

IBM Power Systems Virtual Server Guide for IBM i

Dino Quintero

Jose Martin Abeleira

Adriano Almeida

Marcelo Avalos

Sanjeev Chhabra

Michael Easlon

Luis Ferreira

Ahmad Y Hussein

Diego Kesselman

Gabriel Padilla Jimenez

Bogdan Savu

Deepak C Shetty

Luis Eduardo Silva Viera

Travis Siegfried



 **Cloud**

Power Systems



IBM Redbooks

IBM Power Systems Virtual Server Guide for IBM i

March 2022

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

First Edition (March 2022)

This edition applies to Version:

IBM i 7.2

IBM i 7.4 TR4 RS740-G

IBM i Access Client Solutions 1.1.8.0

Bus4i version 7.1

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

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

Contents

Notices	vii
Trademarks	viii
Preface	ix
Authors	ix
Now you can become a published author, too!	xiii
Comments welcome	xiii
Stay connected to IBM Redbooks	xiii
Chapter 1. Introduction to IBM Power Systems Virtual Server	1
1.1 Overview	2
1.1.1 Potential consumers	2
1.2 Creating a Power Systems Virtual Server	3
1.2.1 IBM Cloud Dashboard	4
1.2.2 IBM Cloud Catalog	5
1.3 Power Systems Virtual Server service	6
1.4 Power Systems Virtual Server instance	9
1.4.1 VSI ready to use	11
Chapter 2. Migration to Cloud with IBM Power Systems Virtual Server	13
2.1 IBM i and data migration	14
2.2 Backup and Restore migration	16
2.2.1 Migration with partial saves and restores	17
2.2.2 Full system backup and restore by using native commands	19
2.2.3 Full-system backup from IBM i Source by using BRMS and IBM Cloud Storage and Restore on IBM Power Systems Virtual Server	39
2.2.4 Solution diagrams	50
2.2.5 Migrating IBM i by using PowerVC to PowerVS	55
2.2.6 Migrating IBM i by using Mass Data Migration	57
Chapter 3. IBM Power Systems Virtual Server in IBM Cloud Network	63
3.1 IBM Power Systems Virtual Server virtual private network connectivity introduction ..	64
3.2 IBM PowerVS network overview	65
3.3 IBM PowerVS network scenarios	67
3.3.1 Nonproduction proof-of-concept scenarios	67
3.3.2 Production scenarios	71
Chapter 4. Backing up IBM i on IBM Power Systems Virtual Server	79
4.1 Backup and restore considerations on IBM Power Systems Virtual Server	80
4.1.1 Overview	80
4.2 Backing up and restoring by using IBM Backup, Recovery, and Media Services	81
4.2.1 IBM BRMS overview	81
4.2.2 IBM Cloud Storage Solutions	82
4.2.3 Cloud solutions for IBM i characteristic	84
4.2.4 IBM BRMS turn-key	85
4.2.5 Backing up the IBM i system to the cloud	86
4.2.6 Full-system backups from the cloud	89
4.2.7 Recovering the IBM i system from the cloud	93
4.2.8 Full-system recovery from the cloud using IBM i as an NFS server	94

4.3	Creating object-level backups	113
4.3.1	Introduction	113
4.3.2	Copying files to the cloud	113
4.3.3	Copying files from the cloud	114
4.3.4	Backup and restore by using save files.	116
4.3.5	Backup by using image catalog and IBM Cloud Object Storage.	121
4.3.6	Sample save and restore IBM i objects to IBM Cloud Object Storage	124
4.3.7	Transferring a save library to the cloud by using IBM BRMS	128
4.3.8	Automatically transferring media to IBM Cloud Object Storage	130
4.4	Dynamic Solutions International Virtual Tape Library	132
4.4.1	Introduction	132
4.4.2	Overview	132
4.4.3	IBM BRMS requirements for DSI VTL iSCSI	133
4.4.4	Backing up IBM BRMS by using DSI VTL iSCSI	134
4.4.5	Recovering IBM BRMS by using DSI VTL iSCSI	136
4.4.6	DSI VTL restore on IBM PowerVS	137
4.4.7	Sizing and ordering DSI VTL	137
4.5	Full-system snapshot, cloning, and restore.	137
4.5.1	Overview	137
4.5.2	Snapshot.	137
4.5.3	Cloning a volume	138
4.5.4	Restore a snapshot.	138
4.5.5	IBM PowerVS API use	139
4.5.6	Use cases	143
 Chapter 5. Hints and tips for IBM i deployments on IBM Power Systems Virtual Server		
	145	
5.1	Overview	146
5.1.1	Connecting to an IBM i virtual machine.	146
5.1.2	Remote access to IBM i by tunneling	151
5.1.3	IBM i 5250 console through LAN adapter.	153
5.2	Using snapshots on IBM i instances	157
5.2.1	Taking snapshots	157
 Chapter 6. Disaster Recovery with IBM Power Systems Virtual Server.		161
6.1	Introduction	162
6.2	Solution components and requirements	162
6.2.1	Solution components	162
6.2.2	Solution requirements.	163
6.3	Operating system level replication use case with PowerHA SystemMirror for i geographic mirroring	192
6.3.1	Introduction	192
6.3.2	Prerequisites	193
6.3.3	Planning	193
6.3.4	Implementation of PowerHA SystemMirror for i geographic mirroring in the Power Systems Virtual Server environment	199
6.3.5	Switching a geographic mirroring environment.	209
6.3.6	Troubleshooting	212
6.4	Logical replication use case with Bus4i.	213
6.4.1	Introduction	213
6.4.2	Prerequisites	214
6.4.3	Planning	214
6.4.4	Implementation of Bus4i for IBM i Logical Replication	222
6.4.5	Switching a Bus4i environment between two IBM i virtual machines in the IBM Power	

Systems Virtual Server Cloud	230
Chapter 7. Reference architectural decisions to migrate IBM i on-premises to IBM Power Systems Virtual Server	243
7.1 Introduction	244
7.2 Use case overview	244
7.2.1 Scope	245
7.2.2 Architectural decisions	248
7.2.3 Architecture overview	254
Appendix A. IBM i licensing considerations	257
IBM i Licensed Program Products and operating systems.	258
Other Licensed Program Products	260
Passport Advantages software	265
Movable IBM i	266
Appendix B. Frequently asked questions about IBM i on IBM Power Systems Virtual Server	269
FAQs about IBM i on IBM Power Systems Virtual Server	270
Miscellaneous FAQs about IBM i on PowerVS.	270
Related publications	279
IBM Redbooks	279
Online resources	279
Help from IBM	280

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

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

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

AIX®	IBM Cloud Pak®	Redbooks®
Aspera®	IBM Garage™	Redbooks (logo)  ®
CICS®	Passport Advantage®	SystemMirror®
Db2®	POWER8®	Think®
DB2®	POWER9™	Tivoli®
FlashCopy®	PowerHA®	WebSphere®
IBM®	PowerVM®	XIV®
IBM Cloud®	Rational®	

The following terms are trademarks of other companies:

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.

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

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.

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

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

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

Preface

This IBM® 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 and 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 available at the [IBM Documentation web page](#) and aligns with the educational materials that are provided by IBM Garage for Systems Technical Education.

Authors

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

Dino Quintero is an IBM Redbooks Project Leader with IBM Systems. He has 25 years of experience with IBM Power Systems technologies and solutions. Dino shares his technical computing passion and expertise by leading teams developing technical content in the areas of enterprise continuous availability, enterprise systems management, high-performance computing (HPC), cloud computing, artificial intelligence (including machine and deep learning), and cognitive solutions. He also is a Certified Open Group Distinguished IT Specialist. Dino holds a Master of Computing Information Systems degree and a Bachelor of Science degree in Computer Science from Marist College.

Jose Martin Abeleira is a Senior Systems/Storage Administrator at DGI (Uruguay Taxes Collection Agency). A former IBMer, he is a Gold IBM Redbooks Author, Certified Consulting IT Specialist, and IBM Certified Systems Expert Enterprise Technical Support for IBM AIX® and Linux in Montevideo, Uruguay. He worked with IBM for 8 years and has 18 years of AIX experience. He holds an Information Systems degree from Universidad Ort Uruguay. His areas of expertise include Power Systems, AIX, UNIX, LINUX, Live Partition Mobility (LPM), IBM PowerHA® SystemMirror®, SAN and Storage on IBM DS line, V7000, HITACHI HUSVM, and G200/G400/G370. He also teaches Infrastructure Administration in the Computer Technologist career that was created by the joint venture between UDELAR (Universidad de la Republica), UTU (Universidad del trabajo del Uruguay), and UTEC (Universidad Tecnologica).

Adriano Almeida is a Senior Power Systems Consultant from IBM Systems Lab Services in Brazil. He has worked at IBM for 20 years. His areas of expertise include IBM AIX, PowerVM®, IBM Power Virtualization Center, IBM PowerHA SystemMirror, Linux, SAP HANA on IBM Power Systems, IBM Power Systems Virtual Server, and IBM Cloud®. He is an IBM Certified Expert IT Specialist and IBM Certified Advanced Technical Expert on IBM Power Systems. He has worked extensively IBM Power Systems Project projects, performing implementations, health checking, performance analysis, and consulting on IBM Power Systems environments, and project leadership. He also conducts IBM Power Systems trainings and is TechU speaker. He holds a degree in computing technology from the Faculdade de Tecnologia em Processamento de Dados do Litoral (FTPDL). He is a coauthor of *Exploiting IBM PowerHA SystemMirror V6.1 for AIX Enterprise Edition*, SG24-7841; *IBM PowerVM Best Practices*, SG24-8062, and *Implementing IBM VM Recovery Manager for IBM Power Systems*, SG24-8426.

Marcelo Avalos is an Architect that works for Kyndryl. Previously, he worked for IBM and DATEC LTDA (an IBM Business Partner). He is responsible for direction and support of a portfolio in North America and Latin America migrating from anywhere-to-anywhere, moves to any target that includes private, public and hybrid cloud. He has experience delivering IBM semiconductors, IBM Power Systems, mid-range, and high-end storage systems, and system software. He holds IBM Certification in Technical Specialist at an expert L2 level. He received his Electronic Systems Engineering degree from Escuela Militar de Ingeniería in Bolivia. He also holds a Postgraduate Coursework of Higher Education Teaching degree from Universidad Mayor de San Andrés and is pursuing a master's degree in Project Management at Universidad para la Cooperación Internacional - Costa Rica.

Sanjeev Chhabra is a Power Systems Technical Specialist (Advisory) with IBM Techline. He has 23 years of experience with IBM Power Systems technologies and IT industry experience that encompasses a wide range of skill sets, roles, and industry verticals. He possesses rich corporate experience in architecting solution design, solutions implementations, and system administration, primarily the use of IBM Power Systems, IBM mainframes, and cloud infrastructure management with applications modernization and sizing. Sanjeev shares his technical computing passion and expertise by leading the Worldwide Solution Design offering also known as Capacity Planning and Sizing for Power Systems (on-premises and on cloud). He is also a Level 1 Certified Specialist. His core area role includes providing technical consultancy in Power Systems technologies, including solution design, capacity planning (IBM i Performance Study of Workloads), hardware sizing for on-premises and hybrid and private cloud, and IBM PowerVS on IBM Cloud, including IBM Power Servers that are running with IBM i, AIX, and Linux operating systems, virtualization, high availability, and other Power hardware and software products.

Michael Easlon is an IT Management Consultant within the IBM Garage™ for Systems. He has 20 years of experience with IBM Power Systems and Storage. His areas of focus include: hybrid cloud, Linux, system management, Red Hat Ansible, Terraform, and Red Hat OpenShift Container Platform.

Luis Ferreira is a senior software engineer at IBM Austin, working on cloud containers, Kubernetes-related products, and cloud computing design and architecture. He holds a Master of Science degree from Universidade Federal do Rio de Janeiro in Brazil. Before joining IBM Austin, Luis worked at Tivoli® Systems as a Certified Tivoli Consultant, IBM Brasil as a Certified IT Specialist, and Cobra Computadores as an operating systems designer and kernel developer for the SOX operating system.

Ahmad Y Hussein is an expert level certified technical specialist working for IBM Systems Lab Services. He Holds a master's degree in engineering, has over 17 years of experience, and today provides technical advocacy to enterprise IBM Power Systems clients in cloud and cognitive solutions. Ahmad is an IBM Redbooks author with publications related to cognitive solutions and best practices for IBM Power Systems.

Diego Kesselman is general manager at Esselware Soluciones, an IBM Business Partner in Mexico. 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 also is 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.

Gabriel Padilla Jimenez is a test architect for hardware on IBM Power System. He was involved in early facets of most of the GAs and working closely with development to achieve best quality products. He previously worked on manufacturing lines at IBM Mexico as test lead for XIV® Storage and Power Systems. Gabriel has a Bachelor Degree as Electronic Engineer and a Master's Degree on Information Technology.

Bogdan Savu is a Cloud Infrastructure Architect for IBM Systems. He has over 13 years of experience in designing, developing, and implementing cloud computing, virtualization, automatization, and infrastructure solutions. Bogdan holds a bachelor's degree in Computer Science from the Polytechnic University of Bucharest. He is a Red Hat Certified Architect, IBM Certified Advanced Technical Expert for Power Systems, and VMware Certified Professional. He holds TOGAF 9 Certified and ITIL 4 Managing Professional Certificate certificates. His areas of expertise include cloud computing, virtualization, DevOps, and scripting.

Deepak C Shetty has over 22 years of IT experience that spans various domains that range from application development to systems programming to OpenStack and cloud development and pre-sales consulting on IBM Cognitive and Red Hat OpenShift, multi-cloud, and hybrid-cloud offerings. He has worked in closed and open source development models. In his current role, he leads the Pre-Sales Technical Consulting for IBM Cognitive Offerings at IBM Techline. His primary role is to provide hardware design configurations, solution sizing, and technical consulting on all IBM Cognitive Offerings for IBM Sellers, Business Partners, and distributors worldwide. He also leads the Red Hat OpenShift focal team at IBM Techline and spends extensive time helping the field teams position, size, and sell Red Hat OpenShift on IBM Power System. Based on his experience working with the field teams, Deepak aims to bring in the "tech seller" persona to this publication. He holds four patents, authored several publications, and received several technical awards. He also is an IBM Certified Technical Specialist.

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.

Travis Siegfried is a solutions architect and thought leader in the IBM Garage. His main goal is designing and deploying product strategic solutions that enable Fortune 500 corporations to achieve positive returns on their capital investments. He uses IBM ecosystem Business Partners to create AI-based solutions that work with IBM Cloud and open source to make better-informed decisions, augment teams' creativity, and produce their best work. Travis is knowledgeable in image model training, user experience design, industry solutions, drones, IoT, and the enterprise infrastructure market. He is skilled in meeting financial objectives and developing long-term strategic business plans and has proven customer relationship skills with a strong technical foundation. He excels in meeting challenging opportunities that contribute to a company's success, growth, and profitability.

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

Wade Wallace
IBM Redbooks, Austin Center

Chris Schätzle, Marc Rohloff
TSP Gesellschaft für Informationssysteme mbH

Sergio Leyva
Esselware Soluciones, IBM Business Partner in México

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:

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

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

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

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

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



Introduction to IBM Power Systems Virtual Server

Until now, moving your Power Systems server workloads to the cloud was something of a Utopian idea. The chance to have your IBM i, AIX, or Linux on Power to a public or hybrid Cloud might seem difficult and costly. However, these challenges can now be addressed with the IBM Power Systems Virtual Server offering.

This chapter introduces the conceptual foundations of the IBM Power Systems Virtual Server service and includes the following topics:

- ▶ 1.1, “Overview” on page 2.
- ▶ 1.2, “Creating a Power Systems Virtual Server” on page 3.
- ▶ 1.3, “Power Systems Virtual Server service” on page 6.
- ▶ 1.4, “Power Systems Virtual Server instance” on page 9.

1.1 Overview

The IBM Power Systems Virtual Server (Power 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.

You can use the Power Virtual Servers to deploy a virtual server (also known as a *logical partition* [LPAR] or a *virtual machine* [VM]) in a matter of minutes. As a result, IBM Power Systems clients who often relied on on-premises only infrastructure can now quickly and economically extend their Power IT resources off-premises.

Power Systems servers are confined from others servers with separate networks and direct-attached storage in the data centers. The internal networks are fenced but offer connectivity options to IBM Cloud infrastructure or on-premises environments. This infrastructure design enables essential enterprise software certification and support because the Power Systems Virtual Server architecture is identical to certified on-premises infrastructure.

Power Systems customers who are interested in modernization can benefit from deploying the workloads to Power Systems Virtual Server instead of moving their applications to a new platform that can be expensive and high risk. You can access a stack of enterprise services from IBM, all with pay-as-you-use billing that helps you to quickly scale up and out. IBM Power Systems Virtual Server enables clients to take full advantage of this trend with the ability to provision AIX, IBM i, or Linux instances that are connected to the cloud.

1.1.1 Potential consumers

The potential consumers of the Power Virtual Server service offering are the IT administrators, Managed Service providers/Cloud Service providers (MSPs/CSPs), independent software vendors (ISVs), and application developers.

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.

Managed Service providers/Cloud Service providers

Service providers are interested in expanding the level of service they can offer their clients. 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. They want to continue developing mission-critical applications and then deploy on-premises.

1.2 Creating a Power Systems Virtual Server

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. IBM 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 Systems Virtual Server instance, you must log in to IBM Cloud, as shown in Figure 1-1.

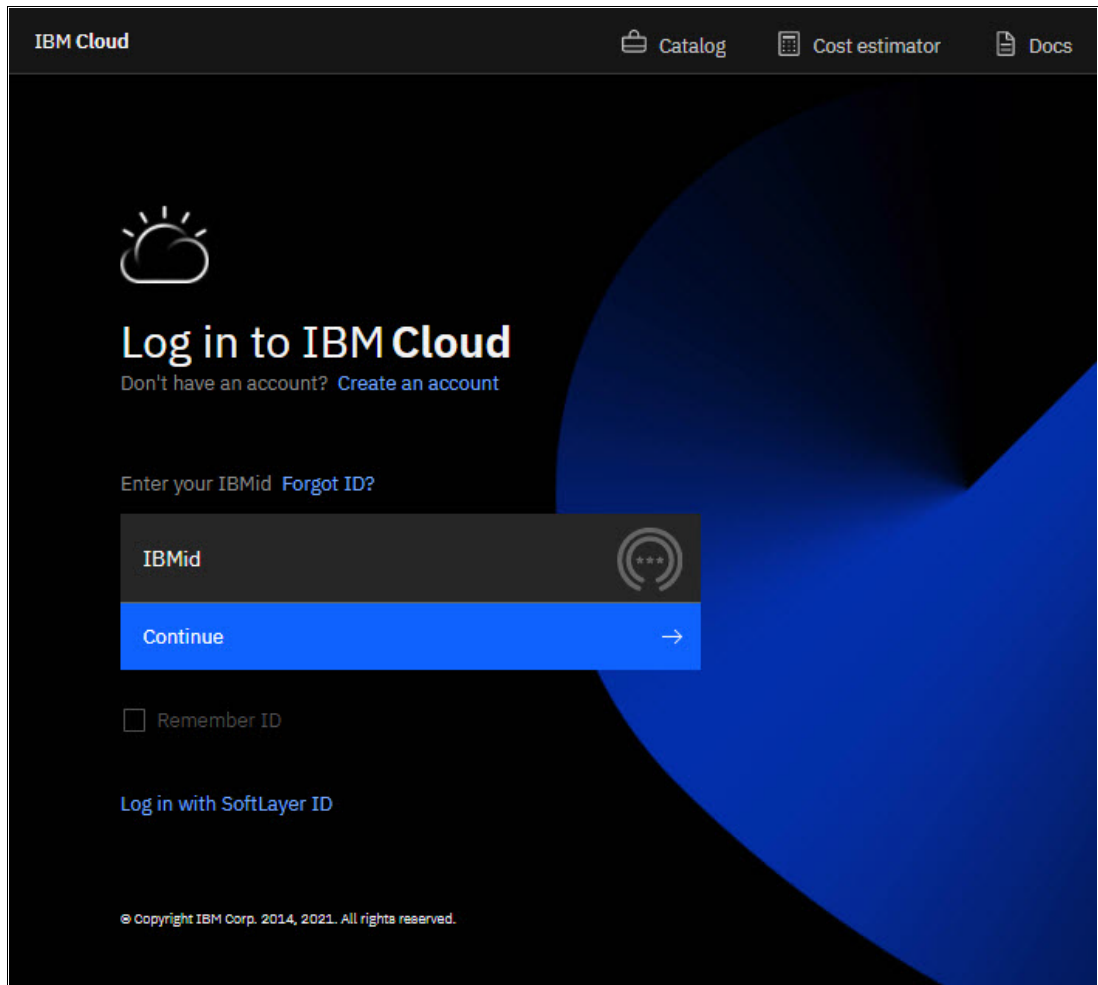


Figure 1-1 Logging in to IBM Cloud

If you do not have an IBM account, you must create one at [this web page](#).

1.2.1 IBM Cloud Dashboard

The Dashboard involves many carefully arranged parts and details and points to useful documentation resources. In addition, many menus, links, and icons help you to fully manage your environment, Virtual Server instances, and associated resources (see Figure 1-2).

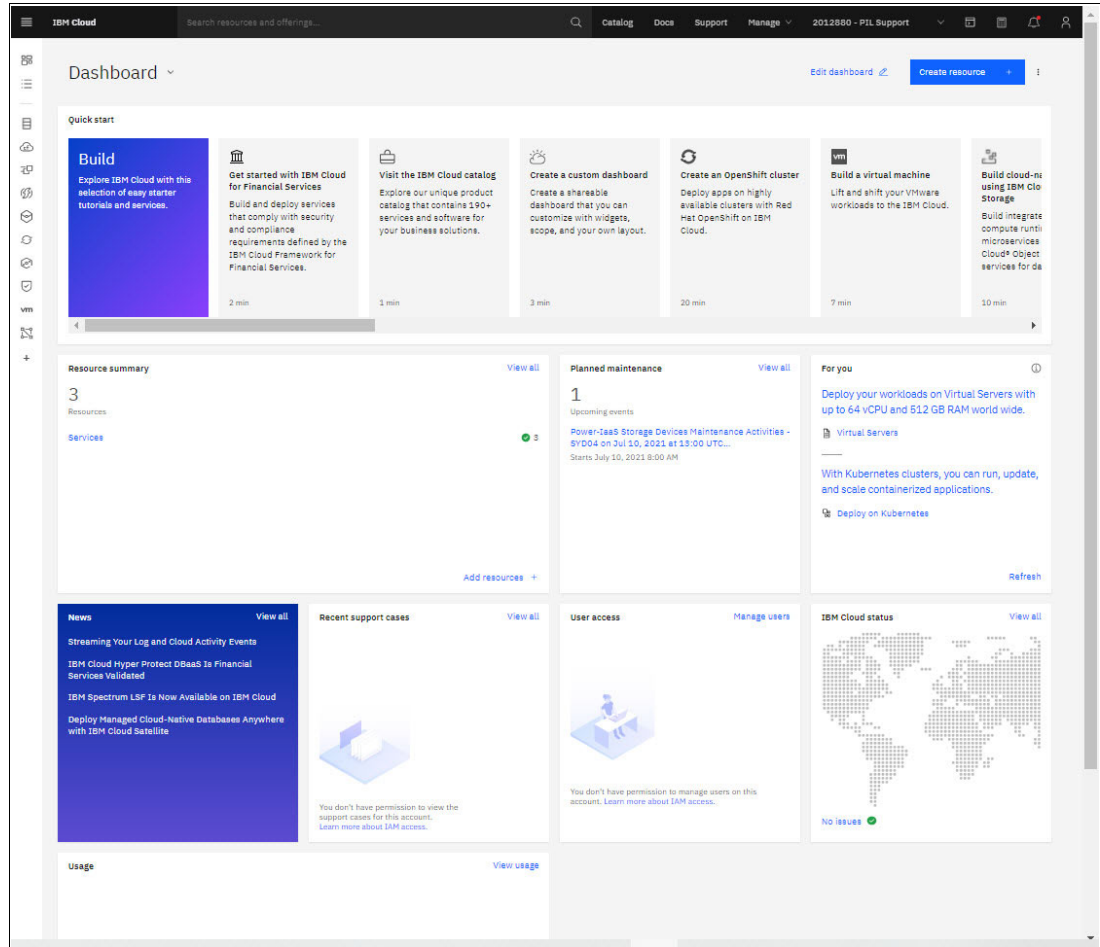


Figure 1-2 IBM Cloud Dashboard

Because dashboards are customizable, you can create a Dashboard that ensures that what is displayed is relevant to you. For example, the dashboards you create can be scoped to specific resources. You also can share the dashboards with users in your account to group resources for specific projects or teams.

You can enable securely authenticate users, control access to Virtual Server resources with resource groups, and allow access to specific resources for a set of users with access groups. This service is based on an Identity and Access Management (IAM) mechanism, which is the one-stop shop for 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 this [IBM Security web page](#).

1.2.2 IBM Cloud Catalog

The IBM Cloud Catalog is an inventory of all offerings that are provided by IBM Cloud, as shown in Figure 1-3.

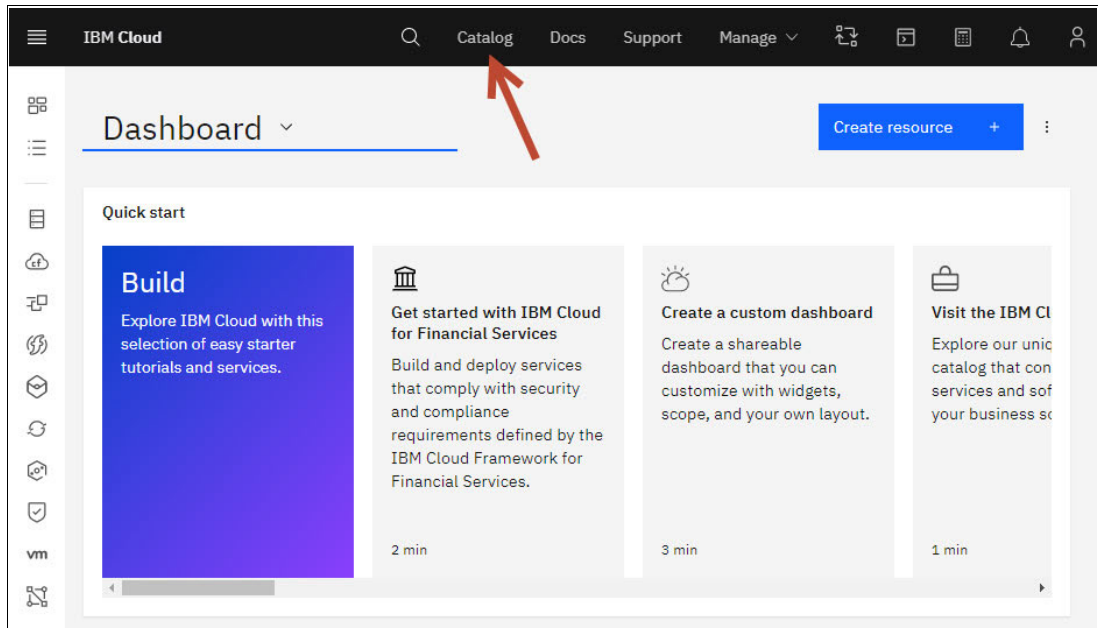


Figure 1-3 IBM Cloud Catalog

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

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

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

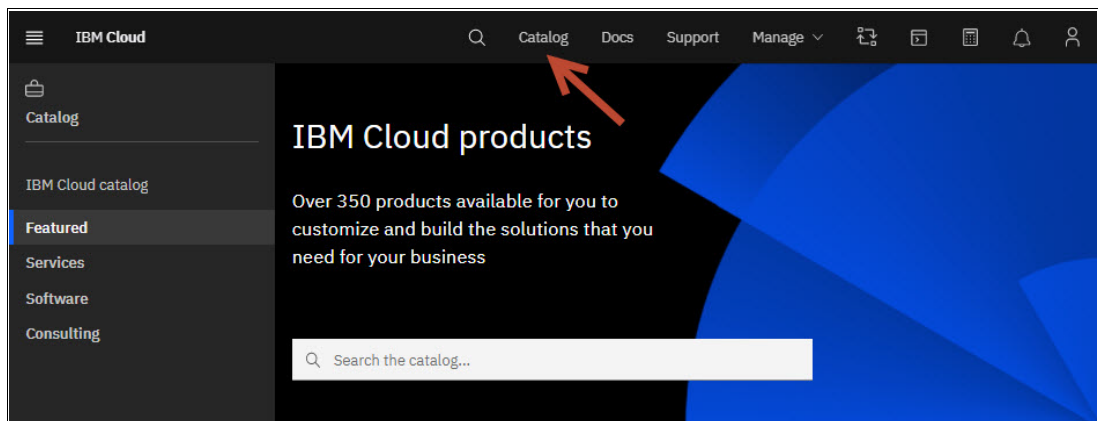


Figure 1-4 IBM Cloud Catalog search

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.

- ▶ Consulting

To get help, from technical to strategic and more, from IBM and our network of partners.

1.3 Power Systems Virtual Server service

To create a Power Systems Virtual Server service, search for Power Systems Virtual Server in the catalog's search field and click **Power Systems Virtual Server**, as shown in Figure 1-5.

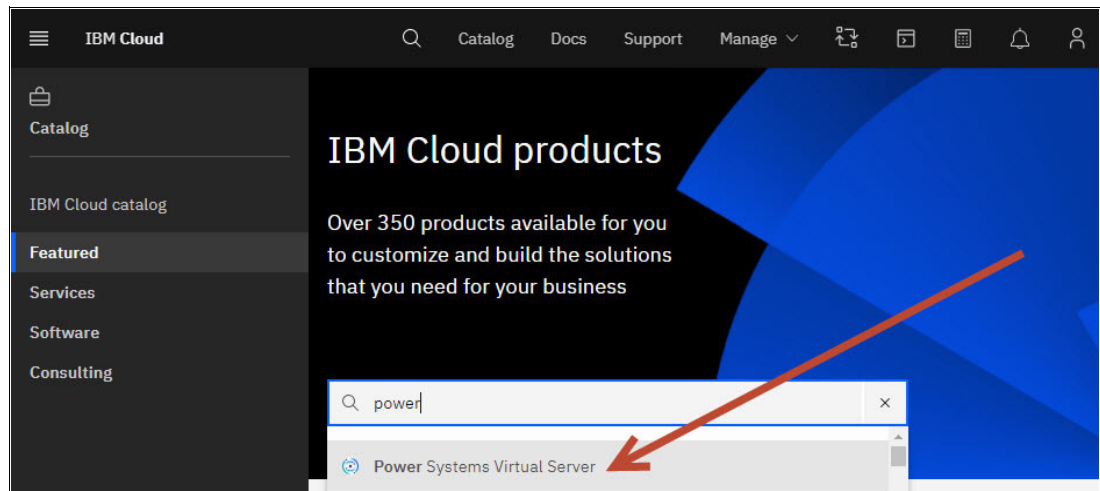


Figure 1-5 Search for Power Systems Virtual Server

A window opens in which you can enter a name for your service and choose where you want to deploy your Power Systems Virtual Server instances (see Figure 1-6).

