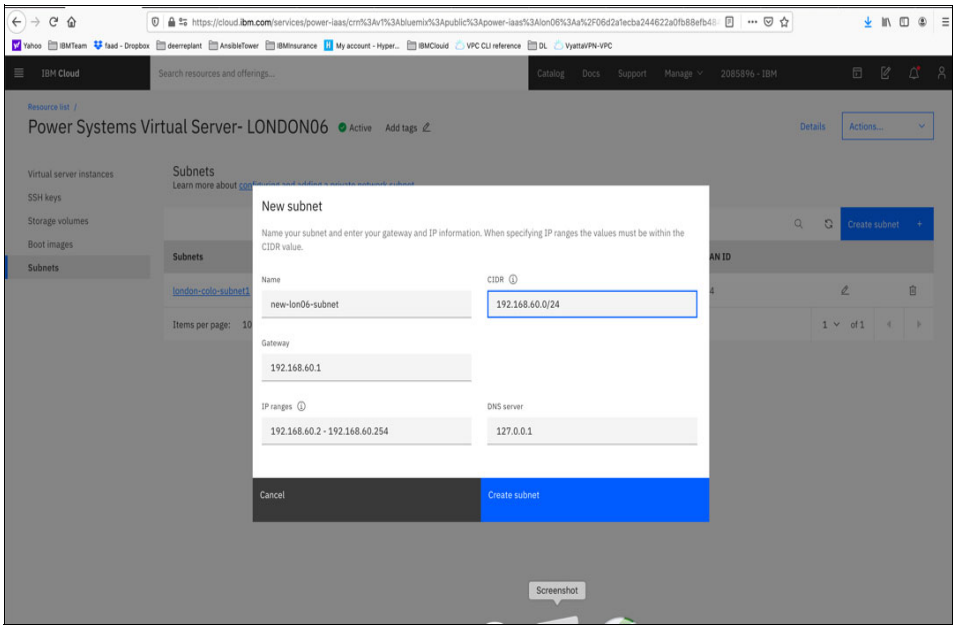

Figure 6-8 Subnet configuration

The rest of the fields are automatically populated based on the CIDR that you provided.

11. Click **Create Subnet**. A VLAN ID is then associated with the subnet, as shown in Figure 6-9.

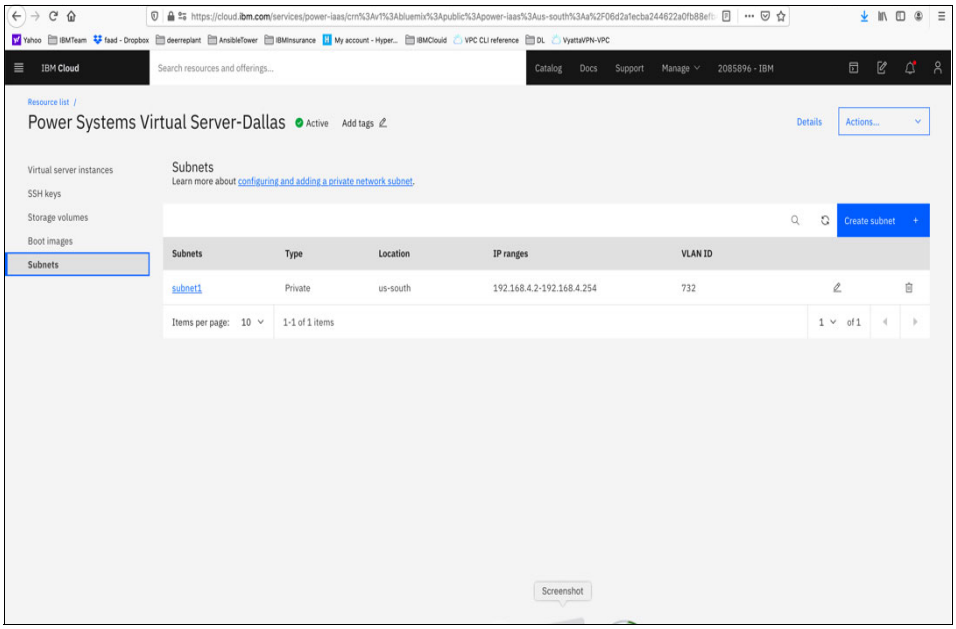

Figure 6-9 Creating a subnet

12. Open a Support Ticket with Power Systems to request that the subnet be configured to allow local communication between any Power Systems Virtual Server Instance you create in this Power Systems Virtual Server location service. Provide your Power Systems Virtual Server location service and your subnet in the ticket.

Without completing this step, the Power Systems Virtual Server Instances that you create cannot ping between each other, even if they are on the same subnet in the same Power Systems Virtual Server location.

Provisioning IBM i VSIs in each Power Systems Virtual Server location

This procedure creates an IBM i 7.4 VSI. The costs that are shown are monthly costs, but they are charged hourly.

Complete the following steps:

1. Go to the IBM Cloud Catalog and click **IBM Cloud** at the upper left of the UI, as shown in Figure 6-10.

b

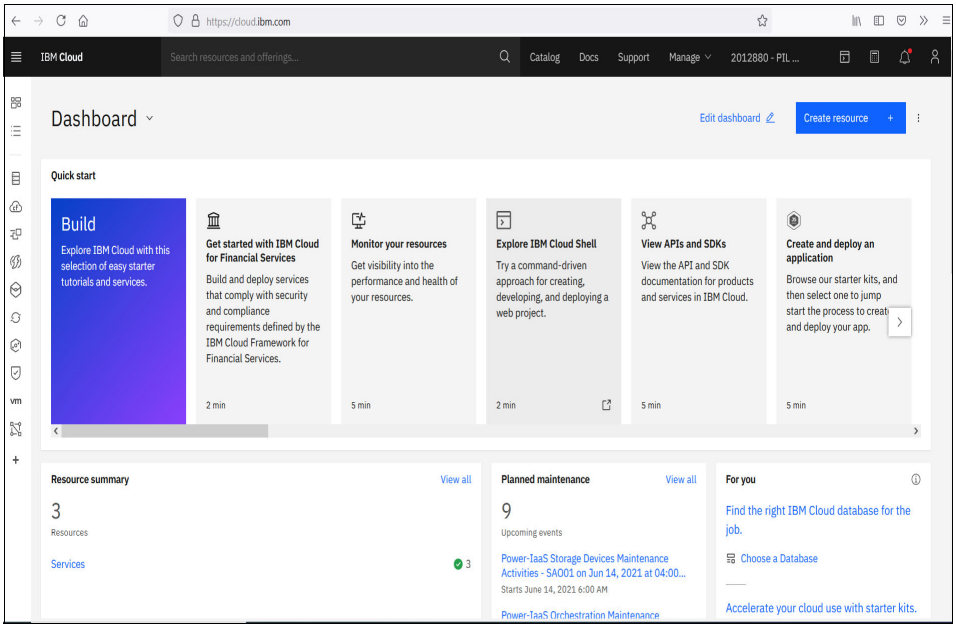

Figure 6-10 Selecting the IBM Cloud option

2. Choose **Services** from the list, as shown in Figure 6-11.

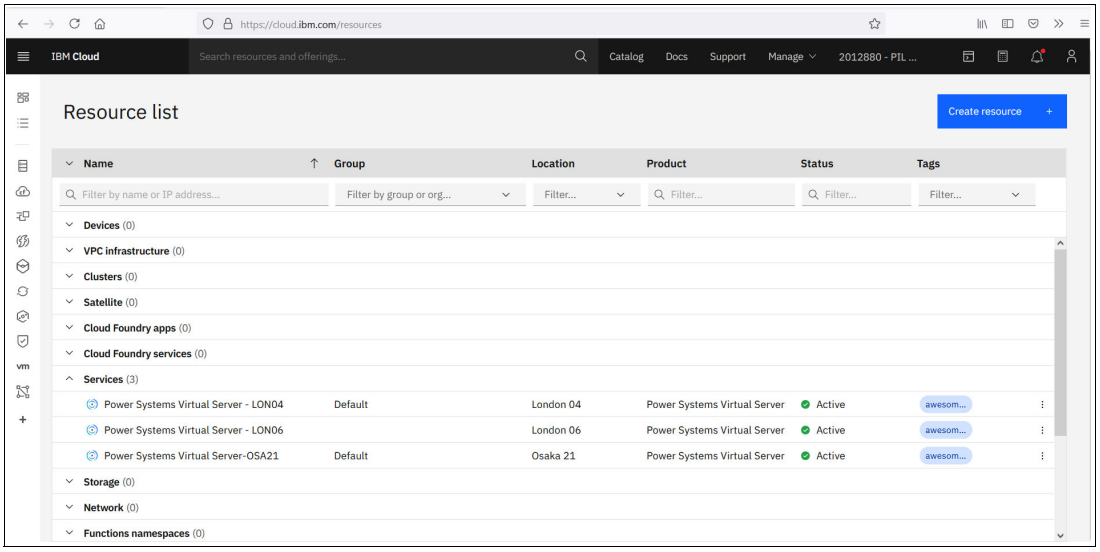

Figure 6-11 Choosing Services option

3. Click the service for data center in which you created a Power Systems Virtual Server location service. In this example, **London06** Power Systems Virtual Server location is selected, as shown in Figure 6-11 on page 222.
4. Because several VSIs are provisioned, they are listed as shown in Figure 6-12. If you are creating VSIs for the first time, the list is empty. Click **Create Instance**.

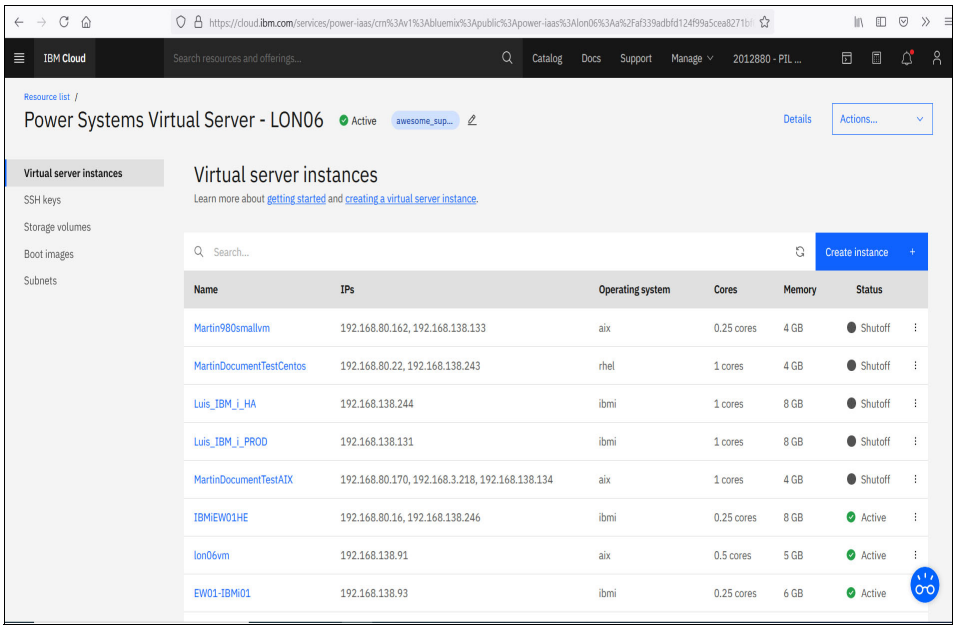

Figure 6-12 London06 Virtual Server instances

It is here that you provision the IBM i VSIs. Consider the following points:

- Enter a name for your VSI (for example, `IBMi-74-Lon06`) and select how many VSIs you must configure. The names of the VSIs are appended with a `-1`, *-2*, and so on, if you select more than one VSI. IBM i gives the LPAR a system name that consists of the first eight characters, and it is best to use only alphanumeric characters for IBM i naming.
- Specify whether this VSI is pinned to the host where it is running by selecting **Soft** or **Hard** pin.
- A public SSH key must be added to securely connect to this VSI, as shown in Figure 6-13 on page 224.

5. Select the operating system IBM i and the wanted image from the image list or any other image that you imported by way of the Boot Image menu.

   Regarding the IBM i Software Licenses, if you implement an operating system Level Replication solution with geographic mirroring, be sure to include IBM i PowerHA. Also, include other IBM i licenses as necessary.
6. For the Tier type, you can choose between the available disk types available: Type 1 or 3. Type 3 is a less expensive option, which is selected in the example.
7. For the Machine type, choose between S922 or E980. S922 is the cheaper of the two options, which is selected in the example.

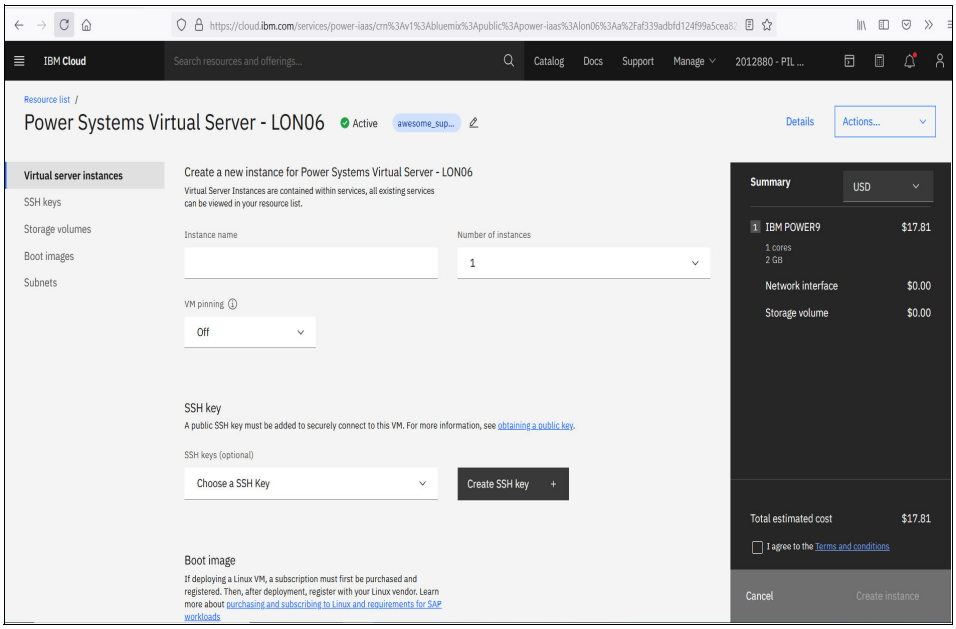

Figure 6-13 Adding SSH key

8. The processor resources can be assigned as Dedicated, Shared, or Shared Capped, as shown in Figure 6-14. Shared is selected because it is less expensive.

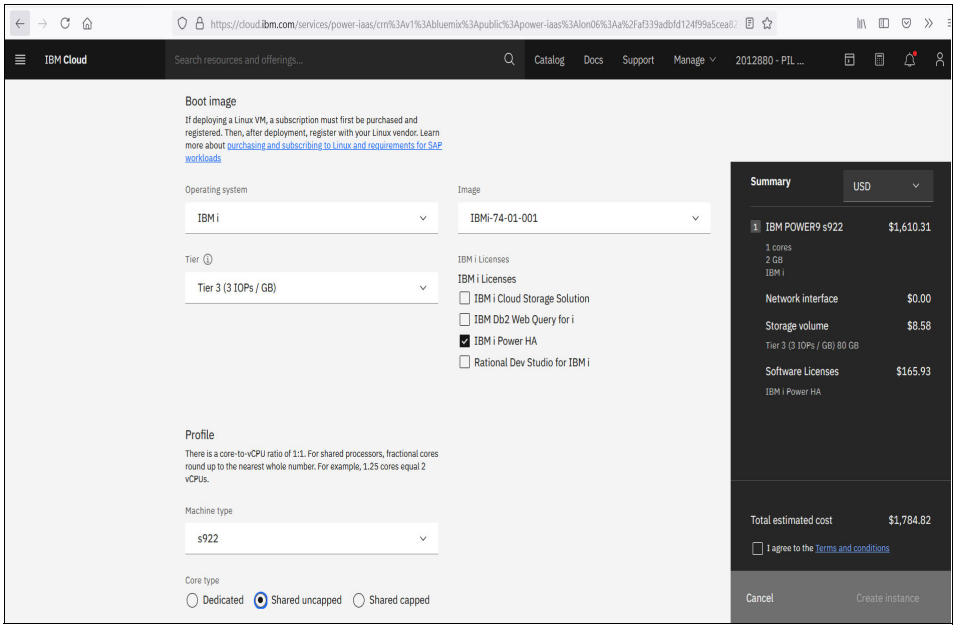

Figure 6-14 Choosing Shared process resource

9. Choose the number of cores and RAM that you need (the minimum core is 0.25). Because IBM i partitions often require at least 4 GB of memory to start, use this minimum value (see Figure 6-15 on page 225).

10. You can also attach more volumes to the VSI. In the example, only the root volume is included. It is highly recommended to add the disks individually after the initial Load Source (LS) volume is created. Doing so can make it easier to track which disk unit ID matches the volume name in the IBM Cloud UI.

11. No matter what operating system version or release that is used, it is always recommended to use the latest PTFs, which can be found at this IBM Support web page.
12. Because PowerHA for IBM i is owned by HelpSystems, see this web page to ensure that the latest HA group PTFs are installed.
13. Scroll down to choose your subnet on which these VSIs are to be provisioned. It is assumed that one or more subnets were created *before* this step.

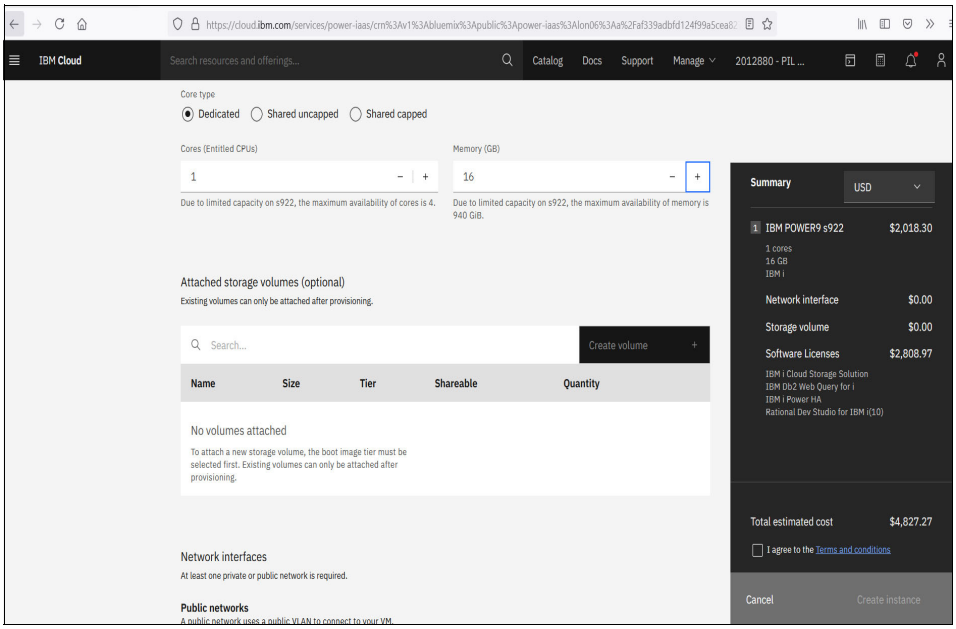

Figure 6-15 Using the minimum memory value

14. Click **Attached Existing** (see Figure 6-16).

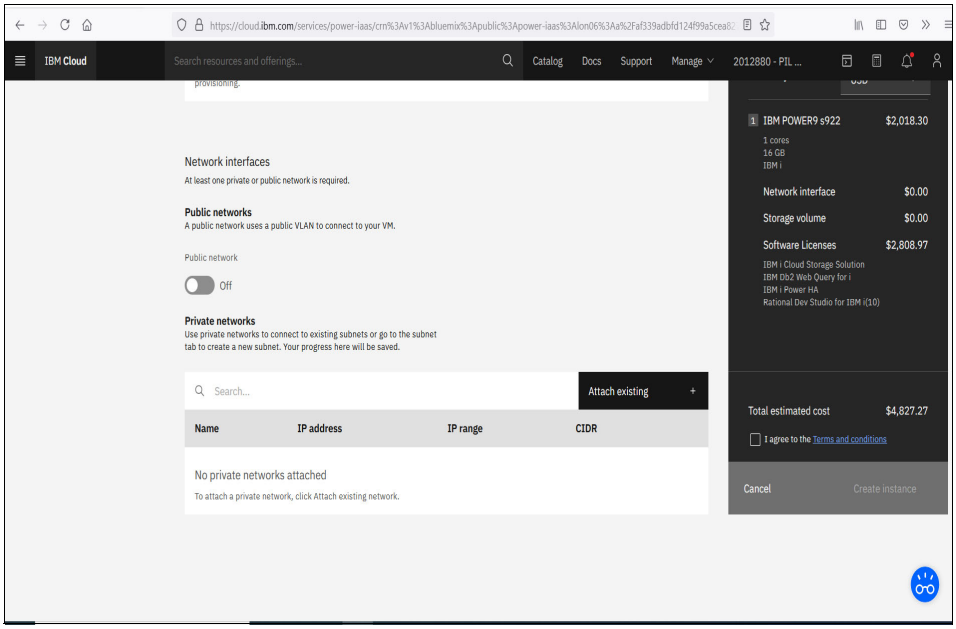

Figure 6-16 Selecting Attached Existing option

15. Choose the subnet that you want to attach and click **Attach**, as shown in Figure 6-17.

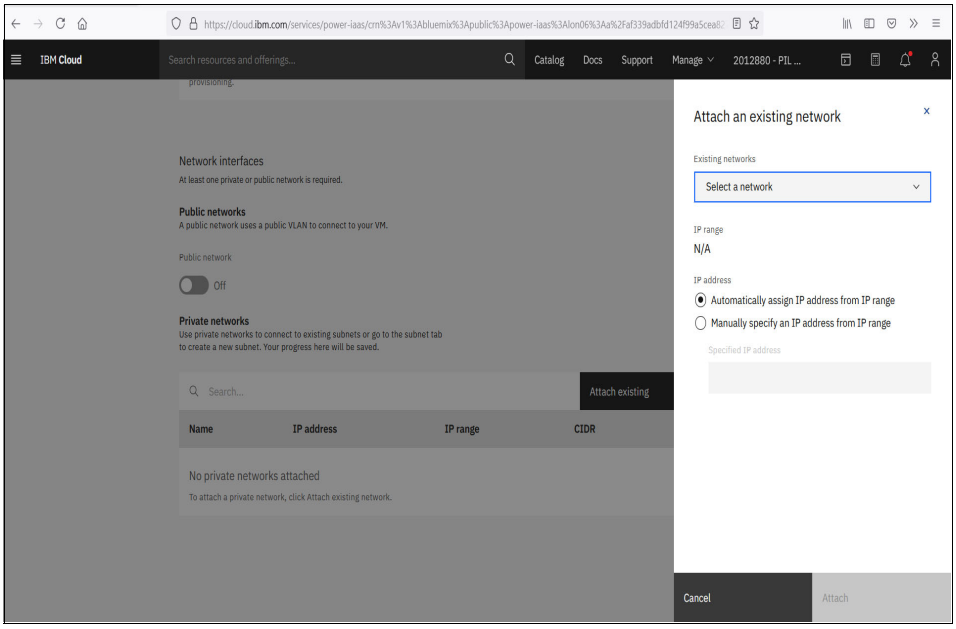

Figure 6-17 Attaching an existing network

16. Select the **I agree to the....** option and click **Create Instance**, as shown in Figure 6-18.

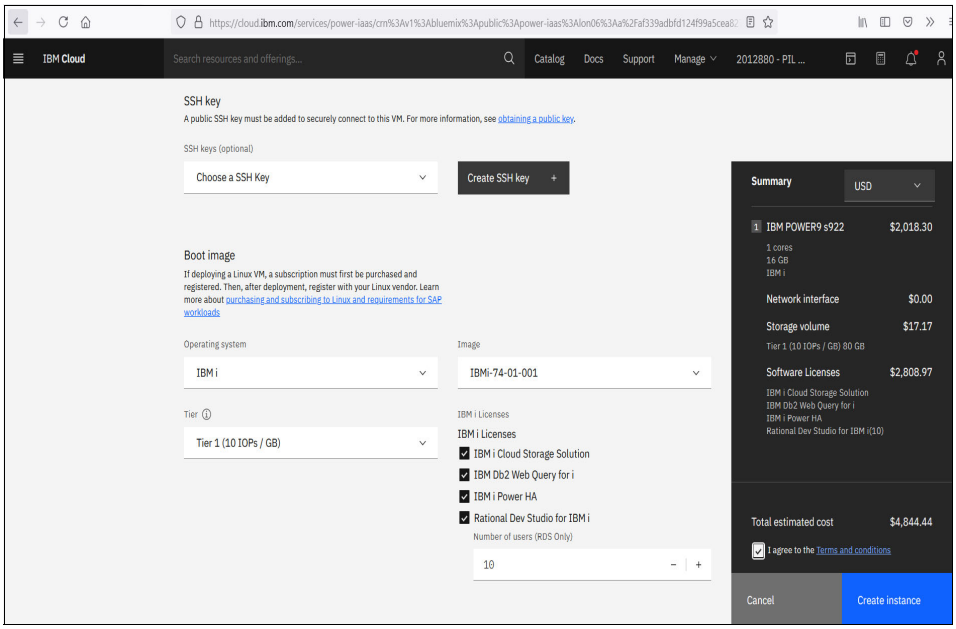

Figure 6-18 Creating an instance

Your VSI is now being provisioned.

Ordering Direct Link Connect Classic to connect Power Systems Virtual

Direct Link (DL) Connect Classic must be ordered so that your Power Systems Virtual Server Instances can communicate with Linux or Window VSIs in IBM Cloud and with all other IBM Cloud services, such as VMware VMs, and Cloud Object Storage.

An order for a DL can take1–2 weeks to complete.

Complete the following steps:

1. Order DL Connect Classic service on IBM Cloud UI.
2. A Support ticket is created. IBM Support sends you a Word document with questionnaires to be completed regarding various DL settings.
3. Complete the questionnaires and upload them to the Support ticket.
4. IBM Support then requests that you create another Support ticket with the Power Systems so that they can complete their tasks for the DL provisioning. Attach any information about the DL in the original ticket to this new ticket.
5. The DL is provisioned, and you are notified when the process is complete.
6. Test connections to any Linux or Windows VSI that is in IBM Cloud and other IBM Cloud services.

Complete the following steps to start the DL order process:

1. Log in to the IBM Cloud UI.
2. Choose **Catalog** and search for `direct`, as shown in Figure 6-19.

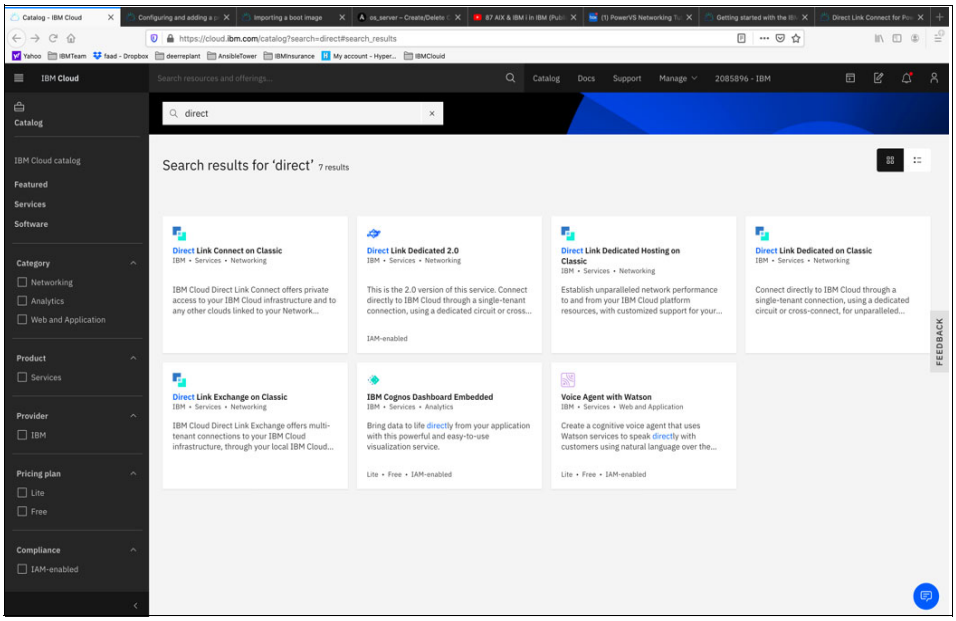

Figure 6-19 Searching for direct

3. Select **Direct Link Connect on Classic** and click **Create**. No options are available to select, as shown in Figure 6-20 on page 228.

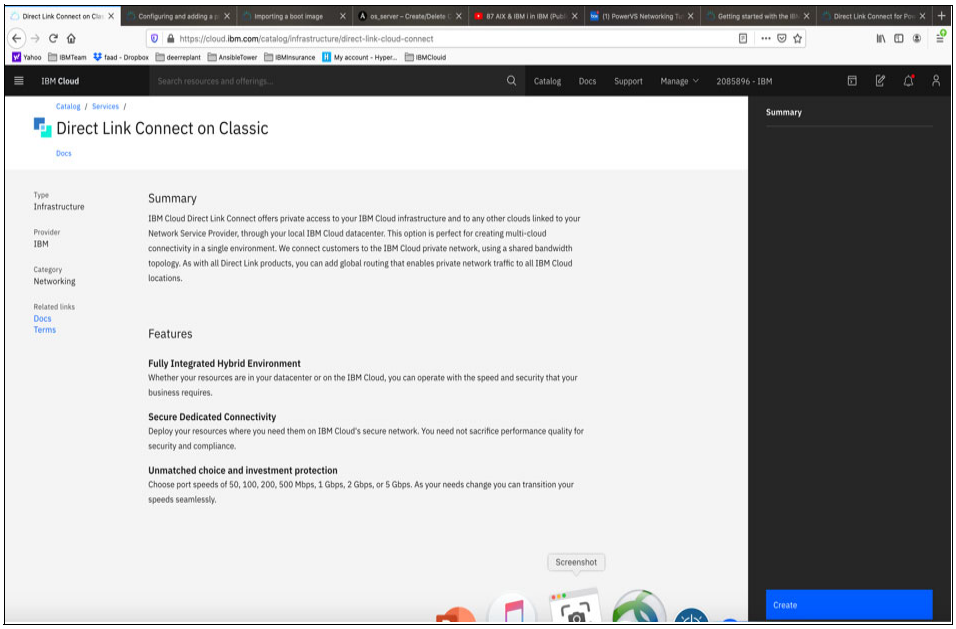

Figure 6-20 Direct Link Connect on Classic

4. Choose **Order Direct Link Connect**, as shown in Figure 6-21.

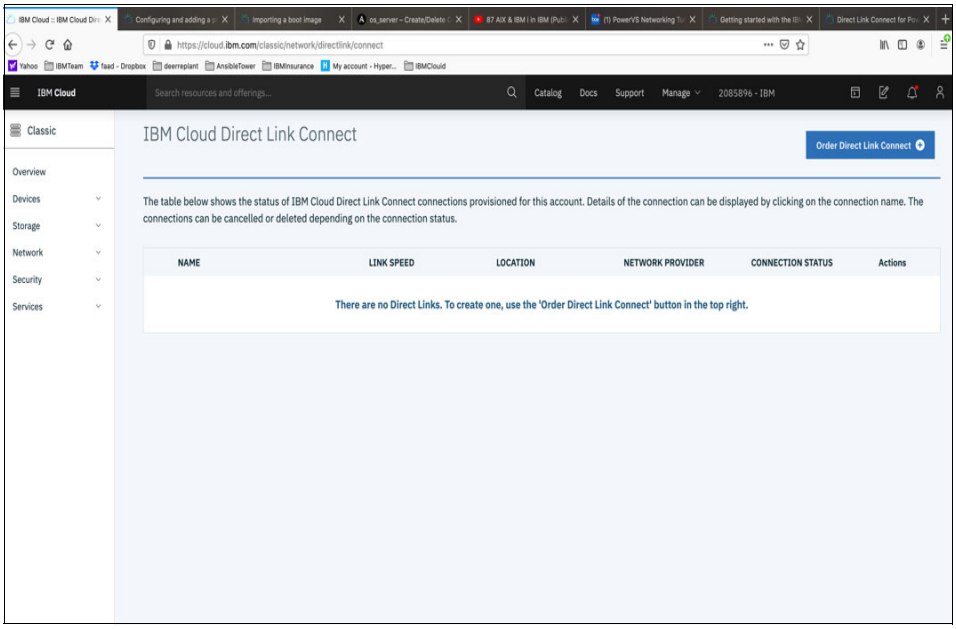

Figure 6-21 Choosing Order Direct Link connection

5. Make the following choices (Figure 6-22 on page 229):
   - Name for the DL.
   - Location for the DL. This location is the same location where you created your Power Systems Virtual Server location service.
   - Link speed (under the Network Provider menu).
   - **Local Routing (no cost)**.

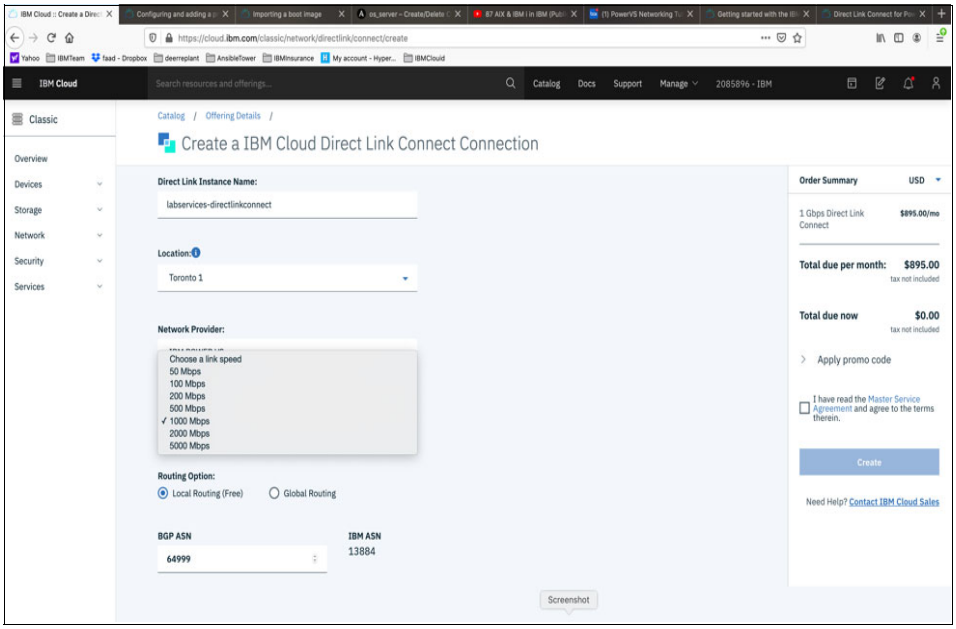

Figure 6-22 Configuring the direct link

Global routing requires extra charges and allows for easier Power Systems Virtual Server location-to-Power Systems Virtual Server location communication. You must also order a Vyatta Gateway Router to complete your Global routing option by way of a GRE tunnel. IBM Support can help you with this part of the order.

In this example, we use Local Routing, ordered a Vyatta Gateway in each Power Systems Virtual Server location, and provisioned a GRE tunnel end-to-end.

6. Select the box to accept the offer and click **Create**.

   A support case is opened that includes the required information, as shown in Figure 6-23.

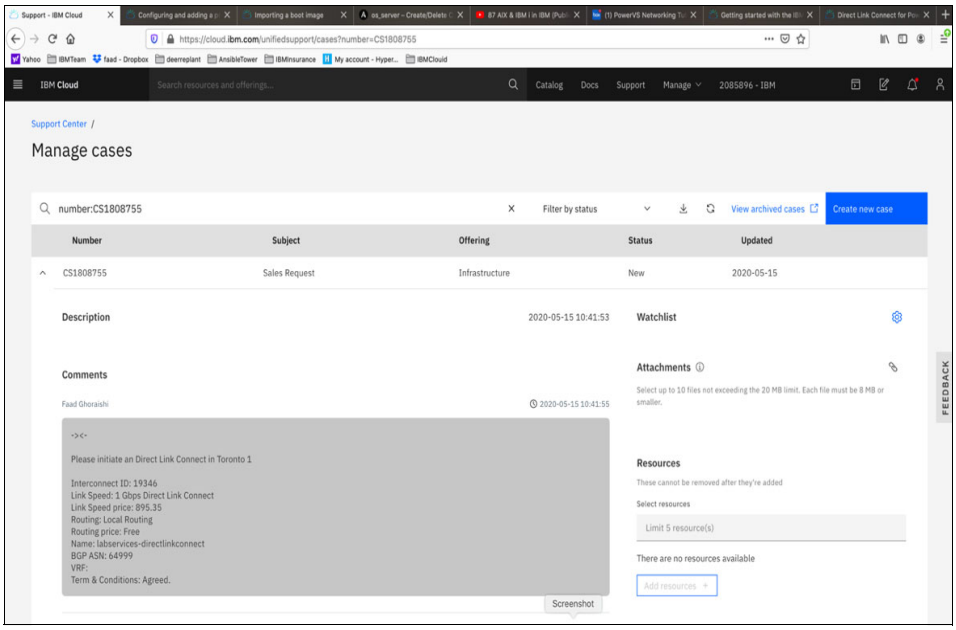

Figure 6-23 Support case window

After this process is complete, you are contacted by IBM Support and requested to complete and answer some questions in an attached document. Return the completed document as an attachment to the same Support ticket.

After this step is complete, IBM Support requests that you open a new Support ticket and address it to the Power Systems. Include the information in the original DL ticket. This new ticket is then sent to the Power Systems Virtual Server location support to configure their side of the DL connection.

This step is the last step before DL communication is enabled. You can test your connection by pinging IBM Cloud Linux or Windows VSI from your Power Systems Virtual Server Instances and vice versa.

Ordering two Vyatta Gateways, one in each data center

To set up communication between the two Power Systems Virtual Server location data centers, you must use Generic Routing Encapsulation (GRE) tunnels. Because the GRE tunnels are provisioned on Vyatta Gateways, you must order one Vyatta Gateway in each Power Systems Virtual Server location.

The example involves ordering one Vyatta in `LON06` and the other in `TOR01` data centers where the Power Systems Virtual Server locations exist.

This example uses two Vyatta Gateways. One in each Power Systems Virtual Server location to provide end-to-end Power Systems Virtual Server location-to-Power Systems Virtual Server location communication by using GRE tunnels.

Complete the following steps:

1. Log in to IBM Cloud and click **Catalog** and search for `Vyatta`, as shown in Figure 6-24.

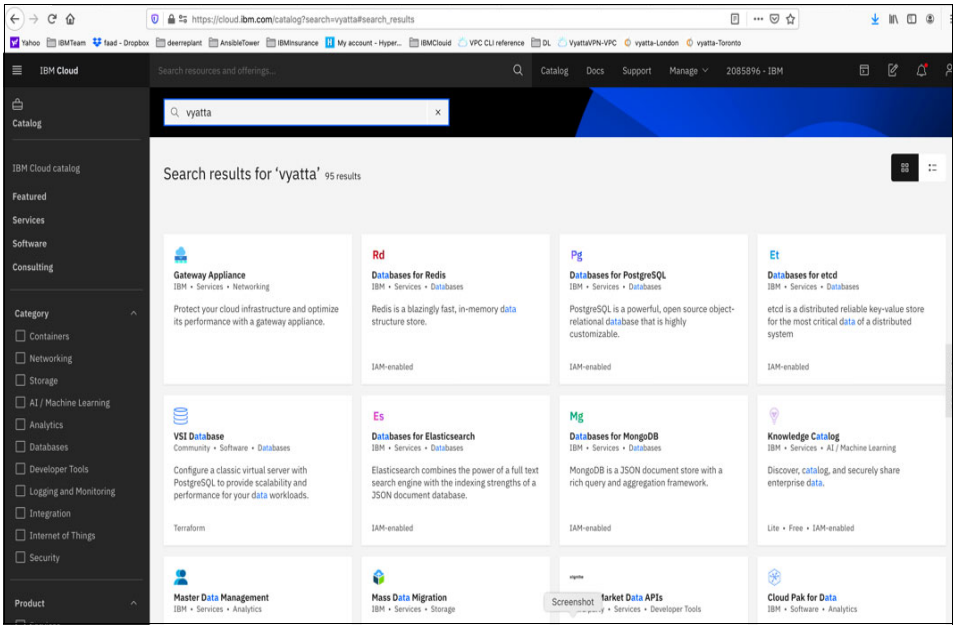

Figure 6-24 Searching for Vyatta

2. Select **Gateway Appliance** and then create it, as shown in Figure 6-25.

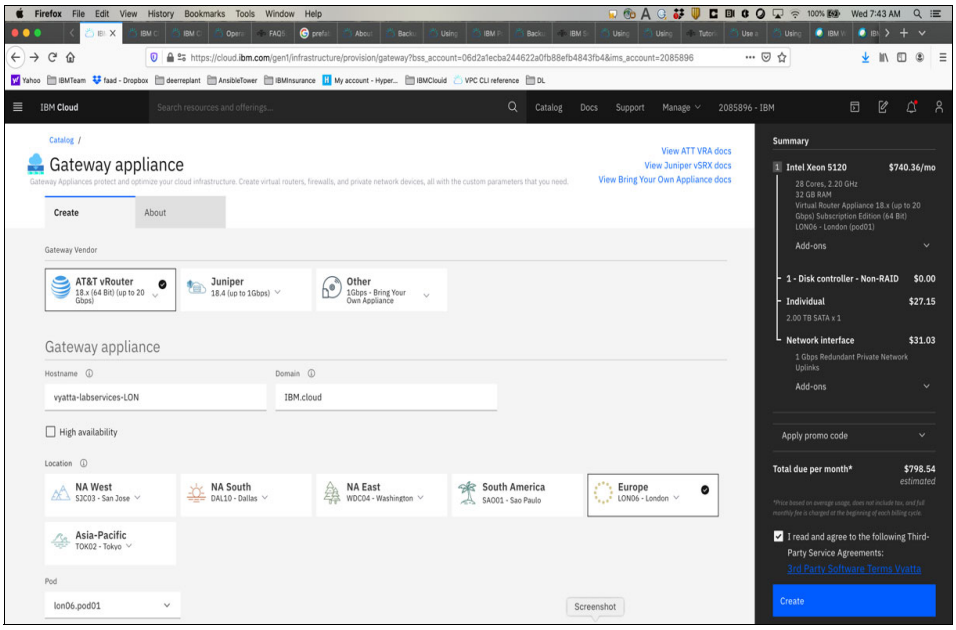

Figure 6-25 Selecting the location

3. Select **AT&T vRouter**, which is the Vyatta Gateway. Other choices of gateways are available, but Vyatta is used in this example.
4. Enter a name for the gateway and include the Power Systems Virtual Server location name in it so that you can distinguish them later.
5. Select a location to match your Power Systems Virtual Server location.

   Choose the following options:

   – Clear the High Availability option unless you want to order one, which means that you order two Vyatta Gateways in each Power Systems Virtual Server location. The option is cleared in this example.

   – Select the location by pressing the arrow key in each location to find the specific data center where your Power Systems Virtual Server locations are available.

   – Choose a POD if several PODs are in the selected data center location.

   – Select the CPU single or dual processor. The example uses a **Single Processor**.

   – Select the amount of RAM that you want and add SSH keys if you want to log in without the use of a password. (This step can be completed later.)

6. Choose **Private network interface** unless you want to use the default setting, which is public/private interface. The example uses a private network interface only (see Figure 6-26).

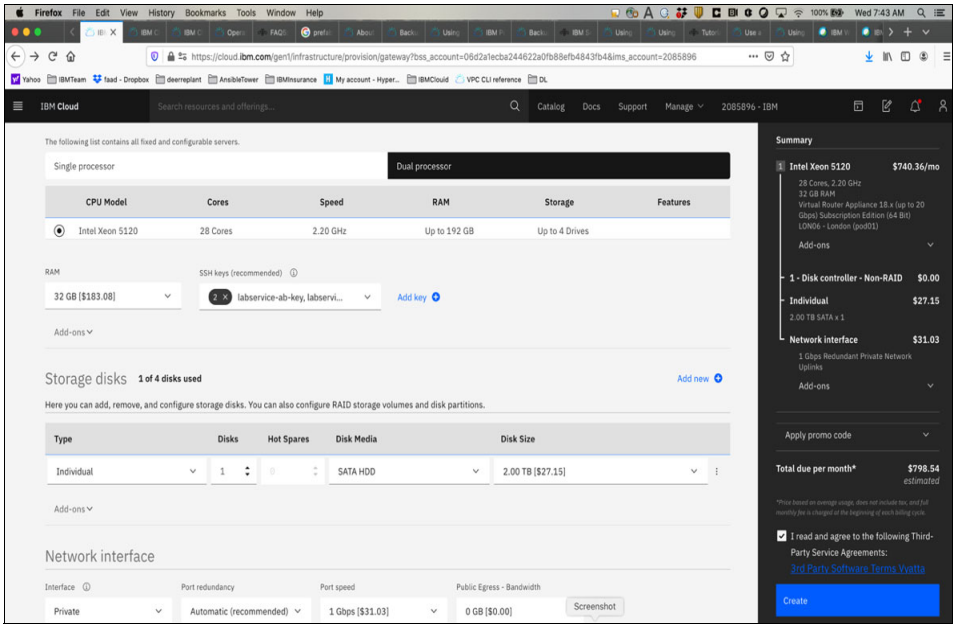

Figure 6-26 Choosing the private network interface

7. Accept the service agreement and click **Create**. The Vyatta gateway is now being provisioned, which can take several hours.

   This process must be done in each of the Power Systems Virtual Server locations.

   After the Vyatta Gateway is provisioned, it is listed under Devices, where you can find your Vyatta and root user passwords (see Figure 6-27).

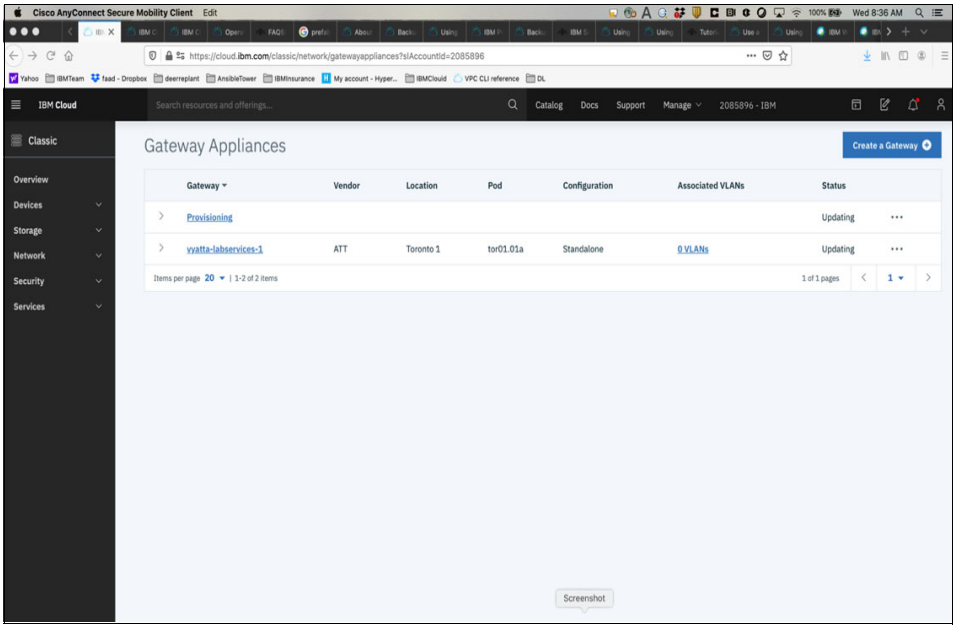

Figure 6-27 Listed gateways

8. Log in to the Vyatta gateway by using a browser and accessing it by using the following link:

   `https://<ip_address_of_the_vyatta_gateway>`

   Use the following credentials, as shown under devices in IBM Cloud UI and password tab on the left (see Figure 6-28):

   – User: `Vyatta`
   – Password: As shown in the GUI

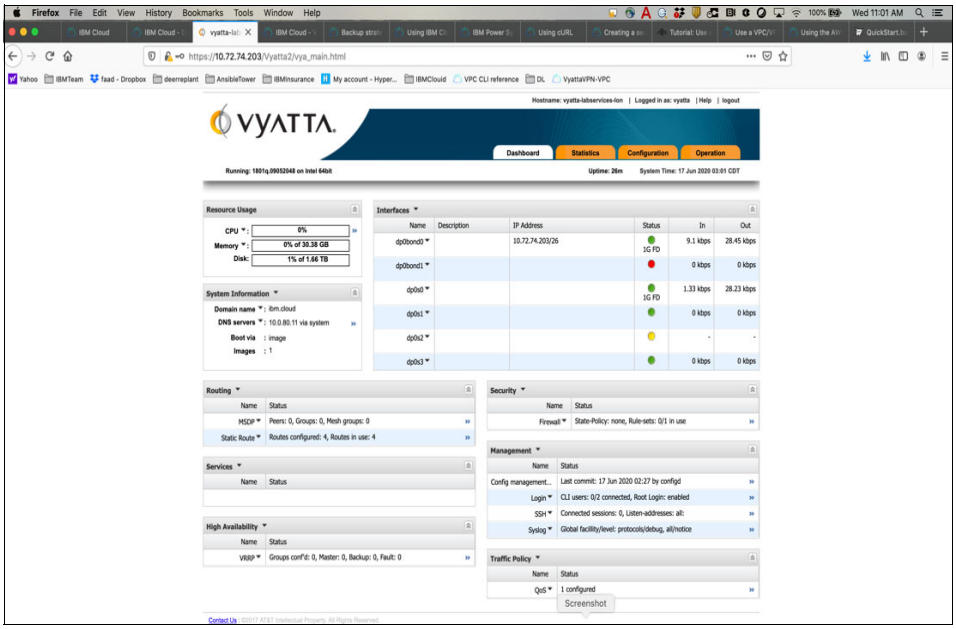

Figure 6-28 Vyatta gateway

Typically, you use a command line to SSH to the Vyatta for more configuration tasks. You use the Vyatta user identification to complete the configurations.

Requesting a GRE tunnel at each Power Systems Virtual Server location

Open a Support ticket with Power Systems and request that a GRE tunnel is provisioned in each Power Systems Virtual Server location. They provision their end of the GRE tunnel and send you the information so that you can continue and provision your end on the Vyatta Gateways.

Also provide the subnets information in each Power Systems Virtual Server location in the ticket.

The Power Systems support team sends you the information for your GRE tunnels (see Figure 6-29) after they finish the request.

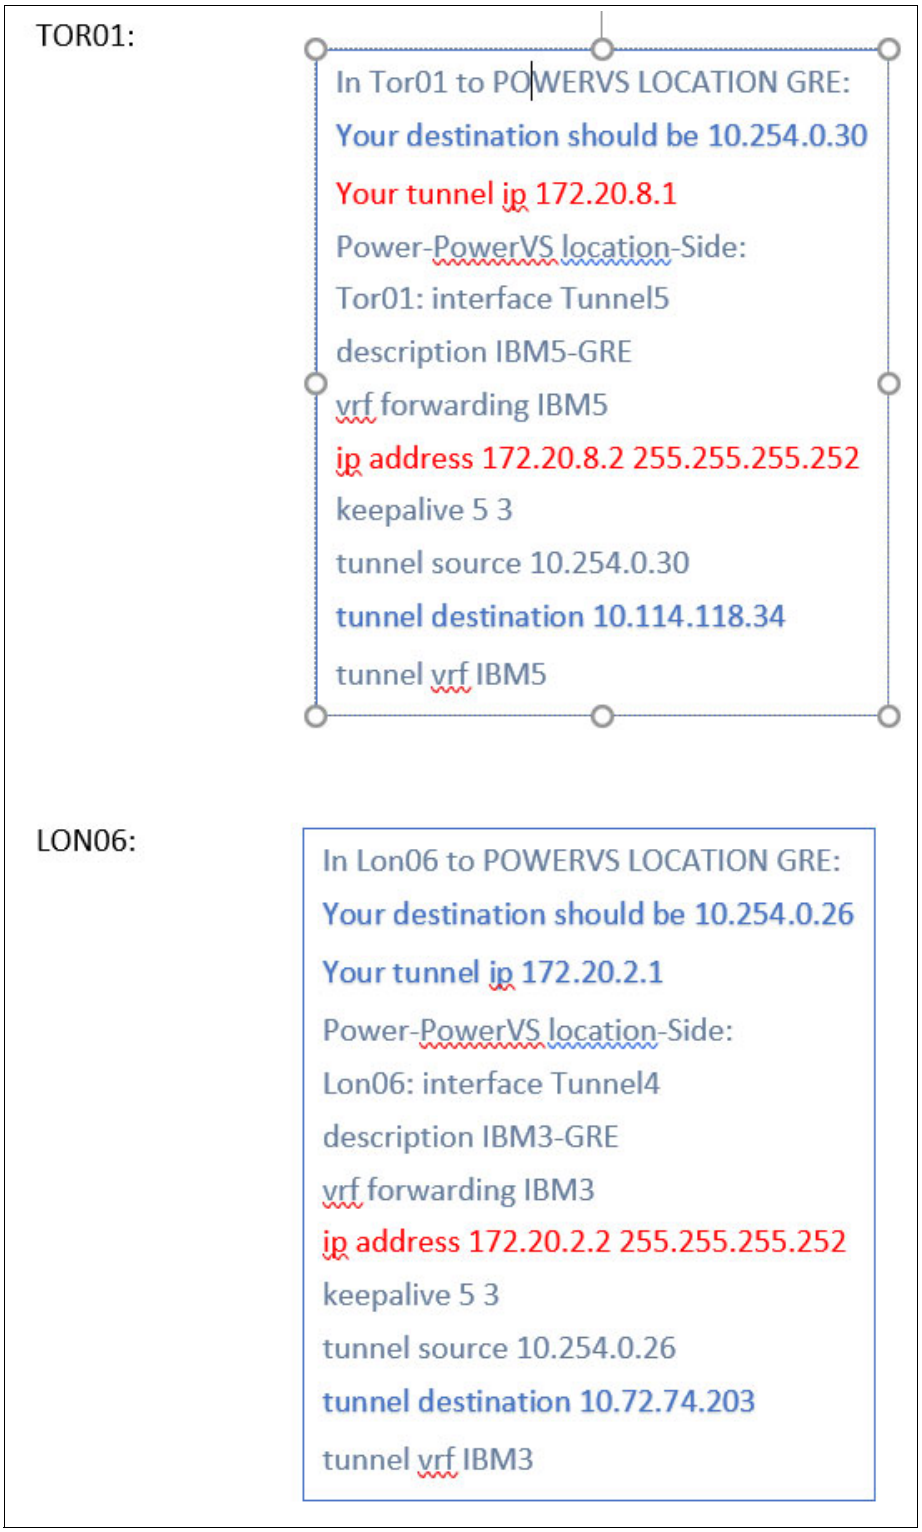

Figure 6-29 GRE information

The items that are shown in red in Figure 6-29 are what must be configured at your end of the GRE tunnel in each Vyatta Gateway.

Consider the following points about your tunnel:

- IP address is `172.20.2.1/30`, where `255.255.255.252` translates to `/30`
- Destination IP is their tunnel source IP
- Source IP is the IP address of the Vyatta gateway

Use the following steps to verify your Vyatta Gateway access:

1. The Vyatta Gateway address can be found in the IBM Cloud UI under Devices.

   Log in to IBM Cloud UI and click **IBM Cloud** on the upper left side. Click **Devices** and choose the Vyatta system that you want to configure:

   - `vyatta-labservices-lon.ibm.cloud`
   - `vyatta-labservices-tor.ibm.cloud`

2. For `LON06`, click **London Vyatta: vyatta-labservices-lon.ibm.cloud** (see Figure 6-30).

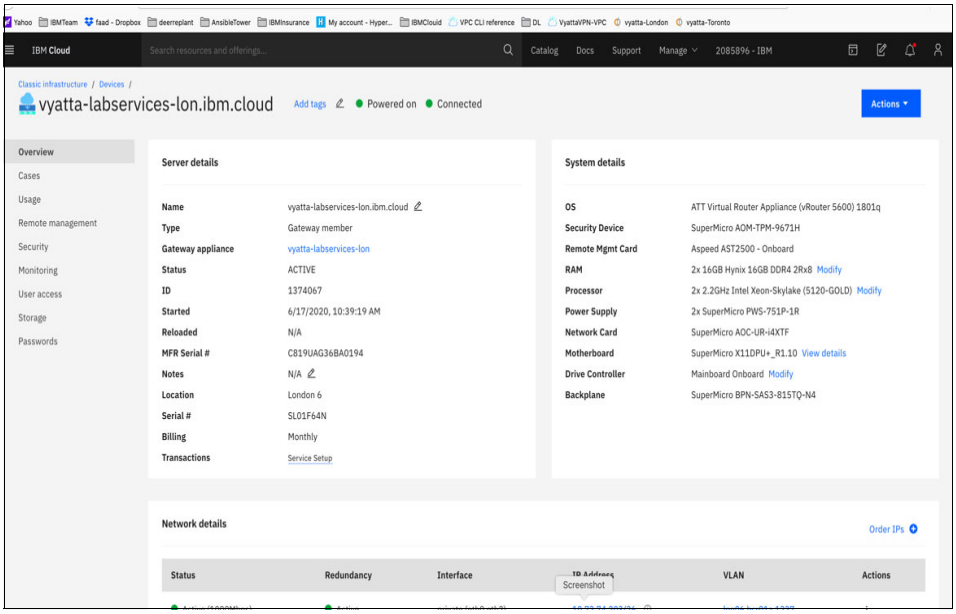

Figure 6-30 London Vyatta

Under Network Details, you see your Vyatta Gateway IP address.

3. Your credentials are under the password menu that is on the left side. Click the icon that is next to the password to see it unencrypted, as shown in Figure 6-31 on page 236.

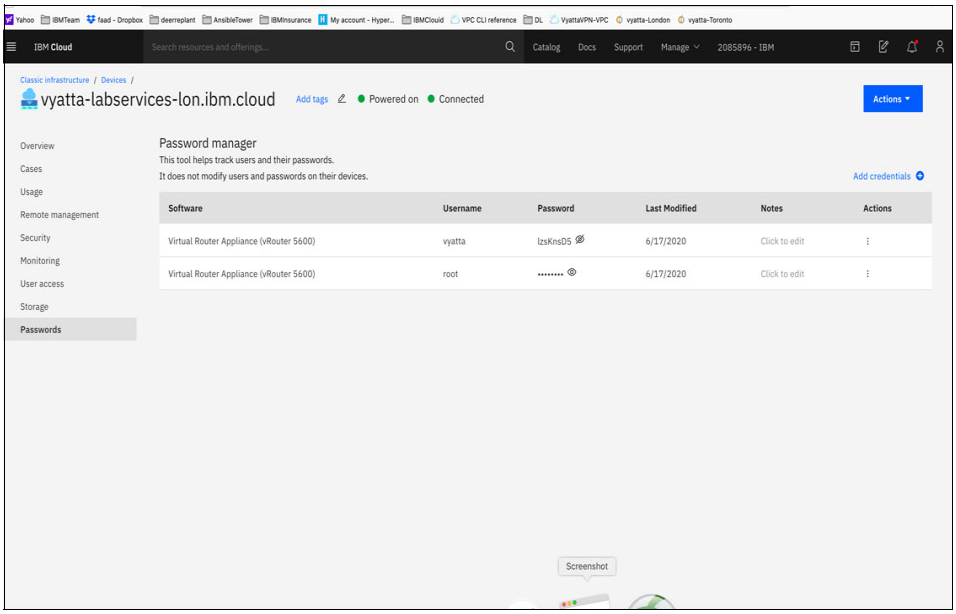

Figure 6-31 Vyatta password

4. Open a browser and log in to the Vyatta Gateway by using the following credentials (see Figure 6-32):
   - User ID: `Vyatta`
   - Password: As shown in the GUI
   - https://10.72.74.203
   - ssh vyatta@10.72.74.203

Note: Before logging in to a 10.x.x.x private IP in IBM Cloud, you must start your MotionPro Plus VPN access, which gives you access to IBM Cloud private IPs.



Figure 6-32 Vyatta login window

5. Log in by using the user ID and password as shown in Figure 6-33.



Figure 6-33 Vyatta interface

Now that you verified you access to the Vyatta Gateways, you can access it by way of SSH to continue your GRE tunnel provisioning.

Configuring GRE tunnels in the Vyatta Gateways

For more information about configuring the GRE, see this IBM Cloud Docs web page.

After the support team finishes configuring the GRE tunnel, you must configure your end of the GRE tunnel on the two Vyatta Gateways.

The following GRE tunnels are required:

- GRE tunnel on Vyatta to end the Power Systems Virtual Server location GRE in `LON06` (or “data center 1”).
- GRE tunnel on Vyatta to end the Power Systems Virtual Server location GRE in `TOR01` (or “data center 2”).
- GRE tunnel across the two Vyatta gateways: one on each side.

For more information about configuring GRE tunnels, see the following resources:

- IBM Cloud Docs: VPN into a secure private network
- Vyatta System’s *Tunnels Refeence Guide*
- IBM Cloud Docs: Configuring connectivity to Power Systems Virtual Server

Open a command window on your notebook.

Note: Before logging in to a 10.x.x.x private IPs in IBM Cloud, you must start your MotionPro Plus VPN access.

Setting up GRE Power Systems Virtual Server location tunnel in LON06

Log in to the Gateway by using the following information (see Figure 6-34):

- User ID: Vyatta
- Password: As shown in the GUI
- ssh vyatta@10.72.74.203
- ssh to LON06 Vyatta Gateway

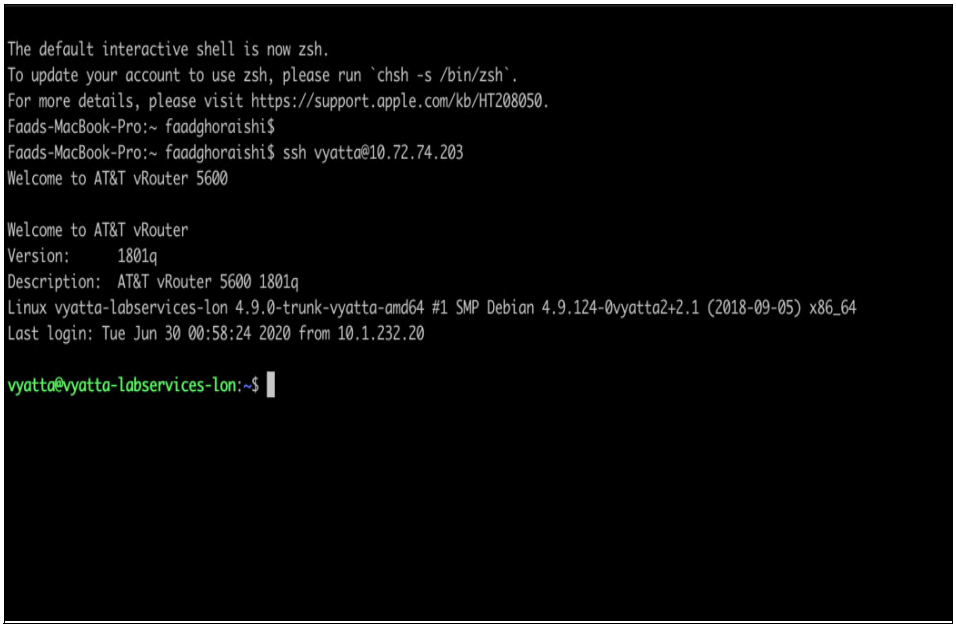

Figure 6-34 Logging in to GRE Tunnel in London06

For this example, use the information that is provided by IBM Support for LON06 GRE, as shown in Figure 6-35.

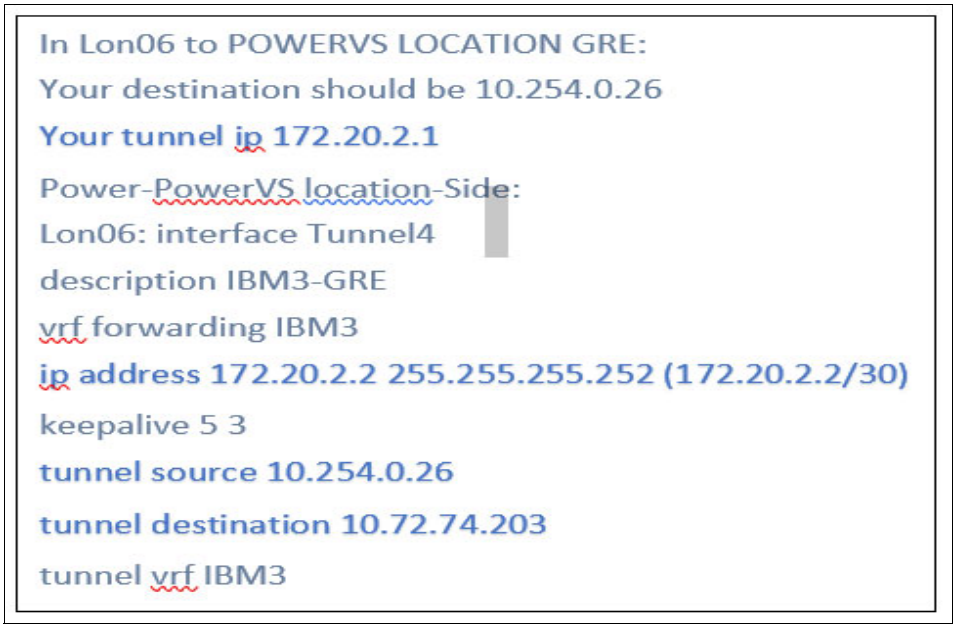

Figure 6-35 London06 GRE information

Run the commands that are listed in Example 6-1. Note that the tunnel is named `tun0` on the Vyatta Gateway.

Example 6-1 Commands to configure the London GRE tunnel

```
configure
set interfaces tunnel tun0 address 172.20.2.1/30
set interfaces tunnel tun0 local-ip 10.72.74.203
set interfaces tunnel tun0 remote-ip 10.254.0.26
set interfaces tunnel tun0 encapsulation gre
set interfaces tunnel tun0 mtu 1300
commit
exit
```

You can verify that your GRE tunnel is set up by running the following commands:

```
configure
show interfaces tunnel
```

For more information, run the following commands:

```
show interfaces tunnel tun0
exit
```

Setting up GRE Power Systems Virtual Server location tunnel in TOR01

Log in to the Vyatta Gateway by using the following information:

- User ID: Vyatta
- Password: as shown in the GUI
- ssh vyatta@10.114.118.34
- ssh to Tor01 Vyatta Gateway

For this example, use the information that was provided by IBM Support for TOR01 GRE, as shown in Figure 6-36.

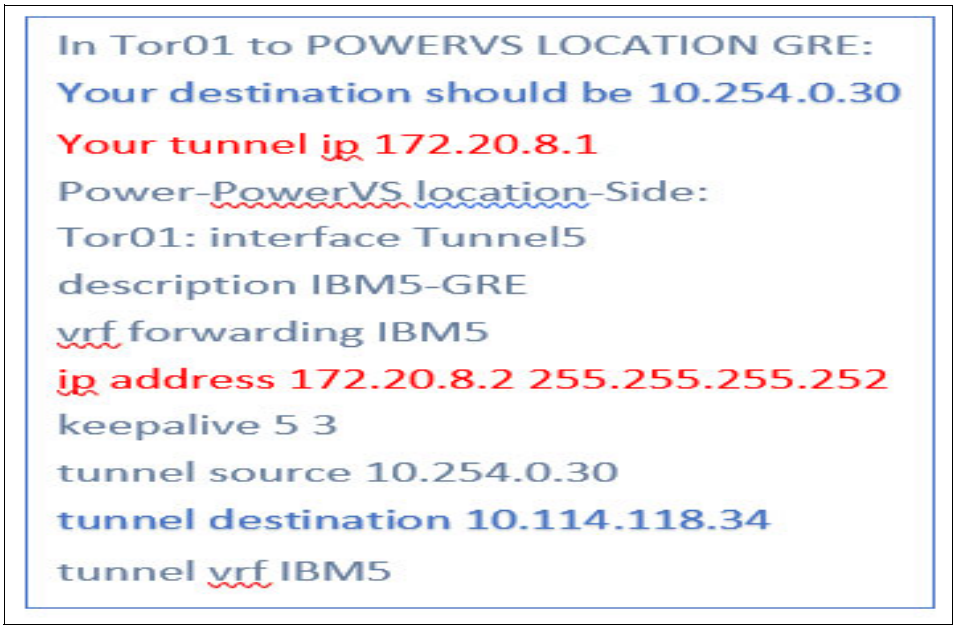

Figure 6-36 GRE Tunnel in Toronto01

Run the commands that are listed in Example 6-2. Note that the tunnel is named `tun0` in the Vyatta Gateway, which is the same name as the other Vyatta Gateway.

Example 6-2 Commands to configure the Toronto GRE tunnel

```
configure
set interfaces tunnel tun0 address 172.20.8.1/30
set interfaces tunnel tun0 local-ip 10.114.118.34
set interfaces tunnel tun0 remote-ip 10.254.0.30
set interfaces tunnel tun0 encapsulation gre
set interfaces tunnel tun0 mtu 1300
commit
exit
```

You can verify that your GRE tunnel is set up by running the following commands:

```
configure
show interfaces tunnel
```

For more information, run the following commands:

```
show interfaces tunnel tun0
exit
```

Setting up GRE tunnel between two Vyatta Gateways

This section describes how to configure a new tunnel in each of the two Vyatta gateways to allow for cross-Vyatta connection through a GRE tunnel.

The example has a specific tunnel address and tunnel source and destination IPs. The tunnel address can be any IP subnet that you choose. In the example, the tunnel is named `tun1` in both Vyatta gateways.

A similar IP as the IPs that are used in the Power Systems Virtual Server location GRE tunnels is selected. A CIDR of /30 is chosed because only two IP address are needed: one in Tor01 and one in Lon06. Consider the following points:

- In `Lon06` Vyatta, the GRE Vyatta-to-Vyatta tunnel address is `172.20.4.1/30`.
- In `Tor01` Vyatta the GRE Vyatta-to-Vyatta tunnel address is `172.20.4.2/30`.
- Your tunnel destination IP is the IP address of the Vyatta gateway in each location.
- Your tunnel source IP is the IP address of the Vyatta gateway in each location.
- We call the tunnels tun1 in both locations.

Run the commands in Example 6-3 to configure TOR01 GRE.

Example 6-3 Commands to configure tun1 in TOR01 GRE

```
configure
set interfaces tunnel tun1 address 172.20.4.1/30
set interfaces tunnel tun1 local-ip 10.114.118.34
set interfaces tunnel tun1 remote-ip 10.72.74.203
set interfaces tunnel tun1 encapsulation gre
set interfaces tunnel tun1 mtu 1300
commit
exit
```

Run the commands in to configure LON06 GRE:

Example 6-4

```
configure
set interfaces tunnel tun1 address 172.20.4.2/30
set interfaces tunnel tun1 local-ip 10.114.118.34
set interfaces tunnel tun1 remote-ip 10.72.74.203
set interfaces tunnel tun1 encapsulation gre
set interfaces tunnel tun1 mtu 1300
commit
exit
```

The final step that is needed is to configure static routes in each Vyatta Gateway to point to the subnets for the Power Systems Virtual Server location to the correct tunnels.

Find the subnets that you created in each Power Systems Virtual Server location in `TOR01` and `LON06` by accessing the services in the IBM Cloud UI for each Power Systems Virtual Server location.

The static routes in the `LON06` and `TOR01` subnets must point to the other subnet.

Configuring GREs to the Power Systems Virtual Server locations

Run the following commands in each Vyatta Gateway after login in by way of ssh using the Vyatta userID:

- In TOR01 Vyatta:

```
configure
set protocols static route 192.168.6.0/24 next-hop 172.20.8.2
set protocols static route 192.168.50.0/24 next-hop 172.20.4.2
commit
exit
```

- In LON06 Vyatta:

```
configure
set protocols static route 192.168.50.0/24 next-hop 172.20.2.2
set protocols static route 192.168.6.0/24 next-hop 172.20.4.1
commit
exit
```

End-to-end connectivity exists, and you can ping between your Power Systems Virtual Server Instances in each Power Systems Virtual Server location. You can also ping from the Power Systems Virtual Server Instance to IBM Cloud services, such as Linux/Windows VSI.

If you cannot ping the IBM Cloud VSIs from the Power Systems Virtual Server location VSIs, open a ticket so that IBM support can address the issue.

Configuring a reverse-proxy Centos VSI

This section describes the procedure to configure a reverse proxy to allow access to a private Cloud Object Storage endpoint.

All access to Cloud Object Storage from Power Systems Virtual Server Instance is by way of this reverse proxy. It is accessed by way of `https://<reverse_proxy_ip>`.

A Centos or Red Hat VSI must be provisioned in the IBM cloud to configure a reverse proxy. This VSI must have public access. After the configuration process is complete, the public access can be made unavailable.

Architecture

The overall architecture of the example deployment is shown in Figure 6-37.