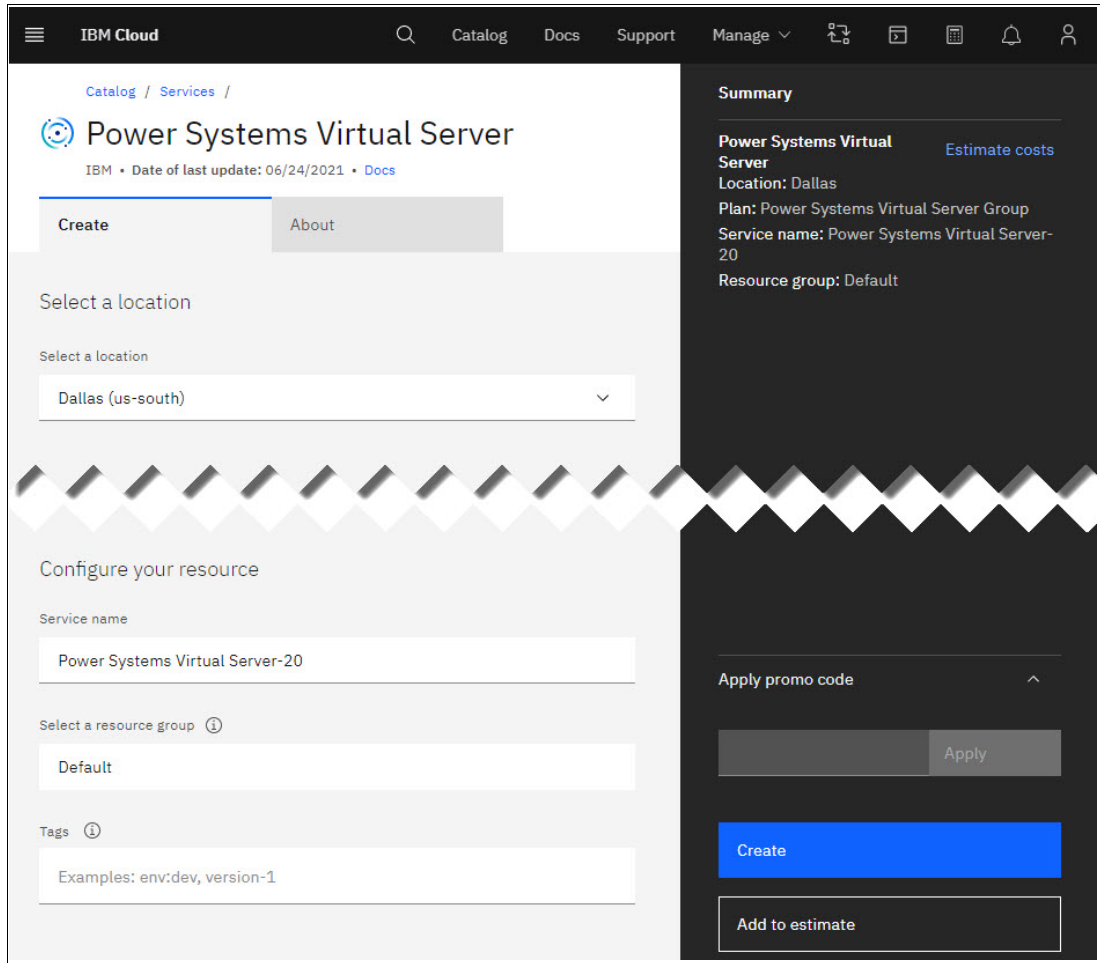


Figure 1-6 Service name and location

Click **Create** and the Resource List window opens, which contains your list of account resources. You can use this window to view and manage your platform and infrastructure resources in your IBM Cloud.

Another way to access the Resource List window is by clicking the navigation menu (the upper-left “hamburger” icon) and then, click **Resource List** in the Dashboard menu (see Figure 1-7).




Figure 1-7 Resource List on Dashboard menu

You can search for resources from anywhere in the IBM Cloud by entering the resource name or tag in the search field from the menu bar. Figure 1-8 shows the result of a Power Systems Virtual Server that were created. All Power Systems Virtual Servers are listed under Services and software resources.

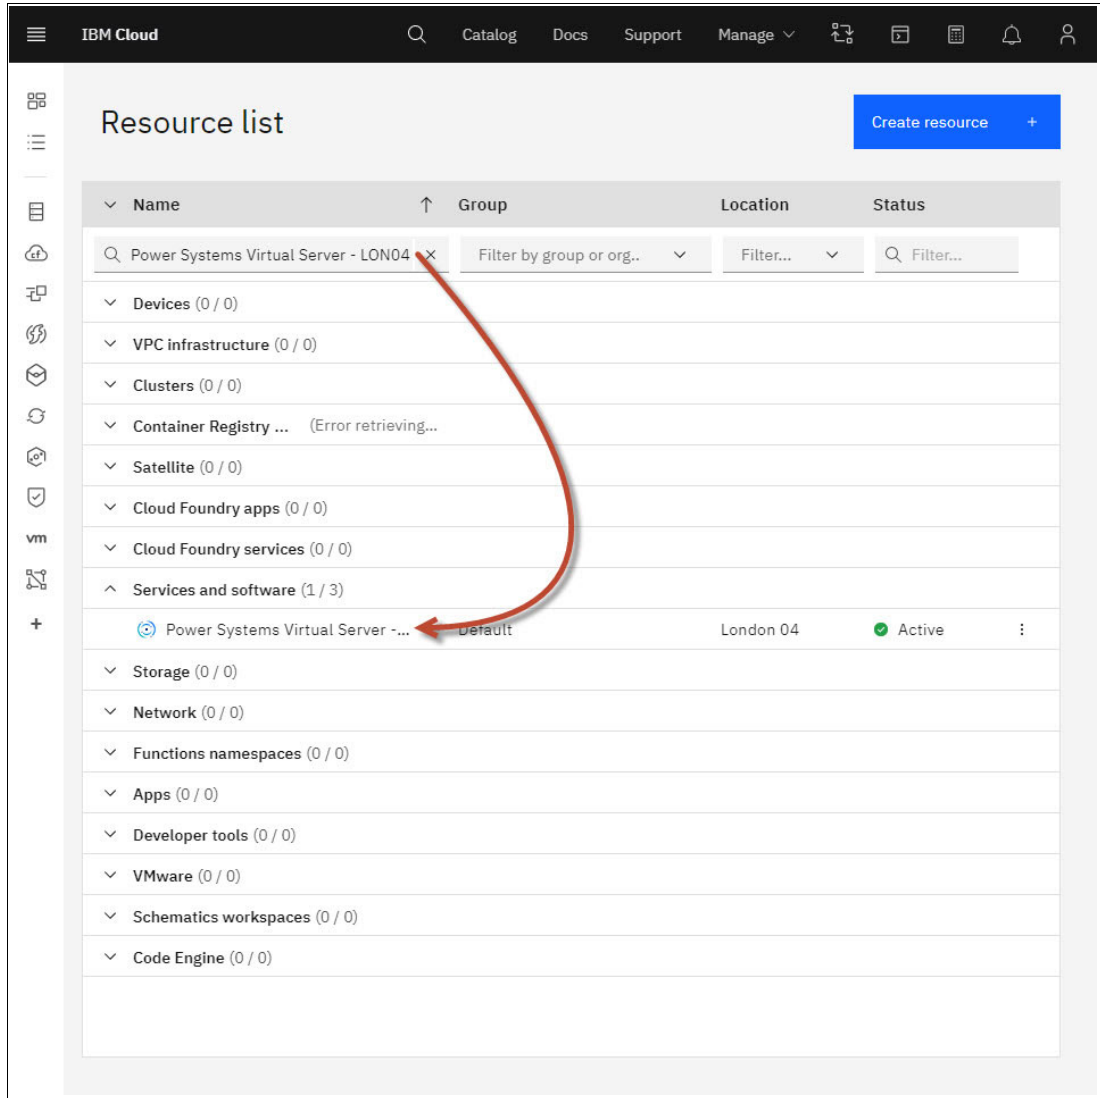


Figure 1-8 Search for resources in the IBM Cloud

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.

1.4 Power Systems Virtual Server instance

You can work with your resources in various ways from your resource list. To directly manage the Power Systems Virtual Server service, click the resource's name to go to the resource details window, as shown in Figure 1-9.

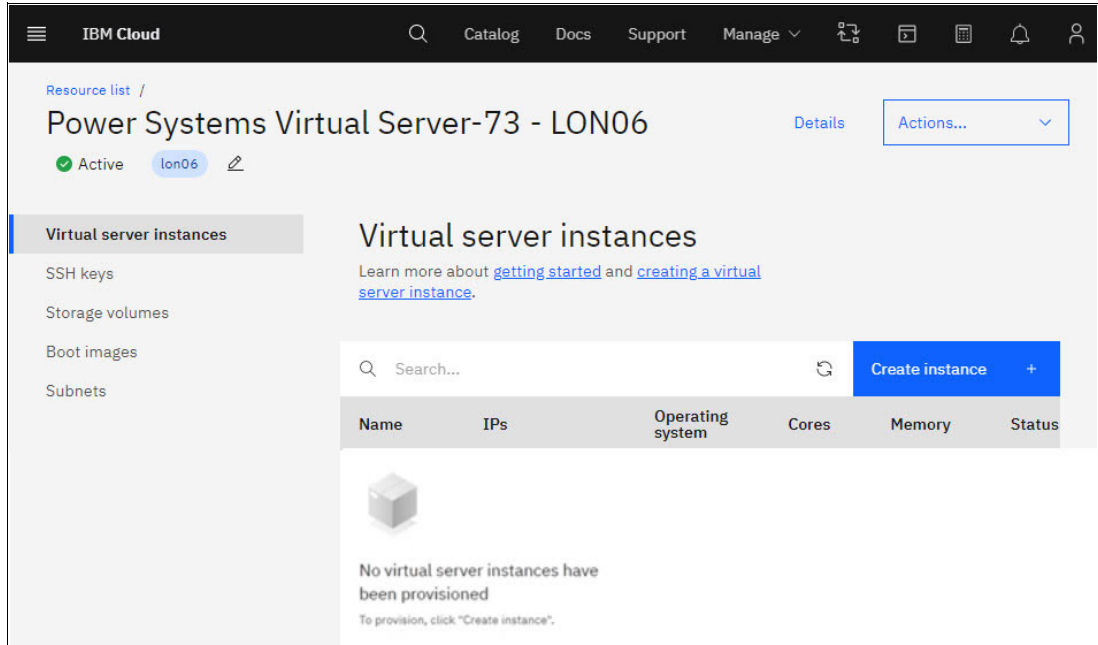


Figure 1-9 Virtual servers instances section

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

The total due per month is dynamically updated in the Order Summary based on your selections, as shown in Figure 1-10. You can easily create a cost-effective Power Systems Virtual Server instance that satisfies your business needs.

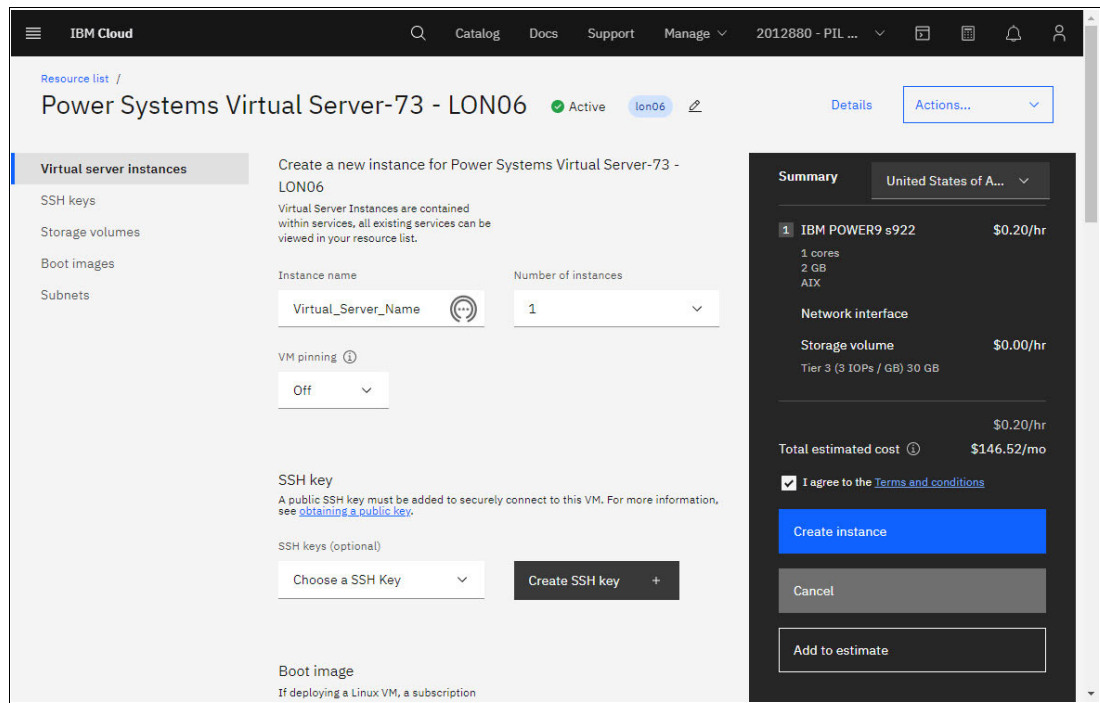


Figure 1-10 Virtual servers instances section and summary cost

In addition to the instance, you also must complete following fields:

- ▶ **Number of instances**
Specify the number of instances that you want to create for the Power Systems Virtual Server. If you specify more than one instance, more 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 host is back to its operating state. The hard pin option restricts the movement of the VM during remote restart, automated remote restart, DRO, and live partition migration.
- ▶ **SSH key**
Choose an SSH key or create one to connect to your Power Systems 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

A core-to-vCPU ratio of 1:1 is used. For shared processors, fractional cores round up to the nearest whole number; for example, 1.25 cores = two vCPUs.
- ▶ Memory

Select the amount of memory for the Power Systems Virtual Server.
- ▶ Boot image

When you select Boot image, the Power Systems Virtual Server user interface allows you to select boot images from a group of stock images or the list of stock images in your catalog.
- ▶ Attached volumes

You can create a data volume or attach an existing volume 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 adding a private network, you can choose a specific IP address or have one auto-assigned.

1.4.1 VSI ready to use

After clicking **Create an instance**, you receive some pop-up messages about your instance, network, and disk space creation. Then, you see the main VSI window, as shown in Figure 1-11. The status sequence of initially Build, Warning, and finally Active is used.

Name	IPs	Operating system	Cores	Memory	Status
Power_H2O	192.168.138.131	rhel	1 cores	8 GB	✔ Active <div style="border: 1px solid orange; padding: 2px; display: inline-block; margin-top: 5px;"> Build, Warning, Active </div>

Figure 1-11 VSI status

Now that you have your VSI running, you see an Actions icon included at the end of the row. By clicking the icon, a pop-down menu opens in which you can shut down, restart, or open a console, or delete your instance, as shown in Figure 1-12. If you click the open console, you a new window opens that includes the login prompt for your system.

Name	IPs	Operating system	Cores	Memory	Status
Power_H2O	192.168.138.131	rhel	1 cores	8 GB	Active
powerH2O	192.168.138.130	rhel	1 cores	8 GB	

Items per page: 10 ▾ Items 1-2 of 2

1 ▾

- OS shutdown
- Immediate shutd...
- Restart
- Open console
- Delete

```

Red Hat Enterprise Linux 8.1 (Ootpa)
Kernel 4.18.0-147.44.1.el8_1.ppc64le on an ppc64le
Web console: https://power-h2o:9090/ or https://192.168.138.131:9090/
power-h2o login: █
  
```

Figure 1-12 Open VSI console

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



Migration to Cloud with IBM Power Systems Virtual Server

This chapter provides information about how to migrate IBM i LPARs on customer's data center to virtual machines (VMs) on IBM Power Systems Virtual Server.

A migration strategy is the most important step in our migration process. Every customer has its own business needs, and the correct strategy can drive the migration project to success.

In this chapter, we discuss the different tools and techniques we can use to transfer our IBM i workloads and the requirements, challenges, and advantages on each method.

This chapter includes the following topics:

- ▶ 2.1, "IBM i and data migration" on page 14.
- ▶ 2.2, "Backup and Restore migration" on page 16.

2.1 IBM i and data migration

When moving workloads to a new system, many issues must be considered based on business rules, technical resources, and the environment.

IBM Cloud Power Virtual Servers features specific requirements and challenges based on its online and self-service nature. The following IBM i constraints also must be considered:

- ▶ We have no access to the hardware layer. No HMC or Storage direct access is provided; therefore, hardware replication cannot be used.
- ▶ We have no access to the data center. A physical tape cannot be sent to perform a restore.
- ▶ No physical tape is available for backup or restore operations.
- ▶ Your system is in a remote facility and communications can have some latency, depending in your location and link type.
- ▶ Most IBM i backup devices use SAS or Fibre Channel connections with a proven reliability and fast transfer rates. Network backups are slow, and even backups to disk can look slow when compared with physical or virtual tape.
- ▶ The amount of third-party backup software for IBM i is not so wide, and a small portion of them use network to transfer data.

At times, you must consider the following business and technical constraints:

- ▶ You do not have 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 shops rely on 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.

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 keep running (see Figure 2-1).

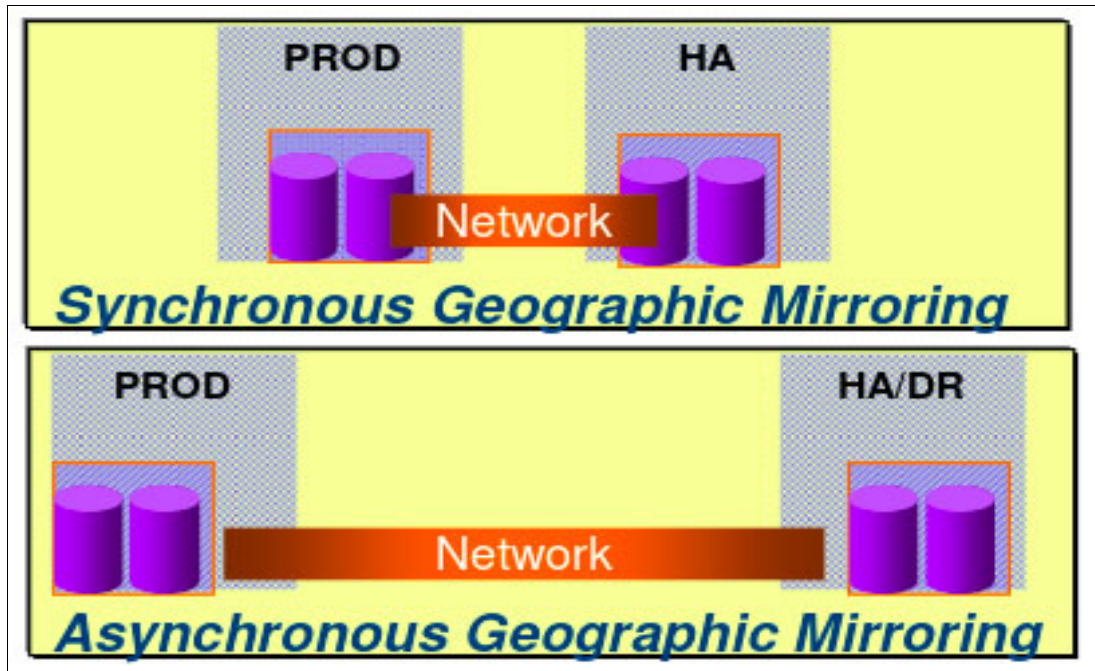


Figure 2-1 From PowerHA SystemMirror for i

But advantages are gained when the instance is running on the IBM Power Systems Virtual Server environment:

- ▶ 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.
- ▶ Instant snapshots can be taken from your IBM i instances.
- ▶ Clones can be taken from our 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.
- ▶ 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 Gb.
 - Most of these features are transparent to the user.
- ▶ A new 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.

- ▶ When all of the data cannot be uploaded to the migration window, IBM Cloud can provide the Mass Data Migration Device, which is a storage device with up to 120 TB and multiple alternatives to save data.

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.

2.2 Backup and Restore migration

Most IBM i shops use and rely on physical tapes for backups of day-to-day operations, and migrating workloads to a new system is not an exception.

Tapes are the main migration method when moving data (and even the operating system) from the system to a new box.

Unfortunately, physical tapes cannot be used with cloud environments; therefore, alternative methods must be used, such as saving data to disk and the network for transfer.

When data is moved to the cloud, data must be uploaded by using the network. At times, the target systems can be accessed, and an intermediate “bucket” is used to reach the target environment. 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 improve the process by reducing its size.

Virtual media can be very large, depending on the amount of data that must be transferred. Consider choosing the correct size for your virtual media so that you can use parallel file transfers when possible and improve transfer times.

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 latest IBM i releases.

For more information about virtual media capacities, see this [IBM Support web page](#).

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 a many features to automate your backup and restore process.

In this section, we discuss how can we migrate data by using the different methods:

- ▶ Full System SAVE and RESTORE:
 - BRMS and IBM i Cloud Storage Solutions
 - Native GO SAVE and GO RESTORE menu and options
- ▶ Partial restores
- ▶ Back up your system when not enough space is available

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 our backups.

A software option is available that is called IBM Aspera®, which improves data transfers with IBM Cloud Object Storage and upload backups by using multiple streams.

When the amount of data to be transferred makes it impossible 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 IBM Cloud network, which avoids direct file transfers from our data centers.

2.2.1 Migration with partial saves and restores

Partial restores can consist of multiple save files, ISO images with virtual optical media, or virtual tape images with which a system can be backed up in parts and then data is 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 that an empty VM instance is created with a working IBM i operating system. Partial restores can be a good option to migrate our LPARs when the operating system images that are provided on IBM Cloud fit a user's needs.

This scenario is easier because the target system already is running. By using a working IP address, communications can be checked. Often, images include the latest resaves and PTFs and some of the following software that is needed to run workloads in the IBM Power Systems Virtual Server are in place (see Figure 2-2):

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

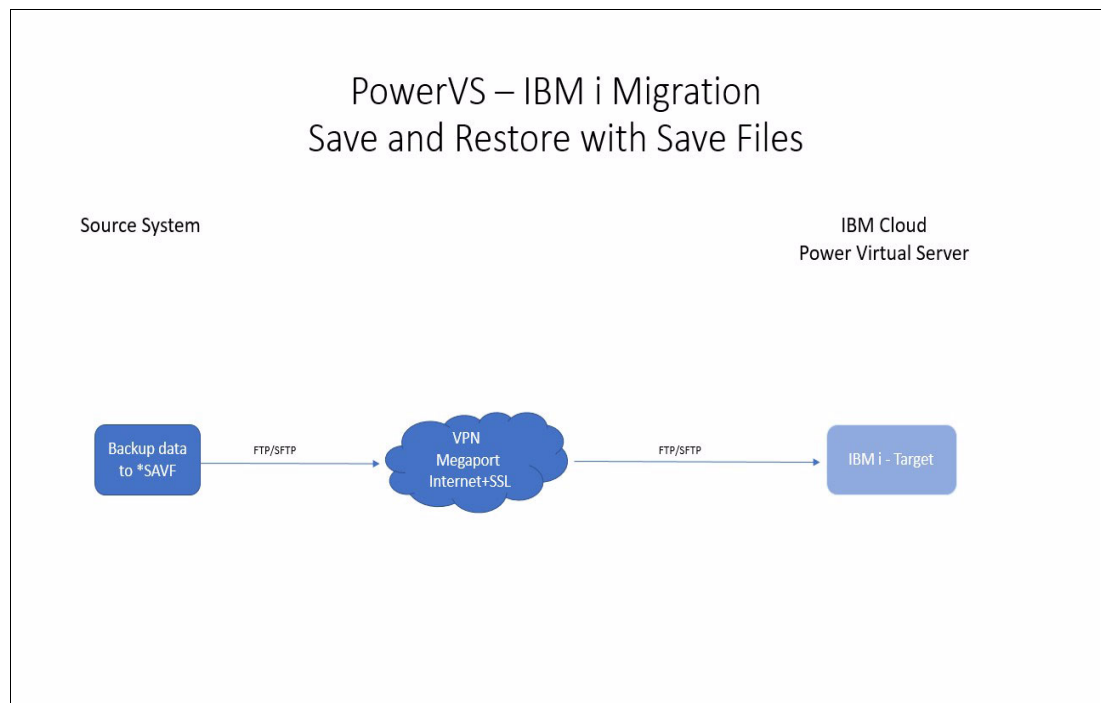


Figure 2-2 Save and restore by using save files

The sample backup and restore scenario that is shown in Figure 2-2 on page 17 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: Correct sizing is recommended. Request such a study from IBM or an IBM Business Partner.

3. Communications must be set up between the local network and IBM Cloud so that the system that was created from our source data center can be reached.
4. Some changes are made on the target server by using the console:

```
CHGSYSVAL SYSVAL(QCTLSBSD) VALUE('QCTL')
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)
```

5. Start *FTP or *SSHD servers to allow file transfer:

```
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 the following example:

```
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)
```

Here, all of the libraries that are needed to migrate are saved:

```
CRTSAVF BKPIBMI/SAVDL01
SAVDLO DLO(*ALL) DEV(*SAVF) SAVF(BKPIBMI/SAVDL01)
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.

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

10. End the subsystems:

```
ENDSBS *ALL *IMMED
```

11. Restore the user profiles and security:

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

12. Restore the configurations:

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

13. Restore all your libraries, one by one:

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

14. 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.

15. Restore the documents:

```
RSTDLO DLO(*ALL) DEV(*SAVF) SAVF(BKPIBMI/SAVDL01) ALWBJDIF(*ALL) SAVASP(*ANY)
```

16. 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)
```

17. Restore the authorities:

```
RSTAUT *ALL
```

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

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

20. Adjust the net name:

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

21. IPL the target system:

```
PWRDWN SYS *IMMED RESTART(*YES) CONFIRM(*NO)
```

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

2.2.2 Full system backup and restore by using native commands

Full system backup saves all of the operating system's configurations, security, and data to media to perform a full system restore from scratch 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 procedures.

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 Cloud 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 operating systems are supported in IBM Cloud Power Virtual Server service:

- ▶ V7R1 with extended support
- ▶ V7R2 with support (but was commercially withdrawn)
- ▶ V7R3 with regular support
- ▶ V7R4 with regular support

An IBM i instance must be created. For more information, see 1.2, “Creating a Power Systems Virtual Server” on page 3.

When restore must be done from this virtual media, special requirements that are based on cloud characteristics are important. The new 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 was used:

- ▶ Target system: IBM Power Systems Virtual Server with IBM i and enough resources for data and workloads.
- ▶ Network installation server: IBM Power Systems Virtual Server with IBM i. At least 1 core and 32 GB, and enough disk Tier1 for backup images.
- ▶ Communications between source and target system: Often, a Direct Link Connect between Power Virtual Servers and IBM Cloud VSI environment, and a Virtual Router Appliance to route traffic between VLANs and enable 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. We recommend the use of any version of Linux (Red Hat Enterprise Linux, Centos, Fedora, SUSE, Debian, and so on). The example that is presented here was made by using Ubuntu Server.
- ▶ A private network on the IBM Power Systems Virtual Server that connected to both servers.

Important: A ticket with IBM Support is needed to make the private network VLAN available so that every system that is connected to this VLAN can see each other.

- ▶ Source System: Enough available disk for backup.

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

Setup process

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

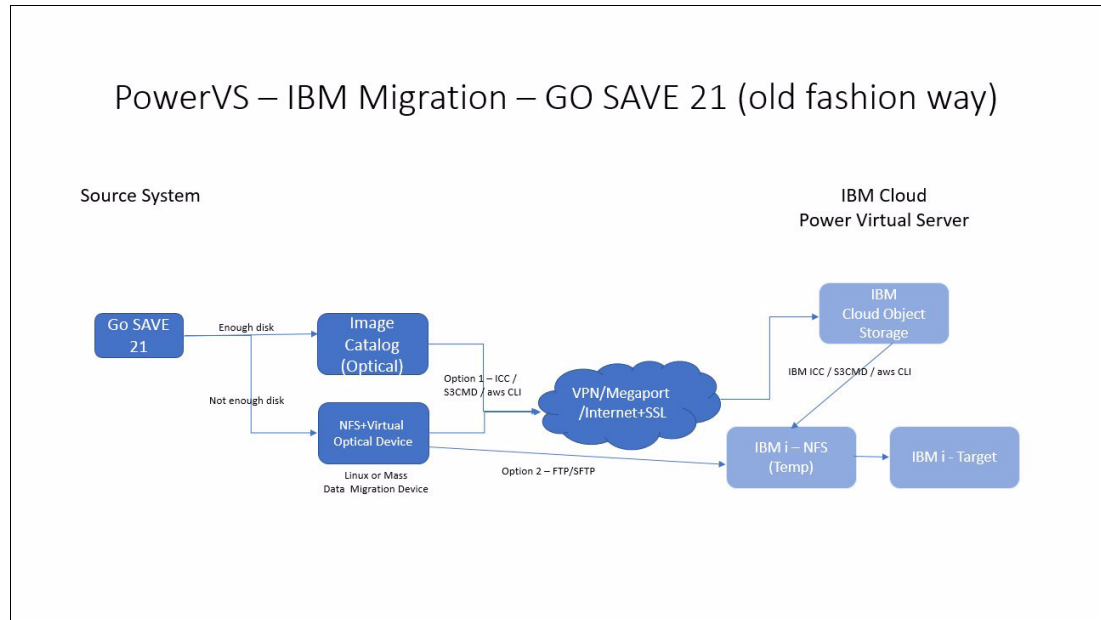


Figure 2-3 IBM i Migration that uses GO SAVE option 21

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

For example purposes, 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 disk space is less than 50%, a full system backup can be performed to the 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. We recommend the use of more than two physical cores and SSDs or SAN-attached storage when possible.
- ▶ The NFS directions are based on this [IBM Support web page](#).
- ▶ IBM i can back up to an NFS server from V7R2 and later. Previous versions also can back up in this way, but with limitations. Our test environment was validated by using V7R2 and V7R4.

- ▶ We recommend installing 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 also must be installed to communicate with IBM Cloud Power Virtual Server installation process. For more information, see the following resources:
 - This [IBM Support web page](#)
 - [RPM pile for IBM i releases in standard support](#)
- ▶ Linux and IBM i must see each other in a 1 GB network or faster connection between them.

Complete the following steps:

1. Install the following tools on your Linux partition (see Example 2-1):
 - Python3
 - NFS Server
 - `awscli` (AWS Client works with most S3 servers)
 - PIGZ

Example 2-1 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. Our source network is IBMi04 with an IP address of 192.168.50.244 and Linux 192.168.50.22, as shown in Example 2-2.

Example 2-2 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 2-3.

Example 2-3 Restart NFS