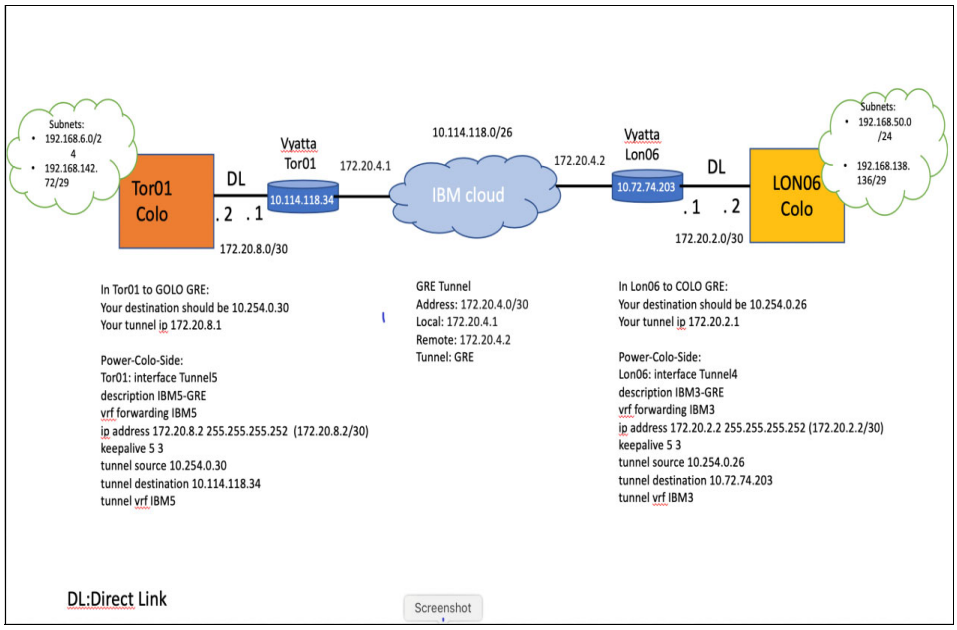

Figure 6-37 Network architecture

6.2 OS level replication with PowerHA SystemMirror for i geographic mirroring

Today, businesses experience increasing demands for application availability, which requires that small and large clients look for a solution that can help eliminate planned and unplanned downtime for their IT services.

Any unplanned outage can include severe implications, especially if the duration of the outage or recovery time exceeds business expectations. Some of these implications can include loss of data, revenue, worker productivity, company reputation, and client loyalty.

PowerHA geographic mirroring offers a straightforward, cost-effective, HA solution for the small to mid-sized client. PowerHA geographic mirroring is a disk solution that runs on all platforms and provides an alternative to solutions that require the extra configuration and management that are associated with an external storage device (see Figure 6-38 on page 243).

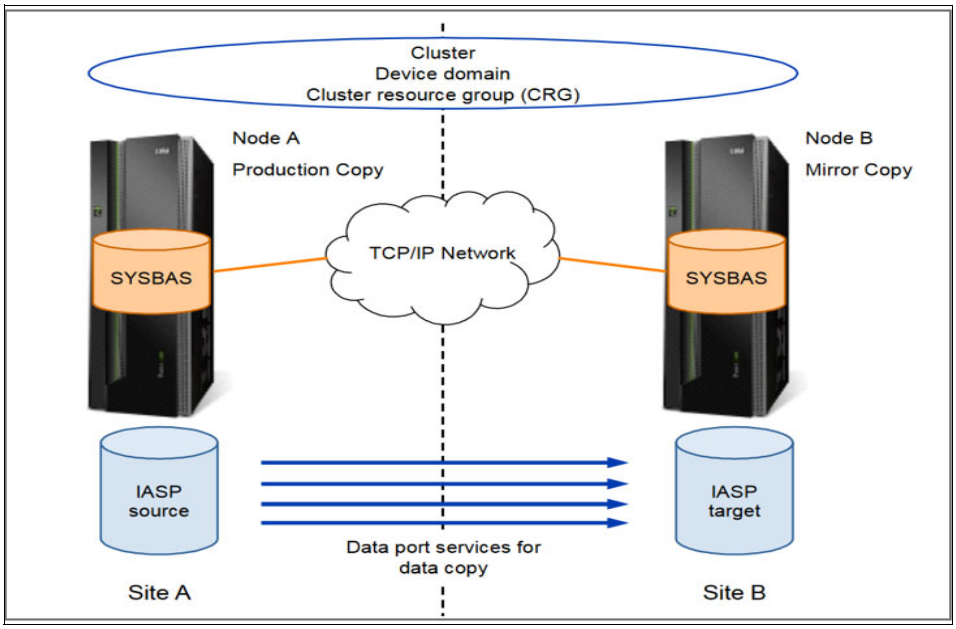

Figure 6-38 PowerHA SystemMirror for i geographic mirroring

Geographic mirroring refers to the IBM i host-based replication solution that is provided as a function of IBM PowerHA SystemMirror for i.

Geographic mirroring is managed by IBM i storage management so that replication is performed on a disk page segment basis. When a page of data is written, storage management automatically manages the replication of this page to the remote system.

Geographic mirroring requires a two-node clustered environment and uses data port services. These data port services are provided by the System Licensed Internal Code (SLIC) to support the transfer of large volumes of data between a source node and a target node.

The processes that keep the Independent Auxiliary Storage Pool (IASP) synchronized run on the nodes that own the IASPs. This transport mechanism communicates over TCP/IP.

6.2.1 Prerequisites

To implement PowerHA SystemMirror for i geographic mirroring, the following products must be installed on both nodes in the cluster:

- HA Switchable Resources, 5770-SS1 option 41.
- IBM PowerHA SystemMirror for i, 5770-HAS *BASE.
- PowerHA for i Standard Edition, 5770-HAS option 2.

If asynchronous mode operation is required, then PowerHA for i Enterprise Edition, 5770-HAS option 1 is also required.

6.2.2 Planning

Geographic mirroring is a host-based replication solution. Because the IBM i manages data transmission between a current production and current backup system, the type of storage that is used can be internal or external.

Geographic mirroring can be configured to use synchronous or asynchronous transmission delivery mode. Depending on the transmission delivery mode, the mirroring mode can also be configured to be synchronous or asynchronous.

Consider the following points:

- *Transmission delivery mode* refers to how the writes are delivered from the production copy system to the mirror copy system from a communications perspective.
- *Mirroring mode* refers to the method of how the data is saved on the mirror copy system.

Various combinations of the synchronous and asynchronous transmission modes are available.

Synchronous transmission mode with synchronous mirroring mode

When geographic mirroring is active in synchronous mode, the disk write operation is sent in parallel to the production copy's disk subsystem (1) and the mirror copy system (2), as shown in Figure 6-39.

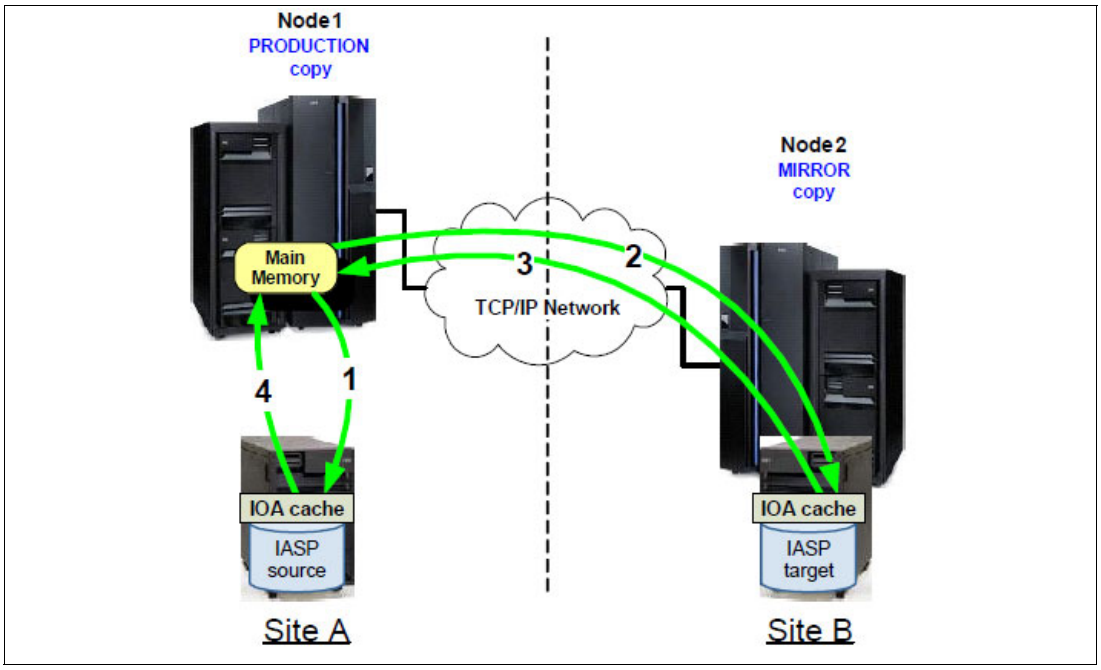

Figure 6-39 Synchronous transmission mode with synchronous mirroring mode

After the data is written to the preferred target system, an acknowledgment operation (step 3 in Figure 6-39) is sent. After the data is written to the production system, an acknowledgment operation (step 4 in Figure 6-39) is sent. After both steps 3 and 4 are completed, an acknowledgment is sent to the storage management function of the operating system of the production copy.

In this configuration, the mirror copy independent auxiliary storage pool (IASP) is always eligible to become the production copy IASP because the order of writes is preserved on the mirror copy system. For this reason, this configuration is preferred, if possible.

However, in many cases, the network infrastructure does not allow this configuration as a practical solution. For example, on lower bandwidth or high-latency communications links, it is not feasible to wait for the acknowledgment of the write being sent to the disk write cache, or even the memory, on the preferred target system.

Synchronous transmission mode with asynchronous mirroring mode

Synchronous transmission mode with asynchronous mirroring mode is similar to sync/sync mode, with one important difference. The communications method of waiting for write acknowledgment from the mirror copy system is still used in this case (see Figure 6-40).

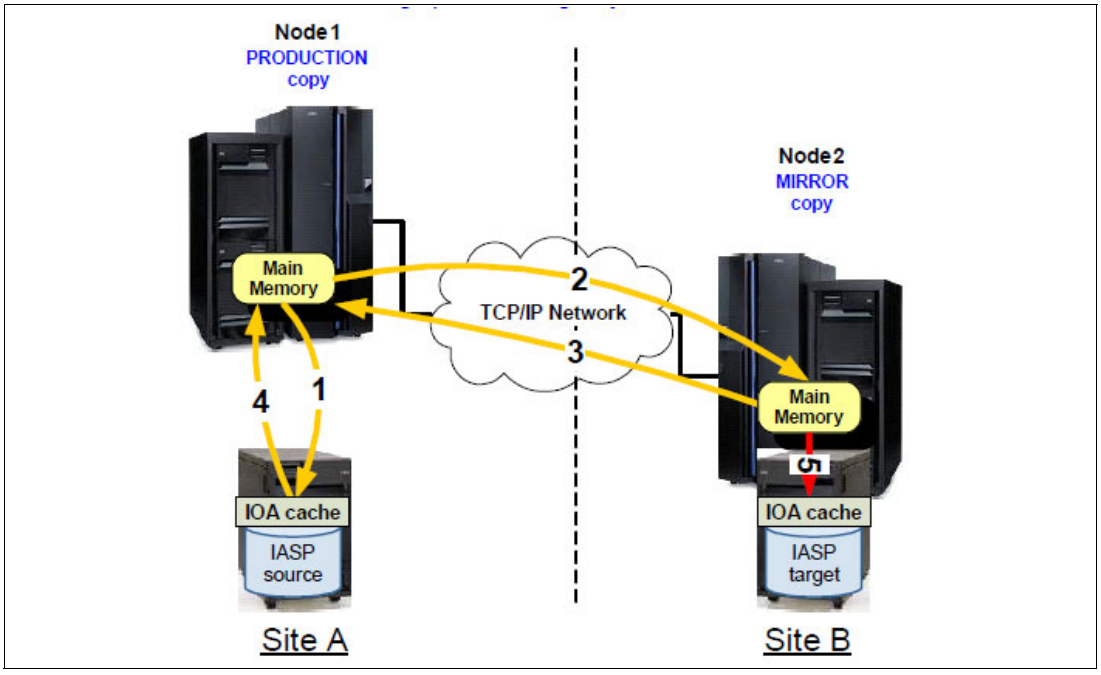

Figure 6-40 Synchronous transmission mode with asynchronous mirroring mode

However, when you use asynchronous mirroring mode, the acknowledgment is sent back from the mirror copy system when that data is in memory on the mirror copy system.

This approach provides a faster acknowledgment because no waiting is necessary for the write to finish on the mirror copy system.

The physical write operation (5) is performed later (asynchronously) to the disk on the mirror copy system. This approach is sometimes referred to as *sync/async mode*.

In this mode, the pending updates must be completed before the mirror copy can become the production copy.

Performance might improve slightly on the production copy system during normal operation. However, switch-over or fail-over times are slightly longer because changes to the backup IASP are still in the main memory of the backup system. These changes must be written to disk before the IASP can be varied on.

Asynchronous transmission mode with asynchronous mirroring mode

Asynchronous transmission mode is available for environments in which synchronous transmission mode is not a feasible solution. In asynchronous transmission mode, the job on the production copy system that performs the write to the IASP does not need to wait for the mirror copy write acknowledgment.

Asynchronous transmission mode with asynchronous mirroring mode specifies that transmission delivery and mirroring mode are asynchronous. This approach is sometimes referred to as *async/async mode* (see Figure 6-41).

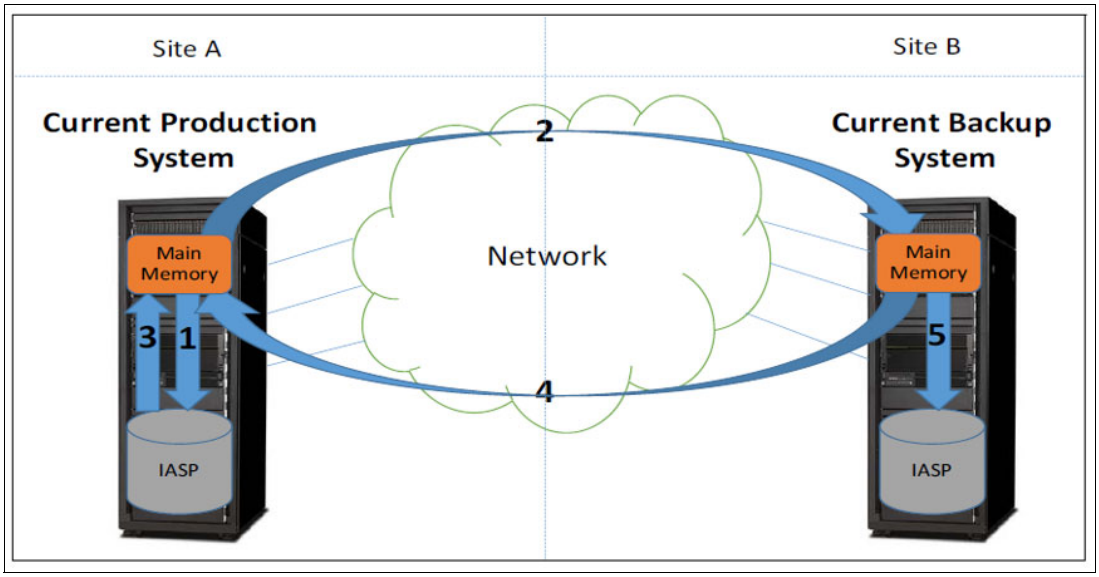

Figure 6-41 Asynchronous transmission mode with asynchronous mirroring mode

Asynchronous transmission delivery allows replication of the IASP beyond synchronous geographic mirroring limits.

The write to disk operation does not wait until the operation is delivered to the mirror copy system. Asynchronous transmission delivery requires asynchronous mirroring mode.

It works by duplicating any changed IASP disk pages in the *BASE memory pool on the production copy system. Likely, page faulting results in the *BASE pool. These duplicated disk pages are sent asynchronously while the write order to the target is preserved.

Therefore, at any time, the data on the target system (although not updated) still represents a crash-consistent copy of the source system.

The asynchronous geographic mirroring option potentially affects performance to system resources, such as processor and memory.

Communication lines with longer latency times might use more memory resources for maintaining changed data. Therefore, the environment must be sized correctly in terms of bandwidth, latency, and system resources.

Other geographic mirroring considerations

This section provides a few considerations when implementing geographic mirroring.

Tracking space

Tracking space enables geographic mirroring to track changed pages while it is in suspended status. Often, with tracked changes, geographic mirroring can avoid full resynchronization after it resumes, which minimizes the exposure of time frames where no valid mirror copy is available.

Tracking space is configured when geographic mirroring is configured. After installation, you can run the CHGASPSSN command to change the tracking space for the auxiliary storage pool session.

Tracking space is allocated inside the IASPs. If you allocate more tracking space, the system can track more changes. The amount of space for tracking can be defined by the user to a maximum of 1% of the total IASP capacity.

Suspend timeout

The suspend timeout in the ASP session specifies how long the application can wait when geographic mirroring cannot be performed. When an error, such as a failure of the communication link, prevents geographic mirroring from occurring, the production copy system waits and retries during the specified suspend timeout before it suspends geographic mirroring, which allows the application to continue.

The timeout value can be tuned by using the command CHGASPSSN. The default value of the suspend timeout parameter, SSPTIMO parameter is 120 seconds.

Network topology

Geographic mirroring can be used in environments over any distance. However, business needs determine the latency that is acceptable for a specific application. Many factors affect communications latency. As a result, these factors might affect geographic mirroring performance.

Consider the following points about communications in a geographic mirroring environment:

- To provide consistent response time, geographic mirroring must have its own redundant communications lines. Geographic mirroring supports a maximum of four communications lines for data port services. If the environment is configured with multiple lines, geographic mirroring distributes the load over multiple lines. Also, multiple communication lines among the nodes provide redundancy to the configuration.
- Geographic mirroring replication must be configured on separate interfaces and lines from the cluster heartbeat interface and line. If the same line is used during periods of heavy geographic mirroring traffic, the cluster heartbeat messaging process can fail between nodes, which can cause a node to go into a Partition status.
- When you use multiple lines, the line speed setting must be the same for each line. If throughput differs, performance is gated by the slowest connection.
- Typically, the throughput for this traffic is determined by the slowest link in the network path between the sites, unless those sites are in the same physical location and local network.
- The communications bandwidth must be sized for the resynchronization process and normal runtime operations that occur in parallel. If geographic mirroring is suspended and a resynchronization must occur, the tracked changes function is performed. During this time, the high priority synchronization tasks are used on the production copy system, which might degrade application performance if the network bandwidth between the sites is saturated.

From an HA perspective, geographic mirroring interfaces that are associated with different Ethernet adapters are considered a preferred practice. The use of redundant switches and routers further improves the overall HA value of the environment (see Figure 6-42).

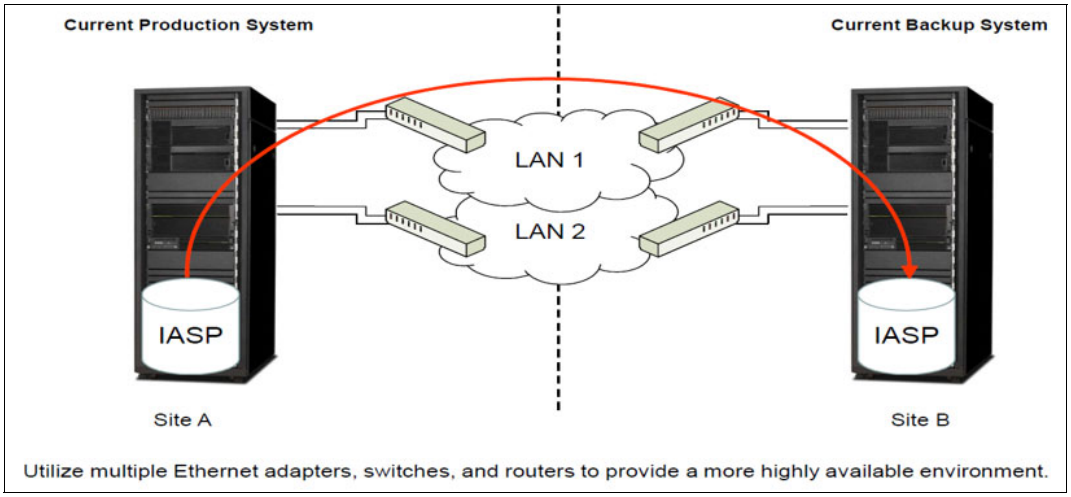

Figure 6-42 Preferred network configuration for Geographic Mirroring

Synchronization priority

One of the attributes that must be considered for geographic mirroring is synchronization priority.

The synchronization priority setting refers to a full or partial synchronization that is performed initially after you set up geographic mirroring, and after the geographic mirror session is suspended or detached.

Synchronization priority determines the importance of how quickly geographic mirroring attempts to synchronize the current production copy of the IASP to the current backup copy of the IASP.

It is always best to synchronize the data as quickly as possible. However, potential drawbacks exist if you simply set the synchronization priority to *HIGH in all configurations.

The following synchronization priority settings are available:

- *High synchronization priority* attempts to perform the synchronization as quickly as possible, without concern for system performance. No artificial delays are introduced into the transmission of the data.
- *Medium synchronization priority* attempts to perform the synchronization with a balanced approach between transmission performance and system performance. An artificial delay of 1 second occurs for every 100 segments of data that is sent.
- *Low synchronization priority* attempts to perform synchronization and ensure that the effect on system resources and performance is minimal. An artificial delay of 1 second occurs for every 10 segments of data sent.

Calculating full synchronization time bandwidth

When you plan for bandwidth requirements, you must determine the bandwidth that is required for a full synchronization.

To determine an approximation for the time that is needed for initial synchronization, take the total space that is used in the IASP times 8 to convert the number of bytes to bits (because throughput is normally measured in bits per second) divided by the effective communications capability of the chosen communications links.

Calculating runtime bandwidth

Writes, which are geographically mirrored to the mirror copy system, and that occur during normal runtime operations to the IASP must also be considered when you determine bandwidth requirements. Runtime bandwidth can be calculated by sampling the collection services data that is on the production copy system.

System performance considerations

With geographic mirroring, IBM i performs the replication. It is important to consider performance when you plan to implement a geographic mirroring solution.

Although asynchronous geographic mirroring allows for more flexibility for the distance between systems, undersizing the source, target, or the network infrastructure between the sites can reduce performance.

Latency is the time that the production system waits for the acknowledgment that the information was received on the target system. Minimizing the latency helps to ensure good application performance.

Backup planning for geographic mirroring

Before you implement HA that is based on geographic mirroring, define and implement a backup strategy to support your business environment.

If the current strategy is to stop or quiesce production applications before starting a back up on the systems, you can continue to run backups as usual if your business does not plan to change this strategy. This strategy ensures that the mirror copy is in sync for the longest period.

If you want backups from the target/mirror copy, geographic mirroring does not allow concurrent access to the mirror copy of the IASP.

This rule has implications that are considered when you perform backups.

6.2.3 Implementation of PowerHA SystemMirror for i geographic mirroring in the Power Virtual Server environment

This section describes the implementation of PowerHA SystemMirror for i Geographic Mirroring in the Power Systems Virtual Server.

Considerations for the creation of the IBM i instances

When you create the IBM i instances by using IBM Cloud Services, note the recommendations that are described in this section in preparation for building the environments.

Creating the Production IBM i instance

When creating the instance, the default size for the load source (LS) volume is 80 GB. Start with *only* the Load Source. Later in the process, add each new volume individually (SYSBAS or IASP) to keep track of each disk unit ID when they are added because this process can otherwise be difficult to map later.

The names of the disks (from IBM Cloud Services) across *all* instances within that server must be unique. These names do not present themselves in the IBM i interface. Instead, the names are visible from the Cloud interface only. Therefore, it is useful to individually keep track of the disk unit ID from the IBM i interface and keep track of the disk name from the Cloud interface so that you can assign your disks to the suitable ASP. Create these disks at a later step in the process.

When choosing a name for the instance, consider that IBM i uses the first 8 characters of that instance name as the system name. Therefore, it is best to choose a unique 8-character instance name as well.

Complete the following steps:

1. In the cloud GUI, choose IBM PowerHA under Software Licenses, which loads and enables the PowerHA Enterprise Edition by default. This choice allows for asynchronous geographic mirroring, which is best for replication between IBM i Cloud instances.
2. After the instance is created and active, open the console and wait for the log-in window. If the default image was deployed, the `qsecofr` password (`QSECOFR`) is disabled and must be changed.

   Note: This window is the operating system log-in. The DST/SST password is changed later.

3. In the Work with software agreements window, use **Option 5 (Display)** for all, and press **Enter**.
4. Press **F15 (Accept ALL)** on each license.
5. After all licenses are accepted, press **F3 (Exit)** to exit to the main menu.

   During this time, the IP interfaces and line descriptions are still being configured. This process can take 5 minutes to complete. If any external IPs were requested, they are not displayed in CFGTCP, Option 1 (Work with TCP/IP interfaces), as shown in Figure 6-43.

```
                        Work with Line Descriptions
                                                              System:   IBMIPROD
Position to  . . . . .  __________   Starting characters

Type options, press Enter.
  2=Change   3=Copy    4=Delete   5=Display   6=Print   7=Rename
  8=Work with status   9=Retrieve source


Opt  Line        Type        Text
  _   CLOUDINIT0  *ELAN
  _   CLOUDINIT1  *ELAN
  _   CLOUDINIT2  *ELAN
```

Figure 6-43 Working with line descriptions

The line descriptions are created as CLOUDINITx names, and the TCP/IP interfaces are assigned to those names automatically. The line descriptions are configured as ONLINE(*YES) and TCP/IP interfaces are configured as AUTOSTART(*YES), but can be changed to fit the needs of the business (see Figure 6-44).

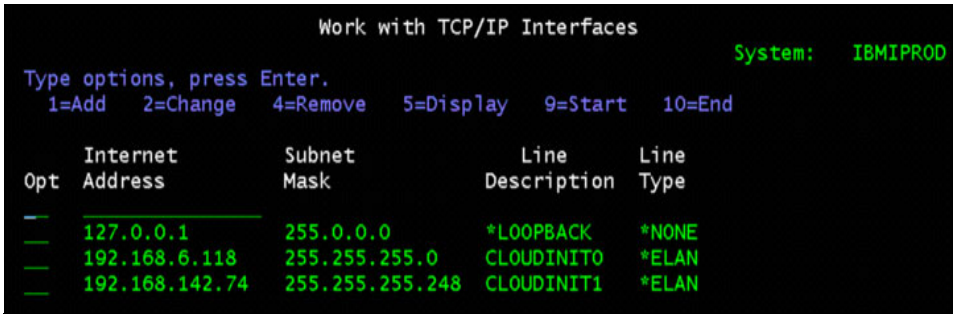

Figure 6-44 Working with TCP/IP interfaces

After TCP/IP is configured, the following changes are recommended:

- CHGSYSVAL SYSVAL(QIPLTYPE) VALUE(1)

  This value is changed back after the DST/SST password is changed upon the next IPL.
- CHGSYSVAL SYSVAL(QLMTSECOFR) VALUE(0)
- CHGSYSVAL SYSVAL(QAUTOVRT) VALUE(100)

  This value can be any number greater than zero.
- CHGTCPSVR SVRSPCVAL(*TELNET) AUTOSTART(*YES)
- CHGTCPSVR SVRSPCVAL(*FTP) AUTOSTART(*YES)
- CHGTCPSVR SVRSPCVAL(*SSHD) AUTOSTART(*YES)
- CHGTCPSVR SVRSPCVAL(*INETD) AUTOSTART(*YES)

  At this point, from the IBM Cloud interface, the extra disks for SYSBAS and IASP can be added.

  It is best practice to add these disks individually, and then record the associated disk unit ID that is presented in the IBM i. This process is useful in determining which disks are assigned to SYSBAS and which are assigned to the IASP.

  Each new disk presents itself as an unconfigured disk. Wait until after the next steps to add or balance any other SYSBAS disks.

  Issue the command PWRDWNSYS OPTION(*IMMED) RESTART(*YES) to shut down and IPL the IBM i instance in Attended mode with DST per the CHGSYSVAL QIPLTYPE.

  When the initial menu is displayed, use DST to change the DST password for `qsecofr` (default is `QSECOFR`).

  Adding or balancing the extra SYSBAS disks can be done now, or as a separate step after the IBM i instance completes the IPL.

  After the IPL and the extra SYSBAS disks are added or balanced, change the IPL type back to Unattended.
- CHGSYSVAL SYSVAL(QIPLTYPE) VALUE(0).

Be sure that all other post-IPL jobs and servers are active as expected.

Creating the HA/DR IBM i instance

Follow the same recommendations that are described for the Production IBM i and ensure that unconfigured disks for the IASP remain unchanged.

Establishing network connectivity between production and HA/DR instances

After both instances are created, it is imperative that both instances can communicate for intended PowerHA Cluster communication and PowerHA Geographic Mirroring replication.

For better performance, it is recommended to use one set of IP interfaces for inter-LPAR communication, such as FTP, remote journaling, and BRMS, and include PowerHA Cluster communication.

A secondary set of interfaces can be reserved for PowerHA Geographic Mirroring so that any spikes in replication traffic have less effect on packet delay for other inter-LPAR communication.

It is also useful to use a separate IP interface for user log-in and administration. This interface can be an internal IP interface, where the user connects to the corporate network, and then accesses the systems by using the internal IP. Optionally, an external IP interface can be assigned to the instance by using the IBM Cloud Services.

Creating IASP

The IASP is created on the Production IBM i instance in preparation for IASP-enablement.

These disks are virtual volumes that are created from external storage, with RAID protection pre-assigned. Therefore, protection does not need to be added to the IASP disks upon creation. To create the IASP, run the following command (see Figure 6-45):

```
CFGDEVASP ASPDEV(<IASP Name>) ACTION(*CREATE) TYPE(*PRIMARY) PROTECT(*NO)
ENCRYPT(*NO) UNITS(*SELECT)
```

```
                    Configure Device ASP (CFGDEVASP)

Type choices, press Enter.

ASP device . . . . . > IASP          Name, *ALL
Action . . . . . > *CREATE       *CREATE, *DELETE, *PREPARE
ASP type . . . . . .   *PRIMARY      *PRIMARY, *SECONDARY, *UDFS
Protection . . . . . .   *NO          *NO, *YES
Encryption . . . . . .   *NO          *NO, *YES
Disk units . . . . . .   *SELECT      Name, *SELECT
             + for more values
```

Figure 6-45 Creating an IASP

Figure 6-46 shows that the user is prompted to select which disks are to be added to the IASP from list the unconfigured disks. After the selections are made, press **Enter** to begin IASP configuration. This process can take several hours, depending on the number and size of the disks.

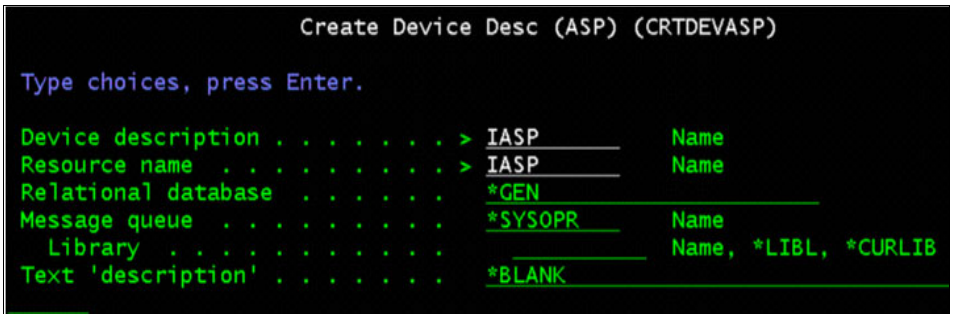

Figure 6-46 Creating an ASP device

After the process is completed, the ASP device description can be varied on, and the resource name shows with the same name as the IASP.

At any time, a matching device description can be created on the HA/DR node with the following command:

```
CRTDEVASP(<IASP Name>) RSRCNAME(<IASP Name>)
```

This information is only the device description. No IASP exists on the HA/DR node now, nor should such an IASP exist. The unconfigured disks for PowerHA Geographic Mirroring are created later.

Note: It is recommended that customers use IBM Lab Services for assistance in migrating a non-IASP environment to an IASP-enabled environment *before* moving forward with any PowerHA-managed solution. Only data that is included in the IASP is being replicated, and switchover or failover activities fail if the partitions were not correctly configured to support IASP.

PowerHA clustering configuration

After the production IBM i system is enabled for IASP, a DR solution with PowerHA Geographic Mirroring can be implemented.

On both instances, run the following commands:

- CHGTCPSVR SVRSPCVAL(*INETD) AUTOSTART(*YES)
- STRTCPSVR SERVER(*INETD)
- CHGNETA ALWADDCLU(*ANY)

Creating the cluster

From the Production IBM i instance, complete the following steps to create the PowerHA Cluster:

1. Run the following command:

   ```
   CRTCLU CLUSTER(<Cluster Name>) NODE((<Prod Node Name> (<Prod Node IP Address>))
   START(*NO)
   ```

2. Run the WRKCLU command and press Enter to verify that the cluster was created.

   From the WRKCLU menu, **Option 6 (Work with cluster nodes)** shows the production node with a status of New.

3. Select **Option 8 (Start)** to start clustering on the production node. The status changes to `Active`.
4. From the same window, select **Option 1 (Add)** to add the DR node to the cluster. Enter the IP address, and specify the Start Indicator as *NO. That node is listed and shows a status of `New`.
5. Select **Option 8 (Start)** to start clustering on the DR node, as shown in Figure 6-47. The status changes to `Active`.

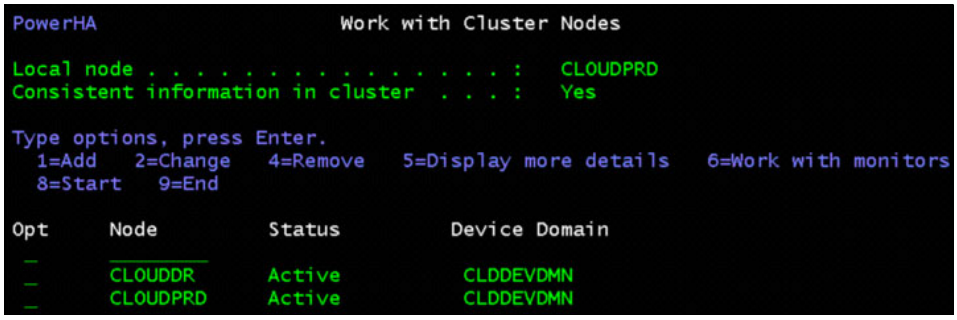

Figure 6-47 Working with cluster nodes

Adding cluster nodes to the device domain

Note that from WRKCLU, Option 7 [Work with device domains], does not show a device domain listed. From the production IBM i instance, complete the following steps to add the PowerHA cluster nodes to the same device domain:

1. Select **Option 1 (Add)** to add a device domain to the cluster, and specify the production node name.
2. The WRKCLU, Option 7 window lists the new device domain with one node (the production node) listed to the right (see Figure 6-48).

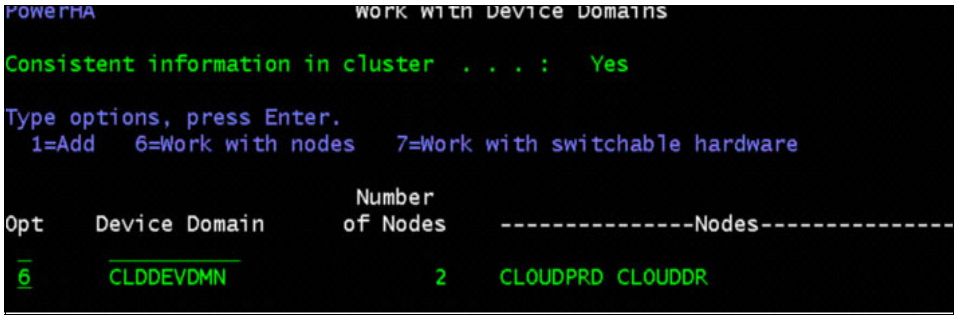

Figure 6-48 Working with device domain

3. Select **Option 6 (Work with nodes)** and select **Option 1 (Add)** to add the DR node to the device domain. The window now lists both nodes as `Active` in that device domain.

The WRKCLU, Option 7 window shows the device domain with both nodes listed, as shown in Figure 6-49.

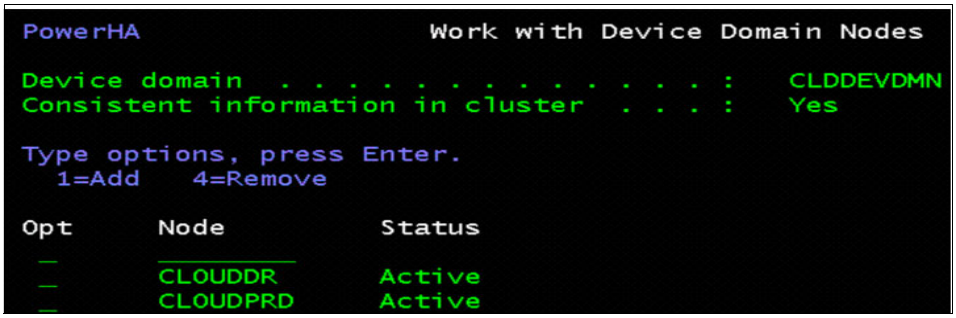

Figure 6-49 Working with device domain nodes

Creating the device cluster resource group

From the Production IBM i instance, complete the following steps to create the device cluster resource group (CRG), which is used for switchover and failover of the IASP (the WRKCLU, Option 9 [Work with cluster resource groups] window is empty):

1. Select **Option 1 (Create)** and enter a name for the CRG that uses 10 characters or less and press **Enter** (see Figure 6-50).

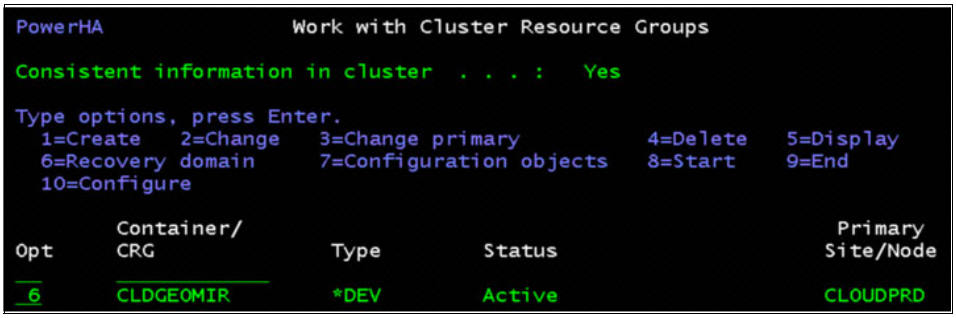

Figure 6-50 Working with cluster resource groups

2. Select **Option 1 (Cluster resource group)** from the pop-up menu and press Enter.
3. Specify `Type=*DEV, Exit Program=*NONE`, and `User Profile=*NONE`.
4. Enter a “+” next to `Recovery domain node list` and press Enter.
5. For the first entry in the recovery domain (see Figure 6-51), use the name of the production node with a node role of *PRIMARY and a site name that is unique from the DR node. Specify at least one IP address to be used by the production node for geographic mirroring (replication) and page down.

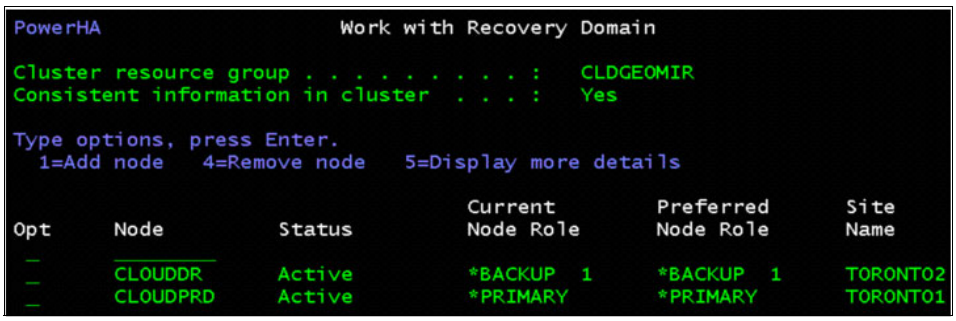

Figure 6-51 Work with recovery domains

6. For the next entry in the recovery domain, use the name of the DR node with a node role of *BACKUP, a sequence number of 1, and a site name that is unique from the Production node. Specify at least one IP address to be used by the DR node for geographic mirroring (replication). Press Enter and page down.
7. Optionally enter a text description for the CRG.
8. Press Enter to create the CRG. The status shows as `Inactive`.

Adding the IASP to the device CRG

From the Production IBM i instance, complete the following steps to add the IASP to the device CRG (WRKCLU, Option 9 [Work with cluster resource groups] shows the device CRG as `Inactive`):

1. Select **Option 7 (Configuration objects)**, as shown in Figure 6-52 to show that objects were not yet added to the CRG.

```
PowerHA                Work with Cluster Resource Groups

Consistent information in cluster  . . . :   Yes

Type options, press Enter.
  1=Create    2=Change    3=Change primary          4=Delete    5=Display
  6=Recovery domain       7=Configuration objects   8=Start     9=End
  10=Configure

         Container/                                                Primary
Opt      CRG               Type        Status                     Site/Node
__       __________
 7       CLDGEOMIR         *DEV        Active                     CLOUDPRD
```

Figure 6-52 Working with cluster resource groups

2. Select **Option 1 (Add)** with the name of the IASP and a type of *DEVD and press Enter (see Figure 6-53).

```
PowerHA                Work with Configuration Objects

Cluster resource group . . . . . :   CLDGEOMIR
Consistent information in cluster  . . . :   Yes

Type options, press Enter.
  1=Add    2=Change    4=Remove    5=Display more details
  6=Configuration status

         Configuration      Object      Device      Device         Vary
Opt       Object Name       Type        Type        Subtype        Online
 _        __________        ______
 _        IASP01            *DEVD       *ASP        Primary        *ONLINE
```

Figure 6-53 Working with configuration objects

3. If you want to set automatic vary-on for the IASP after switchover or failover (which is more common), change the `Configuration object online` to *ONLINE.
4. Specify *NONE for the `Server takeover IP address` because managing an automated switch of a primary IP interface is done more efficiently with IASP exit programs. Press Enter to complete the process.

The IASP is now be listed in WRKCLU, Option 9, Option 7.

Starting geographic mirroring of the IASP

From the Production IBM i instance, complete the following steps to start geographic mirroring on the IASP:

1. Run the following command:

```
CFGGEOMIR ASPDEV(<IASP Name>) ACTION(*CREATE)
SSN(<DR Site ASP Copy>/<Prod Site ASP Copy>/<ASP Session Name>)
DELIVERY(*ASYNC) UNITS(*SELECT)
```

2. Press Enter.

   Consider the following points:

   – DR Site ASP Copy and Prod Site ASP Copy are unique names that are assigned to differentiate which copy of the IASP is referenced. Both copies are the same IASP, but exist on two separate sets of disks (see Figure 6-54).

```
                    Configure Geographic Mirror (CFGGEOMIR)

Type choices, press Enter.

ASP device . . . . . . . . . . . > IASP          Name
Action . . . . . . . . . . . . . > *CREATE       *CREATE, *DELETE
Source site  . . . . . . . . . .   TORONTO1      Name, *
Target site  . . . . . . . . . .   TORONTO2      Name, *
Session  . . . . . . . . . . . .   CLDGEOMIR     Name, *NONE
  Source ASP copy description  .     CLOUDPRD    Name
  Target ASP copy description  .     CLOUDDR     Name
Transmission delivery  . . . . .   *ASYNC        *SYNC, *ASYNC
Disk units . . . . . . . . . . .   *SELECT       Name, *SELECT
               + for more values
```

Figure 6-54 Configuring a geographic mirror

   – ASP Session Name is the label for the geographic mirroring replication, against which actions can be performed; for example, Suspend, Resume, Detach, and Reattach.

3. The next window shows a list of unconfigured disk units from the DR node. Select those units that are to be included in geographic mirroring. Select **F9 (Calculate selection)** to view the capacity that results after the configuration process is complete.

4. Press **Enter** to begin configuration of geographic mirroring. This process can take up to several hours, depending on the number and size of disks used (see Figure 6-55).

```
PowerHA                 Work with ASP Copy Descriptions              IBMIPROD
                                                          07/16/20  09:31:52
Device domain  . . . . . . . . . . . . . :   CLDDEVDMN

Type options, press Enter.
  1=Add copy          2=Change copy       4=Remove copy     5=Display copy
  21=Start session    22=Change session   24=End session    25=Display session

        ASP             ASP             ASP               Session
Opt     Device          Copy            Session             Type

25      IASP01          CLOUDPRD        CLDGEOMIR         *GEOMIR
        IASP01          CLOUDDR         CLDGEOMIR         *GEOMIR
```

Figure 6-55 Working with ASP copy descriptions

5. After the configuration process is complete, verify that `DSPASPSSN SSN(<ASP Session Name>)` shows the status as *RESUMPND.

6. Vary on the IASP device, and recheck the ASP session status to verify that it shows as *SUSPENDED.

7. Select **WRKCLU, Option 9 (Work with cluster resource groups)** and use **Option 8 (Start)** to start the CRG. The status changes to `Active` (see Figure 6-56 on page 258).

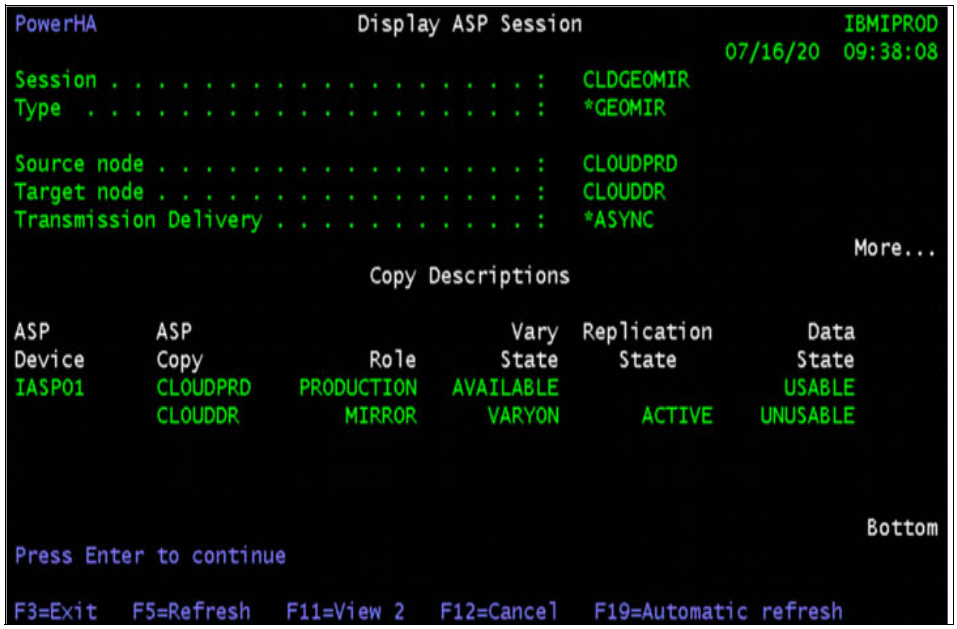

Figure 6-56 Displaying ASP session

`CHGASPSSN SSN(<ASP Session Name>) OPTION(*RESUME)` restarts synchronization of geographic mirroring.

6.2.4 Switching a geographic mirroring environment

This section includes a description of the geographic mirroring synchronization process and shows how to perform a planned switch of a geographic mirroring environment between production and backup nodes. In addition, the following procedures to recover from an unplanned failover are also discussed:

- Synchronization process
- Planned switchover
- Unplanned failover

Synchronization of geographic mirroring

When geographic mirroring is resumed after a suspend or detach, the mirror copy is resynchronized with the production copy.

The production copy can function normally during synchronization, but performance might be affected negatively.

During synchronization, the contents of the mirror copy are unusable, and it cannot become the production copy. If the independent disk pool is made unavailable during the synchronization process, synchronization resumes where it left off when the IASP is made available again.

The message `CPI095D Cross-site Mirroring (XSM) synchronization for IASP` is sent to the QSYSOPR message queue every 15 minutes to indicate the progress of the synchronization and type.

The following types of synchronization exist:

- *Full synchronization* indicates that a complete synchronization occurs. Changes to the production copy were not tracked to apply to the synchronization, or the data status of the mirror copy was not determined.

  First, a full synchronization deletes all data in the backup IASP. Then, the full synchronization copies the current data from the production IASP to the backup IASP.
- *Partial synchronization* indicates that changes to the production copy and mirror copy were tracked while geographic mirroring was suspended or detached. This tracking can shorten the synchronization time considerably because a complete synchronization is unnecessary.

  In this case, when the mirror copy is reattached and geographic mirroring is resumed, only tracked changes must be synchronized. Changes that are made on the production copy (since the suspend or detach was performed) are sent to the mirror copy, and any changes that were made on the mirror copy are overwritten with the original data from the production copy of the IASP.

Important: Any changes that are made on the mirror copy while it is detached are undone, and any tracked changes from the production copy are applied.

Switching a geographic mirroring environment

This section describes how to switch to a geographic mirroring environment.

Planned switchover

Perform a WRKASPJOB to check which jobs use the IASP. End all applications and jobs by using the IASP.

Complete the following steps to perform a planned switch by using IBM i control language (CL) commands:

1. From the Work with Cluster (WRKCLU) menu that is shown, select **Option 9** to work with cluster resource groups.
2. On the Work with Cluster Resource Groups display, specify **Option 3 (Change primary)** to change the CRG primary node.
3. On the Change CRG Primary (CHGCRGPRI) display, press Enter.

   The display shows a status of inhibited while the switch occurs.
4. After the switch is finished, you can run the following Display ASP Session command to confirm that the nodes and ASP copies were reversed:

   ```
   DSPASPSSN SSN (CLDGEOMIR)
   ```

Unplanned switchover

A failover occurs when the source node fails and the backup node takes over. The default failover procedures depend on the cluster failover wait time and failover default action settings.

The following failure scenarios can help to explain failover:

- A primary node failure triggers an automatic failover.

  This scenario can occur by a panic message that is sent to the backup node or cluster monitors in place.
- A primary node or cluster communications failure results in node Partition status.

Each scenario requires different failover and recovery actions.

Important: As with a planned switch, the success of a failover operation also depends on your previous testing and verification to show that the business applications can run on the backup node and users can access those applications.

Primary node failure triggers an automatic failover

The following events can create a panic message that triggers a failover:

- A cluster node is offline.
- TCP is offline.
- All subsystems are offline.
- The system or partition is powered down.

An unplanned automatic failover event can be triggered by a panic message that is sent by the primary node because of ending a cluster node, Transmission Control Protocol (TCP), all subsystems, or powering down the partition or system. The failover can also be triggered by a *Power state change* event that is sent by the Hardware Management Console (HMC) server for a partition failure.

For an automatic failover event, an inquiry message `CPABB02 Cluster resource groups are failing over to node backup-node. (C G)` is sent to the cluster or CRG message queue on the backup node if a failover message queue is defined for the cluster or CRG.

If no failover message queue is defined, the failover starts immediately without posting any message.

The cluster parameters or Failover Wait Time (FLVWAITTIM) and Failover Default Action (FLVDFTACT) determine the next actions.

With the default settings of FLVWAITTIM=*NOWAIT and FLVDFTACT=*PROCEED, an automatic failover begins immediately.

Setting the Failover Wait Time parameter as a duration in minutes or *NOMAX allows a user on the backup node to respond to the CPABB02 inquiry message to proceed with the failover or cancel the failover.

The Failover Default Action parameter determines whether PowerHA proceeds with the failover or cancels the failover processing after the specified failover wait time expires and no response to the inquiry message was entered.

Note: Regardless of the cluster parameter settings, the primary IASP is taken offline by PowerHA for a failover event.

Primary node or cluster communications failure results in node partition status

In the absence of advanced node failure detection, a sudden primary node failure or a cluster communications failure can result in a condition that is known as a *cluster partition*.

The cluster partition condition means that the backup node cannot determine the status of the production node reliably. In this situation, an automatic failover does not occur, regardless of the cluster failover settings. In fact, a failover is not possible unless you perform the following steps:

1. Determine whether the primary node or source node is still in operation.

   If the production workload can continue, you do not need to fail over to the backup node. For a cluster communications failure only, message ID CPDB715 is issued on the primary node.

Mirroring continues and after communications are restored, the cluster software reconnects the cluster nodes automatically.

Otherwise, if the primary node is no longer responsive or available, the status of that node must be changed from Partition to Failed. A status of Partition or Failed allows a failover to proceed to the backup node by running the following steps:

2. Log on to the backup system and verify the node status. If the primary node is in a Partition status, a manual switch or failover cannot be performed. Use Work with Cluster Nodes (WRKCLU) menu Option 6.
3. The status of the primary node must be changed to Failed by using the Change Cluster Node Entry (CHGCLUNODE). This command also varies off the IASP on the target node.
4. The node status changed to Failed. The CHGCLUNODE command also triggers a cluster failover without a failover inquiry message, but it still requires the user to vary on the IASP on the new primary or source node and start the takeover IP interface.
5. Vary on the IASP on the backup (now source) node. The status changes to AVAILABLE. Also, start the takeover address, if required.
6. Run the following Display ASP Session (DSPASPSSN) command to verify that the node ASP copy shows as AVAILABLE:

   ```
   DSPASPSSN SSN(CLDGEOMIR)
   ```

   The backup node now has the production copy.

   Important: The terms can be confusing here. After a switch or failover, the *backup* node (which refers to a physical system or partition) becomes the *source* node and contains the *production* copy of the IASP when it is shown in the DSPASPSSN display.

7. After the original primary system is repaired and started, ensure that the IASP on that node is in a varied-off status and that the cluster node is `Inactive`. Start the repaired node from the current production node and verify that all nodes show Active, as shown in WRKCLU, Option 6.
8. Start the CRG if it is not started by using the following command:

   ```
   STRCRG CLUSTER(CLOUDCLU) CRG(CLDGEOMIR)
   ```

9. Use the Work with Cluster (WRKCLU) menu and select **Option 10** to work with ASP copy descriptions. Select **Option 22** to change session on either ASP's Opt line and press **F4**.
10. On the Change ASP Session display, enter *RESUME.
11. Refresh the display to show that the status changed to RESUMING.

Pressing **F11** displays the progress and the amount of data that is out of sync.

After an unplanned failover, a full resynchronization of the IASP is often required. The full resynchronization can take an extended time.

Consider the backup copy as unusable until the resume process is finished.

Note: It is considered a preferred practice to run the Reclaim Storage (RCLSTG) command on SYSBAS on a failed production node after you perform a failover and before you resume mirroring. Also, you must schedule an RCLSTG of the IASP at the earliest convenience, preferably on the backup node before you switch back to "preferred production".

6.2.5 Troubleshooting

With any new setup, errors can occur. Although some errors can be related to infrastructure issues, others arise because of familiarity with new technology and the differences that using the IBM Cloud presents from traditional solutions.

The following changes that can make alternative methods different than what is expected in non-cloud environments:

- When selecting disk units to assign to the new cloud instances, start by assigning only the initial Load Source (LS) volume, and identifying that volume in the IBM Cloud GUI for easy reference.

  After the instance is created, make note of the disk unit ID in SST/DST. Add subsequent disks individually, and use a naming scheme in the IBM Cloud GUI that easily identifies SYSBAS disks or IASP disks.

  At the same time, keep track of the associated disk unit ID from the IBM i instance. This information is crucial in determining which disk units to add with CFGDEVASP and CFGGEOMIR.
- When selecting a name for the IBM i instances, choose a name that is 8 characters or less, because the IBM i truncates the instance name to 8 characters when assigning a system name.
- As with any geographic mirroring solution, ensuring sufficient throughput on the replication interfaces is key. Separate other traffic onto other subnets (and ideally other adapters) so that the replication traffic is not hindered and subject to auto-suspend issues or lengthy switchover or failover processes.
- It is best practice to keep the PowerHA for IBM i clustering communication on a separate subnet than replication because it requires the “heartbeat” traffic to respond in a timely manner to prevent cluster suspend status (Partition status). This configuration prevents any wanted switchover or failover from being enabled.

For more troubleshooting with Power Systems Virtual Server infrastructure, use the Support option from the IBM Cloud Services GUI to submit a case summary for the IBM team to investigate. These issues are specific to the infrastructure, network, images, servers, disk enlistment, pathing issues, and so on.

For issues that are specific to PowerHA for IBM i geographic mirroring configuration, errors, or performance, start by contacting IBM Support and requesting the High Availability Solutions (HAS) team first. They can then contact development teams or Cloud Support, if necessary.

6.3 Logical replication use case with Bus4i

This section provides a logical replication use case with Bus4i.

6.3.1 Introduction

For many years, logical replication solutions were the only choice for high availability (HA) and Disaster Recovery (DR) for IBM i and its predecessors.

The introduction by IBM of PowerHA and its hardware-based replication solutions (at the operating system level or for External Storage-based) provided a different approach for HA and DR in the IBM i platform.

However, today most HA/DR solution implementations in the IBM i market are still logical replication solutions.

Logical replication solutions work by replicating changes that are made to database files and other IBM i objects in real time, which allows the customer to start a role swap within their HA/DR product if an outage occurs. It also promotes the secondary system into the production system, so the customer can continue operations.

The logical replication products use the IBM i built-in database journaling function to detect changes that are made to database files and the audit journal to detect changes that are made to other types of objects. Later, they use remote journaling or a propriety journal scrape method to get the data out of the journals before sending it across the network of the secondary server, where a journal apply process is used to update the secondary database.

The options that are available for implementing logical replication solutions for IBM i includes the following products:

- Bus4i from T.S.P. Company for Information Systems
- RobotHA from HelpSystems
- From Precisely:
  - Assure QuickEDD
  - Assure MIMIX
  - Assure iTera
- Maxava HA from Maxava
- Ha4i from Shield Advanced Solutions
- iSB-HA from iSam Blue‘s
- hiCluster from Rocket Software

For the purposes of the logical replication use case that is included in this chapter, the Bus4i product is used.

6.3.2 Prerequisites

Bus4i features the following minimum prerequisites:

- 1 GB of main memory
- 20 GB of available disk space
- The system is installed with IBM i 6.1 or later and the TCP/IP services for DDM and REXEC is active

6.3.3 Planning

Unscheduled computer downtimes can be challenging. This issue can be critical if permanent system availability is essential. Situations exist in which downtimes for repairs and maintenance must be avoided at all cost.

The T.S.P. IBM i High Availability System with BUS4i renders possible permanent mirroring of two or multiple complex systems of all IBM Power System models on IBM i basis.

In this process, production databases, objects, user profiles, IFS, authorization lists, spooled files, and control information about batch jobs are permanently mirrored. This mirroring

ensures in this process that operations can be carried out immediately with the secondary system when scheduled or unscheduled downtimes occur Figure 6-57

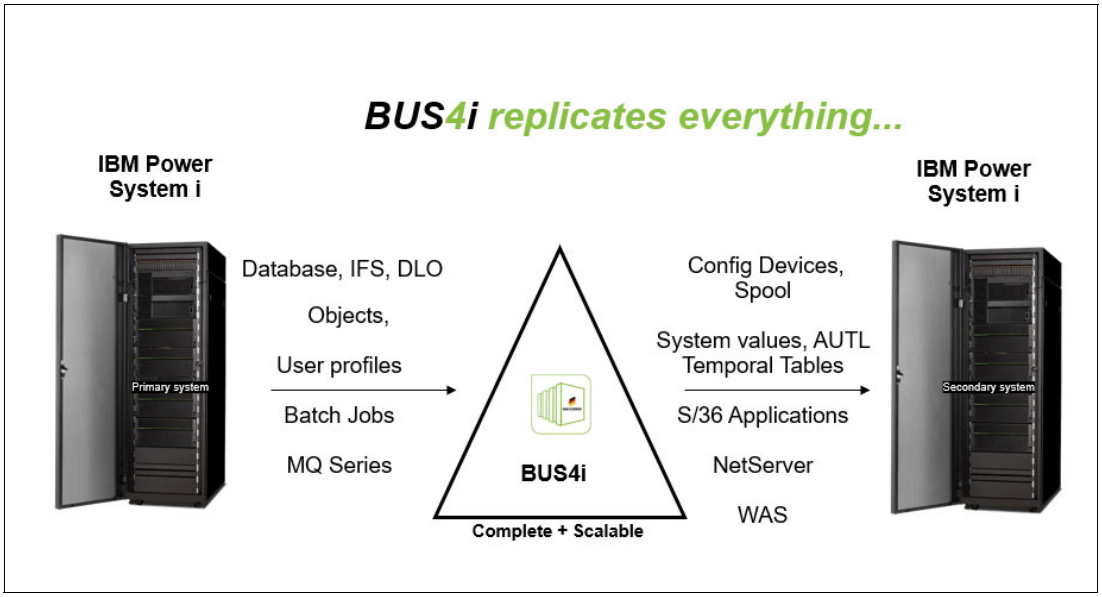

Figure 6-57 BUS4i replicates everything

When the nonfunctional server unit is restored, the systems are resynchronized.

To keep users consistently updated about the mirroring process status, they are notified of malfunctions through SMS or email.

To replicate a system that uses BUS4i, some product characteristics must be considered as described in the following sections.

Full integrated remote journal support

Users can choose between remote journal or local journal. When local journal is used, data transmission to the secondary system is reduced to a minimum by efficient filters and compression. These filters are based on the user requirements on an HA system and the specific application environment.

In various cases, guaranteed transmission of each transaction is required. The use of remote journals might be the best solution in this case because it is also the synchronous transmission and receiving process.

Permanent, high-performance data and object mirroring

Data changes in the databases and object modifications are mirrored in real time to one or multiple secondary systems. To make maximum use of data links, differing transactions are transmitted simultaneously. Data transmission is maximized to minimize data loss.

Intelligent sequence control

The transmitted data is applied in real time on the secondary system. Sequence control ensures that the data is applied in an identical order as in the primary system. Administrative interventions into the primary system are automatically duplicated to the secondary system (see Figure 6-58 on page 265).

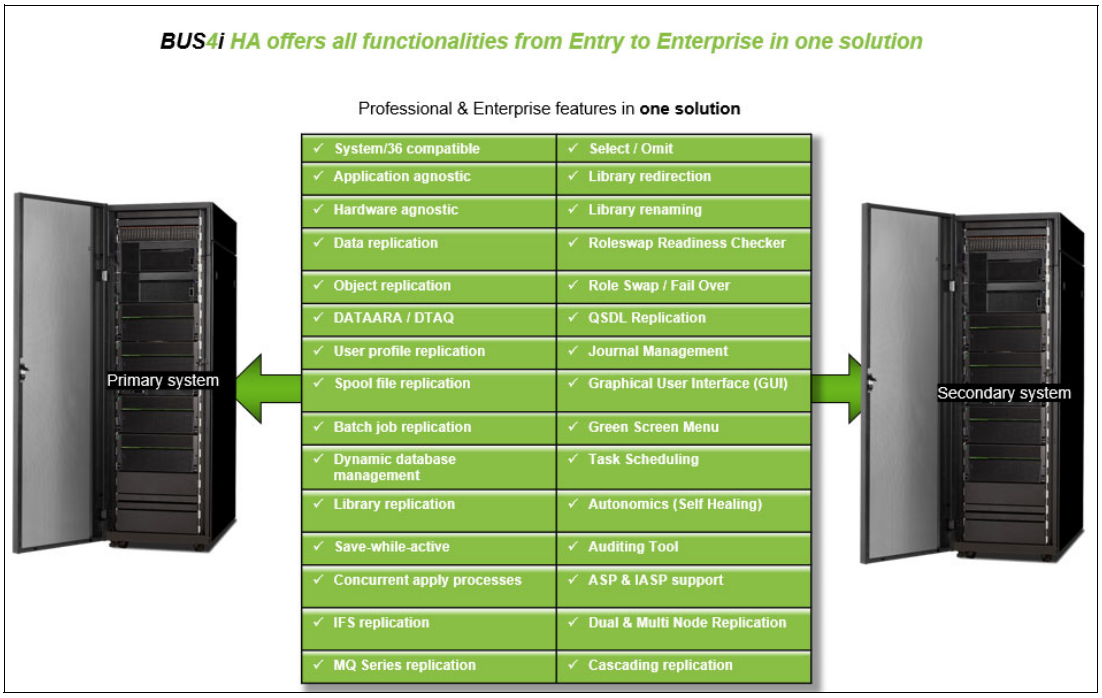

Figure 6-58 Entry to professional and enterprise functions in one solution

Intelligent sequence control restricts the continuous verification of data integrity and consistency to occasional sampling for internal auditing. Administrative costs are reduced to a minimum because all changes, amendments, or deletions are automatically mirrored to the secondary system in real time.

Efficient IFS file mirroring

The files in IFS are regularly updated or are created and never changed. IFS filters check and select only the most important directories for mirroring.

Transmission is restricted to changes, new files, or deletions in critical directories. This restriction leaves the bandwidth of data links available for other user data.

Automatic verification of data integrity

Agents in the primary and secondary systems send failure reports as required per SMS or email distribution lists.

Regular monitoring of systems is reduced to a minimum. At the same time, this safety feature gives administrators and management a sense of security.

Enhanced trigger and constraint support

BUS4i features intrinsic trigger and constraints support. Data doubling is avoided, and constraints are adhered to.

This support ensure that contemporary applications with trigger functions can be entirely mirrored and are fully operational if an emergency occurs.

Integrated LPAR support

The integrated LPAR support allows replication to another partition on the same machine.

You can use mirroring in a test system to verify alterations to the test system, which allows precise definition and minimization of downtimes in the real system.

Autonomy from user applications

It is not necessary to adapt user applications for implementation of HA software BUS4i. Expensive adaptations or special versions are not needed. You work in your familiar environment, as shown in Figure 6-59.

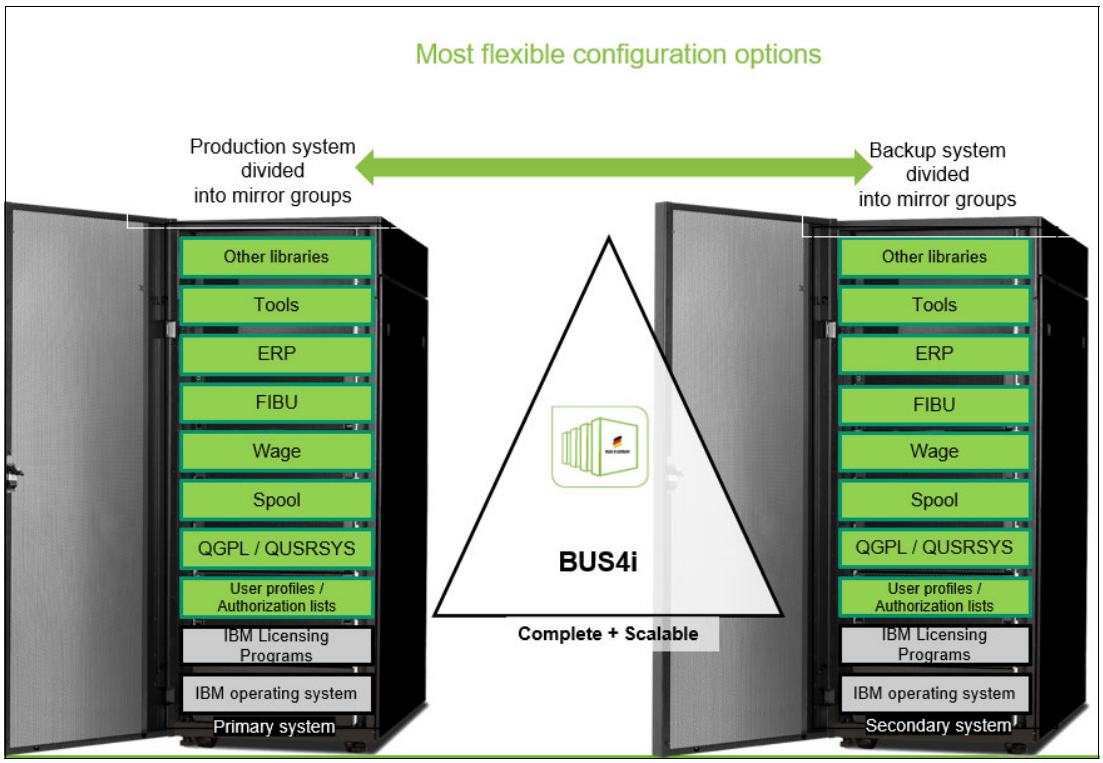

Figure 6-59 Most flexible configuration options

Secondary system with differing hardware models

One or multiple secondary systems are not identical to the original system. The operating system release is identical in only a few functions.

If a new system is installed, the unit that is in place can be used as backup system. A costly identical second system does not need to be installed.

User profile and authorization list

Administrative modifications are immediately mirrored to the secondary system. Changes to user parameters (for example, changing a password) are also mirrored in real time.

Administrative activities can be restricted to the primary system.

Asynchronous or synchronous transmitting and receiving process

In the finance sector or other critical functions, 100% mirroring of data is required. In other sectors, 99%–100% mirroring suffices.

Some applications require guaranteed copies, so synchronized copies are used. The transmission system waits with further processing until the receiving function acknowledges receipt and processing.

Routines that are independent of transmitting and receiving processes

Because of autonomous processes, users have more flexibility when they design their systems and routines (see Figure 6-60 on page 267).

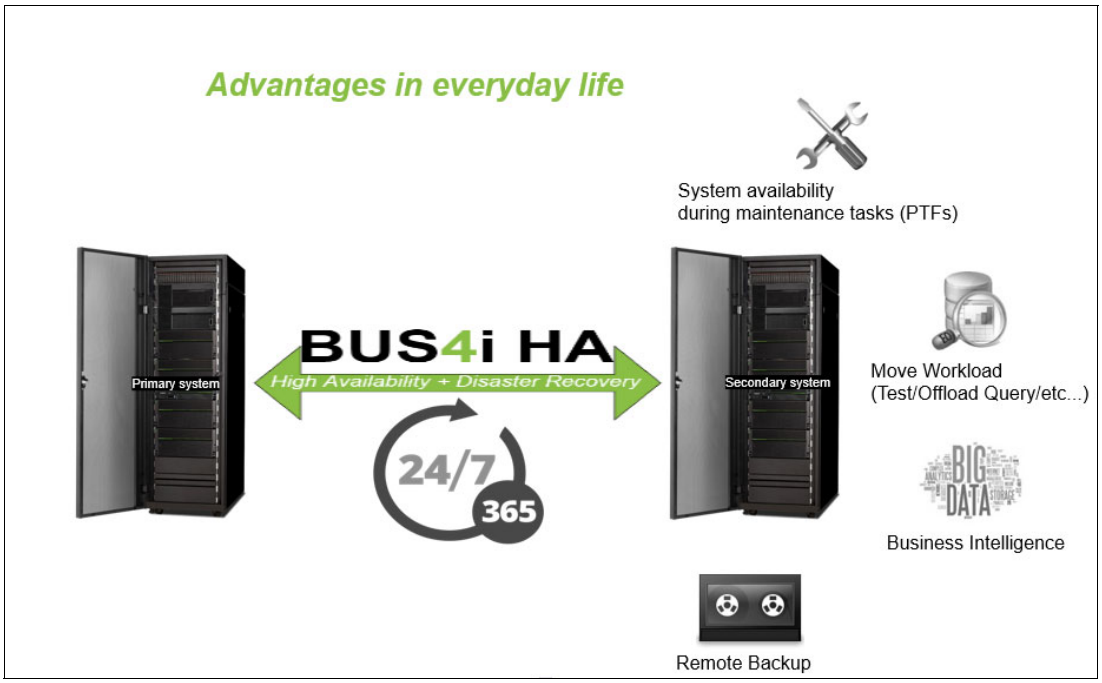

Figure 6-60 Advantages in everyday life

If the processing routine is stopped, consistent data backup on tape is made in the secondary system. Modifications are still accepted by the receiving process. Therefore, the backup window during data protection is minimized.

Spooled files

Some spooled files, such as print-out, cannot be repeated, for example, when printing invoices. The mirrored version in the secondary system is halted to assure only one print-out is made.

If an emergency occurs, print-out can be carried on by the secondary system without any complex recalculation.

Control information about batch jobs

Batch jobs can be of short or long duration. If a switch-over occurs, operators are aware of which jobs were completed and which jobs must still be processed.

When a switch-over occurs, starting new batch jobs on the secondary system is restricted to those jobs that are not yet started in the primary system. It is no longer necessary to look for jobs that must be run.

Integrated process monitoring

To ensure that mirroring is not affected by an inactive process, agents send reports about inactive processes by way of SMS or email, and prevent the asynchronicity of databases.

Your daily inspection of the process list is no longer necessary, and routine activities are reduced to a minimum.