```
# Restart NFS service
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 to the IP address must be in the same network as Linux NFS server:

```
CFGTCP
Option 1
F11
```

- This command shows our network interface and the Line Description name, as shown in the following example:

```
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).

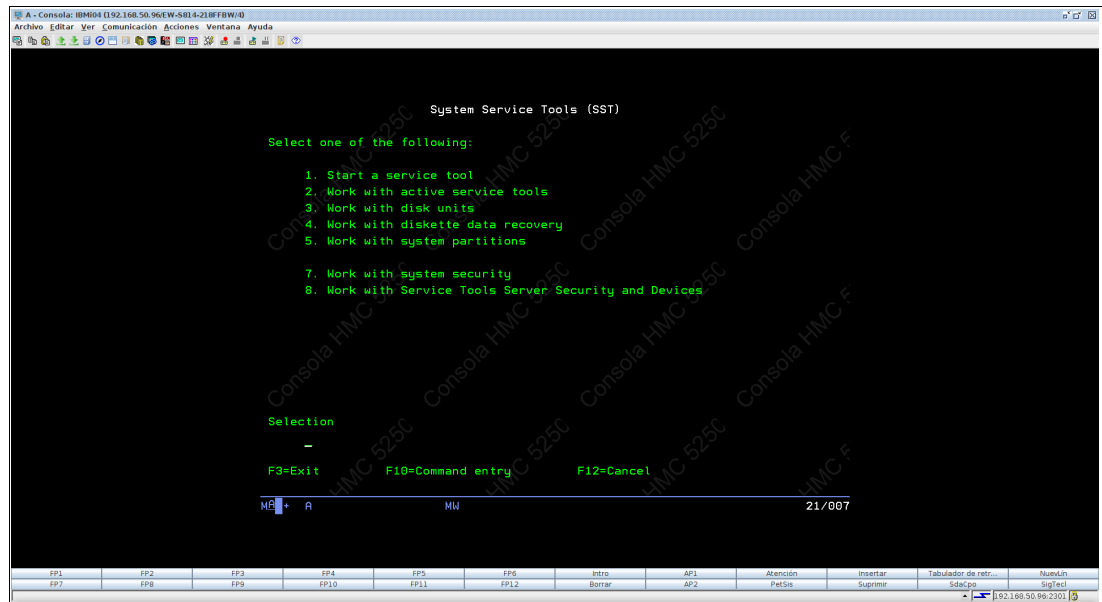


Figure 2-4 SST Main Menu

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

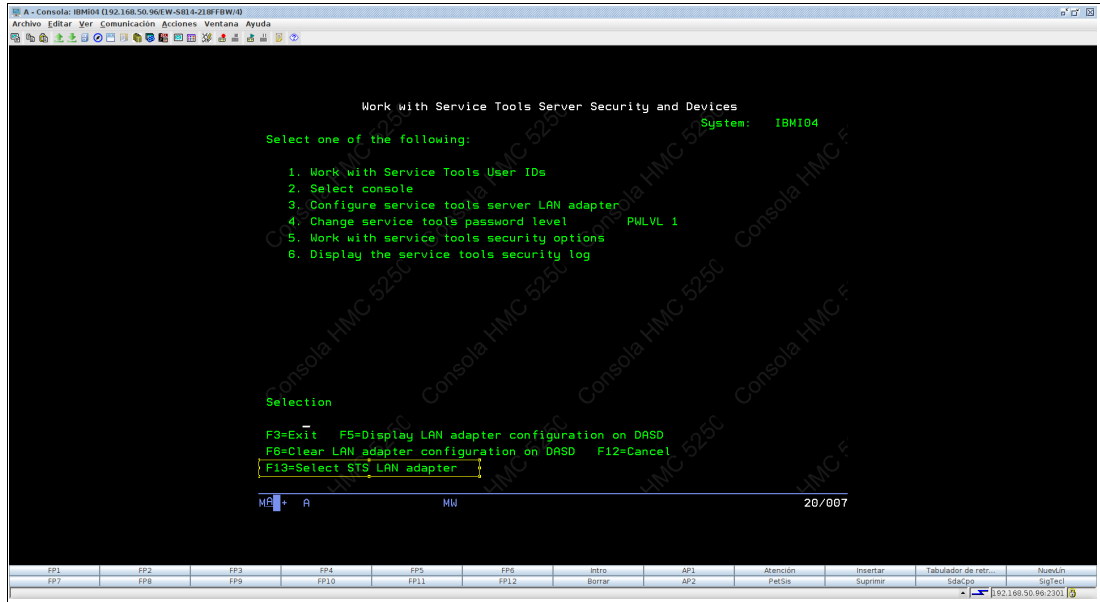


Figure 2-5 SST Work with Service Tools Server Security and Devices

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

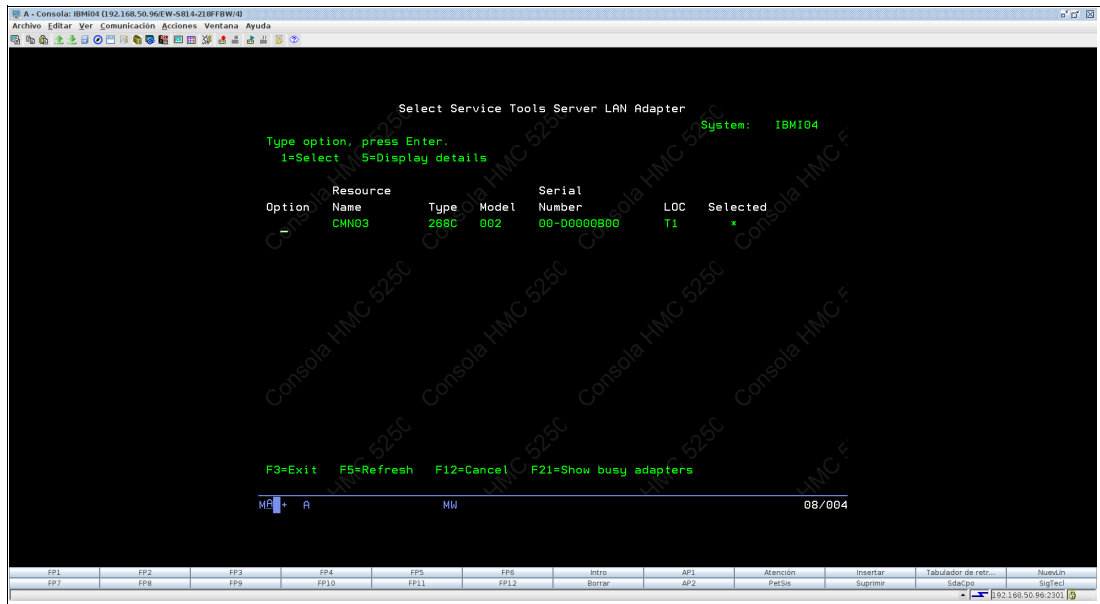


Figure 2-6 Select Service Tools Server LAN Adapter

9. As shown in Figure 2-7 on page 25, set the host name, 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 then, Activate the LAN adapter by pressing **F14**.
11. Press **F3** to exit Service Tools.

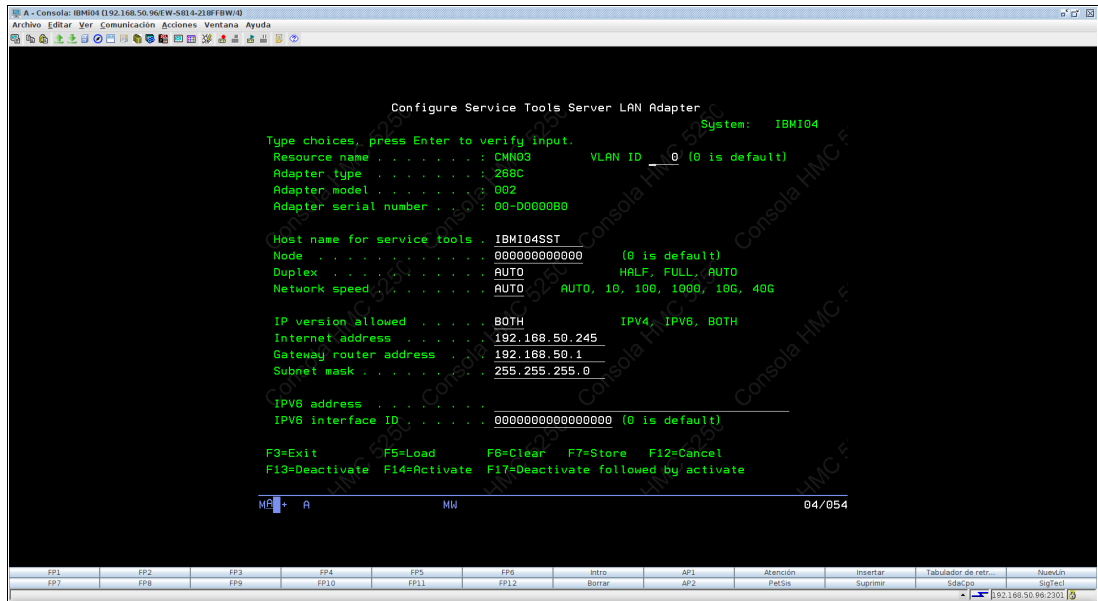


Figure 2-7 Configure Service Tools Server LAN Adapter

12. Restart the system to take the changes:

```
PWRDWN SYS *IMMED RESTART(*YES) CONFIRM(*NO)
```

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

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

14. Mount the NFS share and start creating the virtual media, as shown in Example 2-4.

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

Example 2-4 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 2-5.

Example 2-5 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 2-6.

Example 2-6 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 the, 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 set up the `awscli`. For more information, see [this web page](#).
- Create your Object Storage resource and bucket, as shown in Figure 2-8, Figure 2-9 on page 27, Figure 2-10 on page 27, Figure 2-11 on page 27, Figure 2-12 on page 28, and Figure 2-13 on page 28.

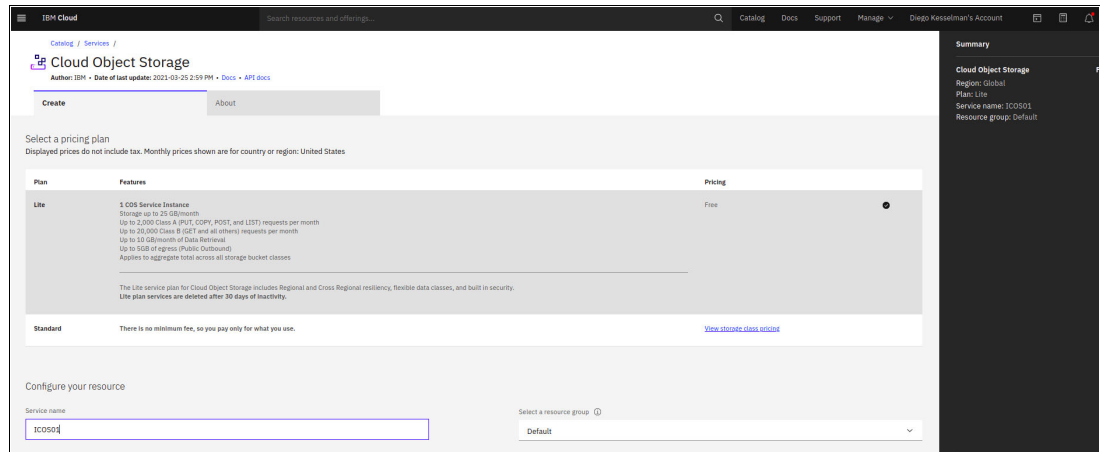


Figure 2-8 Select the resource type: Lite or Standard and create the resource




Figure 2-9 Cloud Object Storage main window

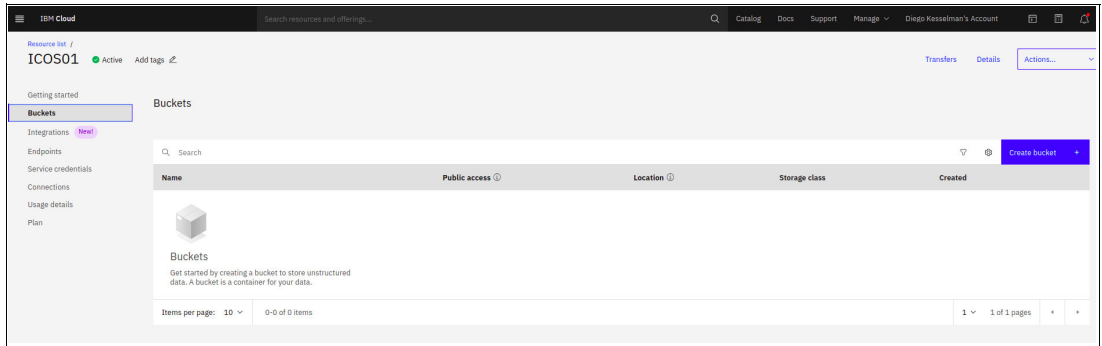


Figure 2-10 Select Bucket and Create the bucket

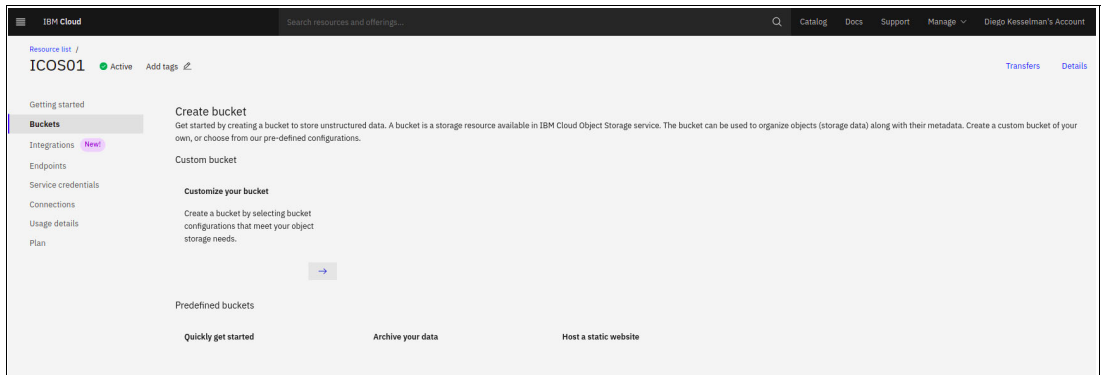


Figure 2-11 Create your bucket

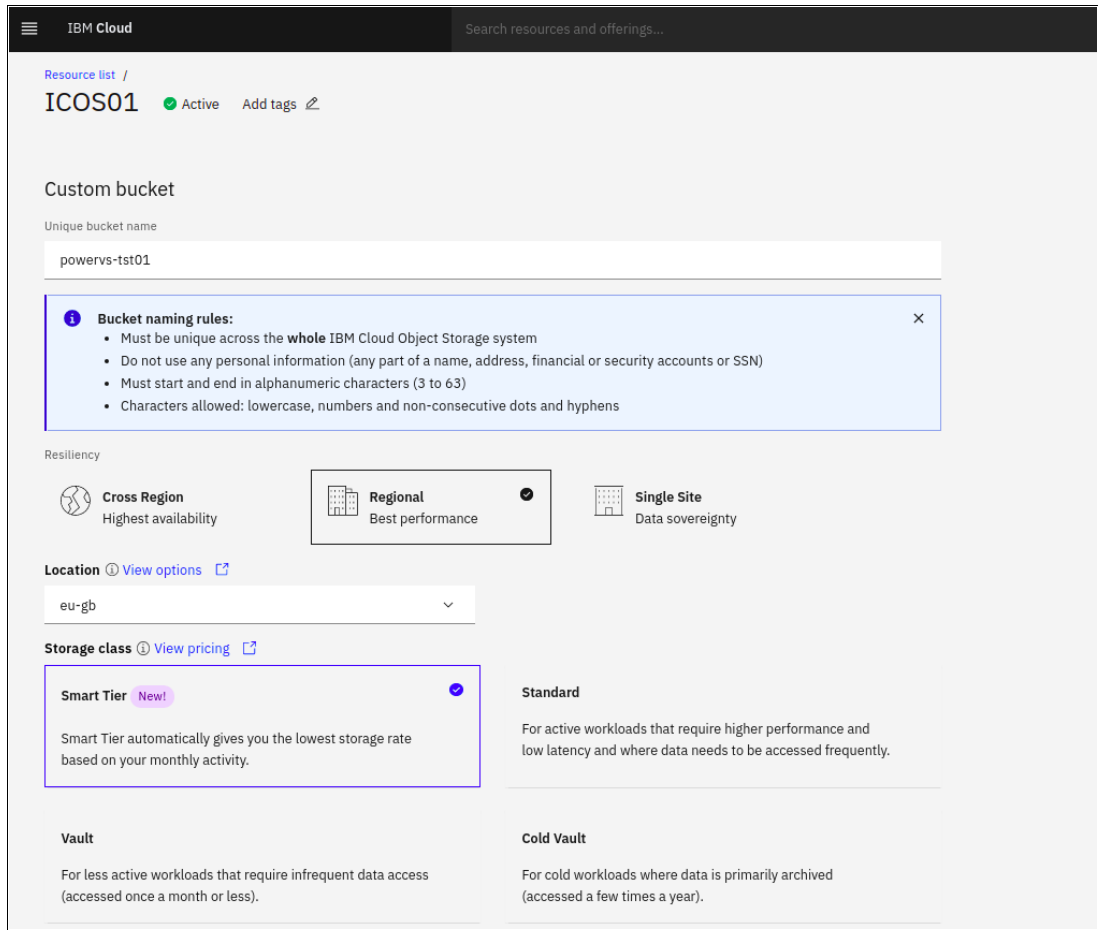


Figure 2-12 Select bucket configuration - Smart Tier is the easy way

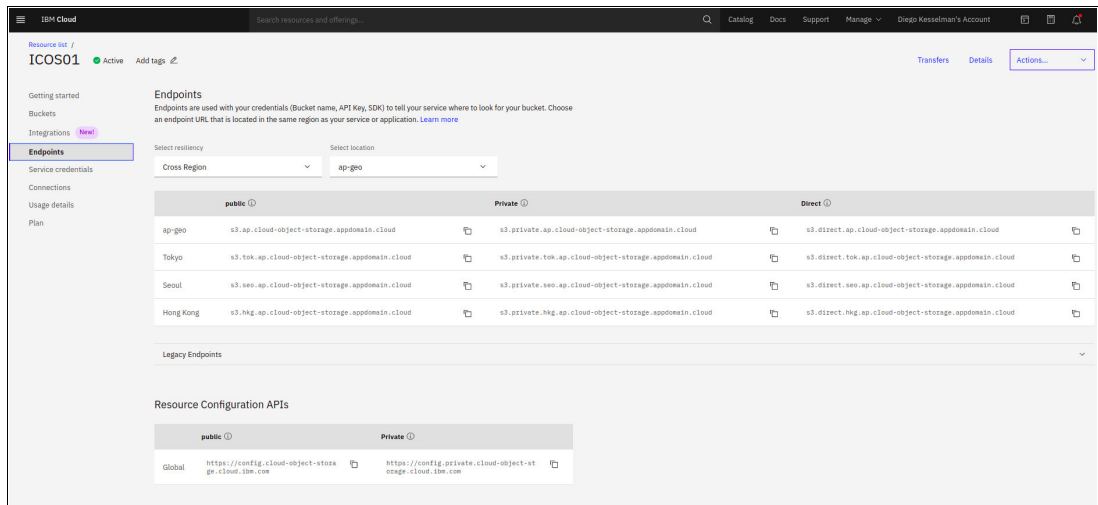


Figure 2-13 The endpoint option shows URLs

22. Take note of the endpoint URI. When connecting by using VPN, select **Private**. When connecting the internet, use **Public** (see Figure 2-14, Figure 2-15, Figure 2-16, and Figure 2-17 on page 30).



Figure 2-14 Service Credentials - Create Credential

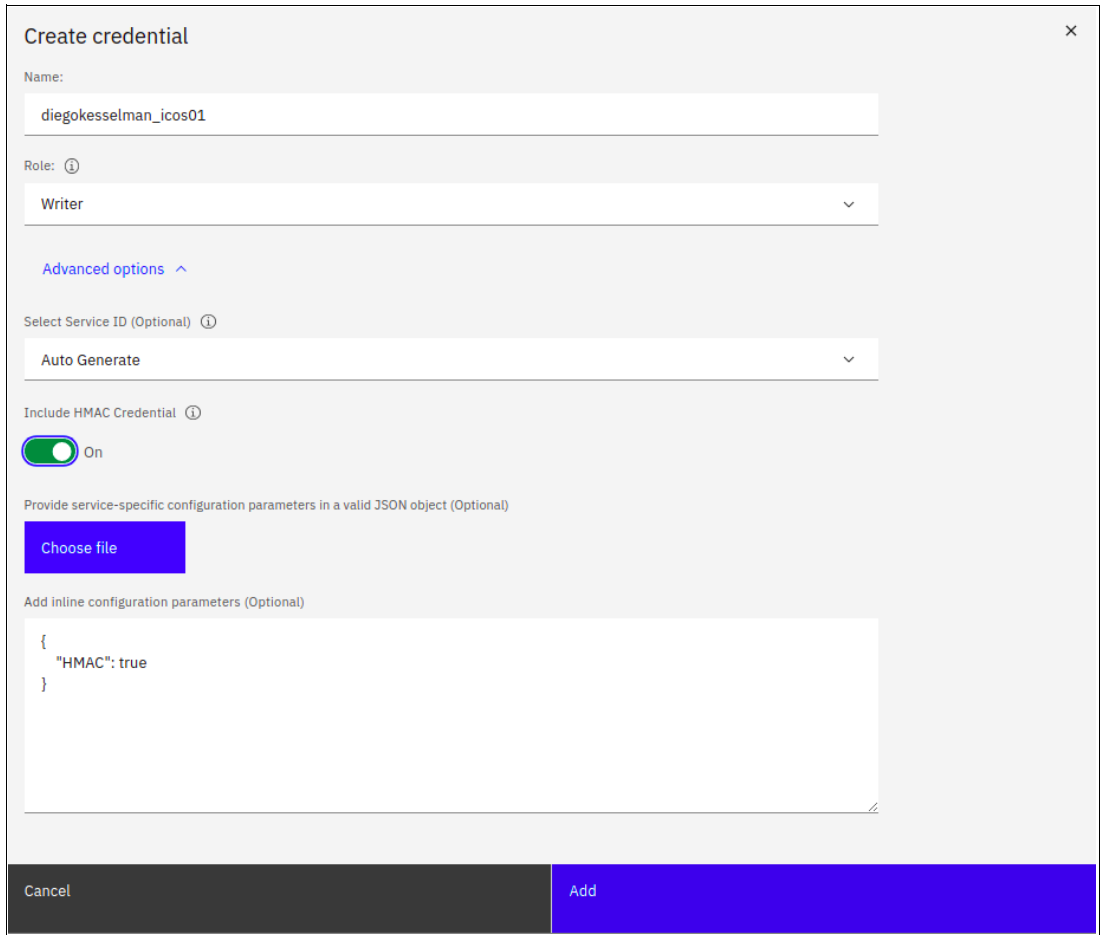


Figure 2-15 Create Credential - Advanced Option - Include HMAC Credential

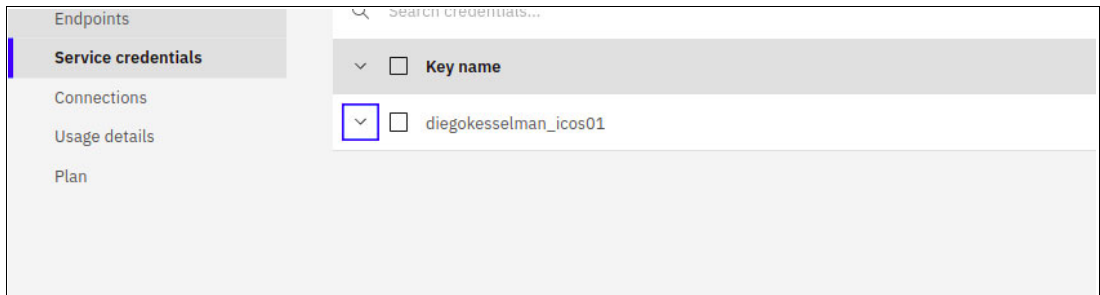


Figure 2-16 List new credential

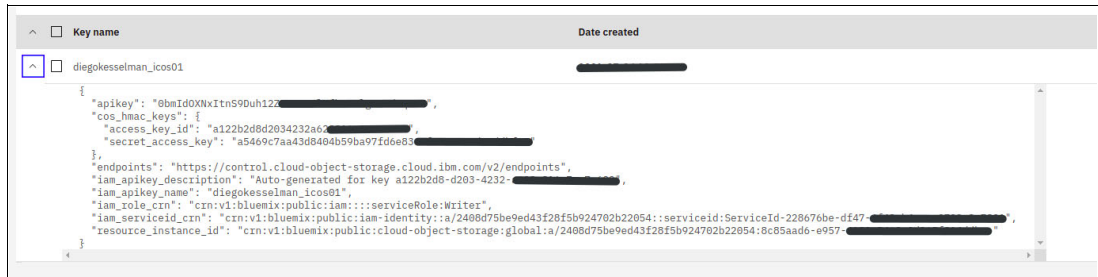


Figure 2-17 Expand credential details

Take note of the Access Key and the Secret Key.

23. Return to Linux:

```
aws configure
```

- Use your credentials to log in.
- For this example, the region is eu-gb and the default output format is **None**.

24. Start compressing and uploading the image files. To avoid the use of more disk space, compress the files as they are being uploaded, as shown in Example 2-7.

Example 2-7 Compressing and uploading image files

```

cd /NFS/IBMi04
cat IMAGE01.ISO | pigz -9 -p40 -c | aws \
--endpoint-url=https://s3.eu-gb.cloud-object-storage.appdomain.cloud s3 cp - \
s3://ibmi-backup/IMAGE01.ISO.gz

cat IMAGE02.ISO | pigz -9 -p40 -c | aws \
--endpoint-url=https://s3.eu-gb.cloud-object-storage.appdomain.cloud s3 cp - \
s3://ibmi-backup/IMAGE02.ISO.gz

cat IMAGE03.ISO | pigz -9 -p40 -c | aws \
--endpoint-url=https://s3.eu-gb.cloud-object-storage.appdomain.cloud s3 cp - \
s3://ibmi-backup/IMAGE03.ISO.gz

cat IMAGE04.ISO | pigz -9 -p40 -c | aws \
--endpoint-url=https://s3.eu-gb.cloud-object-storage.appdomain.cloud s3 cp - \
s3://ibmi-backup/IMAGE04.ISO.gz

```

25. After the process completes, verify whether all files were uploaded to the Cloud Object Storage portal, as shown in Figure 2-18.

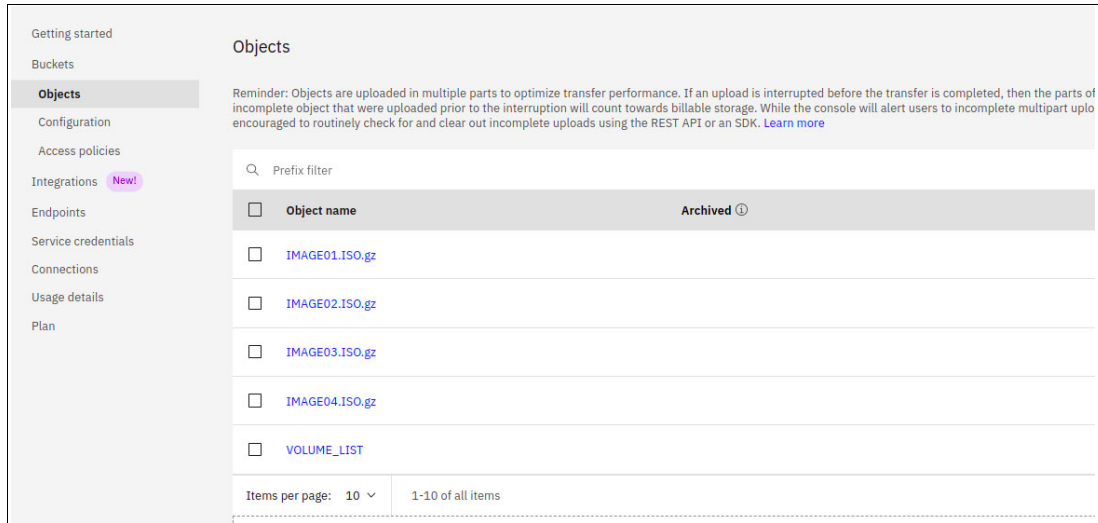


Figure 2-18 View objects in bucket

26. In the command line, use `awscli` and the `ls` command, as shown in Example 2-8.

Example 2-8 Another way using Linux commands

```
aws --endpoint-url=https://s3.eu-gb.cloud-object-storage.appdomain.cloud s3 ls s3://ibmi-backup/
```

27. Return to the IBM Cloud Power Virtual Server environment.

28. Prepare the NFS server and the target system, both running IBM i.

29. 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 save. Consider the following points:
 - This process can be done by using a Direct Link Connect and a reverse proxy. For more information, see this [IBM Systems Lab Services Tutorial](#).
 - The backups also can be downloaded by using the internet and the public access on IBM Cloud Object Storage. however, charges might be incurred on a public egress.

Restore process

Complete the following steps to restore the instance:

1. On the TEMPORARY NFS instance update YUM, as shown in Example 2-9.

Example 2-9 Updating YUM on PASE

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

2. Install some tools by using the ACS or command line, as shown in Example 2-10:

```
pipz gzip gunzip python3 python3-pip
```

Example 2-10 Install other tools from Access Client solutions

```
CALL QP2TERM
```

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

```
yum -y install pipz gzip gunzip python3 python3-pip
```

```
pip3 install s3cmd
pip3 install awscli
```

```
# Then you can configure awscli
aws configure
```

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

3. Download the compressed images, as shown in Example 2-11.

Example 2-11 Downloading compressed image files

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

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

4. Decompress the images, as shown in Example 2-12.

Example 2-12 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) DTAUT(*RWX) SUBTREE(*ALL)
CHGAUT OBJ('/NFS/RESTORE21') USER(QTFTP) DTAUT(*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:


```
CHGTFTP AUTOSTART(*YES) ALTSRCDIR('/NFS/RESTORE21')
```
 - i. Restart the TFTP server:


```
ENDTCPSVR SERVER(*TFTP)
STRTCPSVR SERVER(*TFTP)
```
6. Complete the following steps to set up the environment In the TARGET instance:
 - a. 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.
 - b. Add the NFS instance to the host table:


```
ADDTCPHTE INTNETADR('192.168.80.12') HOSTNAME((IBMiNFS))
```
 - c. 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/*'
```
 - d. 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)
```
 - e. IPL the instance so that the changes can take effect.

- f. Write the following information in your notes (see Figure 2-19, Figure 2-20, and Figure 2-21 on page 35):
- TCP/IP information: IP Address, Gateway, and Hostnames

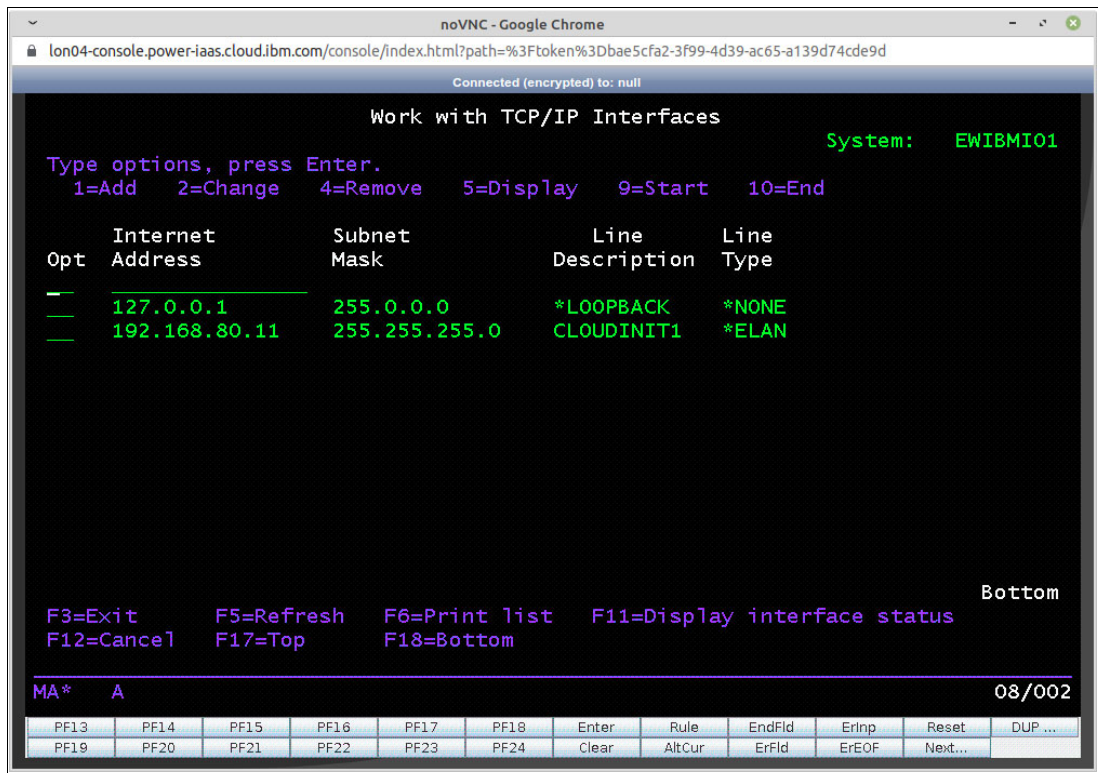


Figure 2-19 Get IP address and associated line description

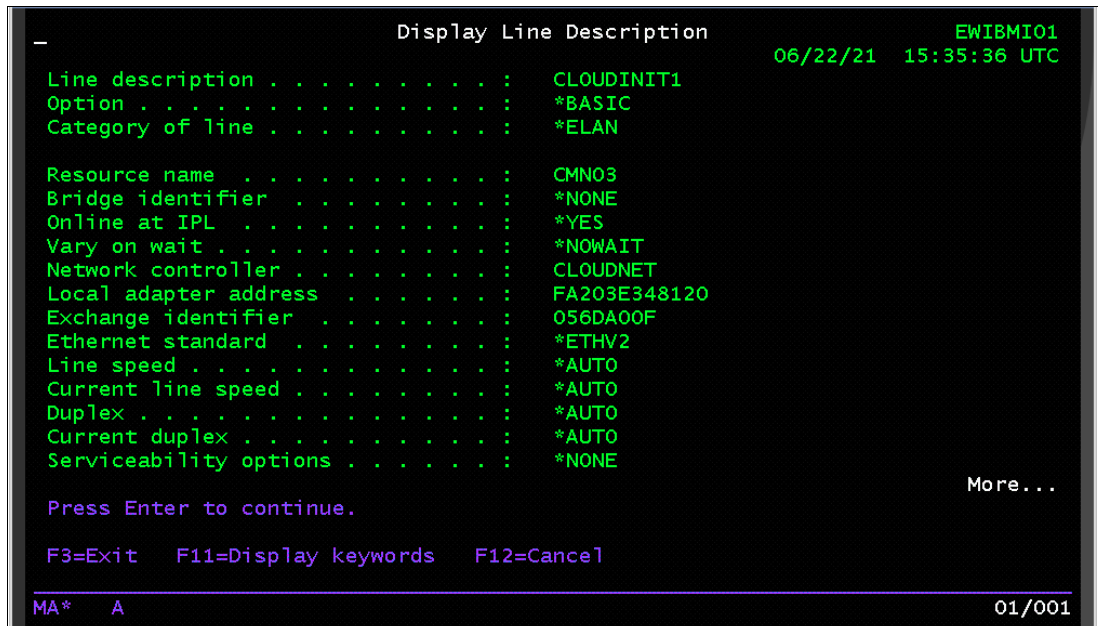


Figure 2-20 Display line description details and get the Resource Name

```

-                               Display Resource Detail                               System:  EWIBMI01
Resource name . . . . . : CMN03
Text . . . . . : Ethernet Port
Type-model . . . . . : 268C-002
Serial number . . . . . : 00-00000
Part number . . . . . :

Location :  U9009.22A.78A2250-V6-C2-T1

Logical address:
SPD bus:
System bus                255
System board              128
System card                2

Press Enter to continue.

F3=Exit  F5=Refresh  F6=Print  F12=Cancel

MA*  A                                                                    01/001

```

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

- License keys: backup licenses to a file and transfer to NFS server:
 - DSPLICKEY OUTPUT(*LICKKEYFILE) LICKKEYFILE(QGPL/LIC01)
 - CRTSAVG QGPL/LIC
 - SAVOBJ OBJ(LIC01) LIB(QGPL) DEV(*SAVF) SAVF(QGPL/LIC)
- g. Transfer the license keys to the NFS instance.
- h. Start the network installation (see Figure 2-22):
 - STRNETINS DEV(NFSRESTORE) OPTION(*LIC) KEYLCKMOD(*MANUAL)

```

Start Network Install (STRNETINS)

Type choices, press Enter.

Network optical device . . . . . > NFSRESTORE      Name
Installation option . . . . . > *LIC              *ALL, *OPSYS, *LNG, *LIC
Keylock mode . . . . . > *MANUAL                *NORMAL, *MANUAL
Language . . . . . > *DFT                       2901-2999, *DFT
Confirm . . . . . > *YES                         *YES, *NO

                                                                    Bottom
F3=Exit  F4=Prompt  F5=Refresh  F12=Cancel  F13=How to use this display
F24=More keys

MA*  A                                                                    08/037

```

Figure 2-22 STRNETINS command

- i. Click **Finish** or confirm the command at TARGET and the IPL process starts (a connection must be established to the console for manual operation).
This process can take several minutes until the DST window is displayed.
- j. From the DST window, complete the following steps to install the License Machine Code (SLIC):
 - i. On the Install License Internal Code Menu, choose **Option 2**, and confirm by pressing **F10**.
The process to initialize the load source disk begins.
 - ii. Select **Install Operating System**. In this example, select the installation device type **Option 5 - Network device**.
 - iii. Press **F10**. The Network Device - Configuration menu is displayed (see Figure 2-23).

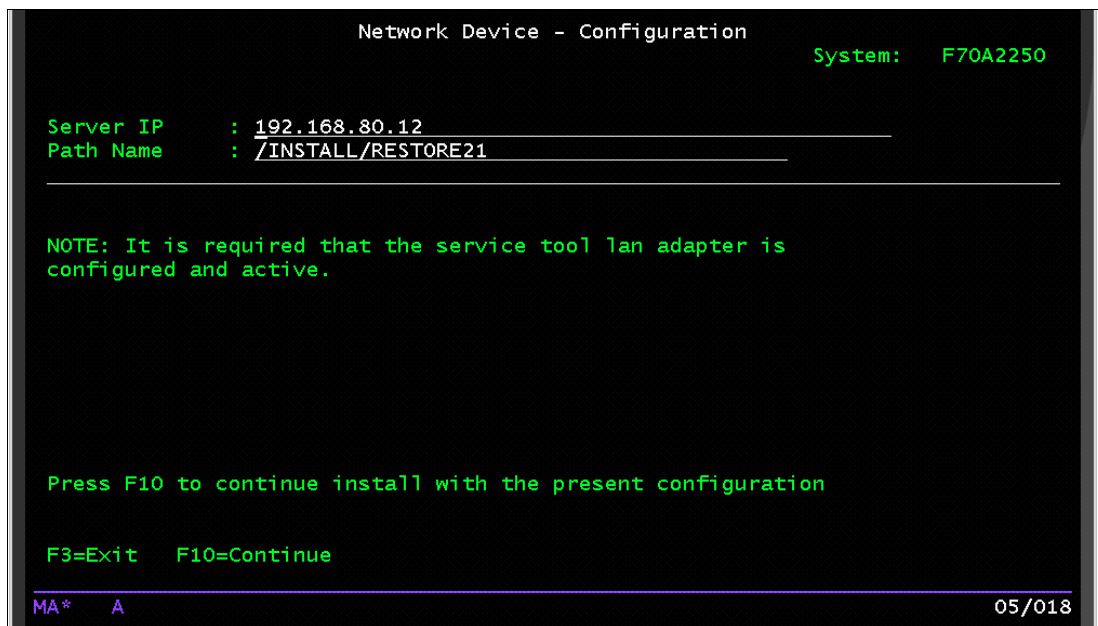


Figure 2-23 Confirm your NFS instance values and continue with F10

- iv. Confirm the installation of the operating system and then, press **Enter**.
- v. Select a language group from the Select a Language Group Menu, (in this example, **2924** is selected).
- vi. If disks must be added, use the Add All Disk Units to the System menu. In our example, option 1 is selected: **Keep the current disk configuration**.

- vii. 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).

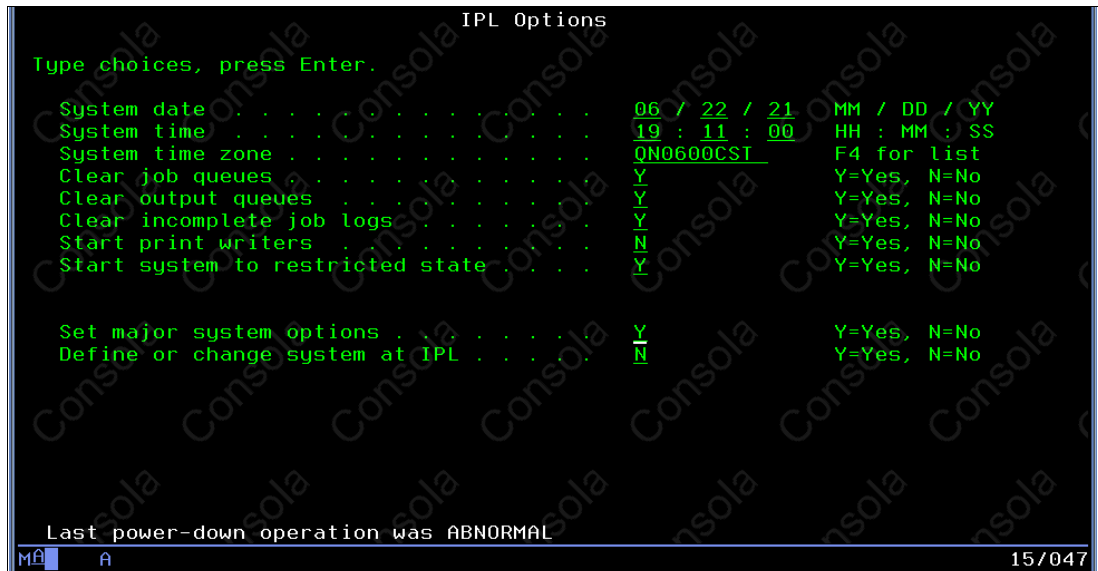


Figure 2-24 Set IPL Options as shown in this image.

- viii. 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.

- ix. After the operating system installation, you are ready to sign on; however, the password must be changed to continue the sign-on request.

7. Continue with setup and restore process:

- a. Check 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) RSRNAME(CMN03)
CRTLINETH LIND(ETH02) RSRNAME(CMN04)
```

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

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

- d. Vary on the Virtual Optical:

```
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 device name.

8. After the restore process completes, complete the following steps to check the messages and system:

a. Check the software products:

```
GO LICPGM
Option 10
```

b. Change system values to fit the new environment:

```
WRKSYSVAL
```

c. Change the network name:

```
CHGNETA SYSNAME(IBMIO4) LCLCPNAME(IBMIO4) LCLLOCNAME(IBMIO4) DTACPR(*ALLOW)
DTACPRINM(*REQUEST) NETSERVER((IBMIO4 *ANY)) NWSDOMAIN(IBMIO4)
ALWVRTAPPN(*YES) ALWHPRTWR(*YES)
```

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

```
CFGTCP
```

e. On the NFS instance, start the FTP service (it is down by default):

```
STRTCPSVR *FTP
```

f. Restore the license keys, as shown in Example 2-13.

Example 2-13 Restoring licensing keys

```
#From 5250 terminal
CRTSAVF QGPL/LIC
#From PASE
FTP '192.168.80.12'
bin
cd QGPL
get LIC (REPLA
quit
# From 5250 terminal
RSTOBJ OBJ(*ALL) SAVLIB(QGPL) DEV(*SAVF) SAVF(QGPL/LIC)
ADDLICENSE LICKEYINP(*LICENSEFILE) LICENSEFILE(QGPL/LIC01)
```

g. Validate license keys:

```
DSPLICKEY
```

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

i. Restart the TARGET instance:

```
PWRDWN SYS *IMMED RESTART(*YES) CONFIRM(*NO)
```

Now, the IBM i TARGET system is 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.

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

This section describes how to perform a save and transfer of 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.

We continue showing how to restore the previously saved IBM i workload in the Power Systems Virtual Server environment.

Solution components and requirements

This section describes the solutions components and requirements.

Required components on the IBM i Local Source LPAR/VM

The following components must be set up 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

The BRMS and IBM Cloud Storage software must be installed. 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

The following program temporary fixes (PTFs) must be installed 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.
- ▶ License Keys can take 5 minutes to appear post deployment:
 - 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 about the minimum required software levels, see this [IBM Cloud Docs web page](#).

- ▶ The following BRMS PTFs must be installed on a stand-alone BRMS system or on all of the systems in a BRMS network:
 - 7.3 - SI61153
 - 7.2 - SI61152
 - 7.1 - SI61151
- ▶ The SI73401 - Cloud Storage Solutions Proxy Support 5733ICC IBM Cloud Storage Solutions for i PTF must be installed 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 post 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 this [IBM Support web page](#).

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 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 required lowest version.
- ▶ The rpm packages can be installed by using yum. For more information, see [this web page](#).

3. Install the following required PTFs and PTF groups:

- 7.2: SF99713, MF61899 (permanently apply), SI62642, SI63163, MF65261, MF65218, and SI72162
- 7.3: MF61942 (permanently apply), SI62642, SI63041, MF63830, and SI72161
- 7.4: SI72142

4. Save and migrate your system.

Note: For IBM i 7.2 and 7.3 on POWER8® and later version, cloud-init is enabled automatically.

For IBM i 7.2 and 7.3 or other IBM i versions that are running on POWER7 or later Power Systems servers, the process that is described in “Installing cloud-init on IBM i” on page 40 *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 machine manually, run this on the virtual machine, as shown in Example 2-14.

Example 2-14 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; then, it is ready to save and migrate.

Downloading the required software

Complete the following steps to download the required software (for more information, see this [IBM Support web page](#)):

1. Click **My Entitled Systems Support** (see Figure 2-25 on page 42).

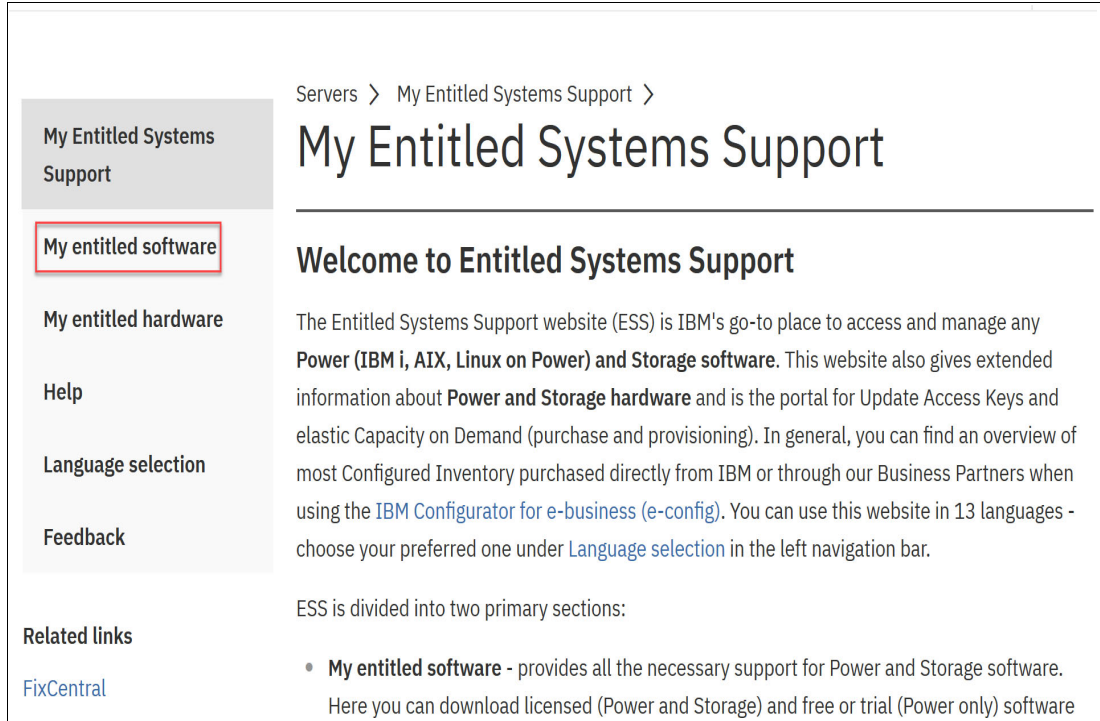


Figure 2-25 Entitled Systems Support (ESS) portal

2. The entitled software downloads are shown in Figure 2-26. Click **IBM i evaluation**. For more information, see [this web page](#).



Figure 2-26 Software downloads option

3. Select a category and group, as shown in Figure 2-27.

The screenshot shows a web interface titled "Software downloads". Below the title is a paragraph of text: "The software available for download is based on all active entitlements associated with all customer numbers registered under your profile. Check [Help](#) for a list of supported product". Below this is a section titled "Step 1: Select a category and group". It contains two dropdown menus: "Category:" with "IBM i" selected, and "Group:" with "7.3 (GA)" selected. At the bottom left of this section is a black "Continue" button.

Figure 2-27 Selecting IBM i and release category

4. Click IBM i product **5770-SS1**. Then, click **Continue**, as shown in Figure 2-28.

The screenshot shows a web interface titled "Step 1: Select product". It has radio buttons for "Entitled" (selected) and "All", and a "Search:" input field. Below is a table with columns "Product" and "Title". The first row is selected, indicated by a red arrow pointing to the checkbox. At the bottom are "Continue" and "Back" buttons, with a red arrow pointing to the "Continue" button.

<input type="checkbox"/>	Product	Title
<input checked="" type="checkbox"/>	5770-SS1	IBM i
<input type="checkbox"/>	5722-WE2	Web Enablement for i5/OS
<input type="checkbox"/>	5770-QU1	Query for i

Figure 2-28 Selecting 5770-SS1

5. Select the main language. Then, click **Continue**, as shown in Figure 2-29.

Software downloads

Start | Step 1: Product | **Step 2: Language** ▶ Step 3: Package | Step 4: T&Cs | Step 5: Method | Step 6: Download

Step 2: Language selection ⓘ

Main language: Lang Grp 1 Americas/Europe ▼

Product	Language
5770-SS1, IBM i	Lang Grp 1 Americas/Europe ▼

Continue **Back**

Figure 2-29 Entitled software downloads - Select language and press Continue

6. Select the package to download. Then, click **Continue**, as shown in Figure 2-30.

Step 3: Product download package selection

Current products ⓘ

Select all Search package title:

Product/Release		Download Package	Size (MB)
5770-SS1	details	<input type="checkbox"/> IBM i	20191
07.03.00	packages	<input checked="" type="checkbox"/> 5837: i7.3 B_GROUP1 v07.03.00,ENU,DVD	14196
07.03.00	packages	<input type="checkbox"/> 5873: i7 Licensed Machine Code v07.03.00,ENU,DVD	531
07.03.00	packages	<input type="checkbox"/> 5973: USB Lic Mach Cde Flash Image v07.03.00,ENU,ESD	2292
07.03.00	packages	<input type="checkbox"/> 6287: IBM i Access 7.1 v07.02.00,ENU,CD-	2484

Figure 2-30 Selecting package for download

7. Save and download the IBM_Cloud_Storage_Solutions_for_i_LCD8_2390_02.udf.zip, as shown in Figure 2-31.

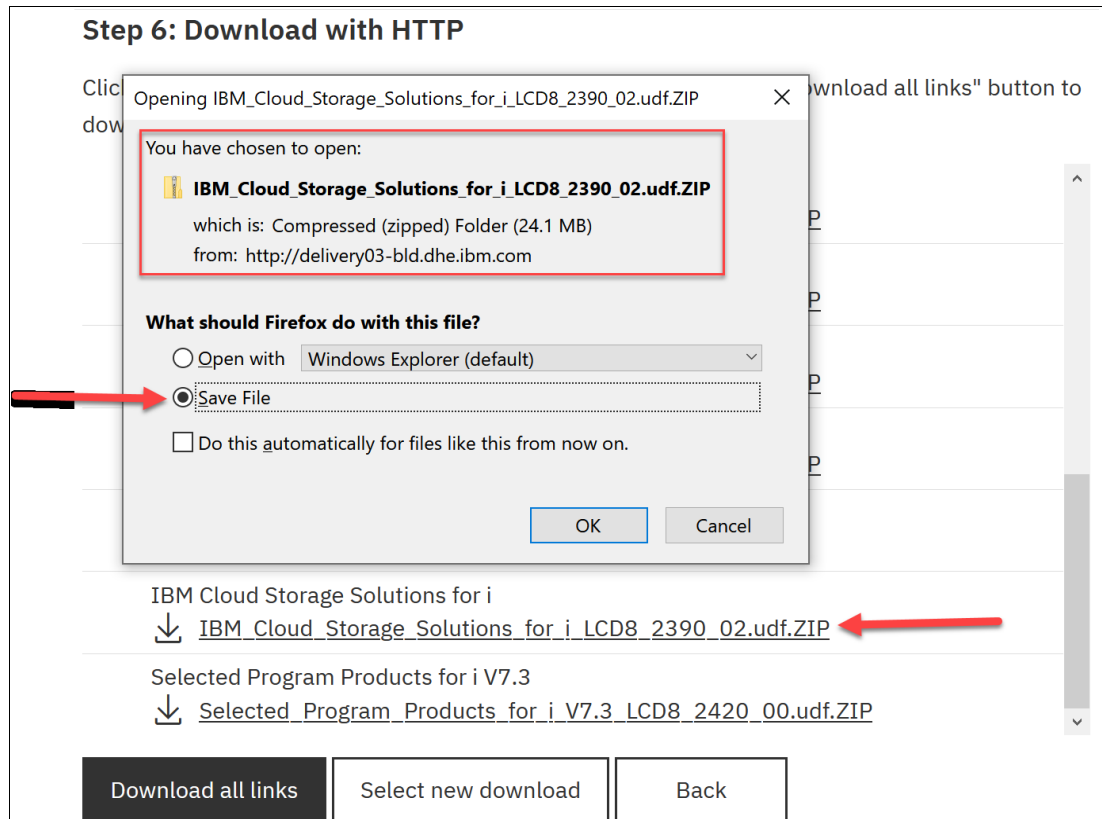


Figure 2-31 Downloading IBM Cloud Storage Solutions for i using HTTP

In our system, we extracted the compressed file content and renamed it to 5733ICC.udf.

Preparing an optical image catalog to install software

In this section, we describe how to create an image catalog and an image catalog entry, add an image catalog entry, and load the image catalog in preparation of performing 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 2-15.

Example 2-15 Create a device description

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

2. Vary on the virtual optical device, as shown in Example 2-16.

Example 2-16 Vary on the Virtual Optical Device

```
VRYPVARY 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 2-17.

Example 2-17 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 at once into your image catalog:
 - Use the **CRTIMGCLG** command, as shown in Example 2-18.

Example 2-18 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 2-19.

Example 2-19 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 2-20.

Example 2-20 Add 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 2-21.

Example 2-21 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 2-22.

Example 2-22 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 2-23.

Example 2-23 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. You also must 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 2-24.

Example 2-24 Verify image catalog

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

To verify that images are added, another method, as shown in Example 2-25.

Example 2-25 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 set up 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 this [IBM Documentation web page](#).

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.

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.

- ▶ 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.

2.2.4 Solution diagrams

This section provides more information about the solution scenarios that are presented in this chapter.

Migrating IBM i to the Cloud using IBM Cloud Object Storage

This section shows a migration of IBM i to the cloud by using IBM Cloud Object Storage. Refer to Figure 2-32.

Consider the following points about migrating IBM i into IBM Cloud Object Storage (see Example 2-32 on page 53):

- ▶ Use 5770-ICC and BRMS to transfer VM to IBM Cloud Object Storage.
- ▶ Requires 2x disk space on-premises.
- ▶ Requires reasonable network connection to IBM Cloud.



Figure 2-32 Migrating IBM i to the Cloud using IBM Cloud Object Storage

- ▶ Run BRMS Control Group QCLDBSYS01, as shown in Example 2-26.

Example 2-26 Run BRMS Control Group QCLDBSYS01

```
STRBKUBRM CTLGRP(QCLDBSYS01) SBMJOB(*N0)
```

- ▶ Run BRMS Control Group QCLDBIPL01, as shown in Example 2-27.

Example 2-27 Run BRMS Control Group QCLDBIPL01

```
STRBKUBRM CTLGRP(QCLDBIPL01) SBMJOB(*N0)
```

Figure 2-33 shows the steps to retrieve your image by using IBM Cloud Storage Solutions for i.

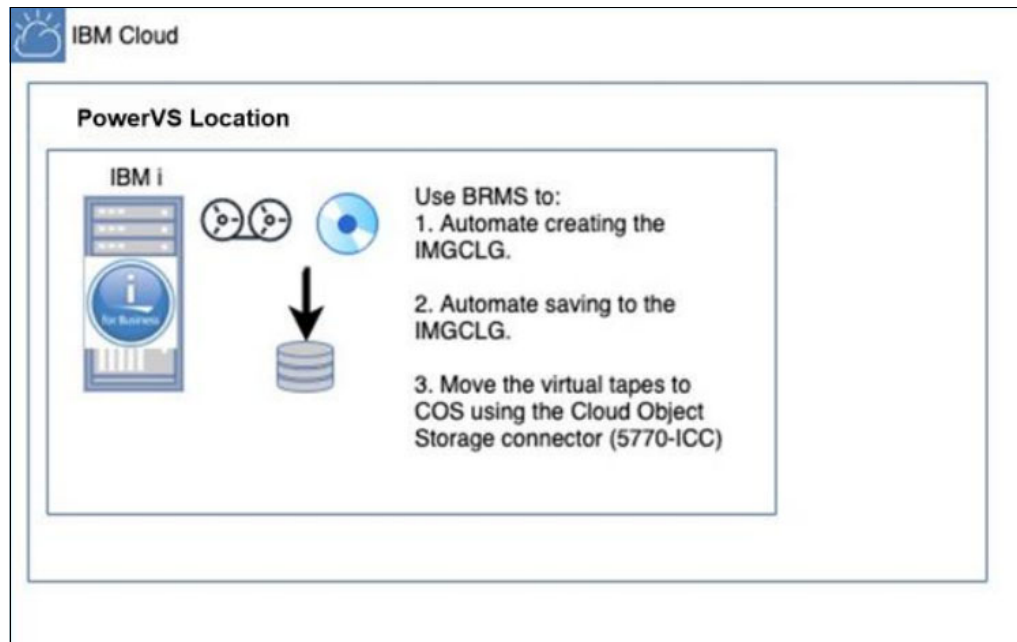


Figure 2-33 Get your BRMS Image Catalog on IBM Cloud Object Storage

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 restricted state as shown in Example 2-28 on page 51.

Example 2-28 Put the system in 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:
 - a. Run the **WRKCTLGBRM** command.
 - b. Find QCLDBSYS01 (see Figure 2-34).
 - c. Select **Option 9=Subsystems** to process.
 - d. Change Restart to *N0 for Seq 10 Subsystem *ALL.

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

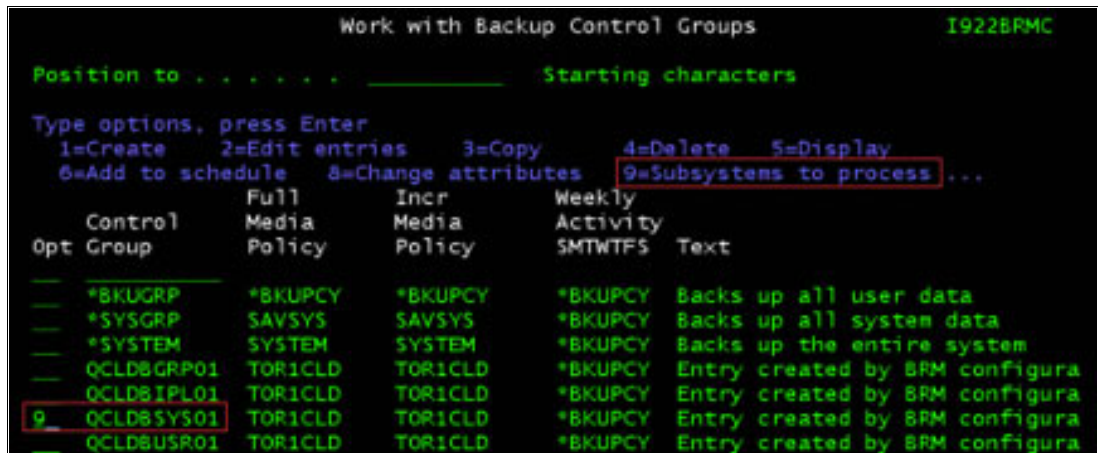


Figure 2-34 WRKCTLGBRM - Find QCLDBSYS01

5. Run the First backup from the console, as shown in Figure 2-29 on page 44.

Example 2-29 Run first backup from console

```
STRBKUBRM CTLGRP(QCLDBSYS01) SBMJOB(*N0)
```

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 completes. You see only subsystem QCTL in a status of RSTD. If not, end all subsystems again, as shown in Example 2-30.