Many replication scenarios

BUS4i can be implemented in many scenarios in a hybrid architecture or in the cloud (Figure 6-61 on page 268):

- Locally with only one Power Systems
- Locally with two Power Systems
- Remotely with on-premises sites

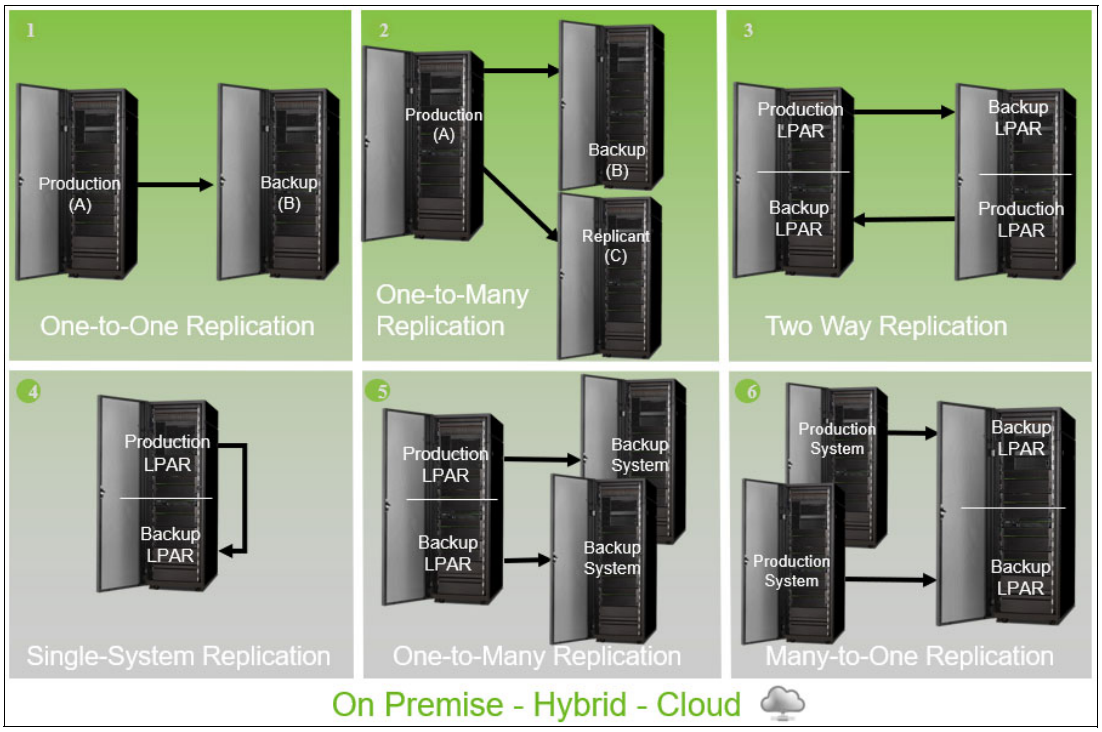

Figure 6-61 Multiple replication scenarios

The T.S.P. system BUS4i is based on the operational principle of journaling. Journals can record all changes to databases and objects. The licensed program BUS4i selects and sends the entries in the journal receiver by way of an IP connection to a second system, which is called a secondary system.

The second system features a copy of the production application, the database, and all affiliated objects. The BUS4i system on the secondary system processes the incoming journal entries and implements the changes to the database and the objects. Any defined commitment boundaries are considered. Therefore, two identical applications are available in autonomous systems.

If a scheduled or unscheduled production system downtime occurs, processes continue to run fast on the secondary system. The application is resynchronized when the production system is active again.

The BUS4i license program is composed of the following autonomous processes:

- Sending
- Receiving
- Applying

After it is configured and started, the sending and receiving processes and the applying processes run automatically. These processes include menu-driven maintenance and monitoring features.

A two-tiers server landscape consists of a production (primary) system and a backup (secondary) system. With BUS4i, a two-tier server landscape is defined with the following characteristics (Figure 6-62 on page 269).

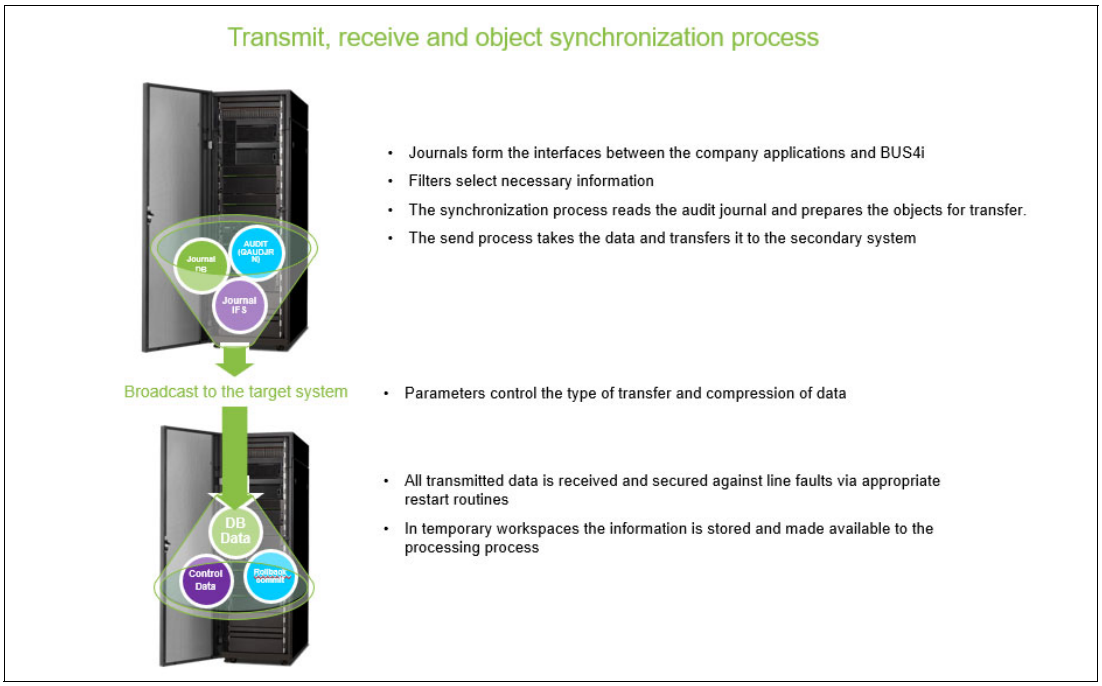

Figure 6-62 Transmit receive, and object synchronization process

An asynchronous sending process or a remote journal process immediately transfers to the secondary system any changes to database and objects.

The receiving process accepts data sets that are sent in the sending process and stores these data sets in a database file.

The applying process in the secondary system is autonomous from the sending or receiving process, as shown in Figure 6-63.

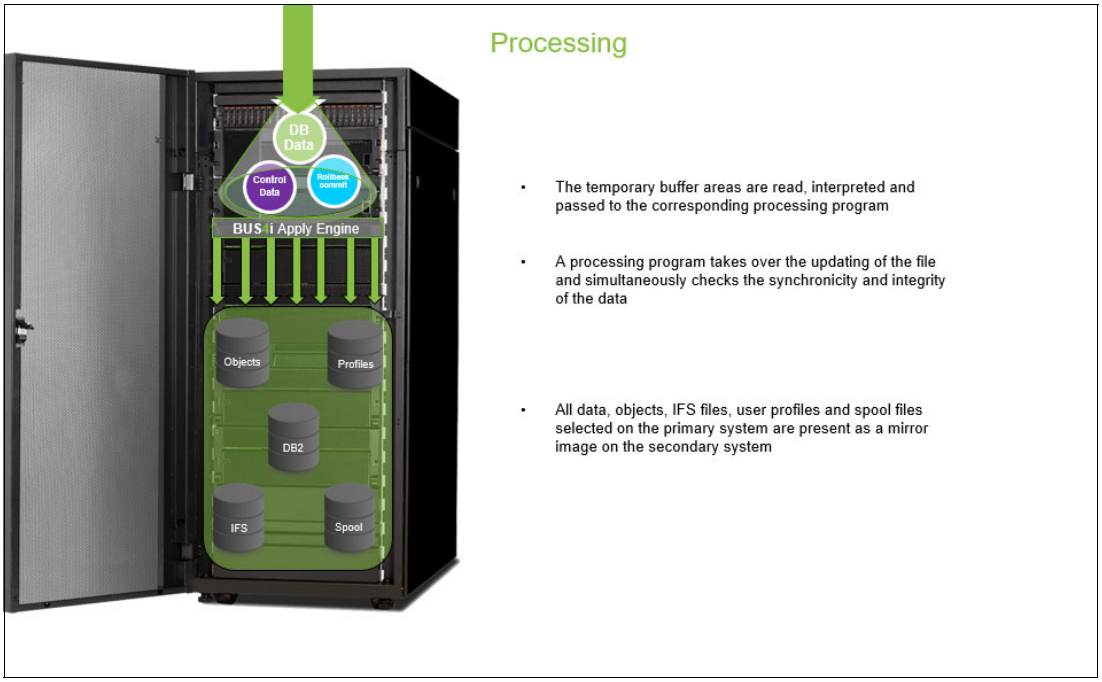

Figure 6-63 Processing process

Graphical user interface

The GUI offers a freely customizable interface, for example, for table columns, order and size. Figure 6-64 shows an overview of the available information.

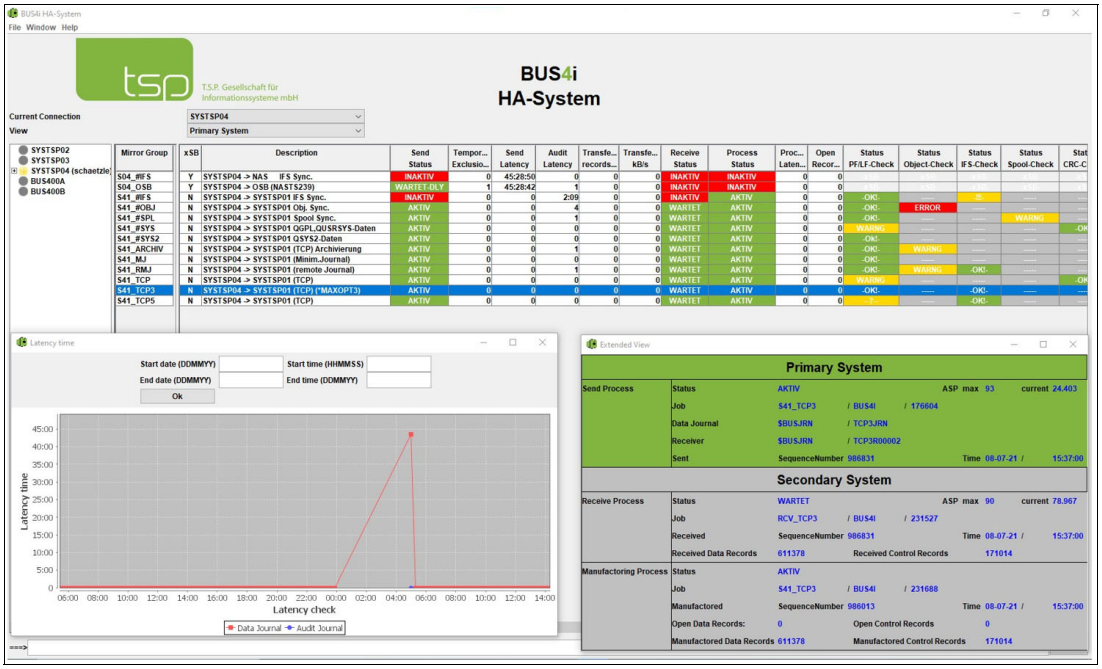

Figure 6-64 Graphical user interface

6.3.4 Implementation of Bus4i for IBM i Logical Replication

BUS4i installation is performed directly by T.S.P. experts or by Business Partners specialists who are certified in BUS4i by T.S.P. Depending on the complexity of the system, a typical BUS4i implementation can take one or two days.

From the Primary System, the BUS4i main menu can be accessed by running the **`bus4i`** command (see Figure 6-65 on page 271).

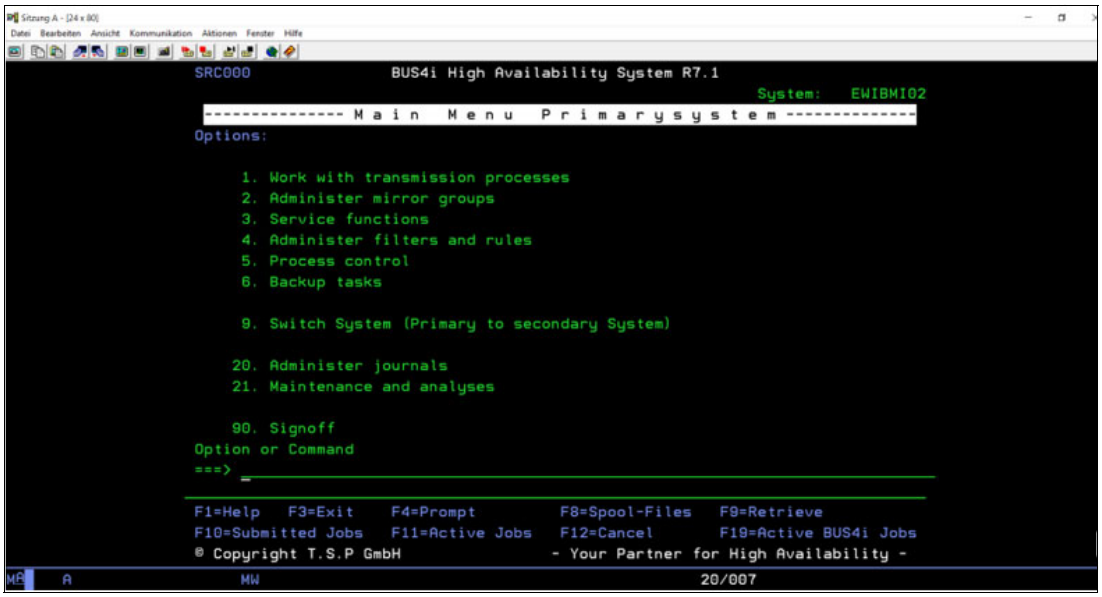

Figure 6-65 BUS4i main menu for Primary System

Normally, BUS4i manages a main menu with the title in white for the primary system environment and another main menu with the title in red for the secondary system environment so that it is easy to identify which environment you are working in (Figure 6-66).

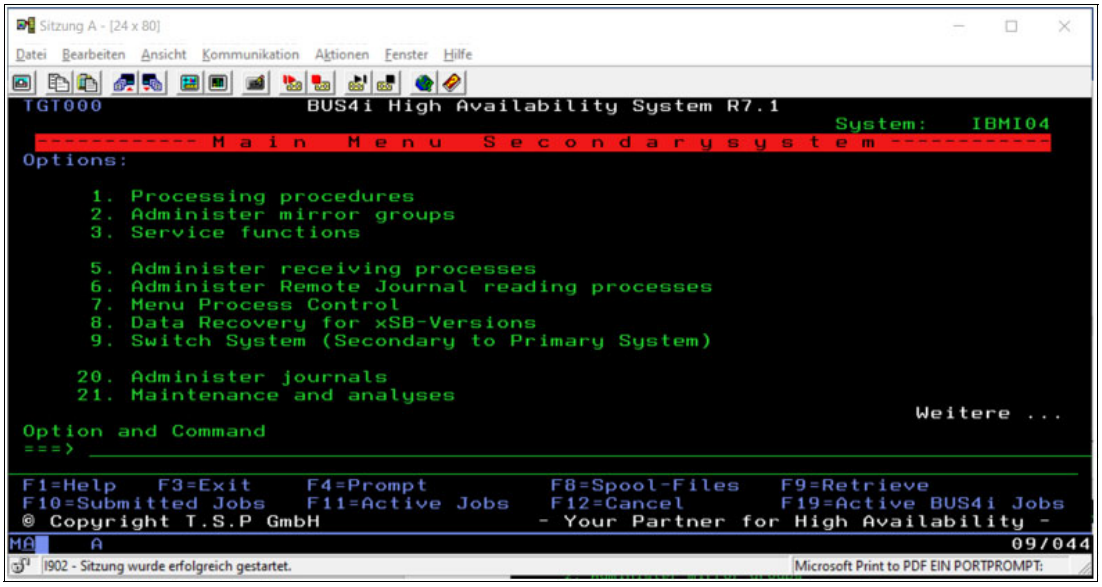

Figure 6-66 BUS4i main menu for Secondary System

BUS4i handles replication between systems by establishing different mirror groups, including the following standard groups:

- P2S_#IFS for the replication of IFS objects
- P2S_#OBJ for non-Database objects
- P2S_#SPL for Spools
- P2S_#SYS for QGPL and QUSRSYS libraries objects
- P2S_#SYS2 for Database catalog objects, such as triggers and constraints

In addition to these standard mirror groups, user mirror groups can be created. For example, the libraries of a specific application can be included, such as the specific mirror groups P2S_ERP and P2S_PAY to replicate the libraries of the Payroll and Human Resources applications.

A mirror group P2S_#NEU is also available that can be used to include it in all the libraries that are created new in the system that were not added to another mirror group.

P2S_ is used to indicate that the mirror group handles the replication from Primary to Secondary for the objects that are contained in it.

In the same way, the S2P_ prefix is used for the corresponding mirror groups in the case of replication from Secondary to Primary, as shown in Figure 6-67.

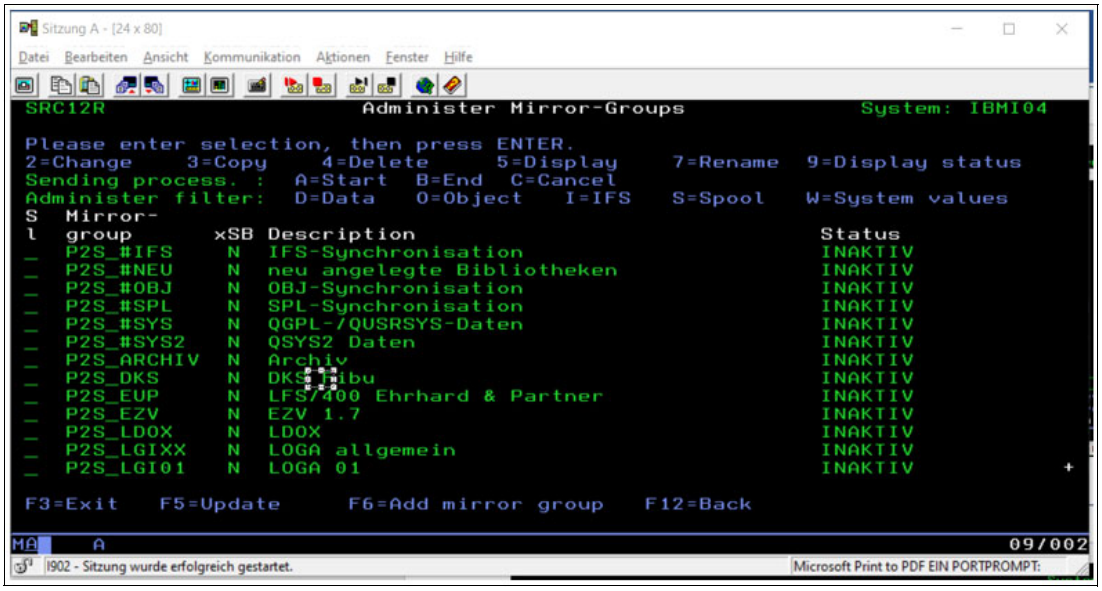

Figure 6-67 BUS4i mirror groups administration

For each configured mirror group, you can customize the different control parameters, such as the Journal and Journal Receiver, that are used to journalize the objects of that group, as shown in Figure 6-68 on page 273.

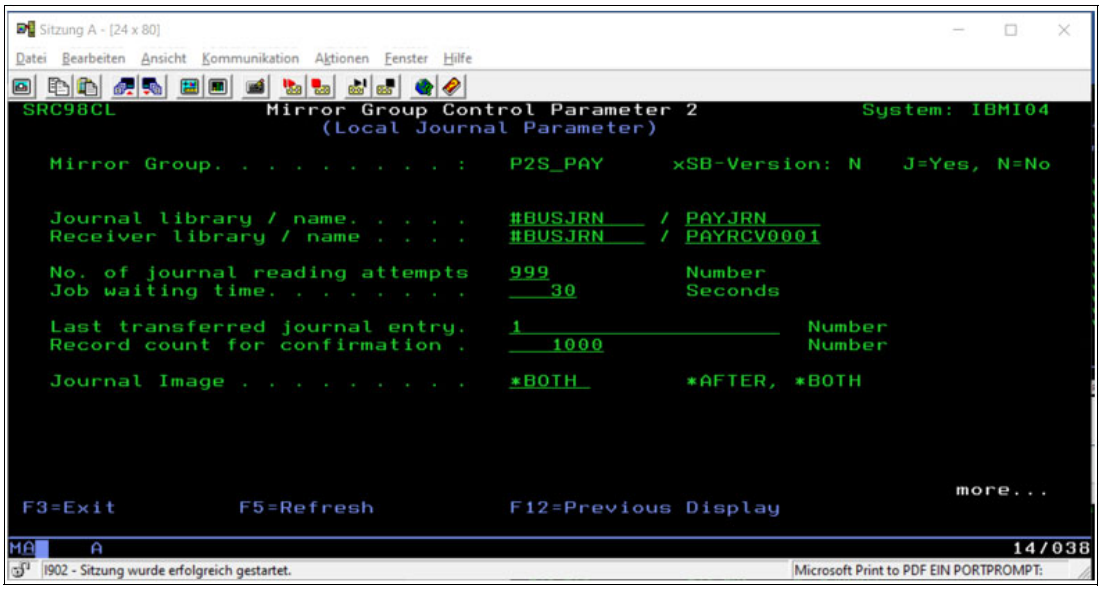

Figure 6-68 BUS4i Mirror Group Control Parameter 2 window

The following details are shown in Figure 6-69:

- IP address of the system that is the target of the replication
- TCP/IP port that is used for the replication of this specific mirror group
- Data compression level
- Version of the operating system on the target system
- Disk occupation threshold on the secondary system from which the process of sending the information stops
- Other parameters that are relevant to the operation of the mirror group

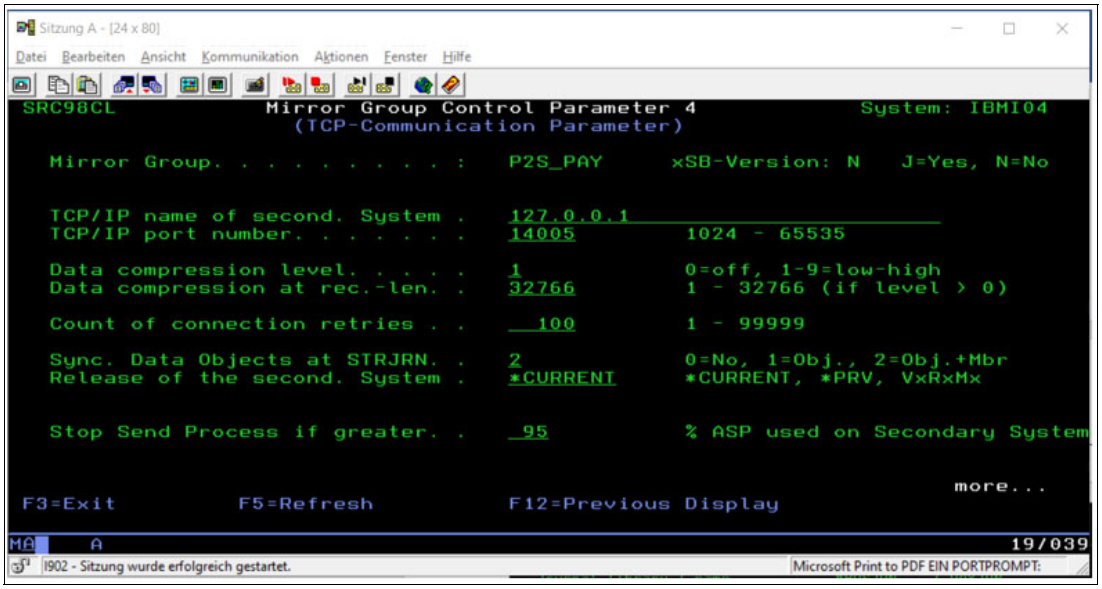

Figure 6-69 BUS4i Mirror Group Control Parameter 4 window

For mirror groups #IFS, #SYS, and QSYS2, the following journals are used (see Figure 6-70 Figure 6-71 and Figure 6-72 on page 275).:

- QAUDJRN system audit journal from the QSYS library
- QAOSDIAJRN from the QUSRSYS
- QSQJRN from QSYS2

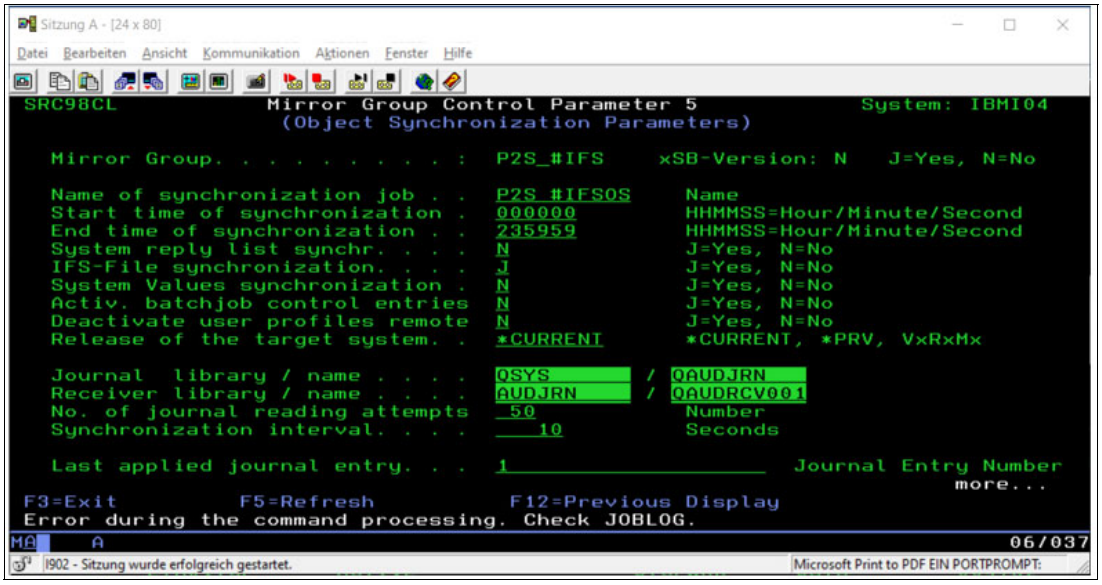

Figure 6-70 BUS4i Mirror Group Control Parameter 5 window

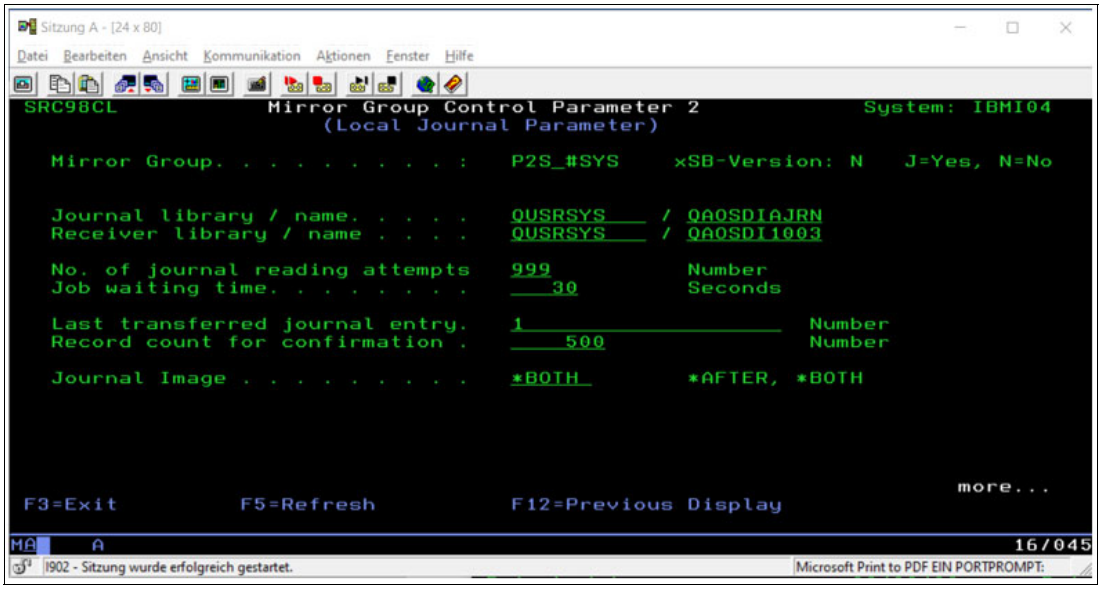

Figure 6-71 BUS4i Mirror Group Control Parameter 7 window

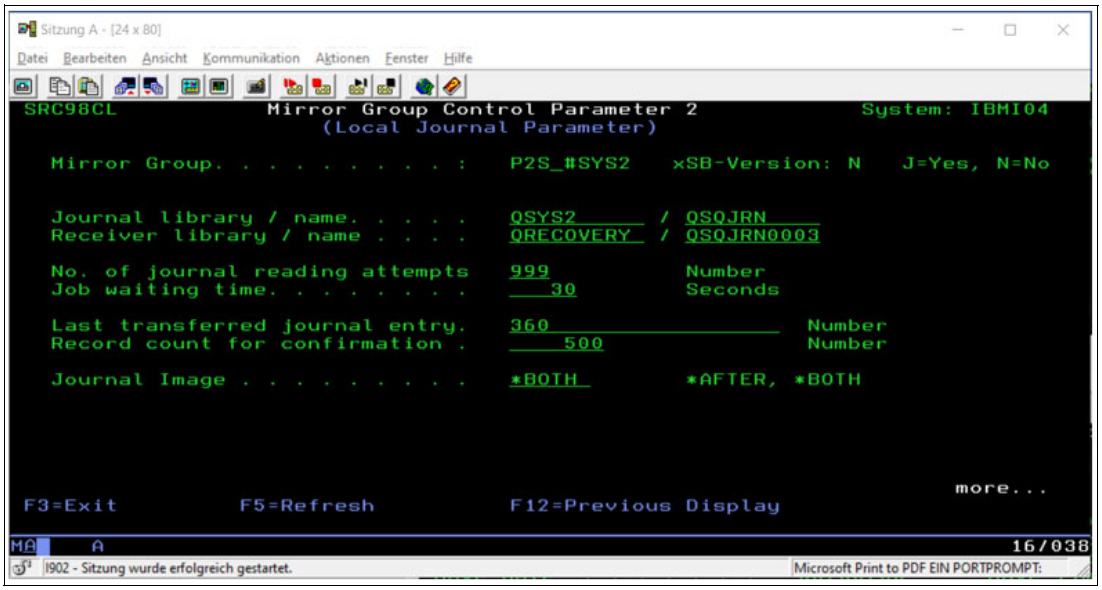

Figure 6-72 BUS4i Mirror Group Control Parameter 2 window

The objects that are included in each of the mirror groups are defined by managing the corresponding filters, as shown in Figure 6-73.

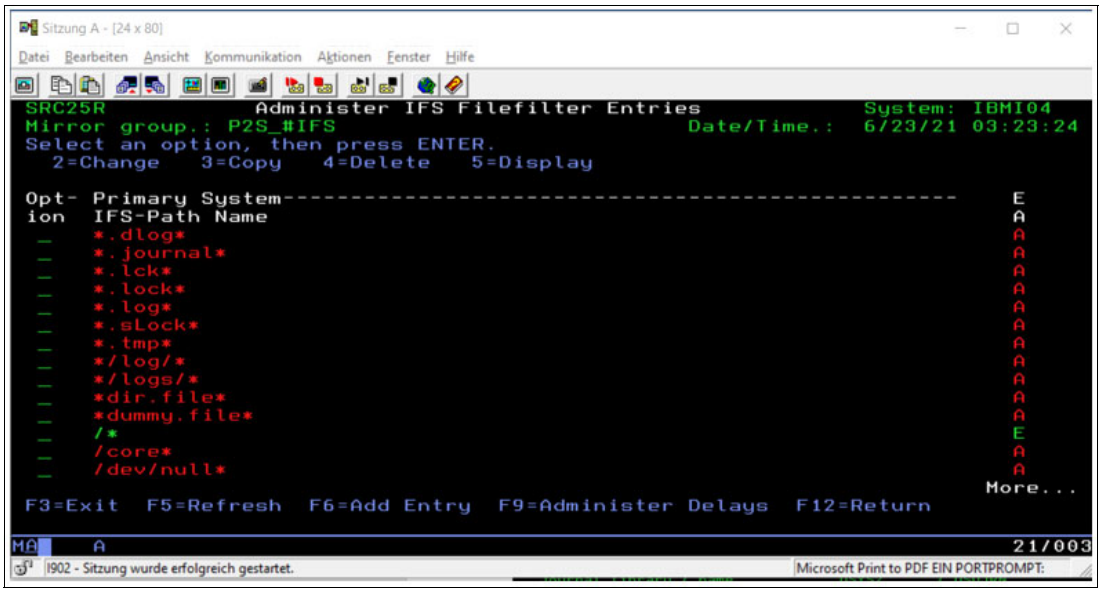

Figure 6-73 BUS4i IFS Filter Entries

Note: In the case of the mirror group P2S_#IFS, which controls the replication of IFS objects from the primary system to the secondary system, you can include a green entry that indicates inclusion for the /* directory. Then, you can include several red entries that indicate exclusion for all of those directories that do not need to be replicated to the secondary system

Similarly, by using the mirror groups P2S_OBJ and P2S_SPL, you can control which objects and spool files are included in or excluded from replication, as shown in Figure 6-74 on page 276 and Figure 6-75 on page 276.

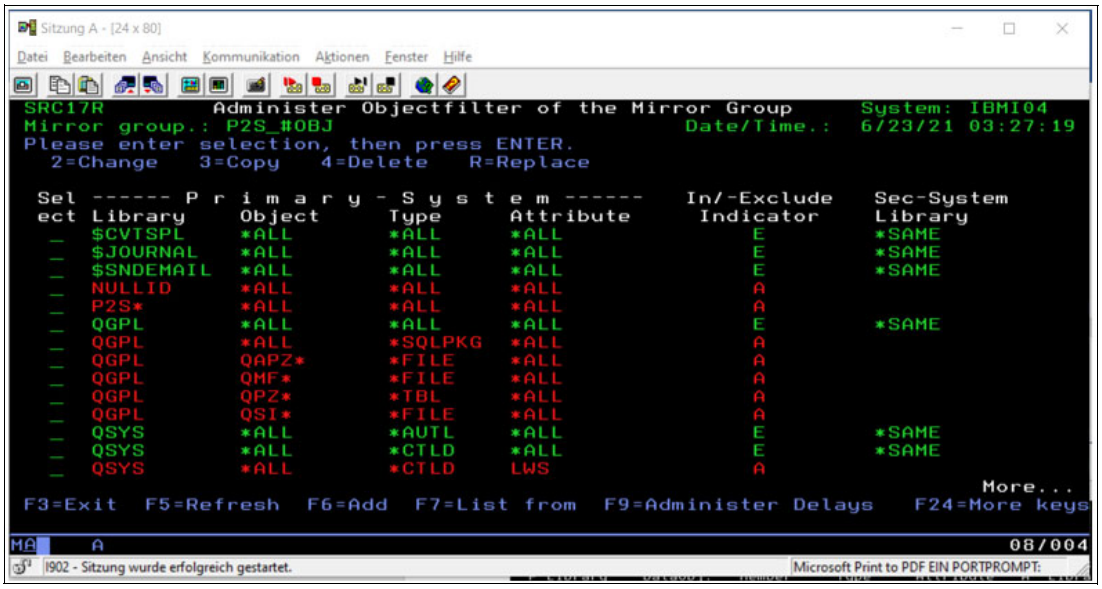

Figure 6-74 BUS4i objects filter entries

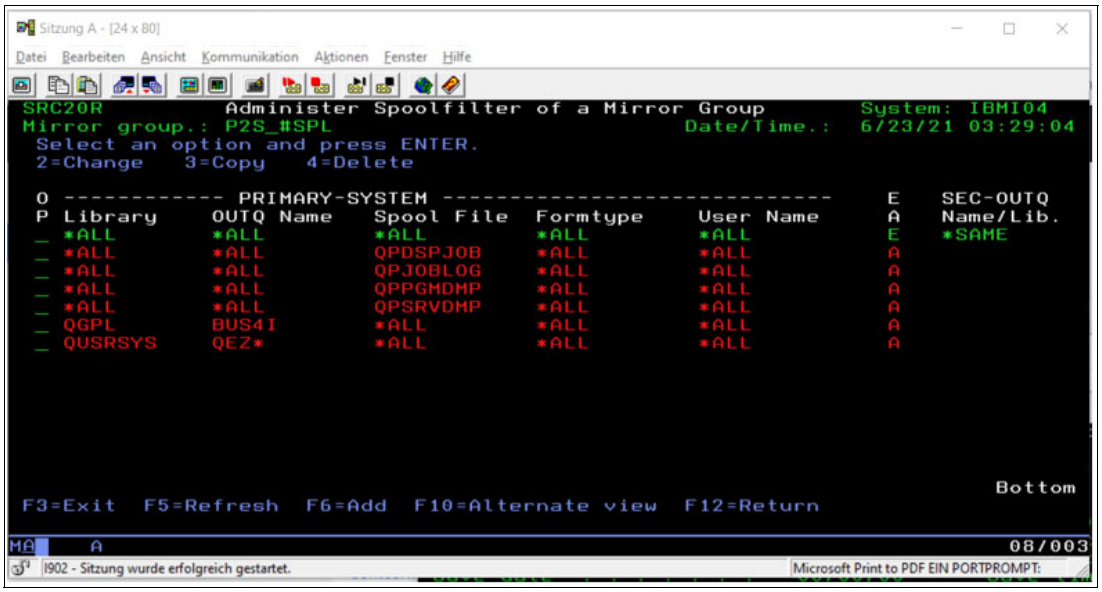

Figure 6-75 BUS4i spool files filter entries

Additionally, by using the mirror groups P2S_#SYS and P2S_#OBJ, you can also control which QGPL and QUSRSYS library objects and system values are included in or excluded from replication (Figure 6-76 and Figure 6-77).

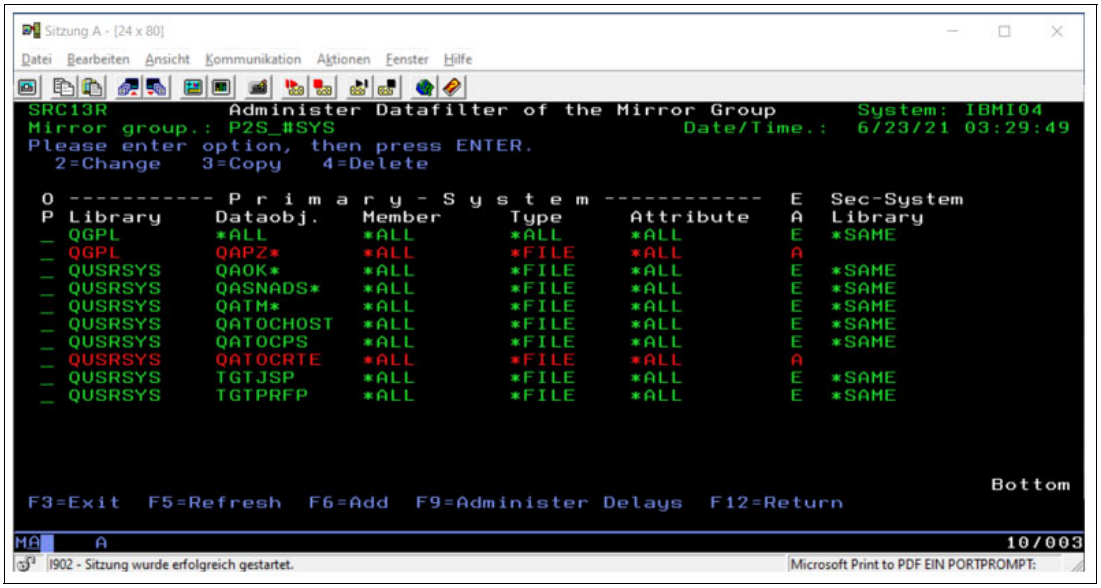

Figure 6-76 BUS4i QGPL and QUSRSYS filter entries

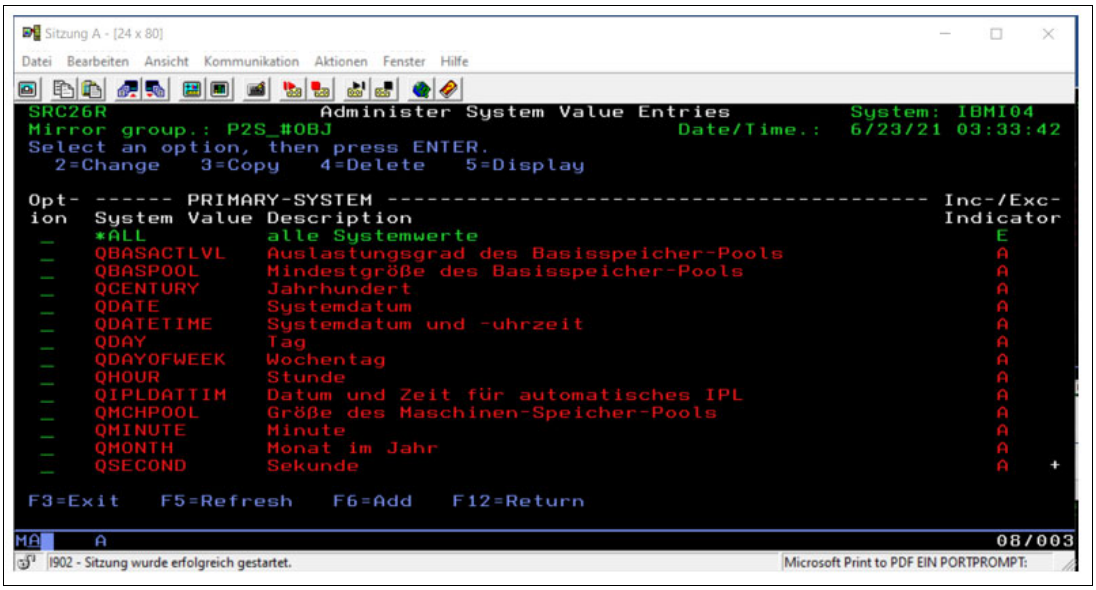

Figure 6-77 BUS4i system value entries

By default, all objects that are in new libraries that were not added to any user mirror group are included in the mirror group P2S_#NEU for replication.

If a specific situation does not require inclusion of some of these libraries or objects to the replication, for example, new libraries with large volumes of redundant data for testing, then you can exclude these libraries or objects by using an entry in red in this same mirror group (Figure 6-78).

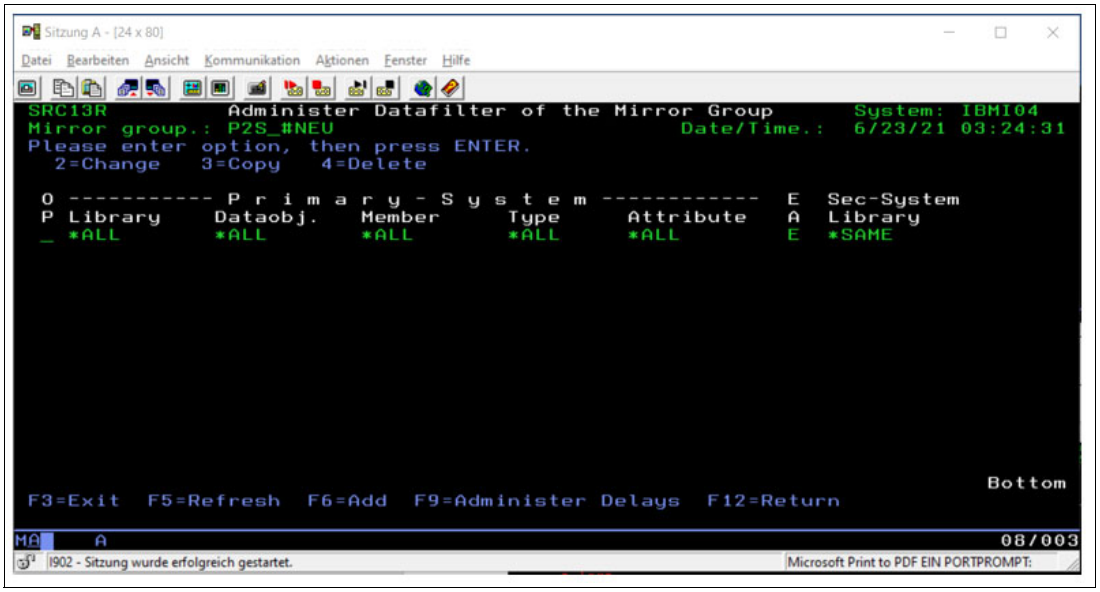

Figure 6-78 BUS4i data filter entries for new libraries

Libraries that belong to a specific application can be assigned to a specific user mirror group for them by individual name or generic name, as shown in Figure 6-79.

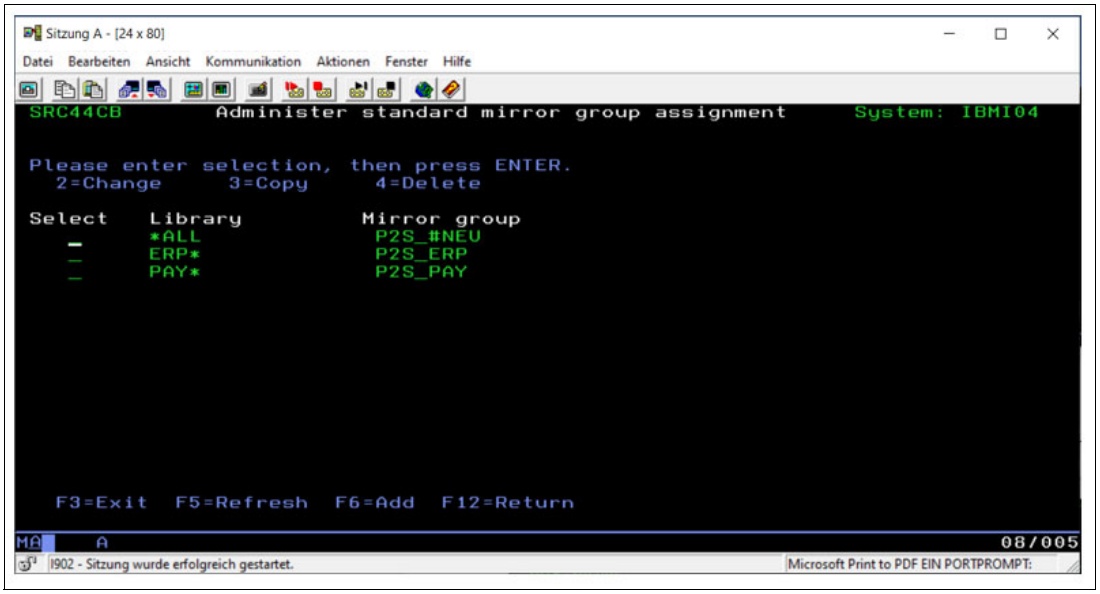

Figure 6-79 BUS4i standard mirror group assignment

By using its Journal Manager (see Figure 6-80), BUS4i can manage the Journal Receivers of all the journals in the system, regardless of whether they are used for replication tasks.

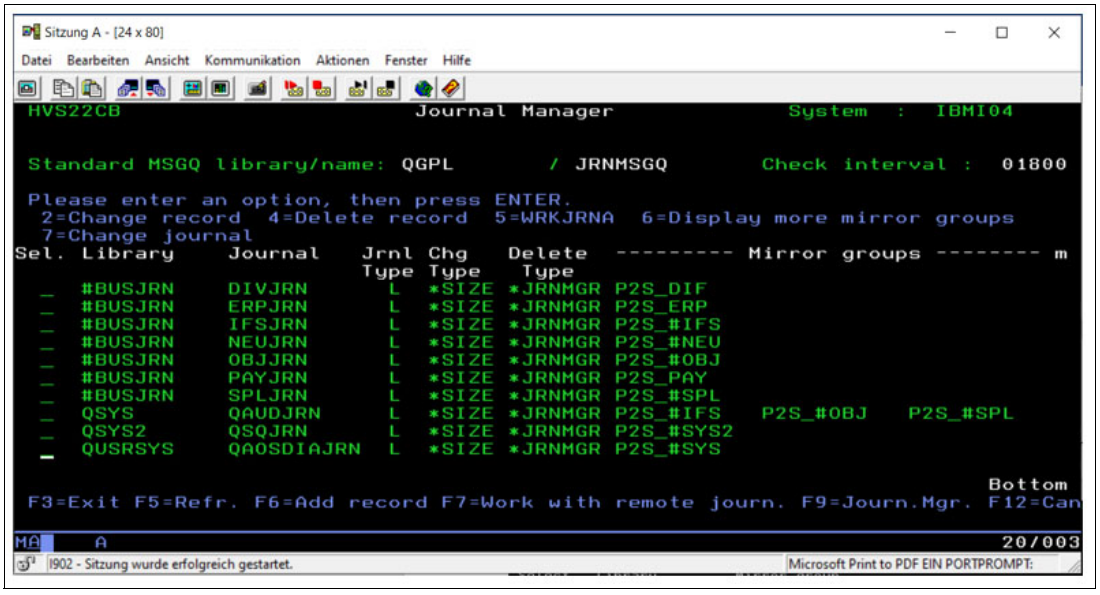

Figure 6-80 BUS4i Journal Manager window

6.3.5 Switching a Bus4i environment between two IBM i virtual machines in the IBM Power Virtual Server Cloud

To document this use case, two IBM i 7.4 virtual machines (VMs) are configured in the IBM Power Virtual Server location `London 04` called `EWIBMI02`, which acts as the primary system and `EWIBMI04`, which acts as the secondary system.

Both IBM i VMs are physically hosted in a Power System S922 and include assigned 0.25 processor units in shared uncapped mode, 8 GB of memory, and a standard disk configuration in their system ASP.

The IBM i VMs are in the same private subnet. The communications between them is done through the ports DDM, REXEC, and 140XX (for the different mirror groups).

Complete the following steps to start the planned switchover:

1. At the command line of the primary system, enter **bus4i** and press Enter to reach the BUS4i main menu for the primary system (Figure 6-81).

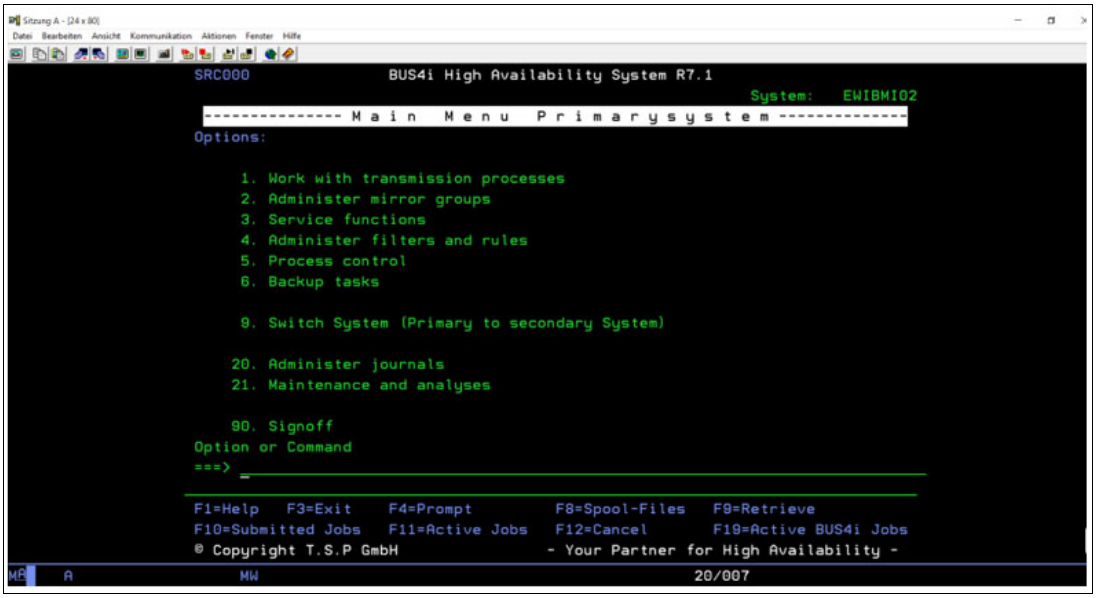

Figure 6-81 BUS4i main menu of the primary system

2. Select **Option 1. Work with transmission processes** to check that all of the mirror groups are defined, in `ACTIVE` status, and no latency exists on them, as shown in Figure 6-82.

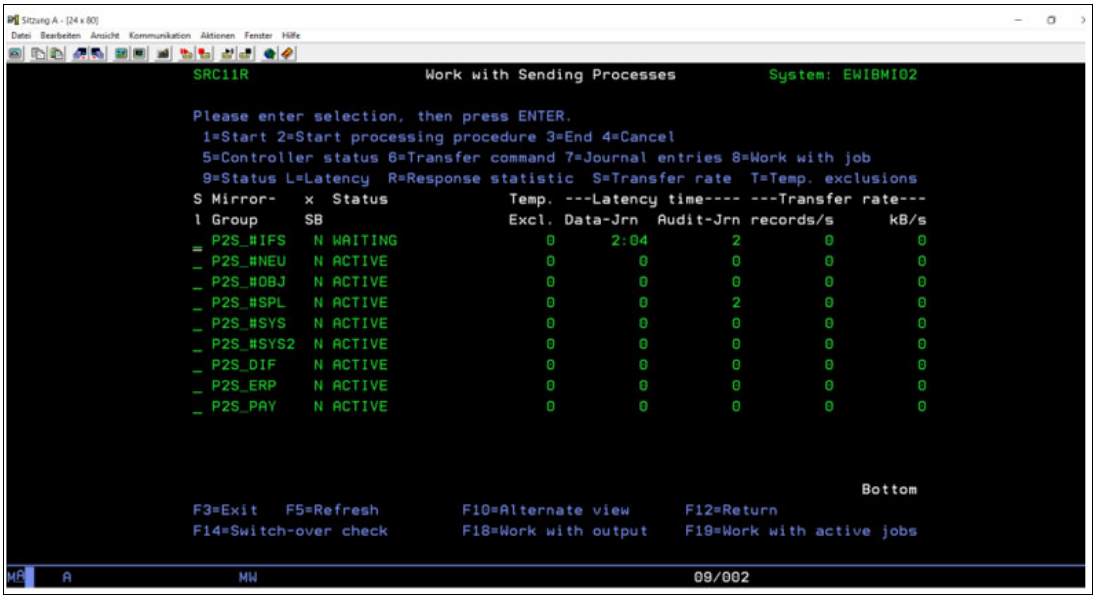

Figure 6-82 BUS4i Working with sending process

In the main menu of the primary system, Option 2. Administer Mirror Groups can also be used.

3. Check the transfer rates for each mirror group by using **Option S. Transfer Rate** (Figure 6-83).

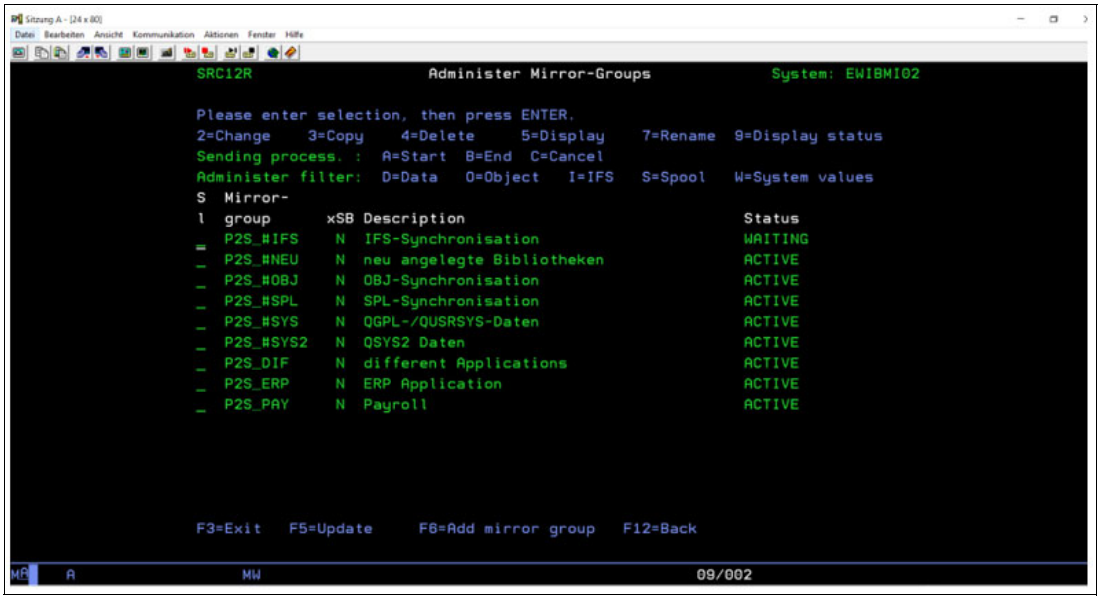

Figure 6-83 BUS4i Administering mirror groups

4. In the main menu of the primary system, select **Option 3. Service Functions** to check the status of the compared data for each mirror group.

   Verify that all of the mirror groups stay in `ACTIVE` status and without any errors or warnings reported as shown in Figure 6-84.

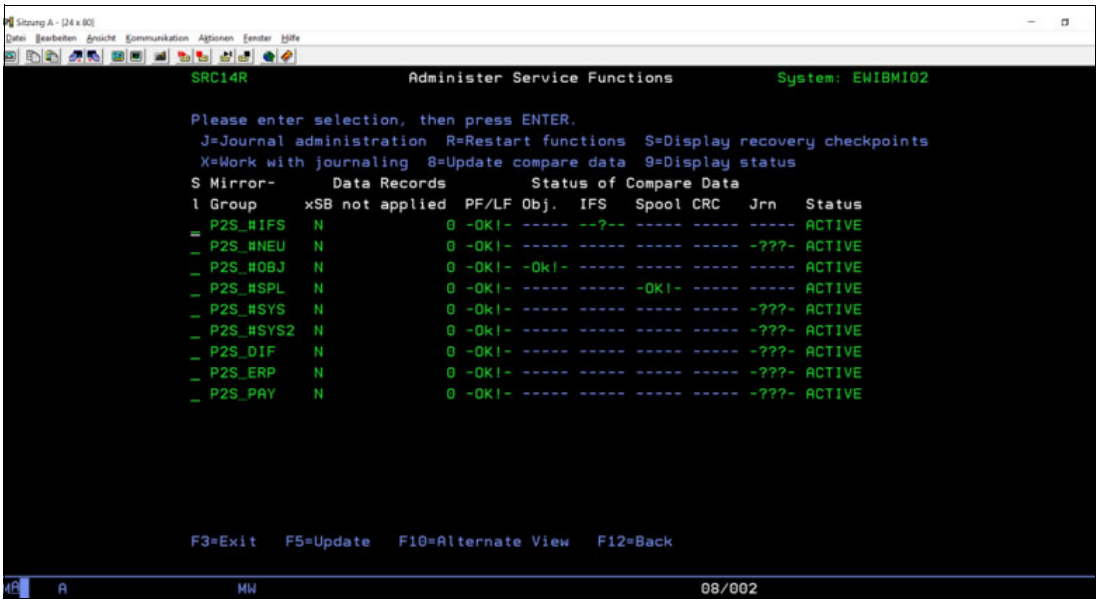

Figure 6-84 BUS4i administering service functions

If errors or warnings exist for some mirror group, they must be investigated and fixed before switchover operation is performed. Errors are displayed in red and warnings are displayed in yellow, as shown in Figure 6-85.

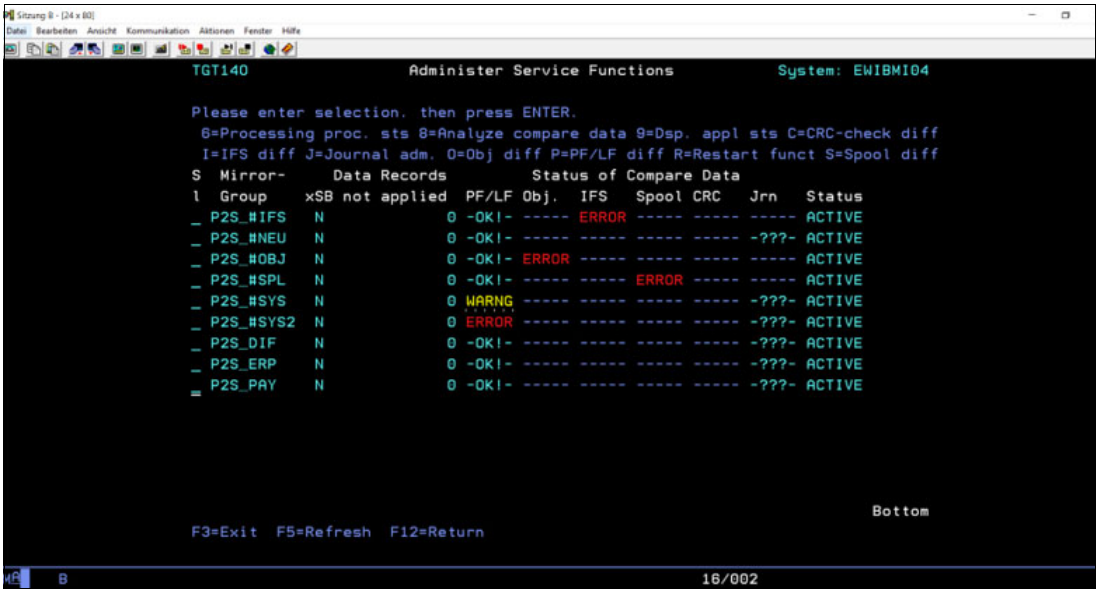

Figure 6-85 BUS4i administering service functions with examples of errors and warnings

5. All of the tasks that are involved in the switchover of the primary system environment are specified in the switchover command. This information can be checked from the main menu of the primary system by selecting **Option 9. Switch System (Primary to Secondary System)** and then selecting **Option 3. Administer Commands** and use the interface to make any necessary adjustments to the switchover (Figure 6-86).

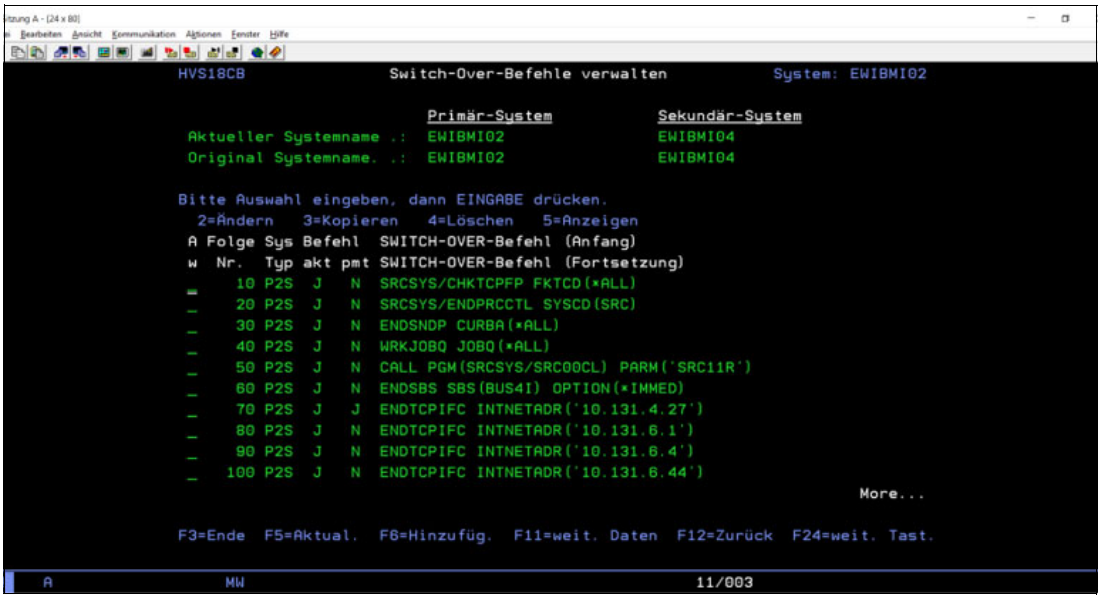

Figure 6-86 BUS4i switchover command

6. If everything is ready, start the planned switchover process from the main menu of the primary system (the system with the title in white), by selecting **Option 9. Switch System (Primary to Secondary System)**.

7. In the switchover menu, select **Option 20, Start Primary to Secondary System**, as shown in Figure 6-87.

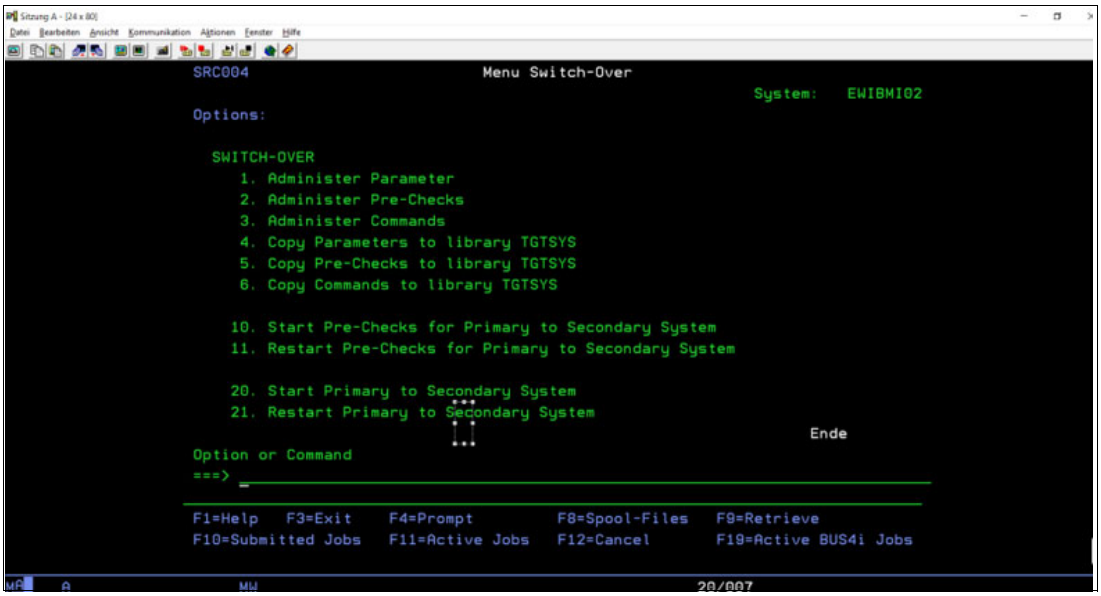

Figure 6-87 Switchover menu

A switchover message appears that indicates you are switching the `PRIMARY` system to the `SECONDARY` system. Answer with **Y=Yes** and press Enter to continue, as shown in Figure 6-88.

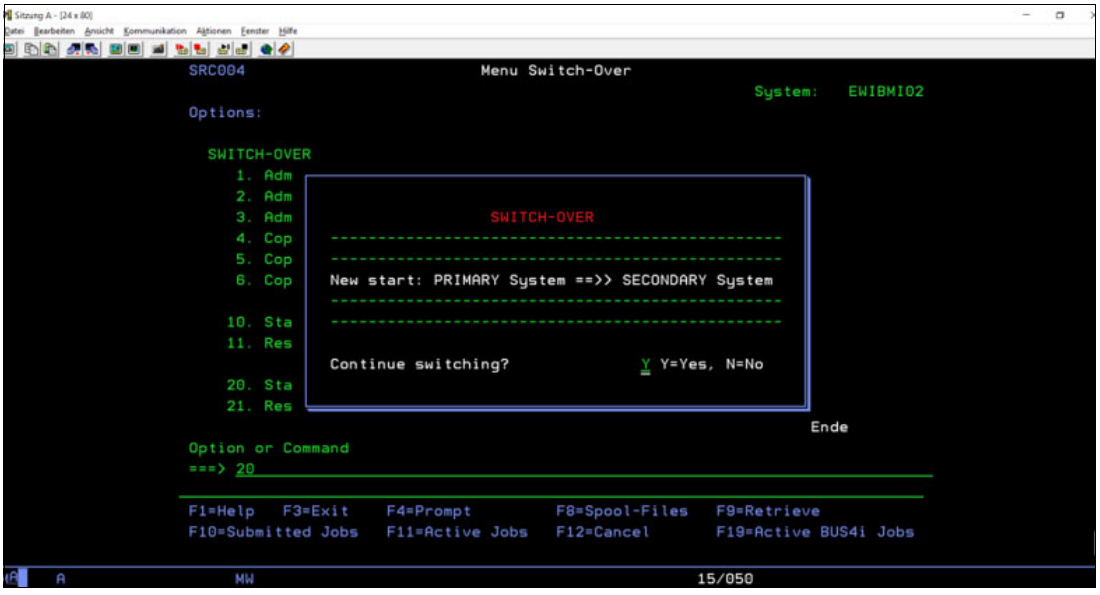

Figure 6-88 BUS4i switchover confirmation message in primary system

8. A warning message appears that indicates a potential problem with duplicated IP addresses. This message is normal and is not an issue. Therefore, answer **Y=Yes** to continue switching as shown in Figure 6-89.

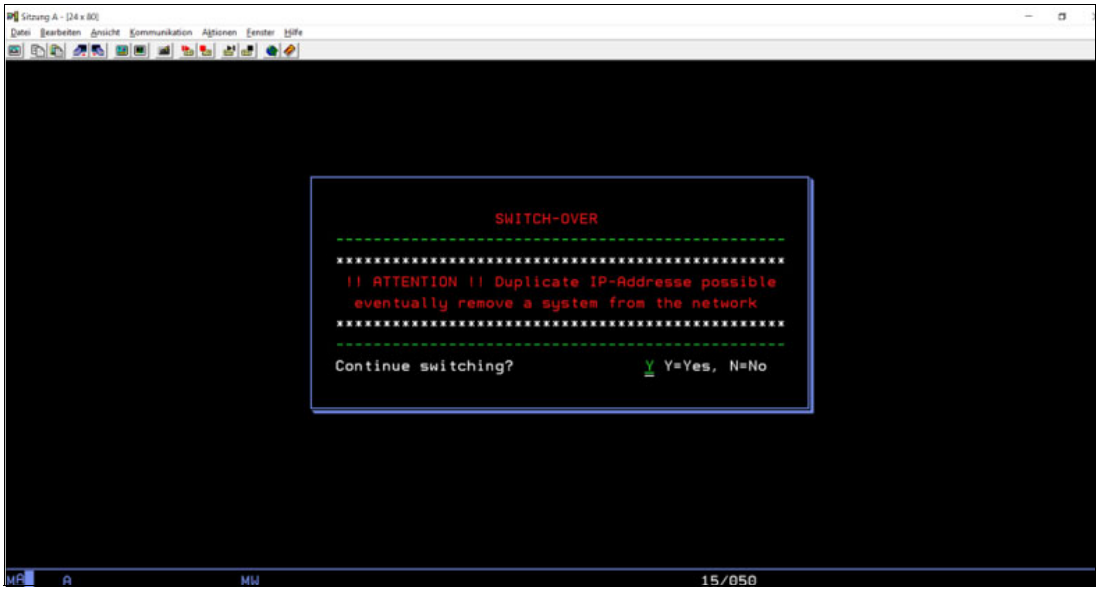

Figure 6-89 BUS4i switchover confirmation message in primary system

9. The Work with All Job Queues window is displayed, as shown in Figure 6-90. Check that all of the application jobs are completed and no active user jobs exist in the system. Press **F3=Exit** to continue with the switchover process.

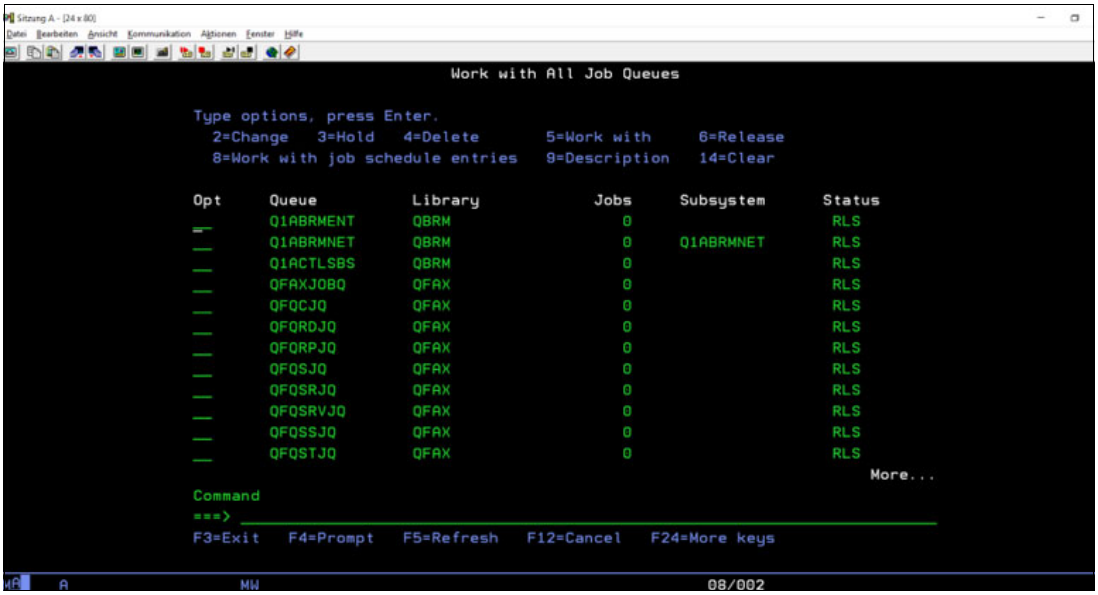

Figure 6-90 BUS4i switchover Work with All Job Queues window

10. The Work with Sending Processes window is displayed, as shown in Figure 6-91. Check that no latencies or temporary excludes appear.