Example 2-30 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 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 2-31.

Example 2-31 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 2-32.

Example 2-32 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 2-33.

Example 2-33 Reset the number of volumes

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

- Run the second backup from the console, as shown in Example 2-34.

Example 2-34 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 2-35.

Example 2-35 Identify volumes transferred

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

10. Transfer the volumes to IBM Cloud Object Storage.

Use the commands that are shown in Example 2-36 to transfer the volumes to IBM Cloud Object Storage.

Example 2-36 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 2-37.

Example 2-37 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 2-38.

Example 2-38 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 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 2-39.

Example 2-39 BRMS - Run daily incremental backups

```
STRBKUBRM CTLGRP(QCLDBUSR01) SBMJOB(*NO)
STRBKUBRM CTLGRP(QCLDBGSRP01) SBMJOB(*NO)
```

Full-system recovery from the Cloud by using IBM i as NFS server

This section describes how to recover a system from the cloud by using IBM i as the NFS server.

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.
IBM i NFS Server is at a minimum on Version 7.2 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 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 that uses the IP address of the NFS server.

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

In this section, we described how to recover from NFS. For more information, see 4.2.8, “Full-system recovery from the cloud using IBM i as an NFS server” on page 94.

2.2.5 Migrating IBM i by using PowerVC to PowerVS

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 need to address the challenges of managing virtual environments, which are far 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- and off-premises. Therefore, you can migrate and deploy IBM i images, as shown in Figure 2-35.

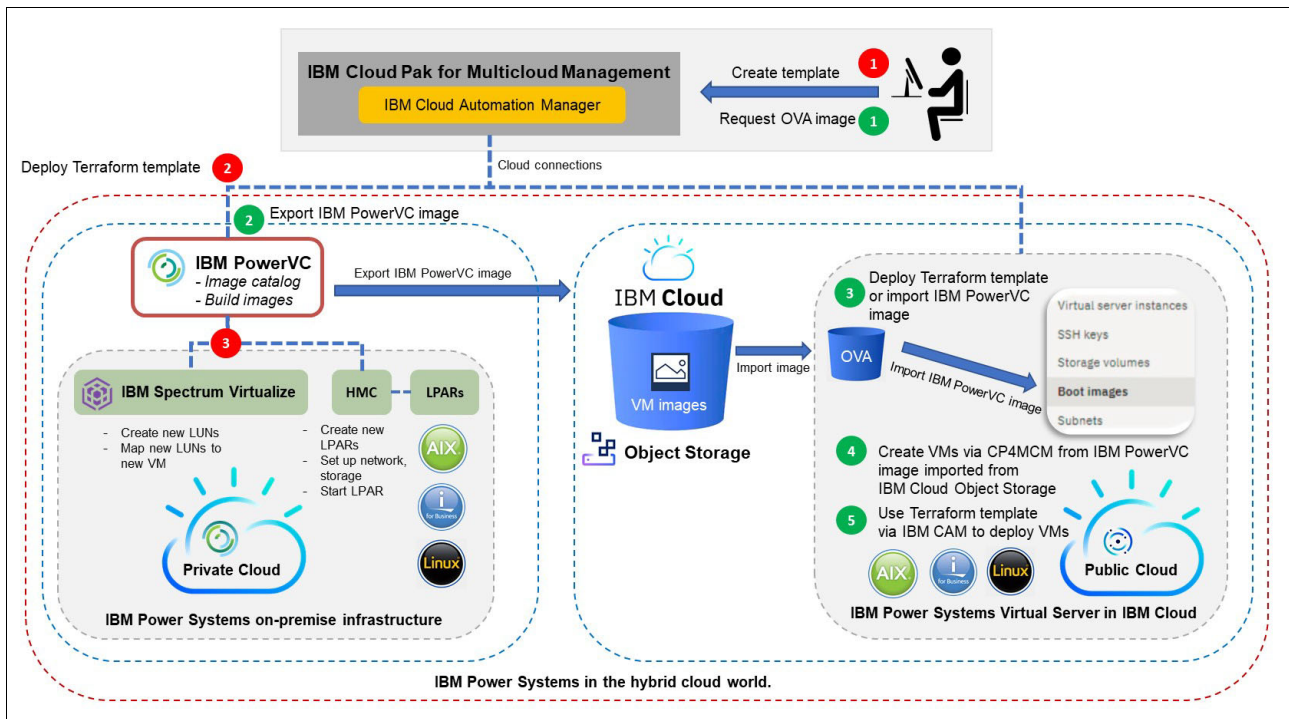


Figure 2-35 Export IBM i PowerVC images to IBM PowerVS

Tip: IBM i images can be deployed on- or off-premises by using Infrastructure-as-Code (IaC) that uses HashiCorp Terraform. You also can 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, and 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.

2.2.6 Migrating IBM i by using Mass Data Migration

Mass Data Migration is a fast, simple, secure way to physically transfer terabytes to petabytes of data to or from IBM Cloud. Whether it is to release on-premises storage space, decommission data centers, or pursue a gateway into IBM Cloud, Mass Data Migration is a versatile physical data transfer solution with which commercial- and enterprise-level customers in every industry can move large amounts of data to IBM Cloud quickly and securely.

Important: Online ordering is available in most regions, except Brazil and India. If the region you are in does not support online ordering, contact your IBM Client Representative to inquire about the use of the service. For more information and support, see this [IBM Cloud Docs web page](#).

Features

Data migration includes the following features:

- ▶ 120 TB usable capacity with RAID-6 configuration
- ▶ Inline AES 256-bit encryption
- ▶ Inline compression
- ▶ Ruggedized, tamper-evident, Trusted Platform Module (TPM)-enabled cases
- ▶ NIST data wipe standards upon completion

Pricing

Ordering a Mass Data Migration device does *not* include any up-front costs. Consider that your account accrues a daily usage charge of USD 50.00.

If you are importing data to IBM Cloud, the charges accrue after the device arrives at your location. If you are exporting data from IBM Cloud, the charges accrue when device is prepared for data copy. The charges stop accruing the day that the device is received by the IBM Cloud data center.

Also, after the data on your device is offloaded into your targeted IBM Cloud Object Storage bucket, your account accrues associated Cloud Object Storage costs. For more information about Mass Data Migration, see this [IBM Technology YouTube video](#).

Migrating IBM i Data to IBM Cloud by using Mass Data Migration

This section describes an alternative method for performing the IBM i on-premises save and transfer of data to IBM Cloud Object Storage (see Figure 2-36).




Figure 2-36 IBM i data save on Mass Data Migration

Note: Restore in IBM Power Systems Virtual Server is performed by using an IBM i VM as an NFS server.

Using the Mass Data Migration device

IBM i data is migrated to the IBM Cloud by using an IBM Cloud Mass Data Migration device. The IBM i data is saved to virtual optical images in a Network File System (NFS) share that is on the Mass Data Migration device at your data center.

You then must send the Mass Data Migration device back to IBM where your data is moved to the IBM Cloud so that it can be accessed through your IBM Cloud Object Storage. the process includes the following steps:

1. The MDM device is mounted to IBM i IFS.
2. An NFS-backed Optical IMGCLG is created.
3. Save the image directly to the MDM device.
4. A transfer rates of ~100 - 120 MBps is observed.
5. The MDM device is sent back to IBM Cloud.
6. The device is uploaded into IBM Cloud Object Storage by IBM Cloud.
7. IBM i is restored from IBM Cloud Object Storage.

Important: For more information, see this [IBM Cloud Docs web page](#).

The tutorial directs you to set up your IBM Cloud account for data migration and request a Mass Storage Migration device.

Preparing IBM Cloud Mass Data Migration device for IBM i data

After you receive the Mass Storage Migration device, review the [IBM Cloud Docs web page](#) and then, use the following checklist to complete the device set up:

Note: IBM Cloud Mass Data Migration devices arrive preconfigured and ready to connect to your network. For more information, see this [IBM Cloud Docs web page](#).

- ▶ After you set up the IBM Cloud Mass Data Migration device for Ethernet connectivity, you can access the device user interface. For more information, see this [IBM Cloud Docs web page](#).
- ▶ You can copy data onto the IBM Cloud Mass Data Migration device by first unlocking and activating the storage pool that is provisioned for the device. More information, see this [IBM Cloud Docs web page](#).
- ▶ Review this [IBM Cloud Docs web page](#) topic to ensure that the shares are set up correctly for the IBM i. The directory (or a parent directory) that contains the virtual optical images must be shared with the following characteristics:
 - The following Internet Protocol (IP) addresses:
 - IBM i client.
 - IBM i client service tools server or the LAN console connection if it is different than the system IP address. For more information, see this [IBM Documentation web page](#).
 - For read and write.

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

Preparing IBM i to use the Cloud Mass Data Migration device

The following requirements must be met to share virtual optical images with an IBM i client through an NFS server:

- ▶ The IBM i NFS client has a Version 4 IP address.
- ▶ The IBM i NFS client has the NFS server share mounted over a local directory during the setup process.
- ▶ The IBM i NFS client has a 632B-003 virtual optical device that uses the IP address of the NFS server.
- ▶ A directory under the NFS server share contains images that are large enough to hold all save data for the IBM i client.

Note: Save operations cannot dynamically create images or dynamically extend the size of images; therefore, understanding of the size of the data that is generated by the save operation is important.

- ▶ A directory under the NFS server share has a volume list file that contains a list of images to be used by an IBM i client. The volume list file must have the following characteristics:
 - Its ID is in the same directory as the image files
 - Is named `VOLUME_LIST`
 - Contains ASCII characters

Consider the following points:

- Each entry in the file is an image file name with access intent or a comment.
- Image file names are limited to 127 characters.
- All characters following a hash symbol (#) are considered a comment and are ignored.
- The order of the image file names in the volume list file indicates the order the images are processed on the IBM i client.

Setting up the IBM i NFS client

Complete the following steps to set up the IBM i client to use virtual optical images that are stored on an NFS server. The following command examples assume that NFS server SERVER01 with an IP address '1.2.3.4' has share /nfs/share01:

1. Create a mount directory for the NFS server:

```
MKDIR DIR('/NFS')
MKDIR DIR('/NFS/SERVER01')
```

2. Mount the NFS server root directory over the IBM i mount directory:

```
MOUNT TYPE(*NFS) MFS('1.2.3.4:/nfs/share01') MNTOVRDIR('/NFS/SERVER01')
```

3. Create image information about the NFS server in the Portable Application Solutions Environment (PASE).

4. Create a device description for a virtual optical device:

```
CRTDEVOPT DEVD(NFSDEV01) RSRNAME(*VRT) LCLINTNETA(*SRVLAN)
RMTINTNETA('1.2.3.4') NETIMGDIR('/nfs/share01/iImages')
```

5. Vary on the virtual optical device:

Note: If the images or VOLUME_LIST file is changed on the NFS server, the virtual optical device must be varied off and then, varied back on to use the virtual images.

```
VRYCFG CFGOBJ(NFSDEV01) CFGTYPE(*DEV) STATUS(*ON)
```

6. Use the device to initialize the image files as optical volumes:

```
LODIMGCLGE IMGCLG(*DEV) IMGCLGIDX(3) DEV(NFSDEV01) INZOPT NEWVOL(IVOL03)
DEV(NFSDEV01) CHECK(*NO)
LODIMGCLGE IMGCLG(*DEV) IMGCLGIDX(2) DEV(NFSDEV01) INZOPT NEWVOL(IVOL02)
DEV(NFSDEV01) CHECK(*NO)
LODIMGCLGE IMGCLG(*DEV) IMGCLGIDX(1) DEV(NFSDEV01) INZOPT NEWVOL(IVOL01)
DEV(NFSDEV01) CHECK(*NO)
```

7. Virtual device NFSDEV01 can now be used for native IBM i save and restore operations. Volume IVOL01 (image file '/nfs/share01/iImages/IMAGE01.ISO') is mounted on device NFSDEV01.

To test the device, perform the following command to verify that a message queue can be saved and restored. Messages from the save and restore commands indicate that one object was saved and one object was restored:

```
CRTMSGQ MSGQ(QTEMP/MSGQ) SAVOBJ OBJ(MSGQ) LIB(QTEMP) DEV(NFSDEV01) CLEAR(*ALL)
RSTOBJ OBJ(*ALL) SAVLIB(QTEMP) DEV(NFSDEV01)
```

The save contents of the virtual optical volume can also be displayed:

```
DSPOPT VOL(IVOL01) DEV(NFSDEV01) DATA(*SAVRST) PATH(*ALL)
```

Saving IBM i data to the Cloud Mass Data Migration device

Perform an IBM i full-system save by using the instructions that are described in “Running a full-system backup from the console” on page 51. However, use the newly mounted NFS device (the MDM device) as a target instead of the local image catalog and IBM Cloud Storage.

Returning the Cloud Mass Data Migration device to IBM

To complete the migration of IBM i data to IBM Cloud, follow the process that is described at this [IBM Cloud Docs web page](#).



IBM Power Systems Virtual Server in IBM Cloud Network

This chapter covers the different Hybrid Cloud Networking use cases for setting up a private network between an on-premises location and IBM Power Systems Virtual Server (PowerVS). It also provides guidance about production and non-production networking scenarios.

This chapter includes the following topics:

- ▶ 3.1, “IBM Power Systems Virtual Server virtual private network connectivity introduction” on page 64
- ▶ 3.2, “IBM PowerVS network overview” on page 65
- ▶ 3.3, “IBM PowerVS network scenarios” on page 67

3.1 IBM Power Systems Virtual Server virtual private network connectivity introduction

A key customer requirement for IBM PowerVS is the ability to connect to cloud-based workloads from an on-premises environment. It is not recommended to create IBM PowerVS workloads with a public IP address for security reasons. Customers need the capability for multiple users and multiple on-premises systems to connect securely to workloads in IBM PowerVS over a private network.

This section discusses the different Hybrid Cloud networking use cases for setting up a private network between on-premises systems and IBM PowerVS. It also provides guidance for production and nonproduction networking scenarios.

IBM PowerVS is colocated (COLO) in IBM cloud data centers with their own dedicated infrastructure, and they are connected to IBM Cloud Infrastructure as a Service (IaaS) by way of Direct Link Connect (DLC), as shown in Figure 3-1.

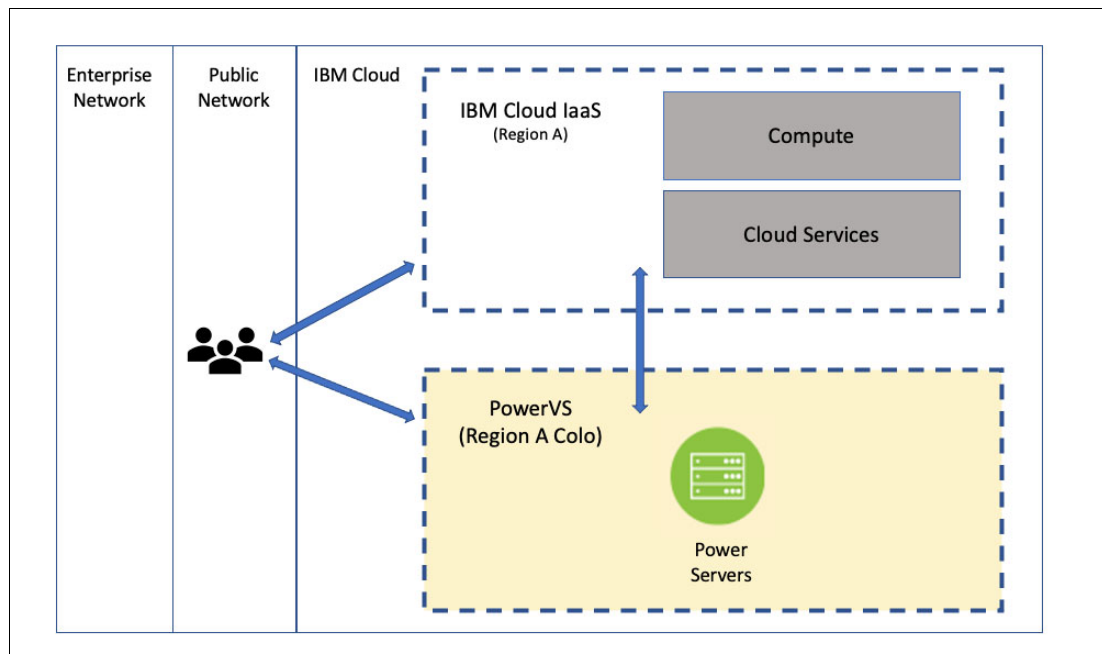


Figure 3-1 Example of the IBM Cloud IaaS and IBM PowerVS connection

DLC enables a key customer benefit in which workloads that are running on IBM PowerVS can integrate with x86-based workloads that are running in IBM Cloud for a single multi-platform business solution. Consider the following examples:

- ▶ An Oracle database that is running in AIX in IBM PowerVS that connects to a Linux application server in an x86 Virtual Server Instance (VSI).
- ▶ A core banking application in IBM i that connects to a point-of-sale application in a VMware-based x86 VSI.

The IBM PowerVS offering includes a highly available link of up to 10 Gigabits per second (Gbps) when Direct Link Connect 2.0 is used to connect IBM Cloud services at no cost for each customer per data center.

For more information about DLC, how to order DLC, and a tutorial about integrating x86 workloads with IBM PowerVS, see the following IBM Cloud Docs web pages:

- ▶ [Direct Link Connect for Power Systems Virtual Servers](#)
- ▶ [IBM Power Systems Virtual Server integration with x86-based workloads](#)

3.2 IBM PowerVS network overview

Customers can configure connections from on-premises systems to an IBM PowerVS location and other locations by using one of the following methods:

1. On-premises-to-IBM PowerVS by way of IBM Cloud IaaS (see Figure 3-2).




Figure 3-2 On-premises-to-IBM PowerVS by way of IBM Cloud IaaS

2. On-premises-to-IBM PowerVS directly, as shown in Figure 3-3.

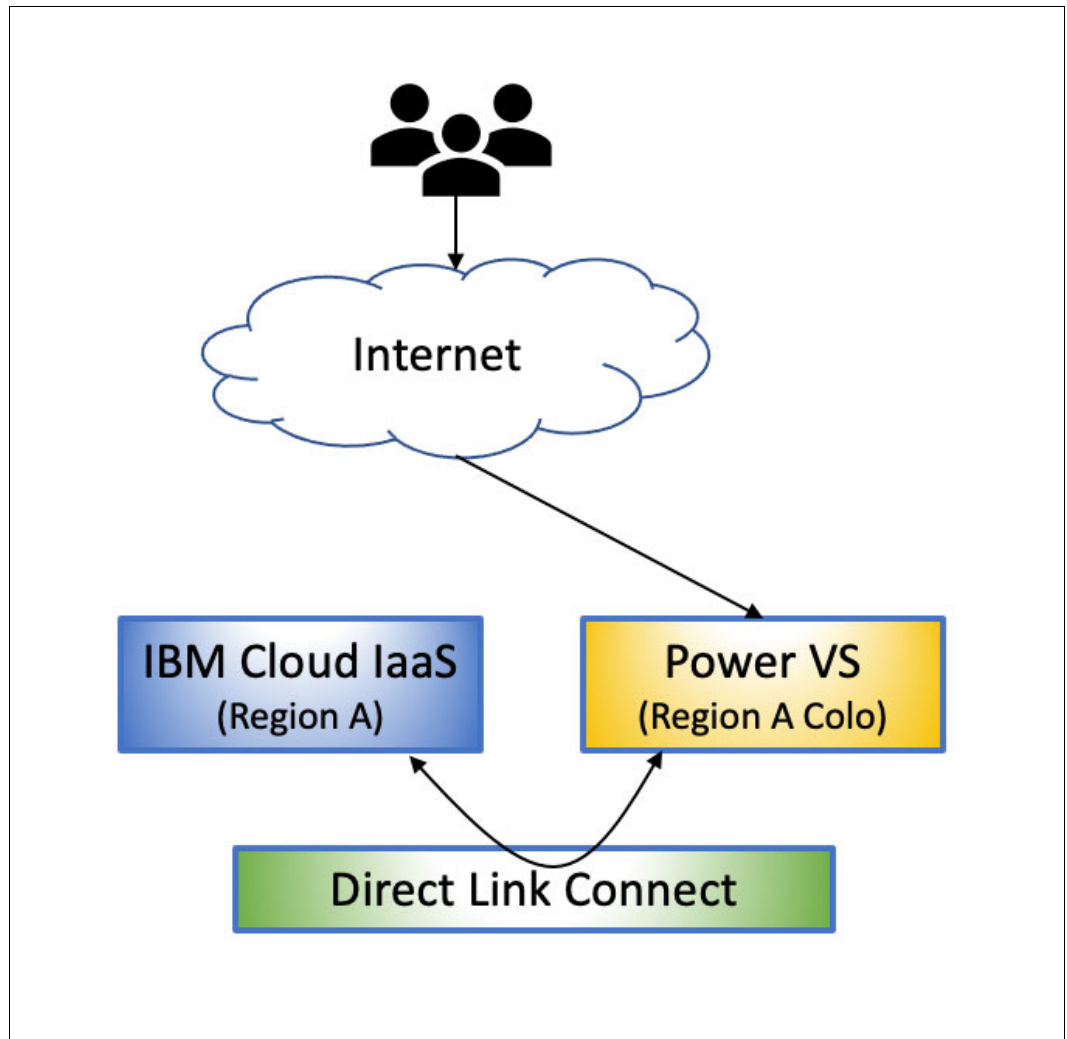


Figure 3-3 On-premises-to-IBM PowerVS directly

3. On-premises-to-IBM PowerVS (multiple COLOs) as shown in Figure 3-4.

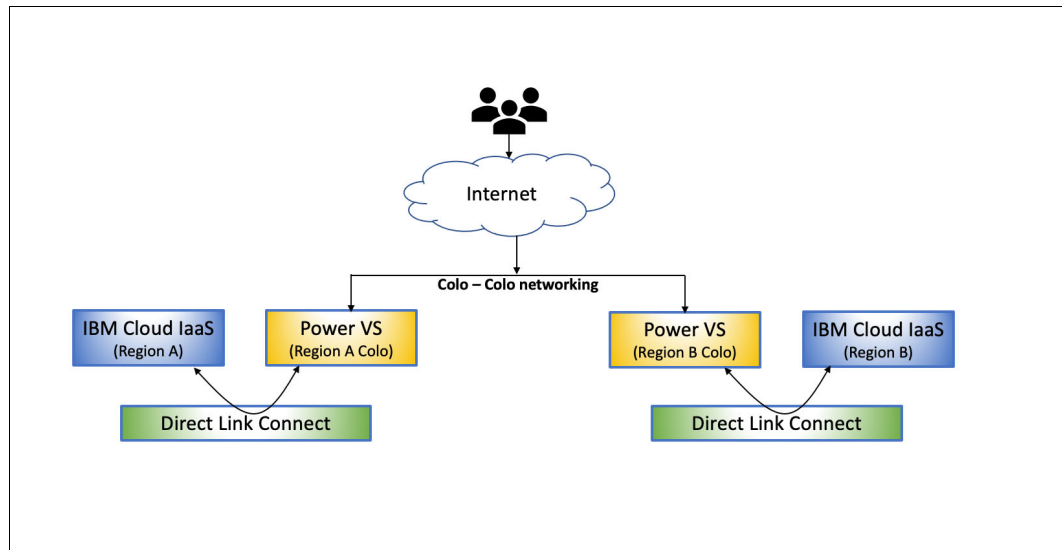


Figure 3-4 On-premises-to-IBM PowerVS (multiple COLOs)

3.3 IBM PowerVS network scenarios

Customers can create a private network connection between on-premises systems and IBM PowerVS by using different methods. Based on the networking technology that is used, required network hops, network latency, network session, and stability considerations, networking scenarios can be classified be as:

- ▶ Nonproduction/Proof-of-Concept (POC) production
- ▶ Production

3.3.1 Nonproduction proof-of-concept scenarios

Nonproduction/POC environments can use the following private connections:

- ▶ SSL + Jump Host + Direct Link Connect
- ▶ IPsec Virtual Private Networking (VPN) + Direct Link + Edge Gateway

These connections are described next.

Private connection by using SSL + Jump Host + Direct Link Connect

This scenario is used for environment management and test and development use cases from the public network. It is *not* recommended for production workloads.

A jump server must be used because (at the time this publication was written) a VPN connection cannot be used to directly connect to the IBM PowerVS instance.

The reference architecture is shown in Figure 3-5.

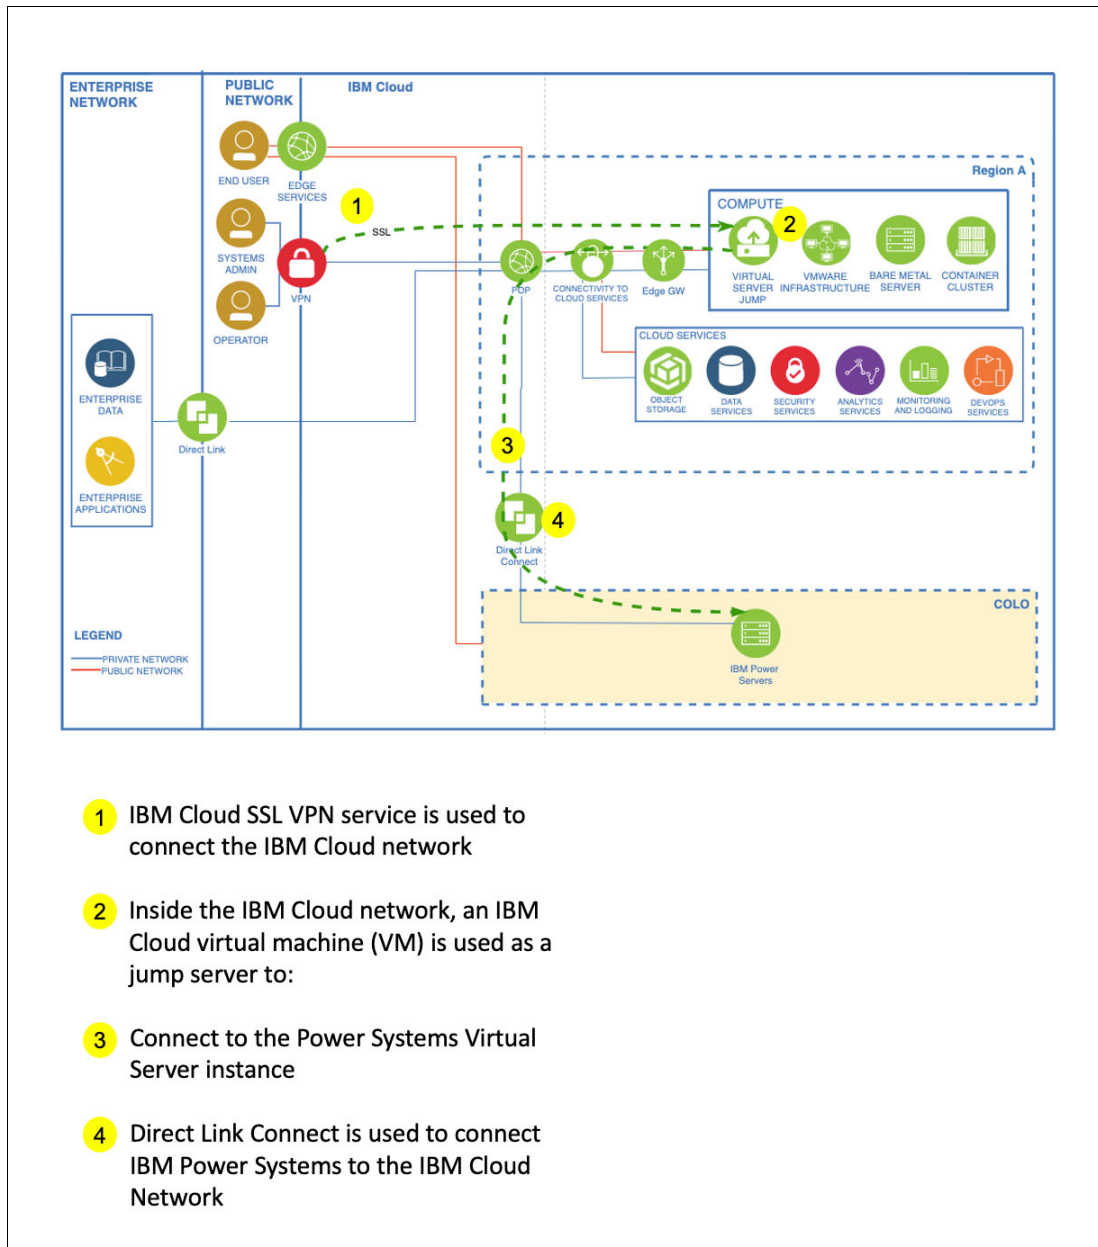


Figure 3-5 Private connection by using SSL + Jump Host + Direct Link Connect

Setting up this connection is a two-step process:

1. Private access is established to a jump server/host (Windows or Linux VM on the IBM Cloud).
2. After you are logged in to the jump host, you can connect to the IBM PowerVS environment over the direct-link connect between IBM PowerVS and IBM Cloud.

This direct-link connect must be established by way of a ticket. For more information about Direct Link Connect for IBM PowerVS, see this [IBM Documentation web page](#).

In addition, you cannot connect directly to the IBM PowerVS environment from on-premises in this setup. You must access the jump host first and then, start the second connection to the IBM PowerVS environment.

That initial access to the IBM Cloud classic private environment is done by using the no-cost SSL VPN client (IBM Cloud SSL VPN service).

For more information about enabling the SSL VPN access, see the following resources:

- ▶ [Enabling SSL VPN Access](#)
- ▶ [Using an SSL VPN](#)

For more information about accessing and downloading the no-cost SSL clients, see the following resources:

- ▶ [Getting Started with VPN Access](#)
- ▶ [Connecting to SSL VPN from MotionPro clients \(Windows, Linux, and Mac OS X\)](#)

Private Connection by using IPsec Virtual Private Networking (VPN) + Direct Link + Edge Gateway

In this scenario, a user can have dedicated and direct access to the Power environment by using an IPsec VPN setup between on-premises systems and a customer-owned gateway appliance setup in IBM Cloud and a DLC setup IBM PowerVS and the IBM Cloud.