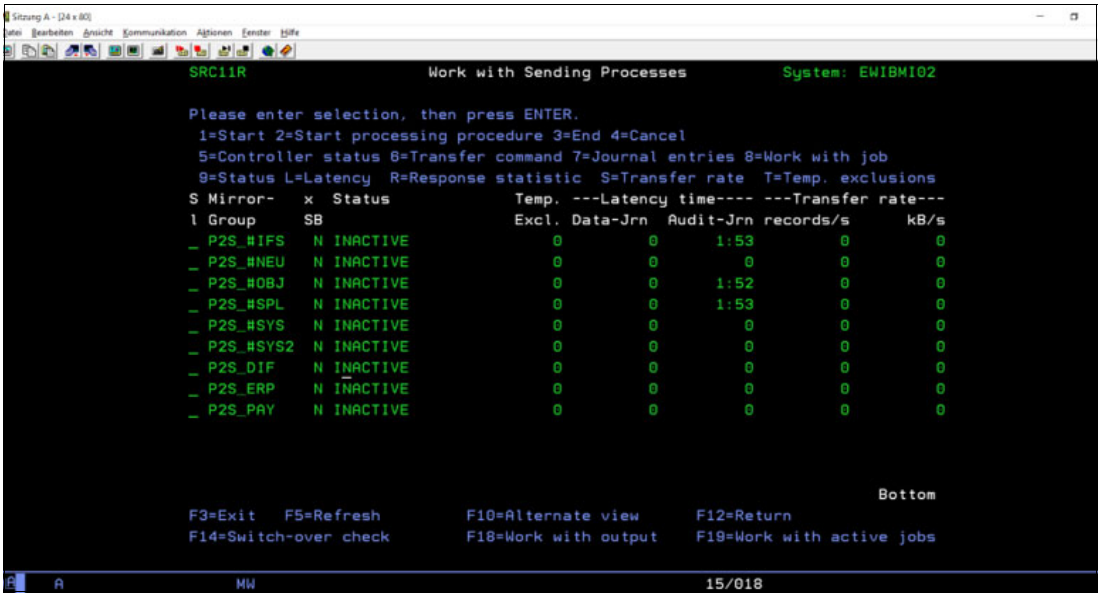

Figure 6-91 BUS4i switchover Work with Sending Processes window

After the switchover process is completed in the primary system, complete the following steps to revert the switchover from the secondary to the primary system:

1. From the command line of the secondary system, enter **`bus4i`** and press Enter to access the main menu of the secondary system, the system with the title in red.
2. Select **Option 9. Switch System (Secondary to Primary System)**.
3. In the switchover menu, select **Option 20, Start Secondary to Primary System** and press Enter and proceed in the same way as you did for the primary system (see Figure 6-92 and Figure 6-93 on page 286.

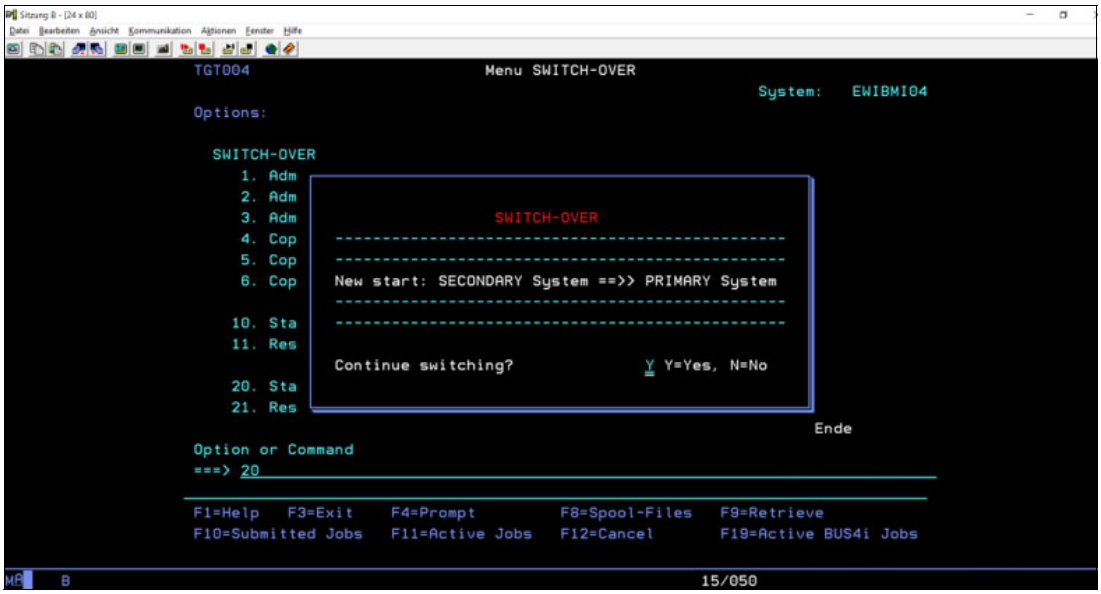

Figure 6-92 BUS4i switchover confirmation message in secondary system

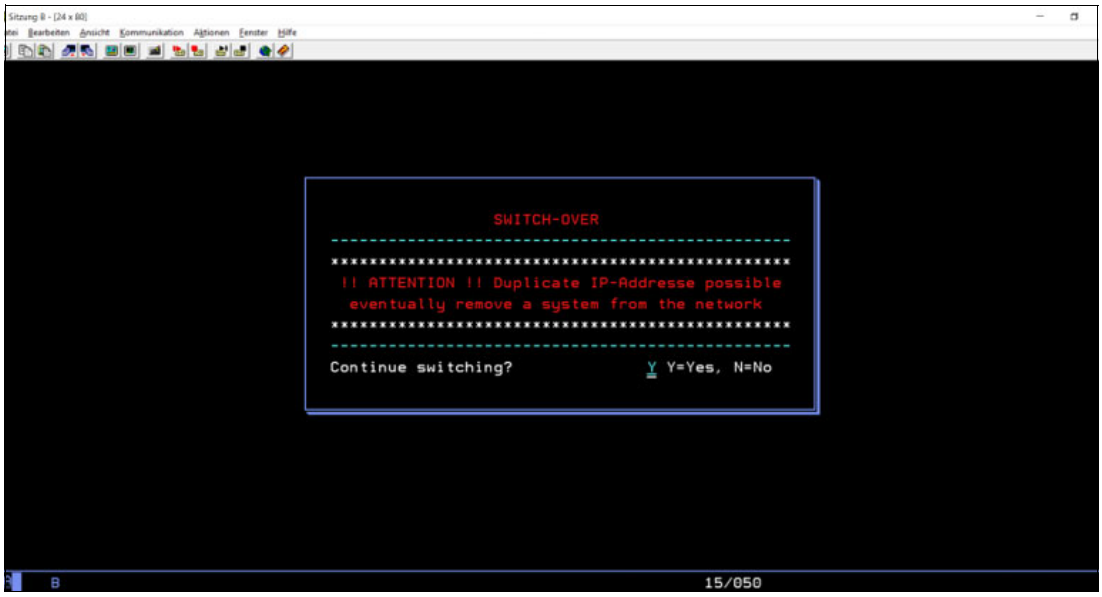

Figure 6-93 BUS4i switchover confirmation message in secondary system

4. After the switchover process is complete on the secondary system, verify the new roles of each system by entering **bus4i** in the command line of each system.

   As shown in Figure 6-94, the preferred target system `EWIBMIO4` shows the main menu for the primary system, which has a title that is printed in white.

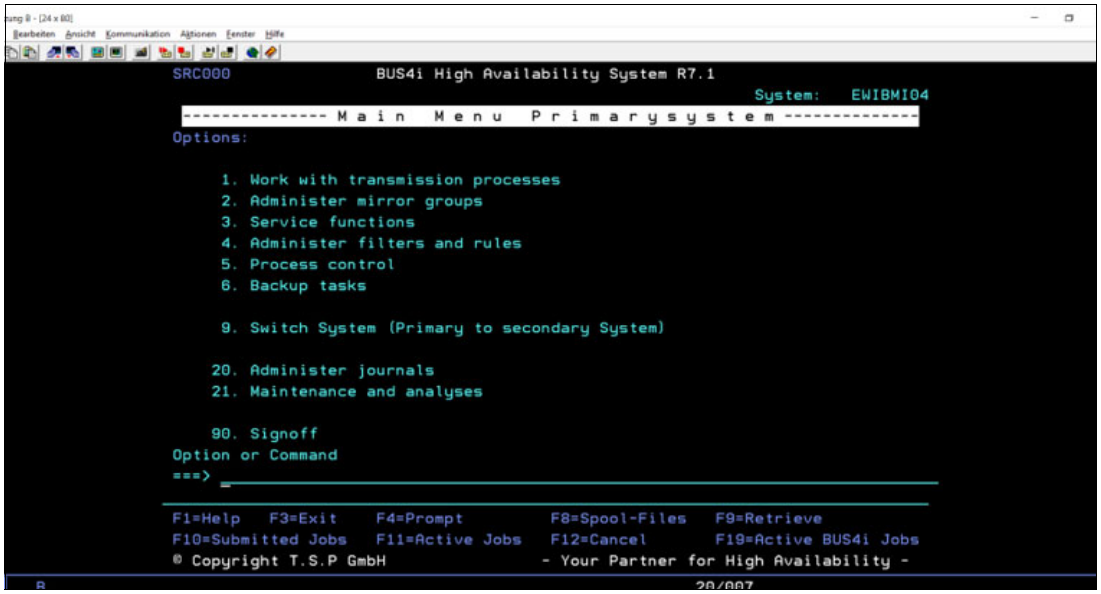

Figure 6-94 BUS4i new primary system after switchover

The preferred source system EWIBMI02 shows the main menu for secondary system, which has a title that is printed in red as shown in Figure 6-95.

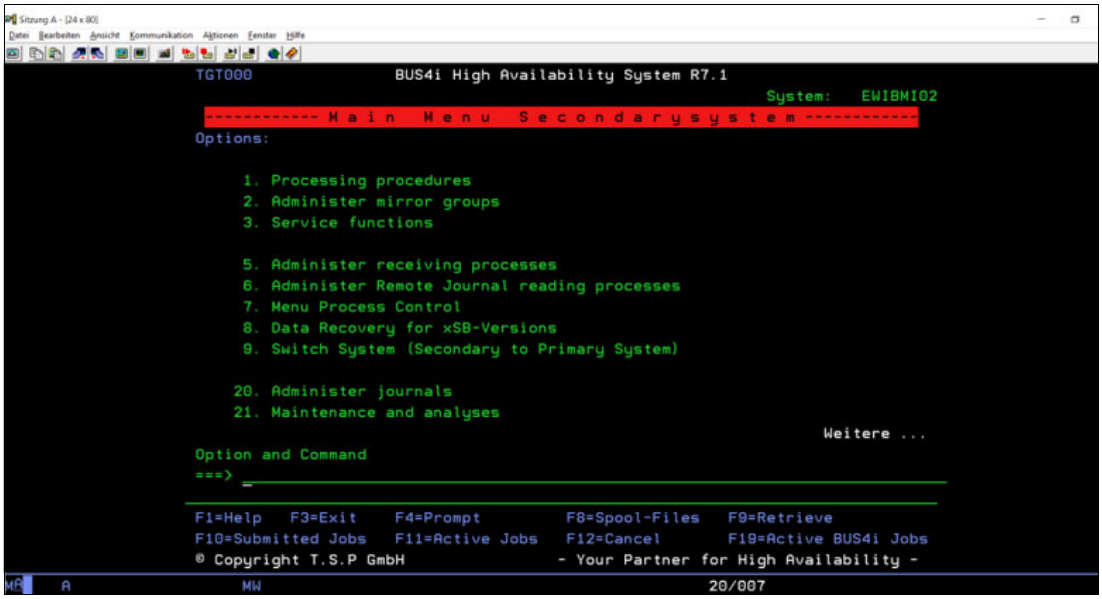

Figure 6-95 BUS4i new secondary system after switchover

From the switchover menu, you can use **Option 1. Administer Parameters** to verify whether the roles for the primary and secondary systems were switched, as shown in Figure 6-96 and Figure 6-97 on page 288.

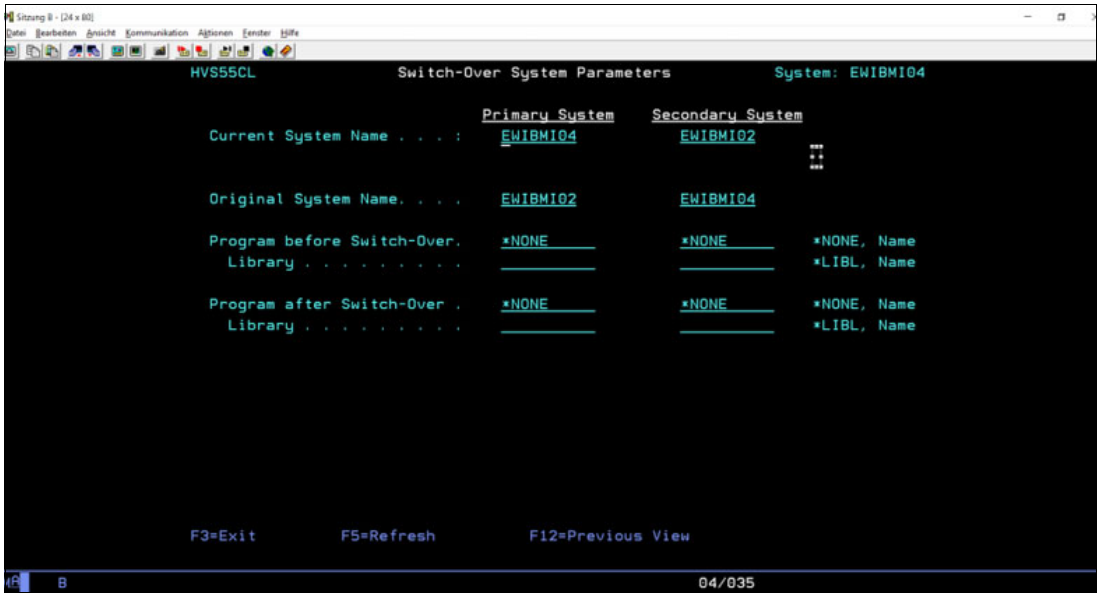

Figure 6-96 BUS4i switchover system parameters on new primary system after switchover

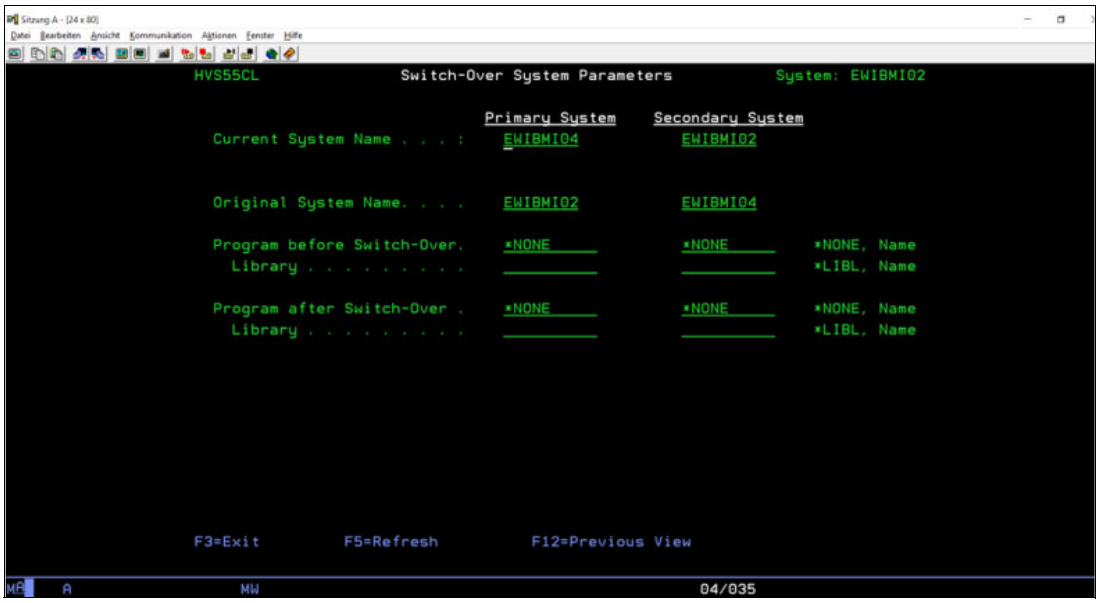

Figure 6-97 BUS4i switchover system parameters on new secondary system after switchover

5. Verify that the switchover process was run correctly in the primary and secondary system environments, from the command line of the new primary system (`EWIBMI04`, which is the former secondary system).
6. Start the replication process from the new primary system to the new secondary system by running the **`bus4i *yes`** command (Figure 6-98).

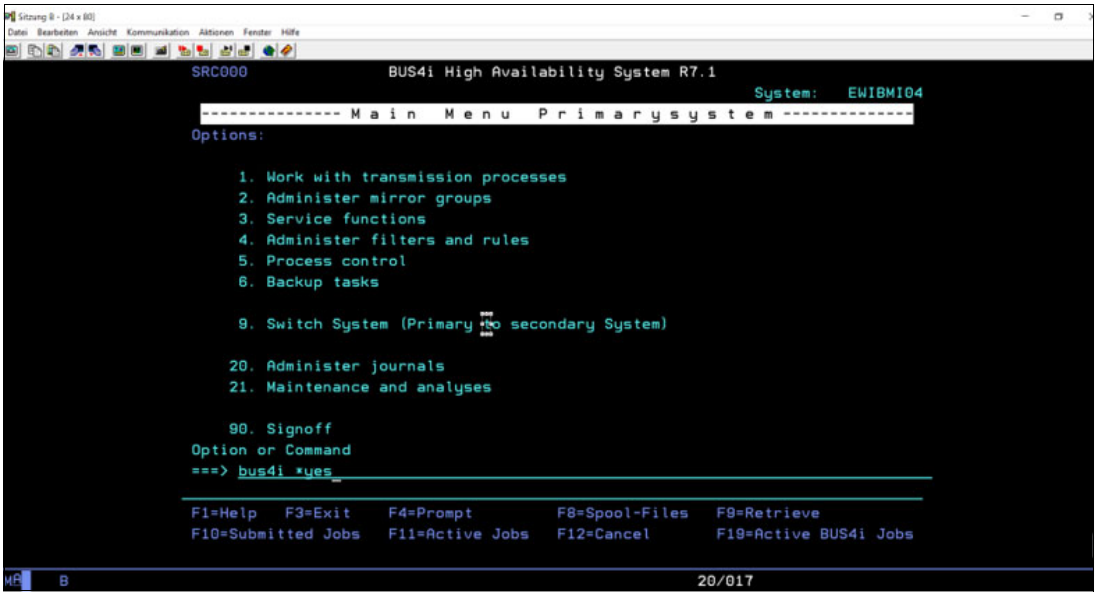

Figure 6-98 BUS4i starting the replication processes

The BUS4I subsystem and all jobs that are related to the replication system tasks must be started on the primary and secondary system, as shown in Figure 6-99 and Figure 6-100.

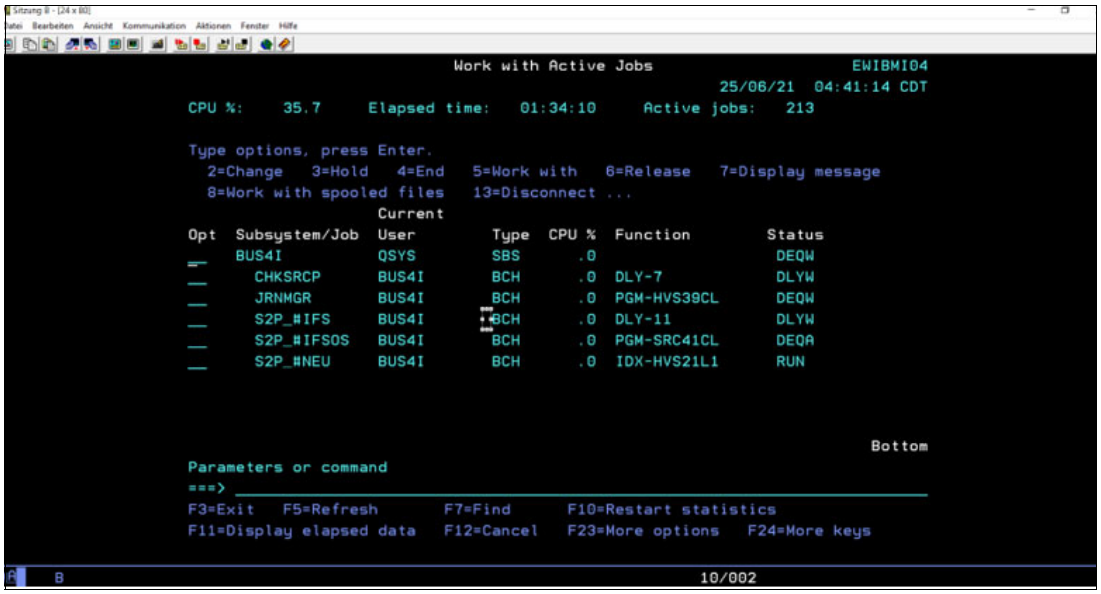

Figure 6-99 BUS4i subsystem in the new primary system

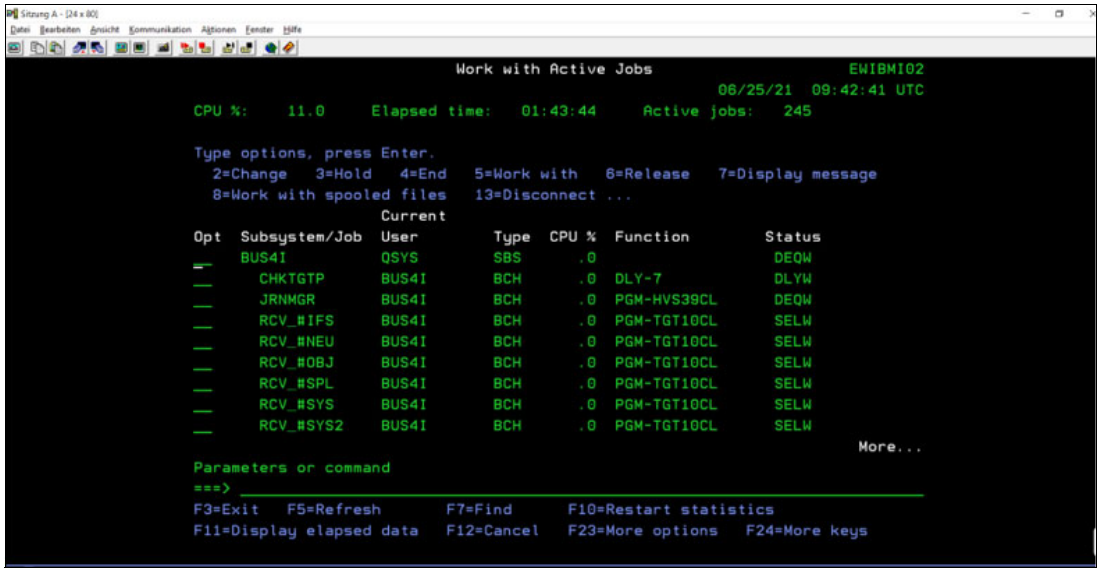

Figure 6-100 BUS4i subsystem in the secondary system

7. In the main menu of the primary system, use **Option 2. Administer mirror groups** to check the status of the different mirror groups, as shown in Figure 6-101.

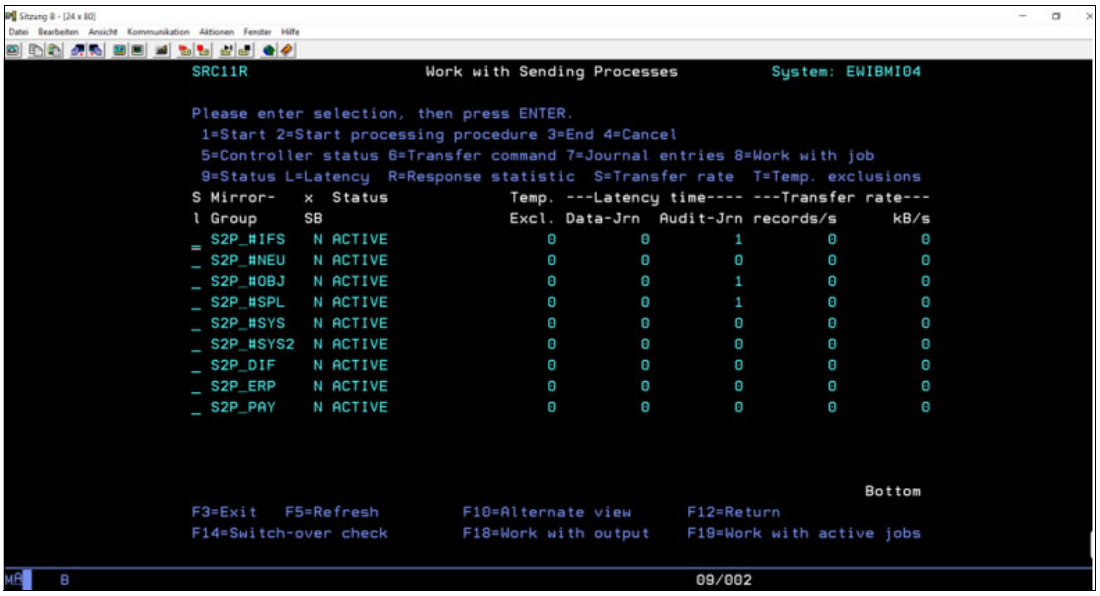

Figure 6-101 BUS4i Work with Sending Process

If you wait approximately 15 minutes, you can use **Option S. Transfer Rate** to see the transfer rates for each mirror group, as shown in Figure 6-102.

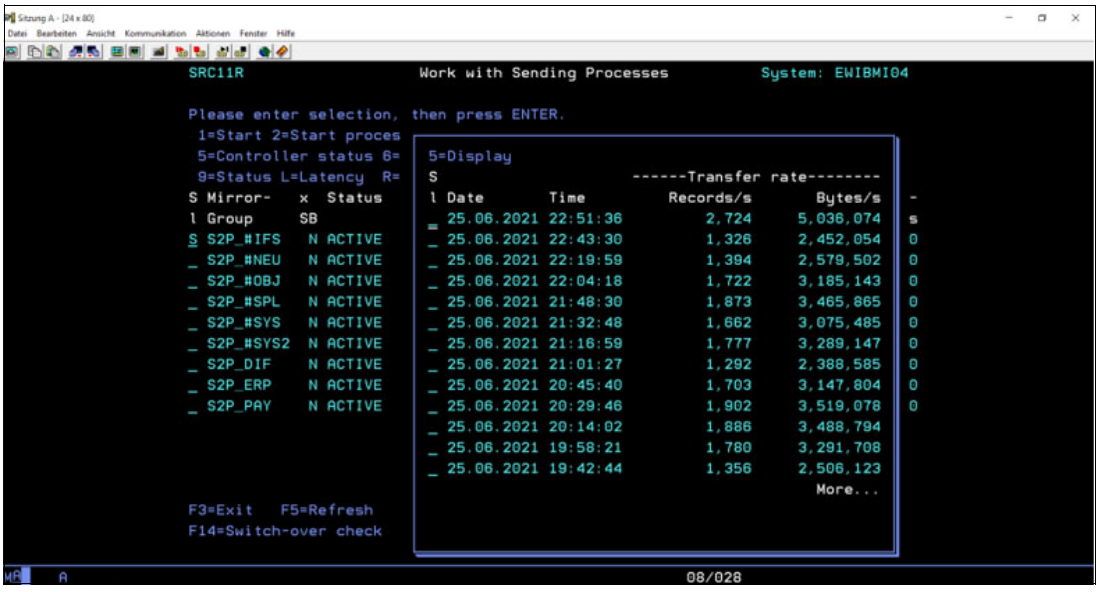

Figure 6-102 BUS4i mirror groups Transfer Rate

Important: If an unplanned switchover or failover is required, the same procedure is applied, but only the steps that correspond to the secondary system are run.

6.3.6 BUS4i ISB Disaster Recovery

A replication environment implemented with Bus4i HA can be combined with Bus4i ISB to, among other things, protect the environment from accidental or intentional deletions of information or from ransomware attacks on the IFS, allowing fast recovery, at any time.

In normal operation, BUSi4 ISB replicates the IBM i system to the ISB NAS device in real time, that device can be located on campus or elsewhere, as shown in Figure 6-103.

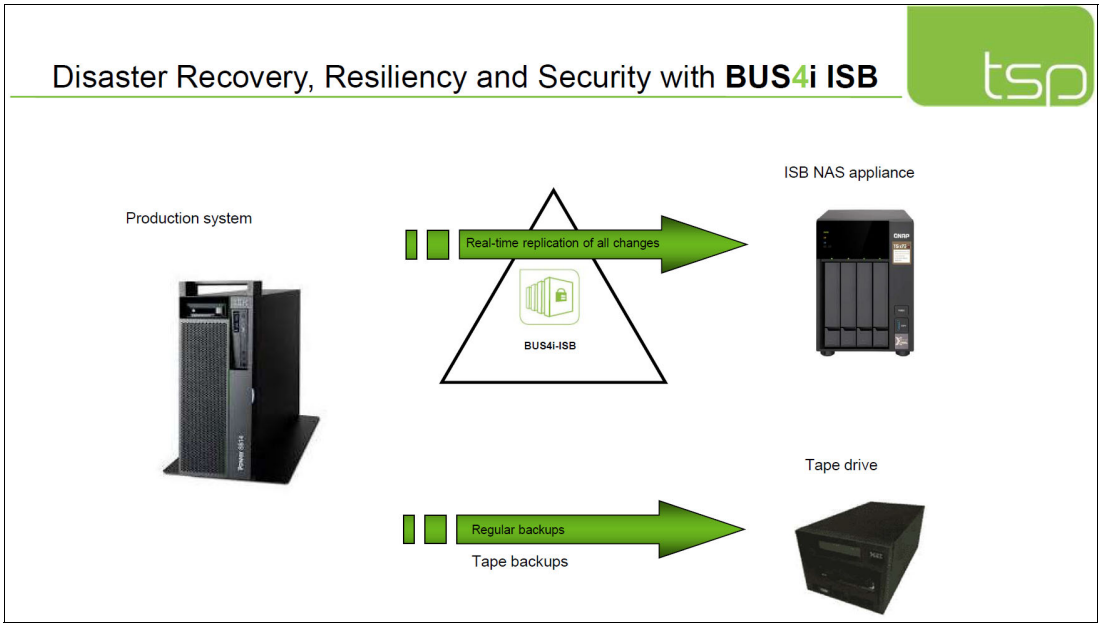

Figure 6-103 Diaster Recovery, Resiliency, and Security with BUS4i ISB

Recovery procedures with daily backups to tape at the end of each day Monday–Friday can result in data loss of 24 hours or more, depending on the case. For example, a crash might occur on Monday midday. If the last backup was on the end of the day on Friday, then all data from Friday until midday Monday is unrecoverable.

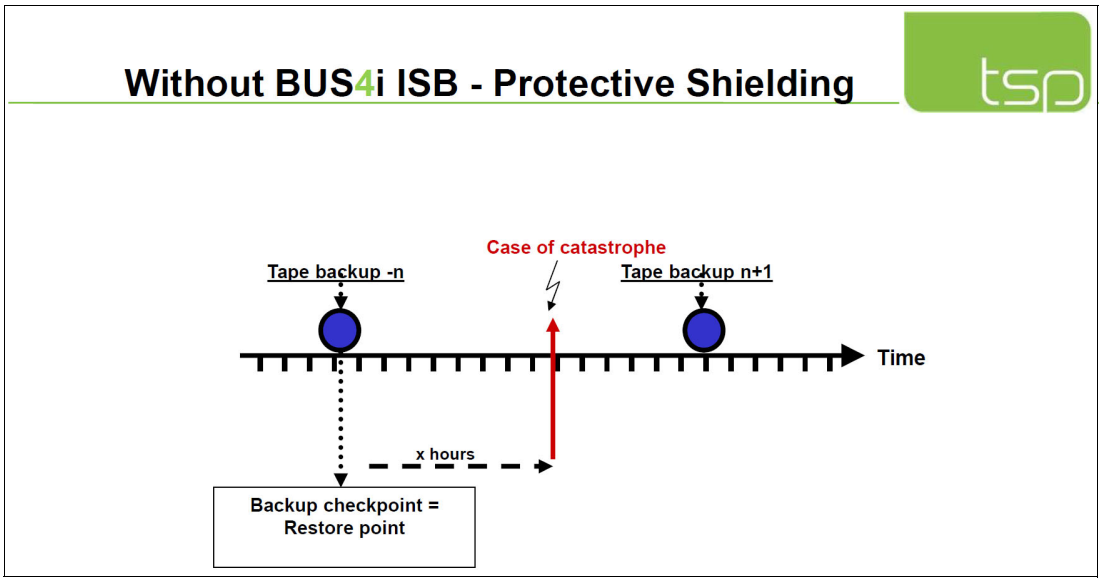

Figure 6-104 Protective Shielding without tool

The restore procedure with daily tape backup and BUS4i ISB protects against data loss, object changes, and printouts (spool files).

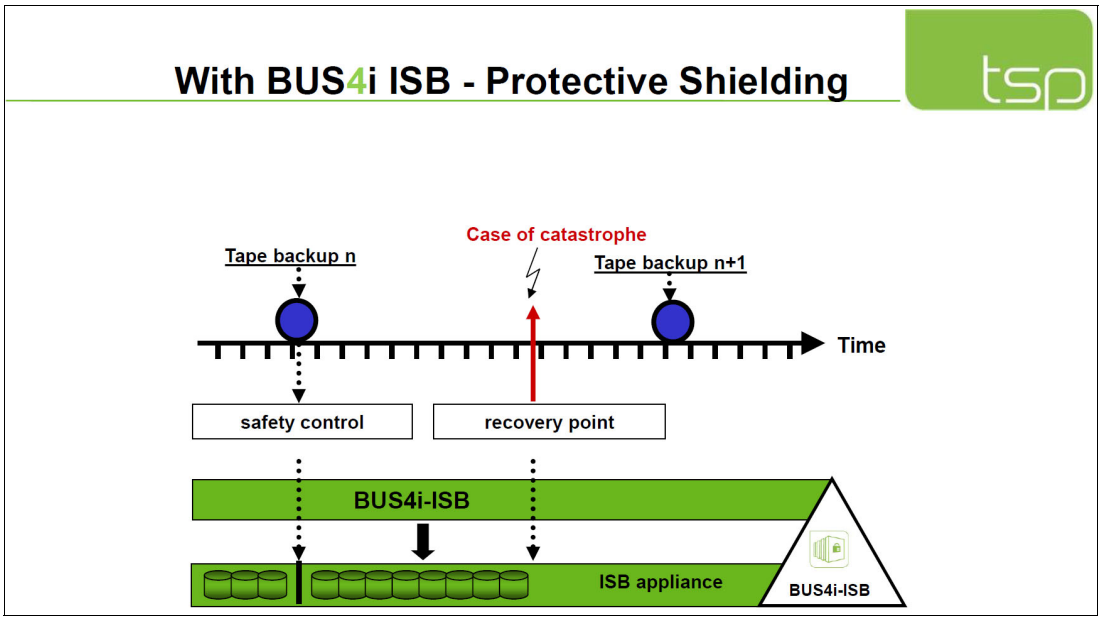

Figure 6-105 Recovery with BUS4i ISB

BUS4i ISB protection includes the following characteristics:

- Permanent backup of data and objects
- Efficient backup of IFS and DLO files
- Backup of user profile and authorization changes
- Immediate backup of printouts (spool files)
- Support for triggers, integrity conditions, and constraints
- Integrated process monitoring with SMS or email message
- Periodic storage of changes in the device
- Integrated safeguarding tasks
- Independence of corporate applications
- Simple and selective retrieval functions

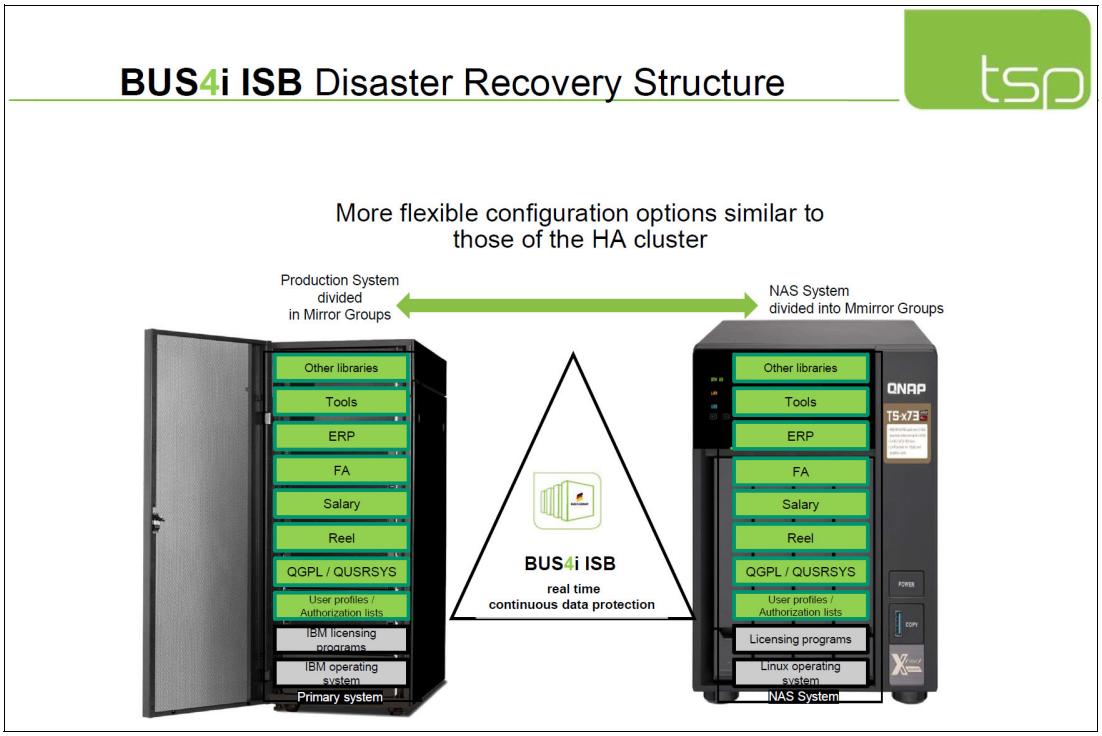

Figure 6-106 Diaster Recovery Structure

The backup process is explained in Figure 6-107.