A reference architecture is shown in Figure 3-6.

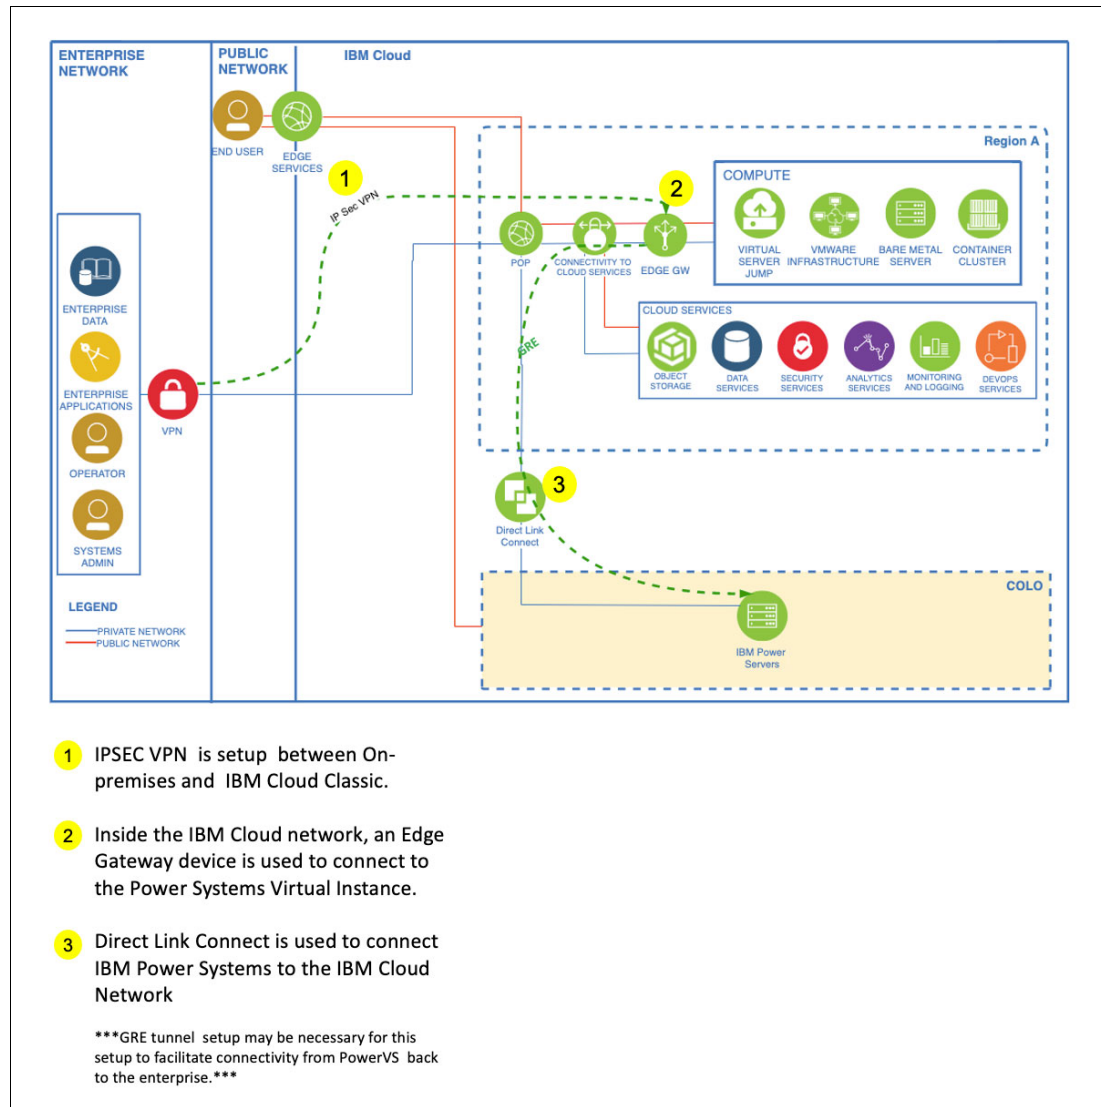


Figure 3-6 IPsec Virtual Private Networking (VPN) + Direct Link + Edge Gateway

In this setup, IBM PowerVS is directly reachable from on-premises. A Virtual Router Appliance (VRA or Edge GW) is needed in IBM Cloud because (at the time this publication was written) a VPN connection cannot be used to directly connect to the IBM PowerVS instance. The customer also needs the same VRA to access IBM Cloud resources and IBM PowerVS.

Various configuration caveats exist to facilitate the connection. For example, because IBM does not advertise the on-premises subnets over the DLC between IBM PowerVS and Cloud, various mechanisms are used to ensure traffic from IBM PowerVS to on-premises systems, including the following examples:

- ▶ GRE setup between the Edge GW and the IBM PowerVS router or firewall (IBM Managed). This way, Edge GW in IBM Cloud is set up as the next hop for any traffic that is intended for on-premises subnets from IBM PowerVS. This set up requires that a support ticket is opened with IBM PowerVS support (by way of the cloud console).

For more information about IBM PowerVS Virtual Private Network Connectivity (which is the preferred option), see [IBM Power Virtual Server Virtual Private Network Connectivity: An IBM Systems Lab Services Tutorial](#).

- ▶ NAT on the Edge GW appliance in IBM Cloud. The Edge GW can *mediate* traffic between on-premises subnets and IBM PowerVS subnet. This process works because IBM PowerVS sends traffic only to IBM Cloud subnets (portable subnets) that can be translated to an on-premises address. This process will only work if only a few on-premises IPs, subnets, or traffic are started from the IBM PowerVS.

3.3.2 Production scenarios

Production environments can use the following private connection configurations:

- ▶ Direct Link + Edge GW + Direct Link Connect
- ▶ Megaport + Direct Link Connect (on-premises to IBM Cloud PowerVS COLO)
- ▶ Megaport - Multi COLO, Multi-Region
- ▶ IBM backbone - Multi COLO

These connections are described next.

Private Connection using Direct Link + Edge GW + Direct Link Connect

Use this scenario for enterprise connectivity when private, dedicated, high-speed access through IBM Cloud is preferred and the customer wants to support workloads in IBM Cloud and IBM PowerVS environments.

The Direct Link option is needed to provide private connectivity to IBM Cloud. It is useful for routing BYOIP addresses through IBM Cloud by using GRE tunnels.

A Direct Link setup (Connect/Exchange/Dedicated) is set up between on-premises systems and IBM Cloud (which typically ends on a customer-owned gateway appliance). A direct-link connection (DLC) setup is established between IBM PowerVS and the IBM Cloud Network.

GRE tunneling is needed in most configurations because IBM Cloud Network does not advertise on-premises subnets over Direct Link and for Bring Your Own IP (BYOIP) considerations.

The reference architecture is shown in Figure 3-7.

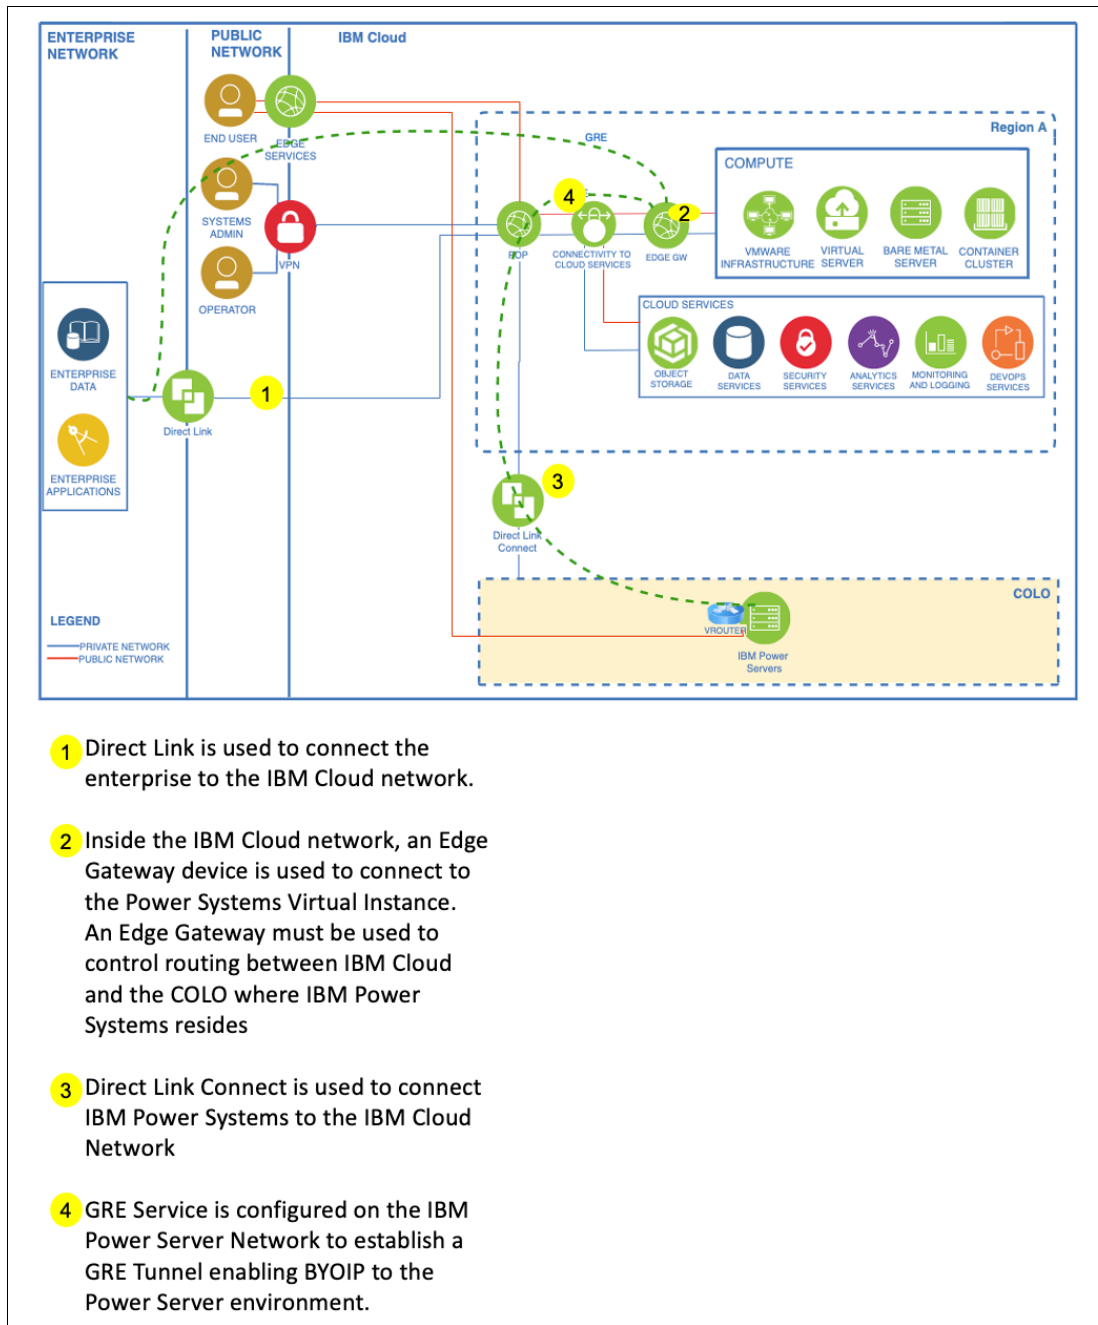


Figure 3-7 Private Connection by using Direct Link + Edge GW + Direct Link Connect

As shown in Figure 3-7, a GRE tunnel is built over the Direct Link Setup between IBM Cloud and on-premises systems. A GRE tunnel also is set up over the Direct Link Connect between IBM PowerVS and IBM Cloud.

This configuration includes several caveats that are needed to facilitate the connection. For example, because IBM does not advertise the on-premises subnets over the Direct Link between IBM PowerVS and Cloud, various mechanisms are used to ensure traffic moves from IBM PowerVS to on-premises systems, including the following examples:

- ▶ GRE setup between the Edge GW and the IBM PowerVS router/firewall (IBM-managed). This way, Edge GW in IBM Cloud is set up as the next hop for any traffic that is intended for on-premises subnets from IBM PowerVS. This set up requires that a support ticket is opened with IBM PowerVS Support (by way of the cloud console).

For more information about IBM PowerVS Virtual Private Network connectivity (which is the preferred option), see [IBM Power Virtual Server Virtual Private Network Connectivity: An IBM Systems Lab Services Tutorial](#).

- ▶ NAT on the Edge GW appliance in IBM Cloud. The Edge GW can *mediate* traffic between on-premises subnets and IBM PowerVS subnet. This configuration works because IBM PowerVS sends traffic to only IBM Cloud subnets (portable subnets) that can be translated to an on-premises address. This process works well if only a few on-premises subnets exist or traffic is started from the IBM PowerVS only.

Private Connection by using Megaport + Direct Link Connect (on-premises to IBM Cloud PowerVS COLO)

In this scenario, the customer wants a secure, low-latency, direct connection between the on-premises system and the IBM PowerVS COLO by using Megaport (a third-party communication provider).

This configuration is ideal for the following use case scenarios:

- ▶ Private and dedicated high-speed, high-bandwidth, low-latency, direct connection to IBM PowerVS environment
- ▶ Replication from enterprise to IBM PowerVS

The customer is responsible for the “last mile” connectivity between the on-premises system and IBM Cloud data center. The customer must order a Direct Link Connection with Megaport to obtain a Service Key, which is required to order or provision Megaport VXC.

In this setup, the customer contacts Megaport to facilitate the connectivity to the IBM PowerVS environment. For more information, see this [IBM Cloud Docs web page](#).

The reference architecture is shown in Figure 3-8.

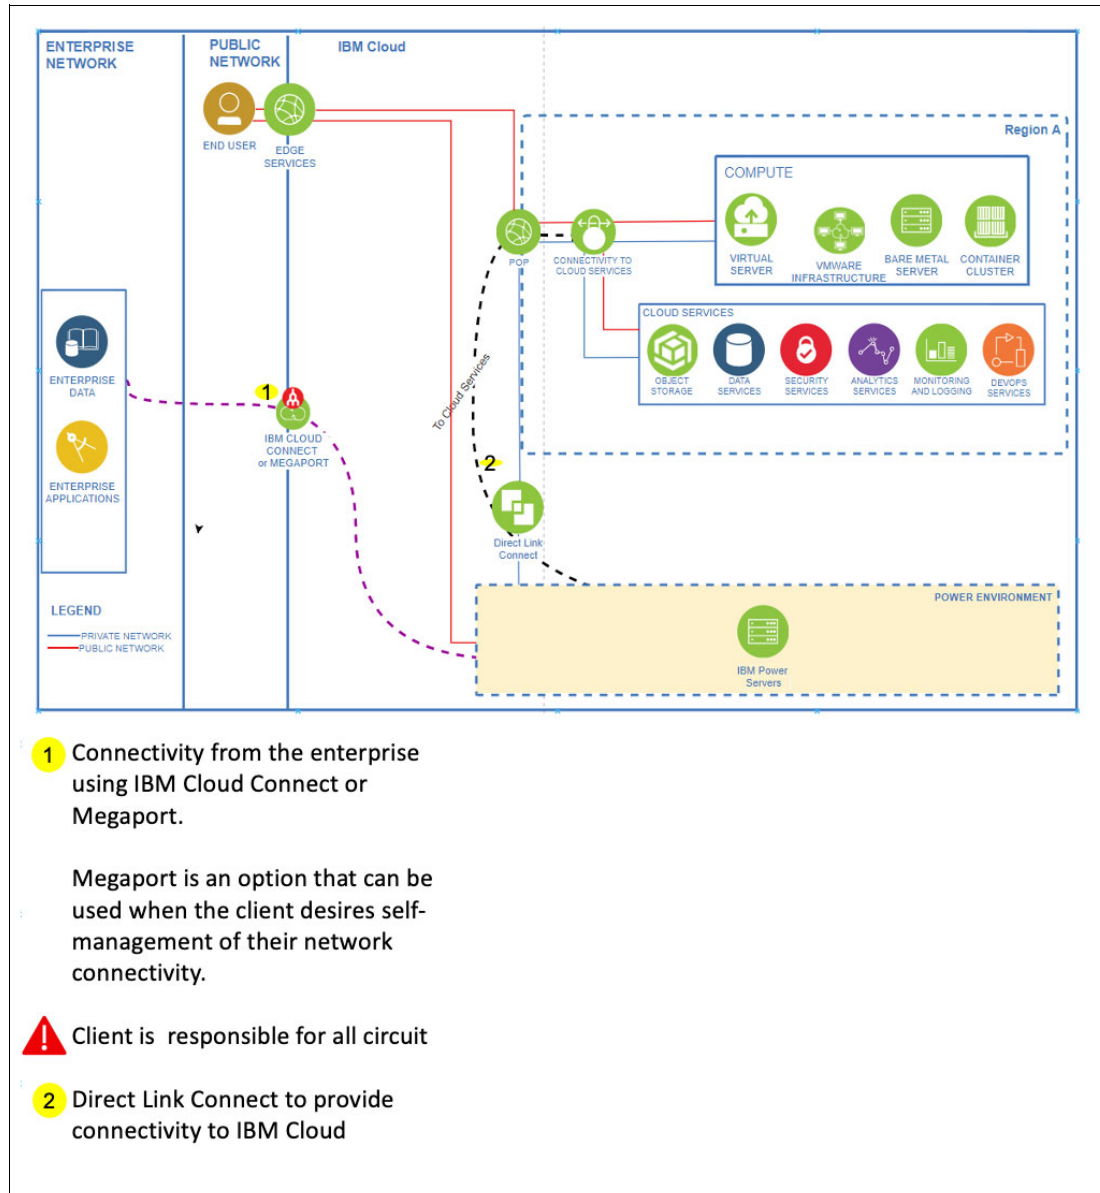


Figure 3-8 Private Connection with Megaport + Direct Link Connect

IBM Cloud Connect is used when the customer wants a fully managed service for connectivity between COLOs. However, IBM Network Services also can provide this as a service with IBM Cloud Connect. IBM Cloud Connect service is an IBM Network service that is wrapped around the Megaport connectivity.

Follow the process that is described at this [IBM Cloud Docs web page](#) before you contact Megaport about this service.

Private connection by using Megaport: Multi COLO, multi-region

In this scenario, the customer wants connectivity between two sites over the IBM Cloud backbone. Because the IBM PowerVS subnets cannot be advertised over the cloud backbone, GRE tunnels are necessary to facilitate the routing between the two locations.

This configuration is ideal for the following use case scenarios:

- ▶ IBM Power Systems that are deployed to more than one COLO and connectivity is required between the two.
- ▶ Replication between IBM sites production and Disaster Recovery systems in IBM Cloud over the IBM backbone.

The customer is responsible for “last mile” connectivity between the on-premises system and IBM Cloud data center. The ordering process requires the customer to order a Direct Link Connection with Megaport to obtain a service key, which is required to order or provision Megaport VXC.

In this setup, the customer contacts Megaport to facilitate the connectivity to the IBM PowerVS environment. For more information, see this [IBM Cloud Docs web page](#).

The reference architecture is shown in Figure 3-9.

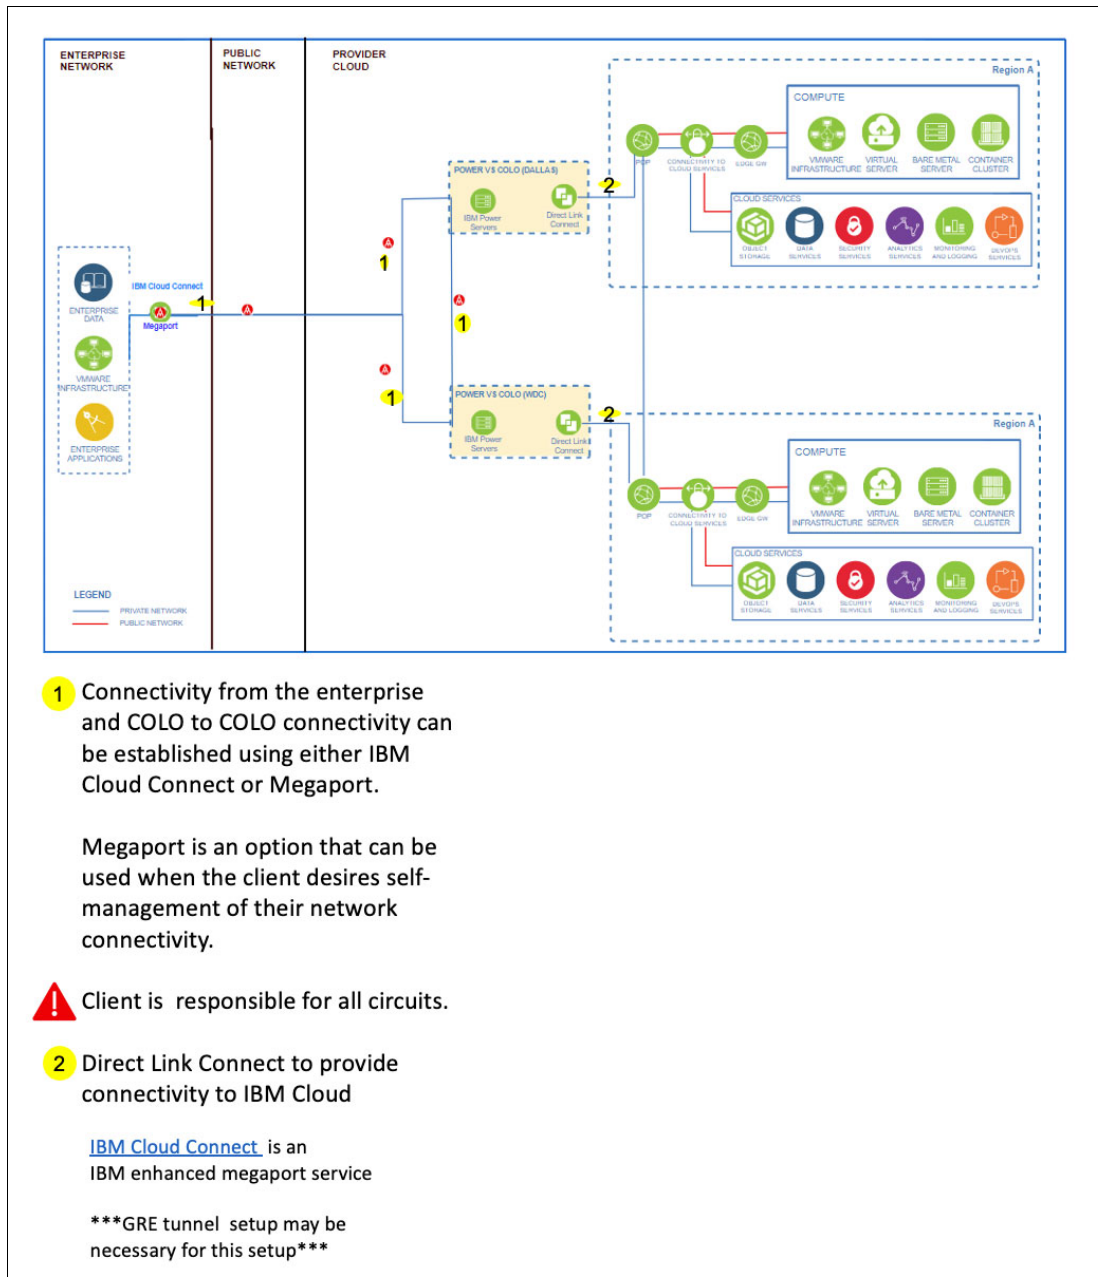


Figure 3-9 Private Connection by using Megaport - Multi COLO, Multi-Region

Private connection by using IBM backbone: Multi-COLO

In this scenario, the customer wants connectivity between two sites over the IBM Cloud backbone. Because the IBM PowerVS subnets cannot be advertised over the cloud backbone, GRE tunnels are necessary to facilitate the routing between the two locations.

This configuration is ideal for the following use case scenarios:

- ▶ IBM Power Systems that are deployed to more than one COLO and connectivity is required between the two.
- ▶ Replication between IBM sites production and Disaster Recovery systems in IBM Cloud over the IBM backbone.

The reference architecture is shown in Figure 3-10.

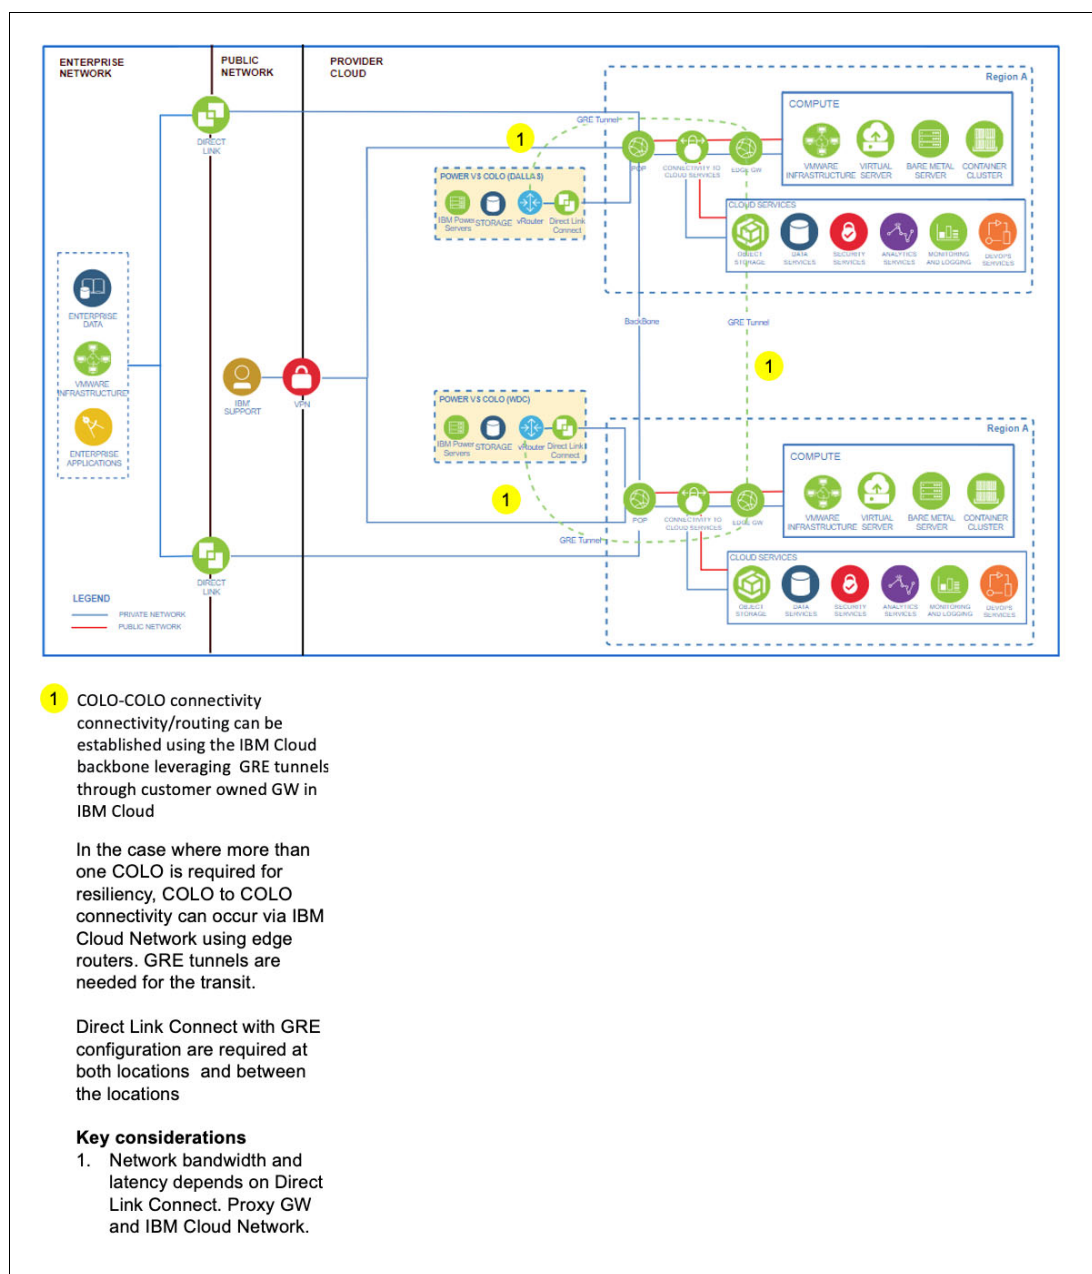


Figure 3-10 Private Connection by using IBM backbone - Multi COLO

The following requirements must be met to set up connectivity:

- ▶ customer-owned gateways (ATT router/Juniper VSRX/FortiGate/BYOGW) are installed at each site.
- ▶ A GRE tunnel exists from IBM PowerVS to the gateway at each location.
- ▶ A GRE tunnel exists between the gateway at Site A and the gateway at Site B.

Note: in this scenario, it is assumed that the Enterprise is connecting to IBM Cloud by way of a Direct Link Setup, as shown in Figure 3-10 on page 77.

However, ISPEC VPN to each site also can be an option, depending on the required bandwidth and use cases.



Backing up IBM i on IBM Power Systems Virtual Server

An effective backup process is a vital part of any data center operation. It insures 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:

- ▶ 4.1, “Backup and restore considerations on IBM Power Systems Virtual Server” on page 80
- ▶ 4.2, “Backing up and restoring by using IBM Backup, Recovery, and Media Services” on page 81
- ▶ 4.3, “Creating object-level backups” on page 113
- ▶ 4.4, “Dynamic Solutions International Virtual Tape Library” on page 132
- ▶ 4.5, “Full-system snapshot, cloning, and restore” on page 137

4.1 Backup and restore considerations on IBM Power Systems Virtual Server

This section describes a few considerations for backing up and restoring data on the IBM Power Systems Virtual Server environment.

4.1.1 Overview

IBM Power Systems Virtual Server includes differences to the regular on-premises environment.

Most on-premises IBM i operations use tape backups because of their capacity, transfer rate, and reliability. IBM provides two solutions: native commands and Backup, Recovery, and Media Services (IBM BRMS) as a complete recovery solution.

Also, only a few software products can back up by using network facilities. Traditionally, tape devices are faster.

When operating on IBM Power Systems Virtual Server environment, only the network can be used to transport backups. Fast connections increase transfer rate, but backup windows must be sized and considered on this environment.

When multiple backups must store a large amount of data, IBM Cloud Object Storage is used. This parallel storage system provides concurrent access from anywhere with an any-to-any-to-any architecture with no single point of failure, which reduces bottlenecks and provides almost endless scalability. IBM Cloud Object Storage can store data in primary or auxiliary storage layers, which allows different service levels for backups and archives, and reduces costs when needed.

When sizing buckets (that is, cloud object storage containers), 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 Cloud Object Storage
- ▶ Needed redundancy

Server-side data compression can use IBM Power Systems Virtual Server hardware compression and uncapped capacity (when possible) to reduce the amount of data that is needed to transfer and store on IBM Cloud Object Storage. This reduction in data also reduces transfer times and storage costs.

IBM i also provides mechanisms to create incremental and differential backups to reduce the amount of data that must be saved and stored. This most used mechanism is effective data is backed up to the cloud, and IBM i is no exception.

By using differential backups, the backup window and size of backup media can be reduced dramatically. However, the time 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 one a week or one a month, and are combined with full disk clones or Disaster Recovery (DR) software.

Today, some Open Source tools can be used to complete backup scripts and compression and encryption capabilities can be used that are included with IBM i Cloud Storage Solutions software to secure and improve transfers.

These pre- and post-processing tasks and the limited capability to start the instance from physical or virtual media creates a challenge to base the disaster recovery strategy only with backup and restore processes.

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 also can be copied by using a gold image and stored on IBM Cloud Object Storage for DR activities, which simplifies and improves the full system backup procedure.

4.2 Backing up and restoring by using IBM Backup, Recovery, and Media Services

This section describes the backup and restore procedures that use IBM Backup, Recovery, and Media Services (IBM BRMS).

4.2.1 IBM BRMS overview

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.

IBM BRMS is the IBM strategic solution for planning and managing save and restore operations on your IBM i product. 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.

By using IBM BRMS, you can manage your most critical and complex save operations, including online backups of Lotus servers. It also supports parallel save operations of a library or 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 Tivoli Storage Manager servers. IBM BRMS enables you 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 then can copy files between the Integrated File System (IFS) on the IBM i computer 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.

Requirements

The following licensed program products must be installed:

- ▶ 5733ICC *BASE IBM Cloud Storage Solutions for i
- ▶ 5733ICC 1 Cloud Storage

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 allow cloud connector resources to be defined for cloud storage providers, such as IBM Cloud, 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 i system. Media also is 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 PowerVS

The following components are required to take IBM i backups on IBM Power Systems Virtual Server on IBM Cloud:

- ▶ 5733ICC *BASE IBM Cloud Storage Solutions for i.
 - Unless double of storage capacity on IBM i virtual machine (VM) to hold the backup images.
- ▶ Direct Link to the Classic Infrastructure (commonly, Direct Link Connect is a good option).
- ▶ An Ngix reverse proxy server. This server can be provisioned on IBM Power Systems Virtual Server instances for bandwidth up to 1 Gbps. If more bandwidth is required, Bare Metal must be used.
- ▶ IBM Cloud Object Storage, with which you can choose the type of speed and resiliency required.
- ▶ IBM BRMS software:
 - 5770-SS1 Option 18: Media and Storage Extensions
 - 5770-SS1 Option 44: Encrypted Backup Ennoblement
 - 5770-BR1 *BASE

- 5770-BR1 Option 1: Network feature
- 5770-BR1 Option 2: Advanced Functions feature
- ▶ If DSI VTL is used, a VM instance that uses Red Hat Enterprise Linux 7 or 6 (can be provisioned on IBM Power System):
 - The required storage resource is typically setup capacity plus 30% more capacity for buffer
 - Network: 1 GbE or 10 GbE
 - Memory: Minimum 16 GB (over 32 GB is suggested for systems that have over 50 TB backup capacity)

Note: Consider the following points:

- ▶ The operating system feature is a prerequisite feature to IBM BRMS. It also is required when HSM dynamic retrieval functions are developed.
- ▶ Media and Storage Extensions provides an API to enable application monitoring and control media usage, including 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 new terminology is important to understand:

- ▶ *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*: 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 changes 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 IBM i characteristic

Solutions in the cloud include the following IBM i characteristics:

- ▶ Creates and maintain IBM Cloud Object Storage and FTP targets.
- ▶ Provides copy to cloud, copy from cloud, and delete from cloud functions for IBM Cloud Object Storage and FTP servers.
- ▶ Supports IBM i and Lynx FTP servers.
- ▶ Manages TCP/IP networking communications to IBM Cloud Object Storage or FTP Servers on IBM i or Lynx.
- ▶ Authenticates resource users with IBM Cloud Object Storage or FTP servers, and seamless updates expired IBM Cloud Object Storage 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 show 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 Cloud Object Storage 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 Cloud Object Storage resource. A resource defines an AWS S3 or IBM Cloud Object Storage 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 Cloud Object Storage resource. A resource defines an AWS S3 or IBM Cloud Object Storage 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 Q1g_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 Cloud Object Storage, 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](#)
- ▶ [How to fill out Request For Enhancement \(RFE\) for Cloud Storage Solutions for i \(5733ICC\)](#)

4.2.4 IBM BRMS turn-key

IBM BRMS turn-key is an automated, pre-scripted solution within IBM BRMS that is simple to set up and run.

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 object for new cloud resources:

- ▶ Run **STRMNTBRM RUNCLNUP(*YES)**
- ▶ Run **INZBRM (*DATA)**

The required IBM BRMS objects are created.

IBM BRMS creates the following objects when it detects the IBM Cloud Object Storage 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:**
 - **QCLDBIPLnn:** Backs up what is minimally needed for a system D-IPL (must be burned to a DVD).
 - **QCLDBSYSnn:** Backs up all system data, except *SAVSYS.
Paired with QCLDBIPLnn.
 - **QCLDBUSRnn:** Backs up all user data.
 - **QCLDBGRPnn:** Backs up what is minimally needed to and from the cloud, except *SAVSYS.
– Paired with QCLDBUSRnn (must be burned to a DVD).

The backup requires the following separate recovery reports:

- ▶ Cloud Connection Recovery (DVD)
- ▶ Recovering user data from cloud

4.2.5 Backing up the IBM i system to the cloud

IBM BRMS stores system backup media in cloud locations the same way 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 for i to 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 names 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 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 can 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 complete, 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 here; 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 `QCLDBGRPnn` and `QCLDBIPLnn` control group data to optical media and storing the discs with other 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.

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 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 are 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 `QCLDCSYSnn` and `QCLDCIPLnn` 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, QCLDUGRPnn, 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, a 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 QCLDU prefix instead adopting a new name with the QCLDC 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 QCLDBSYSnn to a new custom control group and update any entries as needed. Also, rename the control group; 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; 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; 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 control group; 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 IBMid and Password.
2. Go to your Service and select your virtual server instance. Open the console, as shown in Figure 4-1.

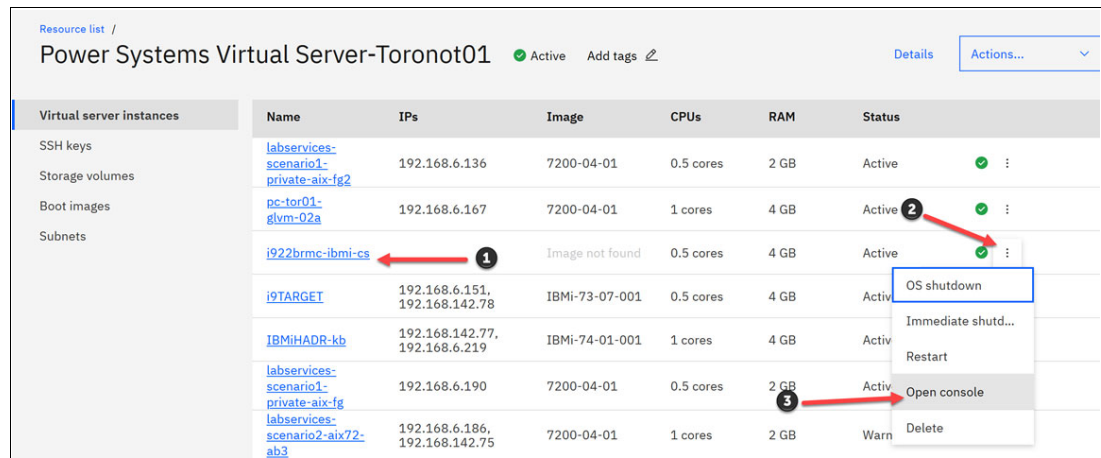


Figure 4-1 Open 5250 console for IBM i

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. You must close your console browser and start a new console connection. For more information, see 5.1.3, “IBM i 5250 console through LAN adapter” on page 153.

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=Subsystem to process**.
 - c. Change the restart to ***N0** for Seq 10 Subsystem *ALL, as shown in Figure 4-2.

```

Subsystems to Process
Use . . . . . : *BKU
Control group . . . . : QCLDBSYS01

Type choices, press Enter.

Seq      Subsystem      Library      End      Delay      Restart
-----  -
  10     *ALL             *ALL         *CNTRLD  120       *NO

```

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).

```

Display ATT Messages
Job . . : DSP01      User . . : QSECOFR      System:  I922BRMC
Number . . . : 224880

234282 blocks processed for sequence 160, volume Q31632, on device
Q1AVTAP001.
52 objects not saved.
185919 objects saved. 1 not saved.
Save of list *LINK completed with errors.
Control group QCLDBSYS01 bypassed automatic save of media information.
The protocol required to support the specified address family is not
available at this time.
Open server connection failed.
Calling exit program Q1ACLEXIT in library QBRM.
An Error occurred. Check the previous messages in the joblog.
COPY FILE TO CLOUD FAILED
MEDIA Q31632 NOT TRANSFERRED TO CLOUD TORICLD.
Control group QCLDBSYS01 type *BKU completed with errors.

Press Enter to continue.
Bottom

```

Figure 4-3 Display possible errors

9. Check the subsystems after the backup completes. 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).

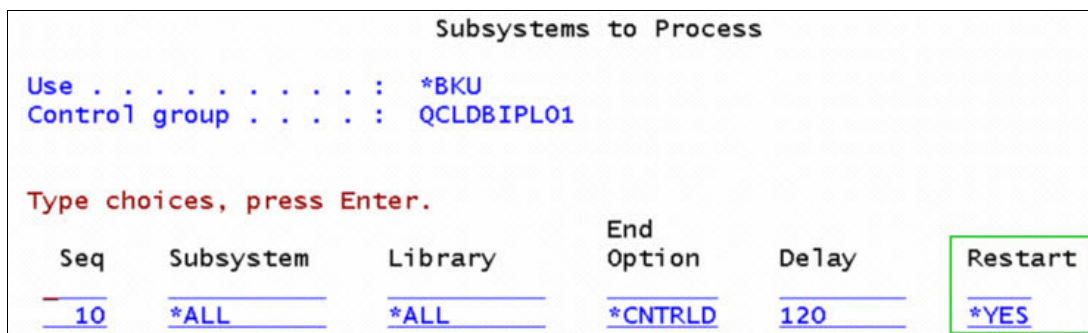


Figure 4-4 Change restart to *Yes on subsystem to process

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 Cloud Object Storage.

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.

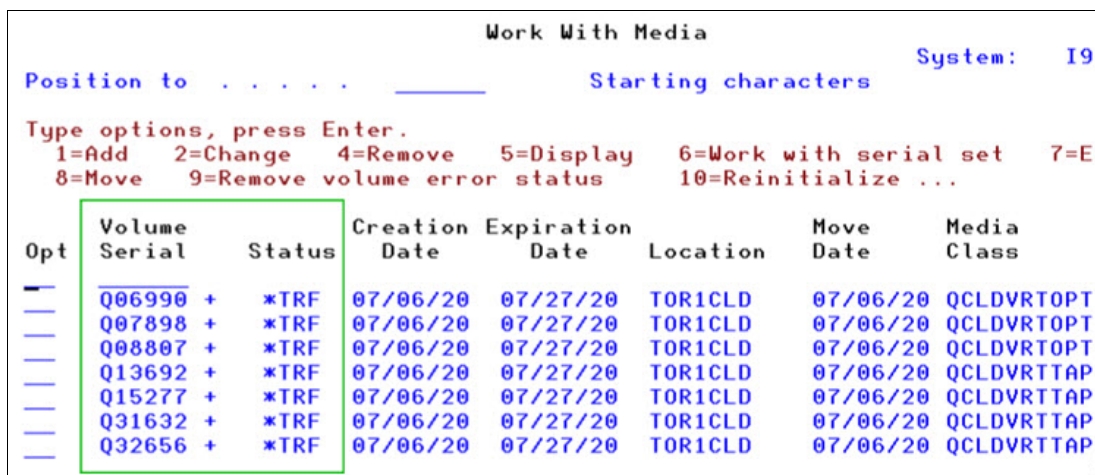


Figure 4-5 Identifying the volumes that were used

16. Transfer the volumes to IBM Cloud Object Storage. Run the **STRMNTBRM** command and the **WRKSTSICC STATUS(*ALL)** command. You see the volume name, status, and complete 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. Then, run the **WRKMEDBRM TYPE(*TRF)** command. 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 Cloud Object Storage.

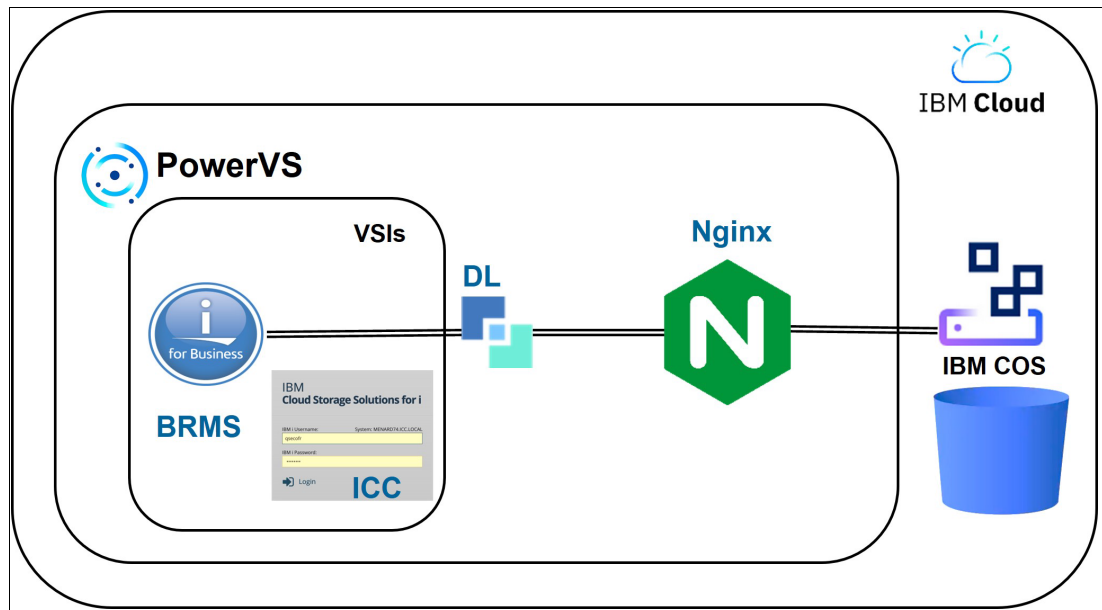


Figure 4-6 IBM i source that uses IBM BRMS and IBM Cloud Storage

Note: Consider the following points about the use of IBM Cloud Object Storage over IBM Cloud Direct link:

- ▶ IBM Cloud customers that purchase IBM Cloud Object Storage and Direct Link can make remote connections to IBM Cloud Object Storage “private” endpoints.
- ▶ HTTPS IBM Cloud Object Storage 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 Cloud Object Storage 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((QCLDBSYSxx 1) (QCLDBIPLxx 2) (QCLDBUSRxx 3) (QCLDBGRPxx 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:
 - 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_XXXXXXX 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 call **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) STRTPPRF(*NO) STRIP6(*YES)
STRTCPIFC INTNETADR('nnn.nnn.nnn.nnn')
```

where 'nnn.nnn.nnn.nnn' is the internet 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

In this section, we describe 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 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 Internet Protocol (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.

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

For more information, see this [IBM Documentation web page](#).

Requirements to create a PowerVS instance for NFS Server

The following requirements must be met to create a Power Systems Virtual Server Instance NFS Server:

- ▶ An IBM i VSI 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 VSI Client in 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 VSI Client Server can ping itself and the IP address of IBM i VSI NFS Server.
- ▶ Verify that the IBM i VSI NFS Server can ping itself and the IP address of IBM i VSI Client Server.

Creating an IBM i VSI for NFS Server

Complete the following steps to create an IBM i IBM Power Systems Virtual Server Instance NFS Server:

1. Create IBM i VSI NFS Server in your Power Systems Virtual Server location.
2. Provision and add your storage.
3. Create a mount directory for the IBM i VSI NFS server:
 - a. Run the **MKDIR DIR('/install')** command.
 - b. Run the **MKDIR DIR('/install/sysip1')** command.
A directory is created in the IFS that is called `/install/sysip1`.
4. From the Cloud, copy the volumes that were created by using the QCLDBIPL01 backup. In this example, three volumes were 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 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_XXXXXXXXX directory, where XXXXXXXXX is the name of the system that performed the backups (see to Figure 4-7).

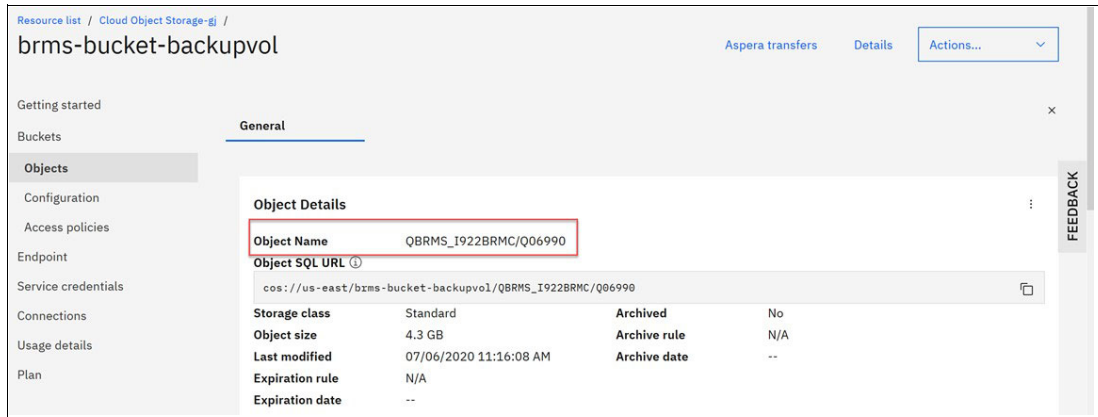


Figure 4-7 IBM BRMS storing media on IBM Cloud Object Storage

6. Run the **CPYFRMCLD** command, as shown in Figure 4-8.

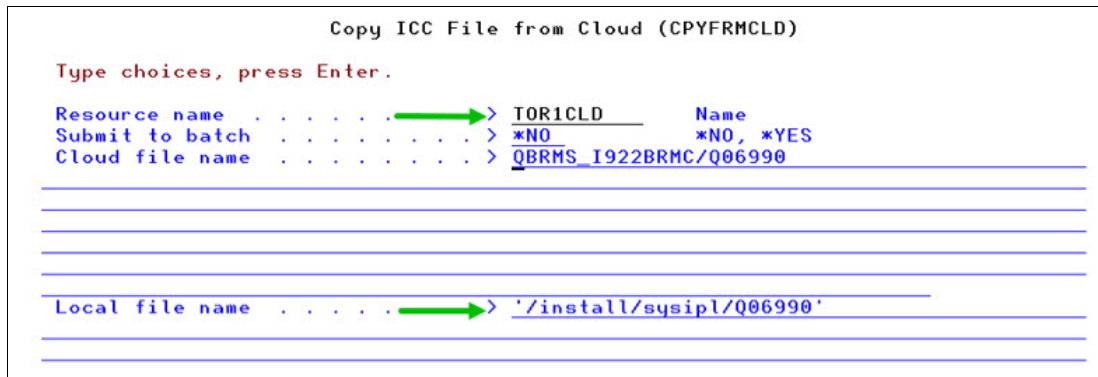


Figure 4-8 Copying IBM Cloud Storage file from cloud

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 virtual optical device on IBM i VSI NFS Server.

Creating virtual optical device on IBM i VSI NFS Server

Complete the following steps to create an optical device on IBM i VSI NFS Server:

1. Create a Virtual Optical Drive that is named "INSTALL" CRTDEVOPT DEVD(INSTALL) RSRNAME(*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 (or example, volume Q06990).


```

Add Image Catalog Entry (ADDIMGCLGE)

Type choices, press Enter.

Image catalog . . . . . > SYSIPL      Name
From optical device, or . . . . .      Name, *VOL
From image file . . . . . > Q06990
-----
To image file . . . . . > Q06990
-----
Image catalog index . . . . . *AVAIL    1-256, *AVAIL
Replace catalog entry . . . . . *NO      *NO, *YES, *INSERT
Text 'description' . . . . . *GEN
-----

```

Figure 4-9 Add image catalog entry

5. Add the next two volumes by using the same name for the TOFILE; that is, the To image file.
6. Load the image catalog, as shown in Figure 4-10.

```

Load or Unload Image Catalog (LODIMGCLG)

Type choices, press Enter.

Image catalog . . . . . > SYSIPL      Name
Option . . . . . *LOAD      *LOAD, *UNLOAD
Virtual device . . . . . > INSTALL    Name
Write protect . . . . . *DFT      *DFT, *ALL, *NONE
Library mode . . . . . *NO       *NO, *YES

```

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; that is, 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) DTAUT(*RWX) SUBTREE(*ALL)** command to change Object Authority.

13. Run the `CHGTFPA AUTOSTART(*YES) ALTSRCDIR('/install/sysip1')` command to change the TFTP Attributes.
14. Specify the alternative source directory where the volumes are stored.
15. Run the `CHGAUT OBJ('/install/sysip1') 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 VSI client server

Complete the following steps to create an IBM i Power Systems Virtual Server Instance Client Server:

1. Create an IBM i VSI Client Server 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**. Then, click **F13** at the bottom of the console window.
6. Choose **Option 1=Select**, as shown in Figure 4-11. Use the same resource that was used to communicate with the IBM i VSI NFS Server.

Select LAN Adapter Port						
Type option, press Enter. 1=Select 5=Display details						System:
Option	Resource Name	Type	Model	Serial Number	LOC	Selected
<u>1</u>	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 (public, direct, or DL) as your IBM i VSI NSF Server. You also MUST End TCP/IP and Vary Off the Line Description by using that CMNxx.

7. Configure Service Tools LAN Adapter (see Figure 4-12):
 - IP version allowed: IPV4.
 - Internet address: The same IP address of the client (TARGET) IBM i VSI can be used.
 - Gateway router address: Use the same gateway on the client (TARGET).
 - Subnet mask: Use the same.

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, which might not be ready immediately.

- d. Select **F3=Exit** to exit the Service Tools window (see Figure 4-12).

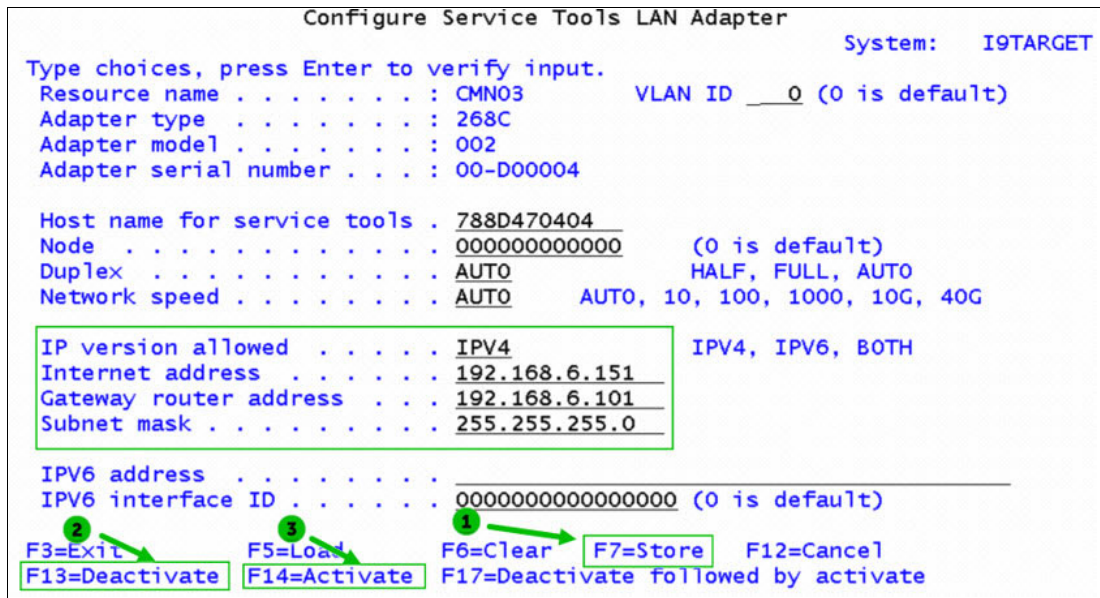


Figure 4-12 Set up TCP IP on LAN Adapter

8. Work with IP IPv4 Connection Status and verify the status of port 3000 status:
 - 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 Service Tools Server.
9. 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/sysip1'.
10. Run the WRKCFGSTS *DEV INSTALL command. Select **Option 1= Vary On for device INSTALL**.
11. Run the WRKIMGCLGE IMGCLG(*DEV) DEV(INSTALL) command to verify that you can access the remote image catalog.
12. Verify the Catalog, Type, and Directory by running the /install/sysip1 command. Verify the status of the volumes (Mounted or Loaded).

Installing LIC and operating system on IBM i PowerVS on client server by NFS server

Complete the following steps to install the LIC and operating system on IBM i PowerVS on client server by using NFS server:

Warning: Before you begin the scratch installation on the (TARGET) IBM i VSI, document all your network information. You must re-create the network information after the installation completes. 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 that are described at this [IBM Documentation web page](#).
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) by using the following settings:

- Server IP: IP address of the SOURCE NFS IBM i VSI Server.
- Path Name: Name of the Directory where the image volumes are located.

Complete the following steps:

- Select **F10 =Continue**.
- Press **Enter** to confirm.

Network Device - Configuration System:

Server IP	: 192.168.6.117	
Path Name	: /install/sysipl	

NOTE: It is required that the service tool lan adapter is configured and active.

Press F10 to continue install with the present configuration

F3=Exit F10=Continue

Figure 4-14 Configure the network device

- Check whether the language group is correct. Then, press **Enter** twice, and confirm.
- 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.
- At the command line, enter GO LICPGM.
- Select **Option 10=Display Installed Licensed Programs**.
The Base IBM i that is installed and the Library QGPL and QUSRSYS are *BACKLEVEL.
- Recover the IBM BRMS product and associated libraries on the IBM i PowerVS 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'
- Run the **WRKCFGSTS *DEV INSTALL** command. Then, select **Option 1= Vary On for device INSTALL**.
- Verify that you can access the remote image catalog by running the following command:
WRKIMGCLGE IMGCLG(*DEV) DEV(INSTALL)
- 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 Item	Save Type	----- ASP Name	----- Number	Save Date	Save Time	Saved	Not Saved	Sequence Number
QBRM	*FULL	*SYSBAS	00001	7/06/20	10:37:08	1184	0	48
		OPTFILE('/IBM BRMSOPTSAV2248801200706103706I922BRMC/QBRM')						
QMSE	*FULL	*SYSBAS	00001	7/06/20	10:37:09	8	0	49
		OPTFILE('/IBM BRMSOPTSAV2248801200706103706I922BRMC/QMSE')						
QUSRBRM	*FULL	*SYSBAS	00001	7/06/20	10:37:41	264	0	58
		OPTFILE('/IBM BRMSOPTSAV2248801200706103708I922BRMC/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(Q06990) ENDOPT(*REWIND) OPTFILE('/IBM BRMSOPTSAV2248801200706103706I922BRMC/QBRM')
RSTLIB SAVLIB(QMSE) DEV(INSTALL) VOL(Q07898) ENDOPT(*REWIND) OPTFILE('/IBM BRMSOPTSAV2248801200706103706I922BRMC/QMSE')
RSTLIB SAVLIB(QUSRBRM) DEV(INSTALL) VOL(Q08807) ENDOPT(*REWIND) OPTFILE('/IBM BRMSOPTSAV2248801200706103708I922BRMC/QUSRBRM')
```

15. Complete Step 005 in your IBM BRMS Recovery Report.

You must 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 Item	Save Type	----- ASP Name	----- Number	Save Date	Save Time	Saved	Not Saved	Sequence 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 BRMSOPTSAV2248801200706103708I922BRMC/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 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.

```

Start Recovery using BRM (STRRCYBRM)

Type choices, press Enter.

Option . . . . . > *SYSTEM      *SYSTEM, *ALLDLO, *ALLUSR...
Action . . . . . > *RESTORE     *REPORT, *RESTORE
Time period for recovery:
  
```

Figure 4-15 Sample of STRRCYBRM

21. Select **F9= Restore Command Defaults** and make the changes that are shown in Figure 4-16.

```

Restore Command Defaults

Type information, press Enter.

Device . . . . . → INSTALL      Name, *MEDCLS
                    _____
                    _____
                    _____

Parallel device resources:
  Minimum resources. . . . . *SAV      1-32, *NONE, *AVAIL
  Maximum resources. . . . . _____ 1-32, *AVAIL, *MIN
  End of media option . . . . . *REWIND *REWIND, *LEAVE, *UNLOAD
  Option . . . . . *ALL          *ALL, *NEW, *OLD, *FREE
  Data base member option . → *ALL      *MATCH, *ALL, *NEW, *OLD
  Restore spooled file data . → *NEW      *NONE, *NEW
  Allow object differences . → *COMPATIBLE *NONE, *ALL, *COMPATIBLE,
                    _____ *AUTL, *FILELVL, *OWNER,
                    _____ *PGP
                    _____
  
```

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 then, press **Enter**.

Note: If you receive a message that reads: Waiting for reply to message on message queue QSYSOPR, select **SysReq** and then, **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 then, press **Enter** (see Figure 4-17).

Select Recovery Items								I922BRMC
Type options, press Enter.								Select action : *ALL
1=Select 4=Remove 5=Display 7=Specify object								Select volume : _____
Opt	Saved Item	Save Date	Save Time	Save Type	Parallel Devices	Volume Serial	File Sequence	Expire 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 then, press **Enter**.

27. Display the remaining objects during the restore (see Figure 4-18).

Select Recovery Items								I922BRMC
- Display Recovery Items								I922BRMC
: Remaining items : 8								17:57:35
: Remaining objects : 2,511								
: Remaining size : 4,967.2540 M 98.7 %								
Saved Item	Date	Time	Save Type	Volume Serial	File Seq	Exp Date	Objects 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 connects 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 Cloud Object Storage (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 then, 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-backupv01

Resource URI . . . . . : s3.us-east.cloud-object-storage.appdomain.cloud
```

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 Cloud Object Storage 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) RSRNAME(*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) STRTPPRF(*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 SBS(DQCTL)
STRTCP
```

34. Move volumes from IBM Cloud Object Storage to the TARGET IBM i VSI Client:

- Run the **WRKMEDBRM** command and then, press **Enter**.