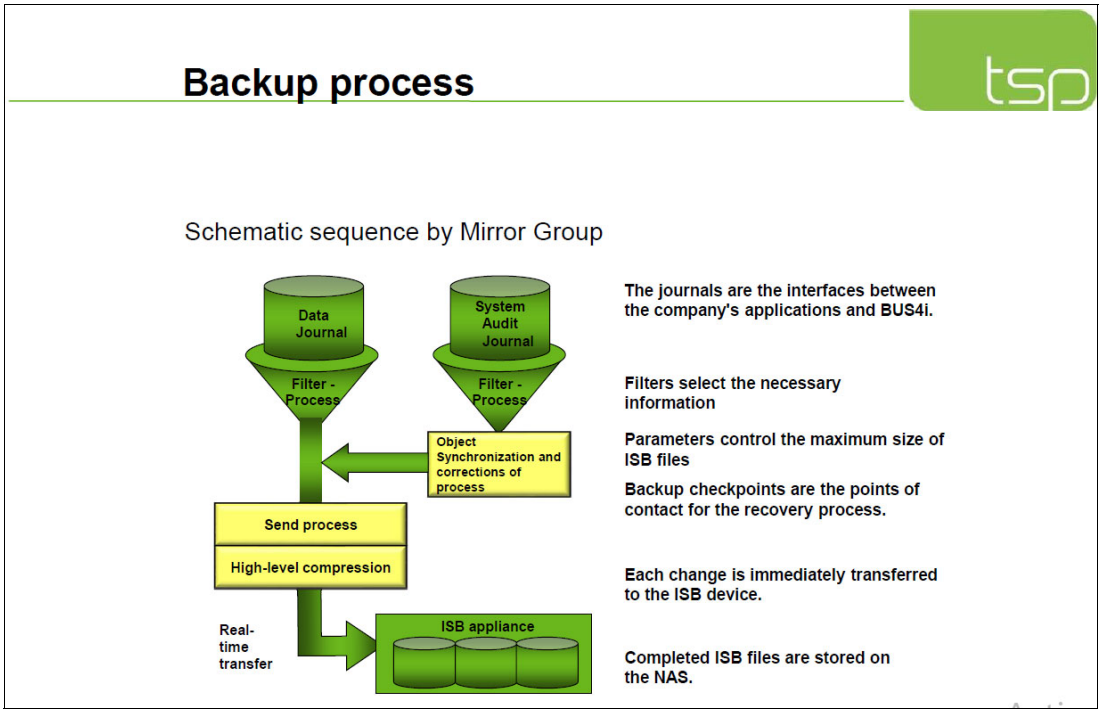

Figure 6-107 Backup process

The recovery process is explained in Figure 6-108 on page 294.

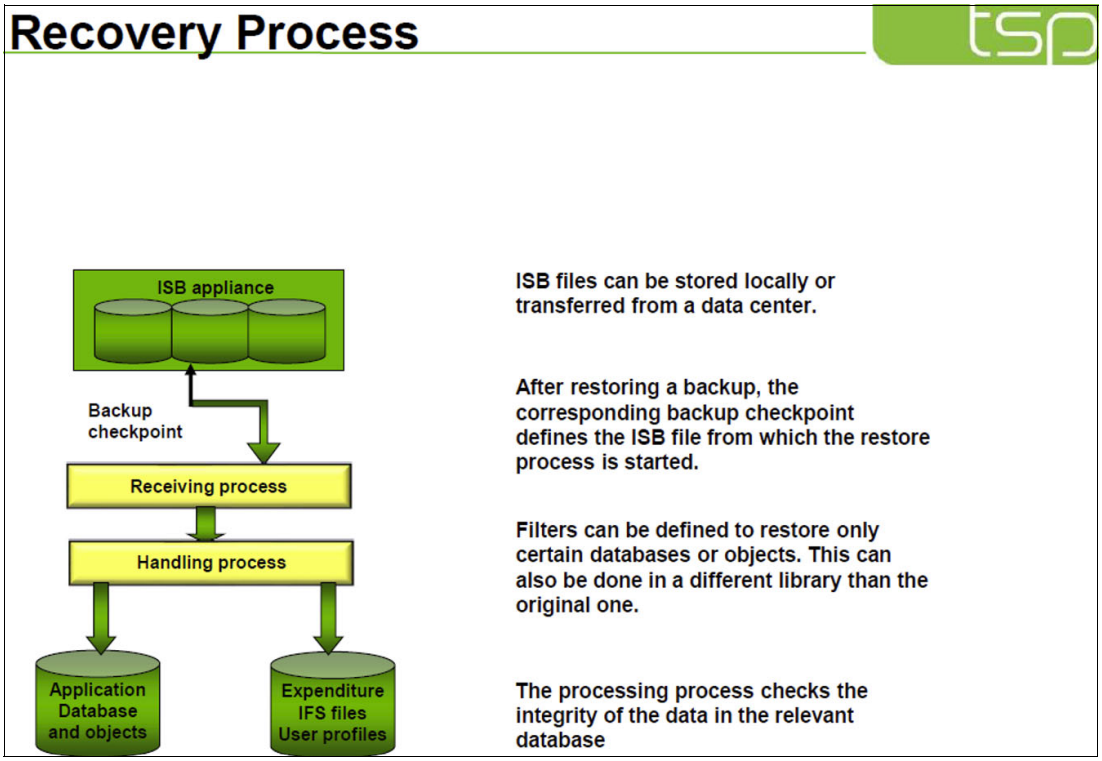

Figure 6-108 Recovery process

The installation of Bus4i ISB includes the following steps:

- Creation or use of the QAUDJRN system journal
- Creation or use of existing data journals
- Installation of BUS4i ISB on the IBM i system
- Configuration of the backup application with the required parameters
- Installation of the ISB Appliance in the internal network
- Start of the permanent automatic backup process in real time

Bus4i ISB includes the following advantages:

- The restore process can be based on one of the previous 9999 backups.
- After a program failure, operator errors, or security incidents, then individual files and objects can be restored in separate libraries.
- IFS files encrypted by ransomware attack can be restored quickly.
- Spool files can be re-created at any time.
- The integrated backup process allows the ISB system to operate with almost no operator intervention.
- Can be combined with BUS4i HA High Availability.

6.4 Disaster Recovery with storage replication

6.4.1 Data replication

Data replication is key to business resiliency because data drives decision-making. Data informs and feeds into mission-critical processes, analytics, systems and ultimately, business insights. Organizations must guarantee that it is constantly available and accessible to users in near real-time. As enterprises develop across geographies and platforms, replication enables them to scale in tandem with their expanding data requirements while maintaining performance.

IBM Power Virtual Server now supports storage level replication between different Power Virtual Server data centers. Global Replication Service on IBM Power Virtual Server provides asynchronous data replication between two sites over long distances. Asynchronous replication updates the backup instance at regular intervals, rather than continuously. This requires less bandwidth and connectivity, is less costly, and uses fewer resources.

6.4.2 Benefits of Global Replication Service

Global Replication Service (GRS) is a valuable feature for high availability and disaster recovery because it keeps your data offsite and available for recovery in a different location if there is a disaster. If the primary instance is destroyed by a catastrophic incident, such as a fire, storm, flood, or other event, then your secondary data instance remains secure off-premises so that you can retrieve your data.

The Global Replication on Power Virtual Server provides the following benefits:

- Maintains a consistent and recoverable copy of the data at the remote site, which is created with minimal impact to applications at your local site
- Efficiently synchronizes the local and remote sites with support for failover and fail-back modes, helping to reduce the time that is required to switch back to the local site after a planned or unplanned outage
- Replicates more data in less time to remote locations
- Maintains redundant data centers in distant geographies for rapid recovery from disasters
- Eliminates costly dedicated networks for replication and avoid bandwidth upgrades

A

IBM i licensing considerations

This appendix describes important considerations about licensing information for IBM i in the IBM Power Systems Virtual Server on IBM Cloud. As cloud changes the computing landscape, customers use cloud technologies to make business decisions regarding how and where to run their compute environments.

IBM has a multi-cloud strategy for IBM i, which is why it is important to know more information about IBM i licensing as part of the journey to the cloud to move IBM i VMs from on-premises to off-premises considering the base operating system and the Licensed Program Products (LPPs).

This appendix includes the following topics:

- “IBM i Subscription Term Licensing”
- “IBM i ISV Licensing in Power Virtual Server”
- “IBM i Support and Power Virtual Server”
- “Other Licensed Program Products”
- “Passport Advantages software”.
- “FAQs about IBM i on IBM Power Systems Virtual Server”.

© Copyright IBM Corp. 2025.

A.1 IBM i Subscription Term Licensing

IBM announces the conversion pricing features for converting from IBM i non-expiring to IBM i Subscription Term.

As announced in AD24-0492, dated 12 March 2024, starting 7 May 2024, the IBM i Processor and User Transfer offering will change for clients who acquire new specific IBM Power10 processor-based servers. This affects the IBM i P05 and P10 software tiers on the following machine types:

- IBM Power S1014 (9105-41B) 4-core processor (P05 software tier)
- IBM Power S1014 (9105-41B) 8-core processor (P10 software tier)
- IBM Power S1022 (9105-22A) (P10 software tier)
- IBM Power S1022s (9105-22B) (P10 software tier)

When purchasing these machine types, the following IBM i licensing options are available on the machine:

- Acquire or renew IBM i Subscription Term licenses.
- Convert IBM i non-expiring licenses to IBM i Subscription Term license at a lower-priced Subscription option. Transfer of IBM i non-expiring licenses is not supported. IBM now provides the conversion pricing features for this lower-priced Subscription option.

A.2 IBM i ISV Licensing in Power Virtual Server

IBM i has thousands of ISVs that provide licenses based on different criteria, including those in the following list:

- Processor group or software tier (that is, P05/P10/P20/P30)
- Serial number
- Partition ID processor feature
- Processor or core count
- Users
- Modules

Other criteria might be used, and licenses might be based on multiple criteria. IBM encourages ISVs to not license to hardware components that can change.

Software tiers are determined by the server model. Power Virtual Server uses S922/S1022, which are P10 tier and E980/E1080, which is P30 tier. A customer on P05 hardware on-prem must license to a higher tier in Power Virtual Server.

For customers that are moving to Power Virtual Server, they get a partition on a system with a different serial number than on-prem. There is no way to preserve an on-prem serial number. Power Virtual Server has a pinning option that can be used to try and keep a partition on a system that they licensed their software against. However, this means that Power Virtual Server cannot dynamically move a partition when planned or unplanned maintenance occurs, which means an outage is required when maintenance is performed.

On the future roadmap, Virtual Serial Number (VSN) support is planned, so that partitions can move around in Power Virtual Server and maintain the same serial number.

In general, IBM software acquired through IBM Passport Advantage® (PPA) can be brought to Power Virtual Server. Examples include IBM WebSphere® Application Server, WebSphere MQ, Db2 Connect, and Lotus Notes.

A.3 IBM i Support and Power Virtual Server

Power Virtual Server supports IBM i 7.2 and later on Power9 S922 (P10 tier) and E980 (P30 tier) servers, IBM i 7.3 and later on Power10 S1022 (P10 tier) and E1080 (P30 tier) servers

IBM i 7.3 moved to extended support 10/1/2023

IBM i 7.2 moved to extended support 5/1/2024

IBM i 7.1 went out of service 5/1/2024 and is NO LONGER SUPPORTED in Power Virtual Server

All customers are recommended to move to IBM i 7.4 or 7.5

Figure A-1 shows the diagram for supported upgrade paths.

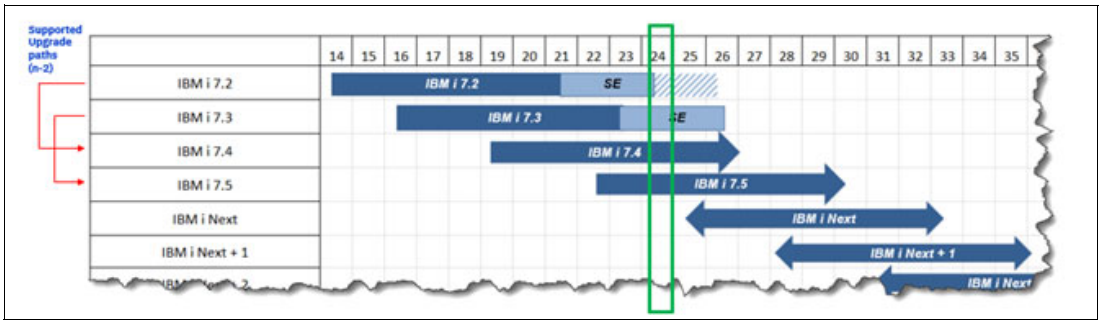

Figure A-1 Support upgrade path

A.4 Other Licensed Program Products

Before you deploy an IBM i on IBM Power Systems Virtual Server on IBM Cloud, you can choose more LPPs. However, a fee payment is required that is shown while creating a virtual machine, and LPPs are individually priced. The selected licensed programs are imported to your IBM i VM by the way of cloud-init. Therefore, the extra LPPs are included in this section.

Each LPP in the package includes all of the features and optional features. For example, the 5770-BR1 solution includes the Network Feature and the Advanced Feature and the base product.

Other LPPs are available for IBM i, which you can include in your VM instance. To include one or more LPPs, complete the following steps:

1. Go to Virtual server instances in the Power Systems Virtual Server user interface and click your instance.
2. Click the **Edit details** option in the Server details window. A menu appears.
3. Select the required licenses that you want to include in your VM instance. As of this writing, you can purchase the following licenses through Power Systems Virtual Server:
   - IBM i Cloud Storage Solutions (5773-ICC).
   - IBM i Power HA (5770-HAS).
   - IBM DB2® Web Query for i Standard Edition (5733-WQS).
   - IBM Rational® Development Studio for i (5770-WDS).

IBM i Cloud Storage Solutions

IBM i Cloud Storage Solutions (5773-ICC) is the IBM cloud solution for backing up IBM i data, as shown in Figure A-2. For more information, see *IBM Cloud Storage Solutions for i User's Guide*.

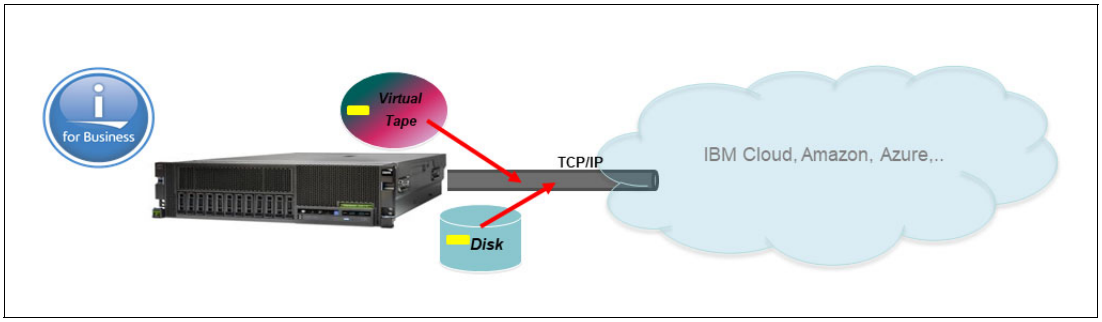

Figure A-2 IBM Cloud Storage Solutions for i

IBM Cloud Storage Solutions for i provides the following features:

- Base plus Advanced feature (encryption and compression).
- Perpetual price per partition. Cloud offering is per, core per month.
- Same price for all system sizes.

Any size of IBM Cloud Storage Solutions for i on IBM Power Systems Virtual Server includes the following list prices:

- Monthly: 71US per core, per month
- Perpetual: 3400US per partition

IBM i PowerHA

IBM i PowerHA geographic mirroring offers a straightforward, cost-effective, high availability (HA) solution for the small to mid-sized client. Typically, PowerHA geographic mirroring is used with internal disk storage, and it provides an alternative to solutions that require the extra configuration and management that are associated with an external storage device.

If the systems are in different locations, it also can provide protection in a site outage for disaster recovery (DR). *Geographic mirroring* refers to the IBM i host-based replication solution that is provided as a function of IBM PowerHA SystemMirror for i.

For more information, see *IBM PowerHA SystemMirror for i: Using Geographic Mirroring (Volume 4 of 4)*, SG24-8401.

Use scenarios

In Scenario 1 (Figure A-3 on page 301), a customer can use IBM i operating system capabilities to set up a DR backup site in Power Virtual Server. The geomirroring capability of IBM i PowerHA allows for IBM i-to-IBM i connectivity across distances. This option applies to only customers that uses Independent ASP (IASPs).

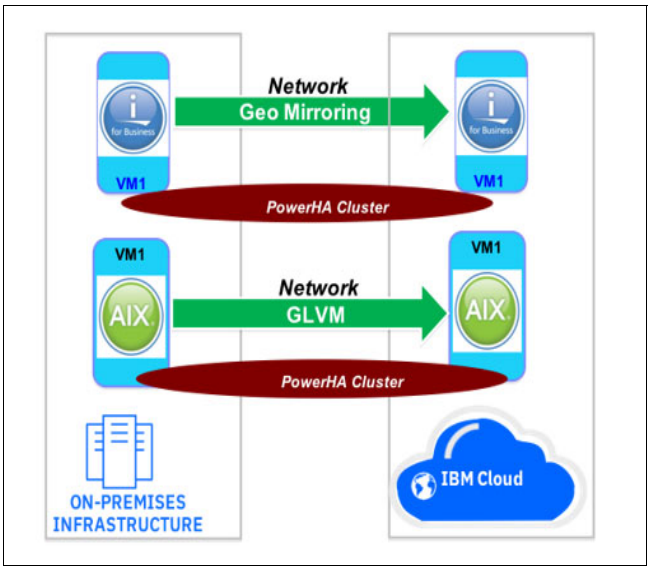

Figure A-3 Scenario 1: Private-to-public cloud DR solution

In scenario 2 (see Figure A-4), a customer can deploy across IBM Cloud regions. For example, PowerHA Enterprise Edition can be used to enable DR across Washington and Dallas.

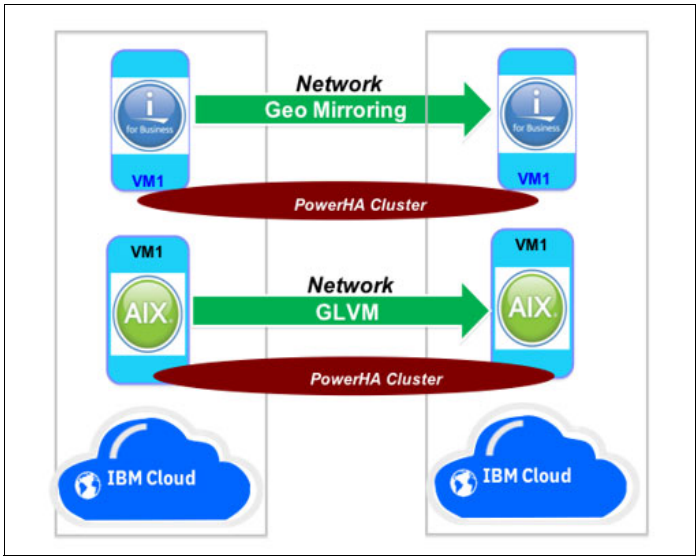

Figure A-4 Scenario 2: Public-to-public cloud DR Solution

Important: IBM PowerHA for i (5770-HAS) Enterprise Edition is available in only the IBM Power Virtual Server on IBM Cloud.

IBM DB2 Web Query for i Standard Edition

IBM DB2 Web Query for i is a key product suite that provides analytics infrastructure and tools to simplify and extend your IBM i data analysis.

By using Db2 for i advanced query optimization technologies, Db2 Web Query provides an on-premises or in the IBM Cloud solution for modernizing traditional reporting. It provides

powerful visualizations for dashboards, and automated report execution and distribution. For more information, see *IBM Db2 Web Query for i*.

IBM Rational Development Studio for IBM i

IBM Rational Development Studio for i is an application development package to help you rapidly and cost-effectively increase the number of e-business applications for IBM i. This package consolidates all of the key IBM i development tools into one offering.

IBM Rational Development Studio-user

The 5770-WDS users are unique users per virtual server, not concurrent. The user entitlement is unique to each developer and cannot be shared.

For example, in one virtual server, the client has two developers during shift 1, and two developers during shift 2, the user count is 4.

Per terms

In each virtual server, a user can start the program from various devices, such as terminals, PCs with terminal emulators, or PCs with IDEs. The licensee must obtain separate, dedicated entitlements for each user who accesses the program by any means directly or indirectly, for example, by way of a multiplexing program, device, or application server. An entitlement for a user is unique to that user and cannot be shared or reassigned other than for the permanent transfer of the authorized entitlement to another user.

IBM Rational Development Studio for i (5770-WDS), the per-user price includes all of the following features:

- ILE Compilers (feature 5101 of 5770-WDS):
  - ILE RPG
  - ILE RPG *PRV
  - ILE COBOL
  - ILE COBOL *PRV
  - ILE C
  - ILE C++
  - ILE IXLC for C/C++
- Heritage Compilers (feature 5102 of 5770-WDS):
  - S/36 Compatible RPG II
  - S/38 Compatible RPG II
  - RPG/4004
  - S/36 Compatible COBOL
  - S/38 Compatible COBOL
  - COBOL/400
- Application Development ToolSet (ADTS) (feature 5103 9f 5770-WDS):
  - Source Entry Utility (SEU)
  - Screen Design Aid (SDA)
  - Report Layout Utility (RLU)
  - Programming Development Manager (PDM)

User licensing information and compiler features are listed in Table A-1 on page 303.

Table A-1 Prices for IBM Rational Development Studio for i (5770-WDS)

Rational Development Studio for i	Monthly license list price	Perpetual license list price
Includes three features (all compilers)	Per user pricing $120 per user	► $3885US per user ► $11,655US for three users

Note: If you lift and shift an IBM i VM from on-premises to off-premises and you licensed one of the extra LLPS, these licenses do not work on IBM Power Systems Virtual Server.

Therefore, you must order a license and approve the service agreement and click **Save edits and order** to complete the instance modification process and accept the price.

A.5 Passport Advantages software

IBM programs that are acquired by way of IBM Passport Advantage can be brought to the IBM Power Systems Virtual Server offering.

One example of an IBM i program that is acquired by way of Passport Advantage is Rational Developer for i (RDi). For RDi, you can bring your current RDi license to the Power Systems Virtual Server offering whether RDi was obtained through the Power Systems hardware channel or through Passport Advantage.

If you do not have an RDi license and need the product, obtain a license by using Passport Advantage. Then, bring that license to the Power Systems Virtual Server offering.

Note: Other examples of IBM software that is acquired by way of Passport Advantage are:

- IBM WebSphere MQ
- Db2 Connect
- Lotus Notes

For more information, see Eligible Public Cloud BYOSL Policy.

A.6 FAQs about IBM i on IBM Power Systems Virtual Server

In this section, the most asked questions about IBM i on Power Virtual Server are addressed.

1. Estimating the amount of IBM Cloud Object Storage that is needed for IBM BRMS backups on IBM i Power Virtual Server, what compression ratio can be expected?

   There is almost no compression by default. You can enable compression on the IBM i side, and it compresses the data by about a third, but it will negatively affect performance of backups.

   However, if IBM Cloud Object Storage performs compression, that process is separate from anything IBM i, IBM Cloud Storage, or IBM BRMS controls.

   As an estimate, if you are saving 2 TB, you need 2 TB of IBM Cloud Object Storage space.

2. What IBM i operating system releases are supported?

   Power Virtual Server supports IBM i 7.2 or later. IBM Power Virtual Server (On-premises) supports IBM i 7.3 or later.

3. Can you choose the image that includes the license from the catalog or must you bring your own image?

   Customers can bring their own customized IBM i image. IBM provides stock images to help customers move their image to Power Virtual Server.

4. Can a QSECOFR user access Power Virtual Server by using SSH?

   The QSECOFR user cannot access Power Virtual Server by using SSH. Therefore, a new user is needed. However, Power Virtual Server is available if you modify the `sshd_config` with **`PermitRootLogin yes`**. This practice is *not* recommended. Therefore, create another user with authority that is similar to QSECOFR.

5. Does Power Virtual Server support virtual serial numbers?

   There is no support for virtual serial numbers in Power Virtual Server. However, that feature is on the roadmap.

6. Can a customer move their IBM i back on-premises?

   The IBM i operating system entitlements stay on a machine serial number by default. If a customer sells the donor machine, the IBM i stays on the machine when it is sold. Therefore, if in the future the customer needs their IBM i entitlement on-premises again, they must keep an on-premises donor machine with IBM i entitlement assigned to it.

   If the customer again uses IBM i on-premises, and the software maintenance agreement (SWMA) lapsed on the donor machine, and the customer needs on-premises support, they must reactivate the SWMA. An after-license fee applies.

7. What is driving the business need for a software-based VTL?

   IBM i customers require backup in the cloud and with no ability to use physical tapes, many opt to use virtual tape, meaning that a software-based VTL is needed.

8. Are there any integrated StorSafe solutions in the IBM cloud catalog?

   FalconStor StoreSafe is available with compatible solutions for both Power Virtual Server cloud and on-premise configurations.

9. Is the customer required to keep SWMA active on-premises?

   If the customer is moving all of their IBM i workload to Power Virtual Server, they do not need to keep IBM i SWMA active. This is covered in Power Virtual Server.

10. What is the expected time to deploy an IBM i VM on IBM Power Systems Virtual Server?

    Typically, under 10 minutes.

11. Are all IBM i licensed programs available in Power Virtual Server?

    A generous set of licensed programs are provided in Power Virtual Server as part of the base install. Optional others are available at additional cost. For more information, see IBM i LIcense Program Products (LPP) and Operating System feature bundles

12. For IBM i BRMS/IBM Cloud Storage backup to IBM Cloud Object Storage. What is the average throughput for 1 TB backup? How do you calculate the amount of time needed?

    For BRMS that uses HAProxy bare metal with private IBM Cloud Object Storage endpoint, approximately 150 MBps without compression, approximately 59 MBps with BRMS compression, and approximately 110 MBps with IBM Cloud Storage compression.

    Important: Consider those averages as guidelines only. The averages were derived from tests and can be different in your scenario or situation.

13. What is the preferred compression solution? IBM Cloud Storage compression or BRMS compression? Which one delivers a better compression ratio?

The following options are available:

- Without any compression, more throughput is realized, but more virtual tapes are used. Therefore, it takes equal time by using IBM BRMS compression.
- With BRMS compression, you need fewer virtual tapes but throughput is less, and the transfer time is the same as without compression.
- With IBM Cloud Storage compression, you need the same number of virtual tapes as without compression (the first option in this list), and you do not get any extra space benefit such as you get with BRMS compression. The time that it takes to transfer to IBM Cloud Object Storage increases by a factor of 5–6, so It is not recommended to use IBM Cloud Storage compression. It is suggested to use BRMS compression because you save extra space in your IBM i LPAR that you can use for retaining virtual tapes on your IBM i LPAR, too. BRMS compression provides approximately 59% compression ratio, and IBM Cloud Storage provides approximately 80% compression ratio.

14. Can IBM Cloud Object Storage be used as backup for IBM i?

    Yes, you can use BRMS with IBM Cloud Storage (5733-ICC) products to perform backups in virtual tapes and move it to IBM Cloud Object Storage bucket. Consider the following points:

    - You can double the attached disk space to stage backup data.
    - Because virtual volumes are transferred to IBM Cloud Object Storage over the network, expect delays depending on how much data is transferred.
    - Full System save is possible. A Full System restore requires more work with an NFS server and network installation.
    - Customers must expect delays for data restores because the virtual volume must be downloaded to a local disk before restore.
    - You need a reverse proxy or Vyatta NAT to reach IBM Cloud Object Storage from an IBM Cloud Private network. It is also possible to perform backups from on-premises directly to IBM Cloud Object Storage by using a public interface whether all requirements are met in on-premises systems. Periodically review data transfer performance.

15. If a customer takes their IBM i image from on-premises to off-premises and that system includes IBM Rational Dev Studio licensed, does it transfer to an instance that is on IBM Power Systems Virtual Server on IBM Cloud? Does IBM Rational Directory Server on Power Virtual Server include any extra cost?

    If you are bringing your IBM i VM to Power Virtual Server by image file or backup or restore, the license key for IBM Rational Directory Server is applied by the IBM Cloud team if you selected a previous version of the VM in the IBM portal and another fee is associated with it.

16. In which regions is Power Virtual Server available?

    Power Virtual Server is now available in 21 data centers around the globe with other regions likely to be added. As of this writing, the data centers are located in Dallas, Sao Paulo, Toronto, Washington DC, Frankfurt, London, Madrid, Sydney, Tokyo, Osaka, and Chennai.

    For more information about data center regional capabilities, see the following web page: `https://cloud.ibm.com/power/overview`

17. How much is Power Virtual Server?

    Pricing is based on the resources that are chosen with options for PAYG, 1-year, and 3-year plans. A cost estimator can be found here at the following web page:

    https://cloud.ibm.com/docs/power-iaas?topic=power-iaas-generating-an-estimate

Abbreviations and acronyms

ACLs	Access control lists
ACS	Access Client Solution
ADTS	Application Development ToolSet
AES	Advanced Encryption Standard
ATL	Automatic Tape Library
BRMS	Backup, Recovery, and Media Services
CA	Certificate Authority
CCSID	Coded character set identifier
CL	Control language
COD	Capacity On demand
COS	Cloud Object Storage
CRG	cluster resource group
DASD	direct access storage device
DL	Direct Link
DR	Disaster Recovery
DRO	Dynamic resource optimization
DoS	Denial of service
ESS	Entitled Systems Support
EULA	End-User License Agreement
FNTS	First National Technology Solution
FTP	File transfer protocol
GCP	Google Cloud Platform
GRE	Generic Routing Encapsulation
GRS	Global Replication Services
HA	high availability
HAS	High Availability Solutions
HMC	Hardware Management Console
HPCS	Hyper Protect Crypto Service
IAM	Identity and access management
IASP	Independent auxiliary storage pool
IFS	Integrated File System
IKS	IBM Kubernetes Service
IOPS	I/O operation per second
IP	Internet Protocol
IPAM	IP address management
IQN	iSCSI Qualified name
ISVs	Independent software vendors
IaC	Infrastructure-as-Code
IaaS	Infrastructure as a service offering
LIC	Licensed Internal Code
LPAR	Logical partition
LPP	Licensed program product
MFA	Multi-factor authentication
MitM	Man-in-the-middle
NAT	Network address translation
NSGs	Network Security Groups
NSP	Network service provider
OTL	Object tracking list
PDM	Programming Development Manager
PER	Power Edge Router
PKI	Public key infrastructure
PPA	Passport Advantage
PSK	Pre-shared keys
PSKs	Pre-shared keys
PTFs	Program temporary fixes
PoC	Proof-of-concept
PoP	Point of Presence
QoS	Quality of service
RBAC	Role-based access control
RDi	Rational Developer for i
RLU	Report Layout Utility
ROKS	Red Hat OpenShift Kubernetes Service
SCC	Security Compliance Center
SCD	System Context Diagram
SDA	Screen Design Aid
SEU	Source Entry Utility
SLIC	System Licensed Internal Code
SWMA	Software maintenance agreement
TCP	Transmission Control Protocol
VM	Virtual machine
VPC	Virtual Private Cloud
VPE	Virtual Private endpoint
VPN	Virtual private network
VRA	Virtual Router Appliance
VSI	Virtual Server instance
VSN	Virtual Serial Number
VTL	Virtual tape library
WAS	WebSphere Application Server

© Copyright IBM Corp. 2024.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- *IBM Power Systems High Availability and Disaster Recovery Updates: Planning for a Multicloud Environment, REDP-5663*
- *IBM PowerHA SystemMirror for i: Preparation (Volume 1 of 4)*, SG24-8400
- *IBM PowerHA SystemMirror for i: Using DS8000 (Volume 2 of 4), SG24-8403*
- *IBM PowerHA SystemMirror for i: Using IBM Storwize (Volume 3 of 4), SG24-8402*
- *IBM Power Security Catalog*, SG24-8568

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Online resources

The following websites also are relevant as further information sources:

- IBM Support portal:

  https://www.ibm.com/support/home/
- Power Systems Virtual Server FAQ:

  https://cloud.ibm.com/docs/power-iaas?topic=power-iaas-power-iaas-faqs
- Power Systems Virtual Server getting started:

  https://cloud.ibm.com/docs/power-iaas?topic=power-iaas-getting-started
- Power Systems Virtual Server network architecture diagrams:

  https://cloud.ibm.com/docs/power-iaas?topic=power-iaas-network-architecture-diagrams
- Megaport ordering considerations:

  https://cloud.ibm.com/docs/dl?topic=dl-megaport
- Megaport portal:

  https://portal.megaport.com/login
- PowerHA SystemMirror for i Product Information:

  https://helpsystemswiki.atlassian.net/wiki/spaces/IWT/pages/163577866/Welcome+to+PowerHA+SystemMirror+for+i

© Copyright IBM Corp. 2025.

- IBM i Migration to Cloud with IBM Power Virtual Server: An IBM Systems Lab Services Tutorial

  https://cloud.ibm.com/media/docs/downloads/power-iaas-tutorials/Power Virtual Server_IBMi_Migration_Tutorial_v1.pdf
- IBM i Backups with IBM Power Virtual Server: An IBM Systems Lab Services Tutorial

  https://cloud.ibm.com/media/docs/downloads/power-iaas-tutorials/Power Virtual Server_IBMi_Backups_Tutorial_v1.pdf
- FALCONSTOR STORSAFE VTL for IBM Power Virtual Server Deployment Guide

  https://falconstor-download.s3.us-east.cloud-object-storage.appdomain.cloud/FalconStor%20VTL%20for%20IBM%20Deployment%20Guide.pdf

For more information, see the following resources:

- Are you saving the right stuff?
- The BRMS Graphical User Interfaces (GUI)
- BRM1744 Issued During a BRMS Full Systems Save or SAVSYS
- BRMS and MONSWABRM or Save-While-Active (SWA)
- BRMS Messages
- BRMS Network Problem Determination Q & A
- BRMS Networking: Adding, Removing, or Renaming a System
- BRMS Networking: Relational Database Directory Entries
- BRMS Networking: Secure DDM
- BRMS Parallel-Parallel versus Multiple-Library Parallel Concurrent Saves
- BRMS Systems Appear Inactive in a BRMS Network
- Factors That May Influence File Size in QUSRBRM library
- Sequence of Events when Running a BRMS Control Group
- Using *EXITs in BRMS Control Groups
- BRMS Save Files
- Using BRMS to Duplicate Media
- Using BRMS to Encrypt Data
- IBM i - BRMS : Welcome to IBM Backup, Recovery & Media Services (BRMS) for i
- BRM1881 when using virtual optical for save
- Replace media devices on BRMS
- Welcome to IBM Backup, Recovery & Media Services (BRMS) for i
- How ADDMLMBRM and Movement Works in a BRMS Network Environment
- How BRMS handles control group saves when there is an error during the backup activity
- How BRMS Move Policies Work
- How to Set Up a Parallel Backup in BRMS
- How to Set Up BRMS Control Group to IPL System After Backup Completes
- Integrated File System Files Created by BRMS to Be Used for Diagnosing BRMS Problems
- Recovering Systems When Using BRMS Object Lists

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Redbooks

Power Virtual Server for IBM i

SG24-8513-01

ISBN

IBM

(0.5” spine)
0.475”<->0.873”
250 <-> 459 pages



SG24-8513-01

0738461962

Printed in U.S.A.

Get connected