Find your volumes that is listed your IBM BRMS Recovery Report.

If you volume has a plus (+) to the right, it is part of a serial set.

- Select **Option 6=Work with serial set**, as shown in Figure 4-20.

Notice that the volumes are in the IBM Cloud Object Storage 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 ...

Opt	Volume Serial	Status	Creation Date	Expiration Date	Location	Move Date	Media Class	Dup 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 then, press **Enter**.

Work with Serial Set									
Type options, press Enter.									
2=Change 5=Display 7=Expire 8=Move 9=Remove volume error status									
10=Reinitialize 11=Mark for Label print									
System:	I922BRMC								
Opt	Volume Serial	Volume Seq	Status	Expiration Date	Location	Move Date	Media Class	Dup Sts	
8	Q13692	0001	*ACT	07/27/20	TOR1CLD	07/09/20	QCLDVRTTAP		
8	Q32656	0002	*ACT	07/27/20	TOR1CLD	07/09/20	QCLDVRTTAP		
8	Q15277	0003	*ACT	07/27/20	TOR1CLD	07/09/20	QCLDVRTTAP		
8	Q31632	0004	*ACT	07/27/20	TOR1CLD	07/09/20	QCLDVRTTAP		

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.

Move Media									
Type storage location or container to receive volumes.									
Storage location . .	*HOME	*SAME, F4 for list							
Container	*SAME	*SAME, *NONE, F4 for list							
Exempt from movement	*SAME	*SAME, 0 - 9999							
Opt	Volume Serial	Status	Creation Date	Expiration Date	Location	Move Date	Media Class	Dup Sts	
8	Q13692	+ *ACT	07/06/20	07/27/20	TOR1CLD	07/09/20	QCLDVRTTAP		
8	Q32656	+ *ACT	07/06/20	07/27/20	TOR1CLD	07/09/20	QCLDVRTTAP		
8	Q15277	+ *ACT	07/06/20	07/27/20	TOR1CLD	07/09/20	QCLDVRTTAP		
8	Q31632	+ *ACT	07/06/20	07/27/20	TOR1CLD	07/09/20	QCLDVRTTAP		

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 Cloud Object Storage:

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 then, 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.

Select Recovery Items							I922BRMC	
							Select action	: *ALL
							Select volume	: _____
Type options, press Enter.								
1=Select 4=Remove 5=Display 7=Specify object								
Opt	Saved Item	Save Date	Save Time	Save Type	Parallel Devices	Volume Serial	File Sequence	Expire Date
4	QICC	7/06/20	10:37:12	*FULL		Q07898	50	7/27/20
4	QUSRICC	7/06/20	10:37:12	*FULL		Q07898	51	7/27/20
4	QIWS	7/06/20	10:37:13	*FULL		Q08807	55	7/27/20
4	QSYSDIR	7/06/20	10:37:13	*FULL		Q08807	56	7/27/20
4	QTCP	7/06/20	10:37:13	*FULL		Q08807	54	7/27/20
-	#CGULIB	7/06/20	9:23:31	*FULL		Q13692	1	7/27/20
-	#COBLIB	7/06/20	9:23:31	*FULL		Q13692	2	7/27/20
-	#DFULIB	7/06/20	9:23:32	*FULL		Q13692	3	7/27/20
-	#DSULIB	7/06/20	9:23:32	*FULL		Q13692	4	7/27/20
-	#RPGLIB	7/06/20	9:23:32	*FULL		Q13692	5	7/27/20
-	#SDALIB	7/06/20	9:23:33	*FULL		Q13692	6	7/27/20
								More...
F3=Exit F5=Refresh F9=Recovery defaults F11=Object View								
F12=Cancel F14=Submit to batch F16=Select								

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 Cloud Object Storage:

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 then, 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.

Select action : *ALL
Select volume : _____

Type options, press Enter.
1=Select 4=Remove 5=Display 7=Specify object

Opt	Item	Save Date	Save Time	Item Type	Objects Saved	Objects Not Saved	Control Group	Encrypted
—	QS36F	7/06/20	9:25:14	*LIB	1	0	QCLDBSYS01	
—	QUSRDIRCF	7/06/20	9:25:14	*LIB	3	0	QCLDBSYS01	
—	QUSRDIRDB	7/06/20	9:25:14	*LIB	191	0	QCLDBSYS01	
—	QUSRIJS	7/06/20	9:25:14	*LIB	104	0	QCLDBSYS01	
—	SYSIBM	7/06/20	9:25:16	*LIB	66	0	QCLDBSYS01	
—	SYSIBMADM	7/06/20	9:25:16	*LIB	17	0	QCLDBSYS01	
—	SYSPROC	7/06/20	9:25:16	*LIB	3	0	QCLDBSYS01	
—	SYSTOOLS	7/06/20	9:25:16	*LIB	60	0	QCLDBSYS01	
4	QUSRSYS	7/06/20	10:37:08	*LIB	1680	0	QCLDBIPL01	
4	QSYS2	7/06/20	10:37:12	*LIB	343	0	QCLDBIPL01	
4	QGPL	7/06/20	10:37:13	*LIB	165	0	QCLDBIPL01	

Bottom

Figure 4-24 Sample of item to be removed created by QCLDBIPL01

45.Select the saved items:

- Review the list again.
- Select **Option 1=Select for each item** or **F16= Select (this will select all the items for you)**.
- 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 Cloud Object Storage, 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.

							Select action : <u>*ALL</u>
							Select volume : _____
Type options, press Enter.							
1=Select 4=Remove 5=Display 7=Specify object							
Opt	Saved Item	Save Date	Save Time	Save Type	Parallel Volume Devices Serial	File Sequence	Expire Date
<u>1</u>	*ALLDLO	7/06/20	9:25:34	*FULL		Q32656 159	7/27/20

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 Cloud Object Storage.

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:

```
PWRDWN SYS 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 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 Cloud Object Storage.

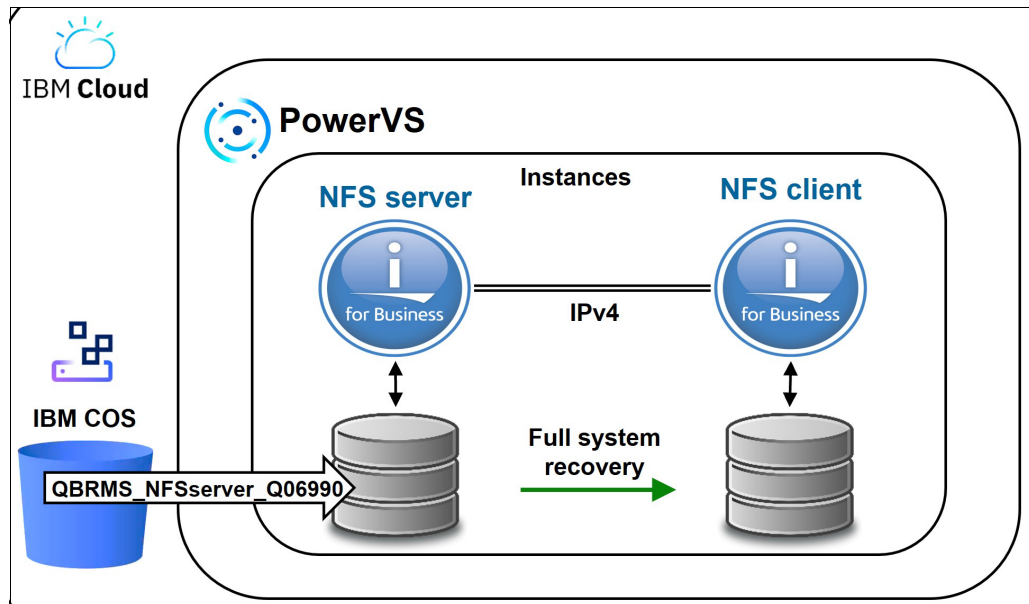


Figure 4-26 Full system recovery from the cloud by using IBM i

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 our 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 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, “Dynamic Solutions International Virtual Tape Library” on page 132).

This virtual media that is produced must be stored and archived. IBM Cloud Object Storage is the native solution to save backups with almost endless capacity.

Communicating with IBM Cloud Object Storage can be done by using the IBM i Cloud Storage Solutions software offering, or the use of the Open Source capacity in the IBM i operating system.

In this section, we discuss how to back up data to a save file or an image catalog by using native commands. Also discussed is how to save and restore data from IBM Cloud Object Storage. Moreover, a library can be saved and transferred to the cloud by the way of IBM BRMS, and the movement to the cloud can be automated.

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 Read authority 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; that is, 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 much time 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; for example, by scheduling 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 then, press **F4**.
2. Complete the required fields as listed in Table 4-1 and then, press **Enter**:

Table 4-1 CPYTOCLD required fields

Field	Action
Resource Name	Enter the name of a resource.
Submit to Batch	enter *N0 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	<p>Enter a path and name for the cloud copy of the file; for example, dir1/dir2/file.txt.</p> <p>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.</p> <p>The container, bucket, or root directory that is defined in the resource must exist before you copy files to the cloud.</p> <p>When overwriting a file, in most cases 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 /Q0pensys file system.</p>

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. |

If 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 Cloud Object Storage 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; that 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 a long time for large files to finish copying before running other commands.

Also, the IBM i facilities can be used to work with asynchronous jobs; for example, by 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 then, press **F4**.
2. Complete the required fields as listed in Table 4-2, and then, press **Enter**.

Table 4-2 CPYFRMCLD required fields

Field	Action
Resource Name	Enter the name of a resource.
Submit to Batch	Enter *N0 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	<p>Enter the cloud path and the name of the file to be copied; for example: dir1/dir2/file.txt.</p> <p>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.</p> <p>In most cases, 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 /Q0pensys file system.</p>
Local file name	<p>Enter the IFS path and file name of the file being copied, for example: /home/user/jdoe/file.txt.</p> <p>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.</p>

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.

This method is simple and includes 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 used to restore License Machine Code (LIC or SLIC).
- ▶ The maximum size for a save file on IBM i V7R1 and later is 2 TB.
- ▶ The 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 backup after the backup completes by using the **JAR** command or Open Source tools.
- ▶ Data in save files can be uploaded to physical tapes or virtual by using **SAVSAVFTA** command.
- ▶ Parallel backups can be performed, which save to multiple save files simultaneously.
- ▶ Backups to multiple save files cannot be spanned by using one command. For example, if 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 security information are saved in two different Save Files, the data is compressed, and then uploaded to IBM Cloud Object Storage, assuming communications exist with Direct Link Connect and CentOS reverse proxy.

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 gzip gunzip python3 python3-pip readline curl
pip3 install awsccli
```

5. Configure `awsccli` to communicate with IBM Cloud Object Storage by using `aws configure`.
6. Use your credentials from your IBM Cloud Object Storage bucket to complete the configuration.

For more information about how to create your IBM Cloud Object Storage resource, bucket, and credentials and configure AWS CLI, see 2.2.2, “Full system backup and restore by using native commands” on page 19 or this [IBM Cloud Docs web page](#).

7. Compress the files and upload to IBM Cloud Object Storage, as shown in Example 4-10.

Example 4-10 Saving and compressing data to Cloud Object Storage 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/
```

Displaying sample script on a menu

Note: This sample code is hosted on an external site to IBM. IBM assumes no responsibility on this sample code. Code is provided as-is with no warranty and can vary over time.

This sample code from GitHub features a sample menu system that was written in BASH. IBM i images on IBM Power Systems Virtual Server include a preinstalled YUM package that enables this code to be run. For more information, see this [GitHub web page](#).

Complete the following steps:

1. Sign in to your system by using an *ALLOBJ capable user profile.
2. Configure OpenSSH environment with BASH:
 - a. Edit `sshd_config` file, as shown in Example 4-12.

Example 4-12 Edit the configuration file sshd_config

```
EDTF '/QOpenSys/QIBM/UserData/SC1/0penSSH/etc/sshd_config'
```

- b. Add the lines that are shown in Example 4-13 to the end of the file.

Example 4-13 Add lines on sshd_config file

```
ibmpaseforienv PASE_USRGRP_LIMITED=N  
ibmpaseforishell=/QOpenSys/pkgs/bin/bash
```

3. Uncomment and change the following line (optional):

from:

```
#UseDNS yes
```

to:

```
UseDNS no
```

4. Update the YUM environment and install some tools, as shown in Example 4-14.

Example 4-14 Update YUM and install tools

```
CALL QP2TERM  
mkdir /home/$LOGNAME  
TERM=aixterm  
export TERM  
PATH=$PATH:/QOpenSys/pkgs/bin  
export PATH  
yum -y update  
yum -y install p7zip pigz gzip gunzip python3 python3-pip readline curl  
pip3 install awscli
```

5. Configure `awscli` to communicate with IBM Cloud Object Storage by using `aws configure`.
6. Use your credentials from your IBM Cloud Object Storage bucket to complete the configuration.

For more information about how to create your IBM Cloud Object Storage resource, bucket, and credentials and configure AWS CLI, see 2.2.2, “Full system backup and restore by using native commands” on page 19 or this [IBM Cloud Docs web page](#).

7. Create the `.ssh` directory and some basic configuration files, as shown in Example 4-15.

Example 4-15 Generate a SSH key

```
ssh-keygen
#Enter to all requests
cd
echo 'export PS1="\u@[\e[32m\]\H[\e[m\]:\w>"' > .bash_profile
echo 'export TERM="aixterm"'>> .bash_profile
echo 'export PATH=$PATH:/QOpenSys/pkg/bin'>> .bash_profile
echo 'export LC_CTYPE=ES_MX'>> .bash_profile
echo 'export QIBM_MULTI_THREADED="Y"'>> .bash_profile
echo 'export LC_ALL=EN_US'>> .bash_profile
chmod +x .bash_profile
#Exit pression F3
```

8. Restart the *SSHD server, as shown in Example 4-16.

Example 4-16 Start and stop TCP/IP Server on IBM i

```
ENDTCPSVR *SSHD
STRTCPSVR *SSHD
```

9. Follow the guidelines that are included in the README.md file.
10. Download the source code to your system, as shown in Example 4-17.

Example 4-17 Git clone on QP2TERM

```
CALL QP2TERM
/QOpenSys/pkg/bin/git clone
https://github.com/dkesselman/IBMi_Cloud_Backup.git
```

11. Connect to your system by using SSH and the *ALLOBJ-capable user profile. You can use PuTTY or OpenSSH client on your system.

On PuTTY, click **Window** → **Translation** → **Remote Character set** → **Add Use font encoding**, as shown in Figure 4-27.

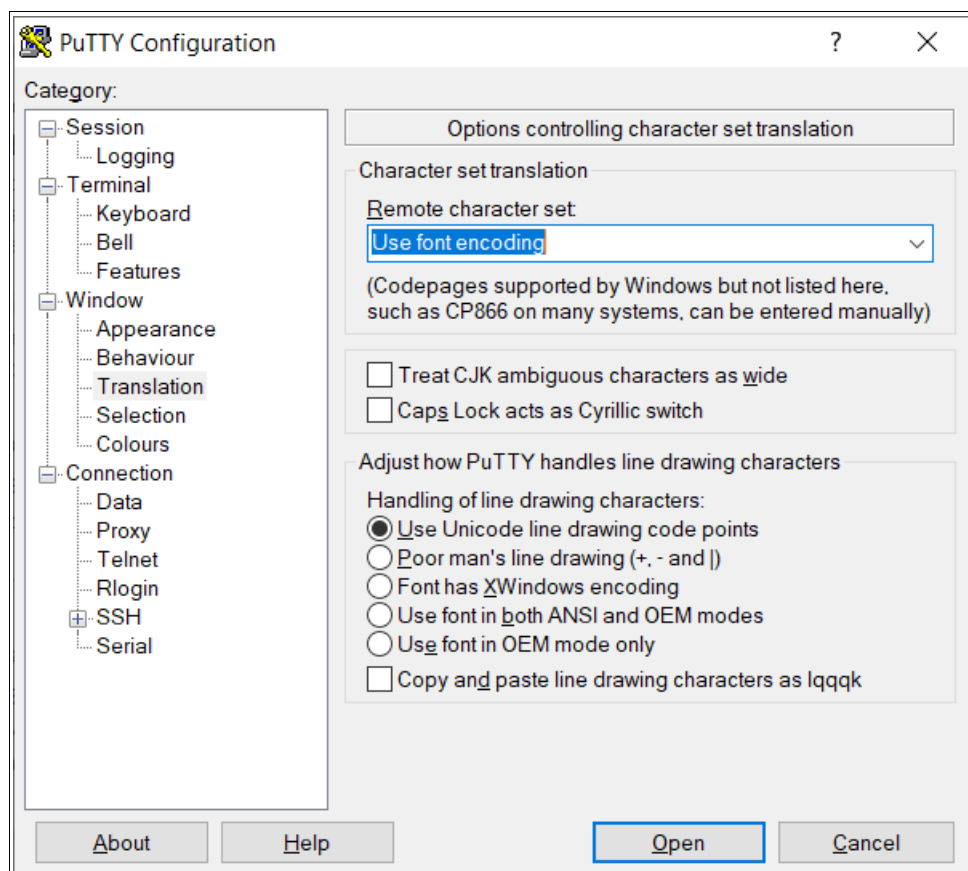


Figure 4-27 Character set translation on PuTTY

12. Edit variables in your environment by using the vi editor, as shown in Example 4-18.

Example 4-18 Edit using vi

```
cd IBMi_Cloud_Backup
vi mnus3_const.sh
# Constants:
# IFSPATH=Temporary space for compressed files
# BUCKETDFT=The name of your bucket
# S3CMD=The command we run. The default endpoint is using US-SOUTH region. You can
change this based on your Endpoint information
```

13. Change the execution permission to **chmod +x mnus3.sh**.

14. Open the Backup menu or the IBM Cloud Object Storage (COS) Backup Menu (see Figure 4-28): `./mnus3.sh`.

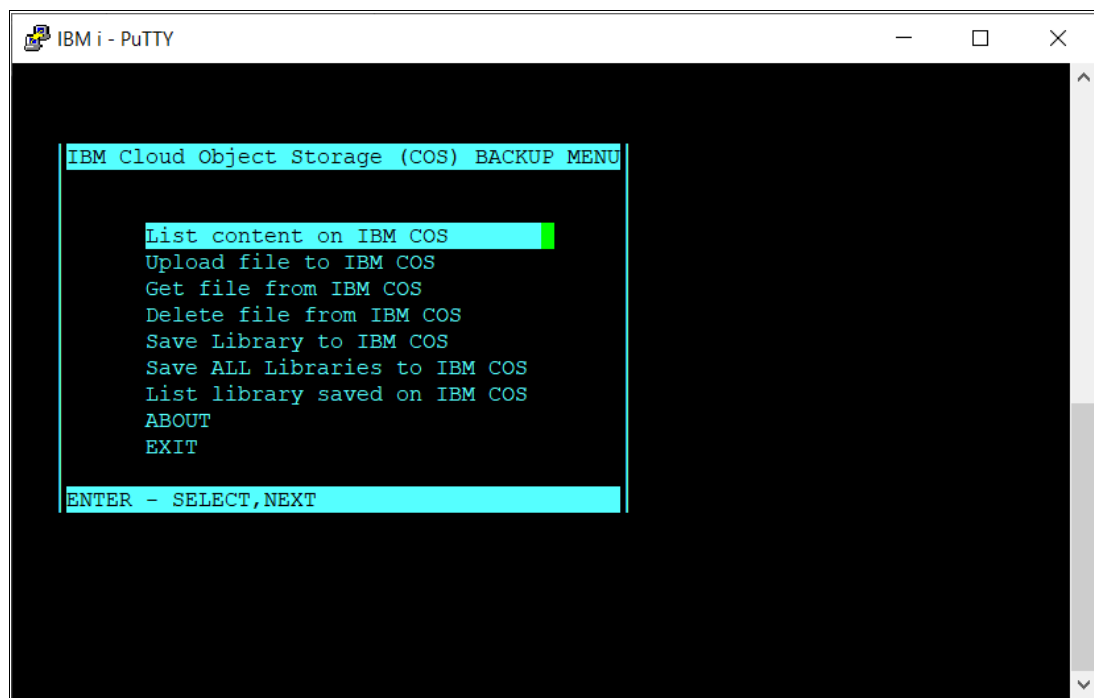


Figure 4-28 Sample tool main menu

The menu includes the following options:

- List content on IBM COS: List buckets and their is content.
- Upload file to IBM COS: Select a specific file to upload from IFS.
- Get file from IBM COS: Retrieve the file.
- Delete file from IBM COS: Delete the file from bucket.
- Save library to IBM COS: Create a Save File, and then, compress and upload the file.
- Save ALL Libraries to IBM COS: This option creates a list of user libraries and starts the backup, compression, and upload process.
- List library saved on IBM COS: An experimental option, it lists the Save File content from the output in CSV format.

Note: This solution is process of a backup solution. Instead, it is a sample script or menu to demonstrate IBM Power Systems Virtual Server capabilities to save on IBM Cloud. In this example, the public endpoint was used. No restore option is available. The file is available from IBM Cloud Object Storage only.

4.3.5 Backup by using image catalog and IBM Cloud Object Storage

The image catalog is a feature on IBM i as an image device from V5R1. From V5R3, you can save to your disk by using optical images; from V5R4, you can create Virtual Tape images.

This useful feature can be used to install PTFs, perform system upgrades, and save our 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 this [IBM Support web page](#).

When backing 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 is created.
- ▶ Managing multiple backup sets can be complicated. Consider the use of IBM BRMS to create backups and 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 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 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 4-19.

Example 4-19 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(10240) TEXT('Optical Backup Image 1 - 10GB') 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 4-20.

Example 4-20 Create virtual optical device

```
CRTDEVOPT DEVD(VOPTBK) RSRNAME(*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 4-21.

Example 4-21 Save security data to the virtual optical device

```
SAVSECDA DEV(VOPTBK)
```

When the virtual optical media must be transferred to IBM Cloud Object Storage, use IBM Cloud Storage Solutions for i or the same procedure, as shown in Example 4-22.

Example 4-22 Upload file with IBM Cloud Storage Solutions for i

```
CPYTOCLD RESOURCE(ICOS01) LOCALFILE('/IMGCLG/OBACKUP/OBKP01') CLOUDFILE(OBACKUP01)
```

The file is transferred to resource ICOS01.

Using Example 4-22 with the configured AWS CLI, upload the file with **gzip** compression and list content to verify the correct file transfer, as shown in Example 4-23.

Example 4-23 Transferring file to IBM Cloud Object Storage 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 OBKP01 | pigz -9 -p40 -c | aws
--endpoint-url=https://s3.eu-gb.cloud-object-storage.appdomain.cloud s3 cp -
s3://ibmi-backup/OBKP01.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 QSHHELL and include `awsc1i` in the CLs to upload data to IBM Cloud Object Storage. 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 OBKP01 | pigz -9 -p40 -c | aws
--endpoint-url=https://s3.eu-gb.cloud-object-storage.appdomain.cloud s3 cp -
s3://ibmi-backup/OBKP01.gz" ')) JOB(S3CP)
```

4.3.6 Sample save and restore IBM i objects to IBM Cloud Object Storage

In this section, we introduce IBM Cloud Object Storage. You also learn how to copy a save file directly from your IBM i partition to your own bucket in IBM Cloud Object Storage.

This process can be useful not just as a backup solution, but also as a simple migration strategy; especially to get started and move a few applications and many small databases to IBM PowerVS to get started with your own PoC or testing.

Therefore, it can be easy to 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 we use the public interface to the IBM Cloud Object Storage Bucket in this example.

On IBM i, and in this sample, we 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 Cloud Object Storage and the IFS. These IFS files can be used with image catalogs to automate the process (as does IBM BRMS).

Although Backup, Recovery, and Media Services (and IBM BRMS Network in an HA/DR scenario) is not used in this example, if you can use IBM Cloud Storage without IBM BRMS, IBM BRMS is fully integrated with IBM Cloud Storage and can be of great 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 Cloud Object Storage.
3. Verify the bucket contents.
4. Drop the file to the IBM Cloud Object Storage bucket.

Backing up IBM i object to IBM Cloud Object Storage

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 Cloud Object Storage 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 then, continue to complete other information (see Figure 4-29).

Create ICC AWS S3 Resource (CRTS3RICC)		
Type choices, press Enter.		
Resource name	CO SRSC	Name
Resource description	CO SRSC	
Access key id	bea445a4d41f40d597ba859bbb3401d6	
Secret access key	ef786ab2c173f24f43e14cd75f0b4ab35526358788f00e49	
Use compression	*NO	*NO, *YES
Use encryption	*NO	*NO, *YES

Figure 4-29 Create an IBM Cloud Object Storage resource by way of CRTS3RICC

4. Enter your Bucket Name and your Resource URI, as shown in Figure 4-30. Press **Enter**.

Bucket	> 'cos-bucket-ibmi01'
Resource URI	> 's3.us-east.cloud-object-storage.appdomain.cloud'

Figure 4-30 Bucket name and resource URI

The resource is created successfully, and the following message received:

```
Resource CO SRSC was created for resource type AWSS3.
```

5. From the console, run the **ENDTCP** command.

6. Begin over again and run:

```
STRTCP
STRTCPSVR *TELNET
STRTCPSVR *SSHD
```

Note: You must restart your SSH Tunnel connections.

7. Start the Cloud Object Storage Subsystem:

```
STRSBS SBSD(QICC/QICCSBS)
```

8. Copy the file to your IBM Cloud Object Storage:

```
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 4-31. Browse to **Buckets** and **Objects**.

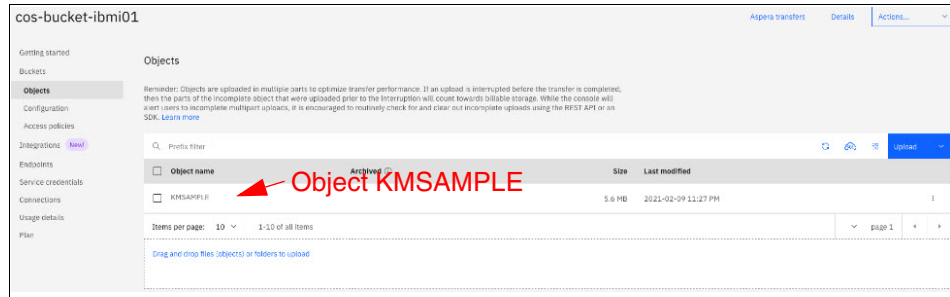


Figure 4-31 IBM i object on IBM Cloud Object Storage

Restoring IBM i object from IBM Cloud Object Storage

In this section, we describe restoring from IBM Cloud Object Storage. Combined with 4.3.6, “Sample save and restore IBM i objects to IBM Cloud Object Storage” on page 124, you now have your first application migration strategy.

You save your applications and data into a save file, move to IBM Cloud Object Storage and then, restore from IBM Cloud Object Storage. As a prerequisite, complete the steps that are described in “Backing up IBM i object to IBM Cloud Object Storage” on page 125.

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 IBM Cloud Object Storage bucket.

The objective is to use an IBM Cloud Object Storage service 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 it works, you are encouraged to create your own IBM Cloud Object Storage service from the IBM Cloud Catalog, your first bucket and credentials, and get started with your first backup and restore.

Complete the following steps to transfer and restore a library from IBM Cloud Object Storage to the Integrated File System:

1. Run the **CPYFRMCLD** command (see Figure 4-32) and complete 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 Cloud Object Storage Bucket.
In our example, `acmeair.file` is the name of the save file that is in the IBM Cloud Object Storage service. It was sent to IBM Cloud Object Storage 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; for example, your home directory `/home/QSECOFR/acmeair.file`. If the file exists, it is overwritten.

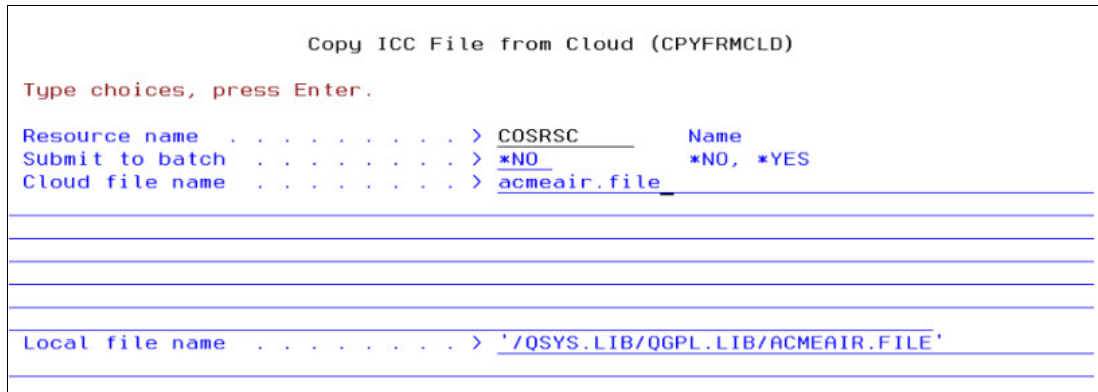


Figure 4-32 Copy IBM Cloud Storage file from cloud

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 4-33:

```
DSPSAVF FILE(QGPL/ACMEAIR)
```

Display Saved Objects						
Library saved : ACMEAIR						
Type Options, press Enter.						
5=Display						
Opt	Object	Type	Attribute	Owner	Size (K)	Data
—	ACMEAIR	*LIB	PROD	QSECOFR	112	YES
—	QSQJRN	*JRN		QSECOFR	12	YES
—	BOOKI00001	*PGM	CLE	QSECOFR	164	YES
—	AIRP000001	*FILE	PF	QSECOFR	184	YES
—	BOOKING	*FILE	PF	QSECOFR	216	YES
—	CUSTOMER	*FILE	PF	QSECOFR	3260	YES
—	CUST000001	*FILE	PF	QSECOFR	348	YES
—	FLIGHT	*FILE	PF	QSECOFR	5460	YES
—	FLIGH00001	*FILE	PF	QSECOFR	228	YES
—	PUSHM00001	*FILE	PF	QSECOFR	200	YES
—	SYSCHKCST	*FILE	LF	QSECOFR	36	YES
—	SYSCOLUMNS	*FILE	LF	QSECOFR	108	YES
—	SYSCST	*FILE	LF	QSECOFR	72	YES

Figure 4-33 Display the objects restored from IBM Cloud Object Storage

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:


```
CRTDEV TAP DEVD(IBM BRMSCLDTAP) RSRNAME(*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 from the cloud.

4.3.8 Automatically transferring media to IBM Cloud Object Storage

IBM BRMS provides specific control groups that can be used to automatically save data to media in the cloud. Then, 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 Cloud Object Storage are shown in Figure 4-34.

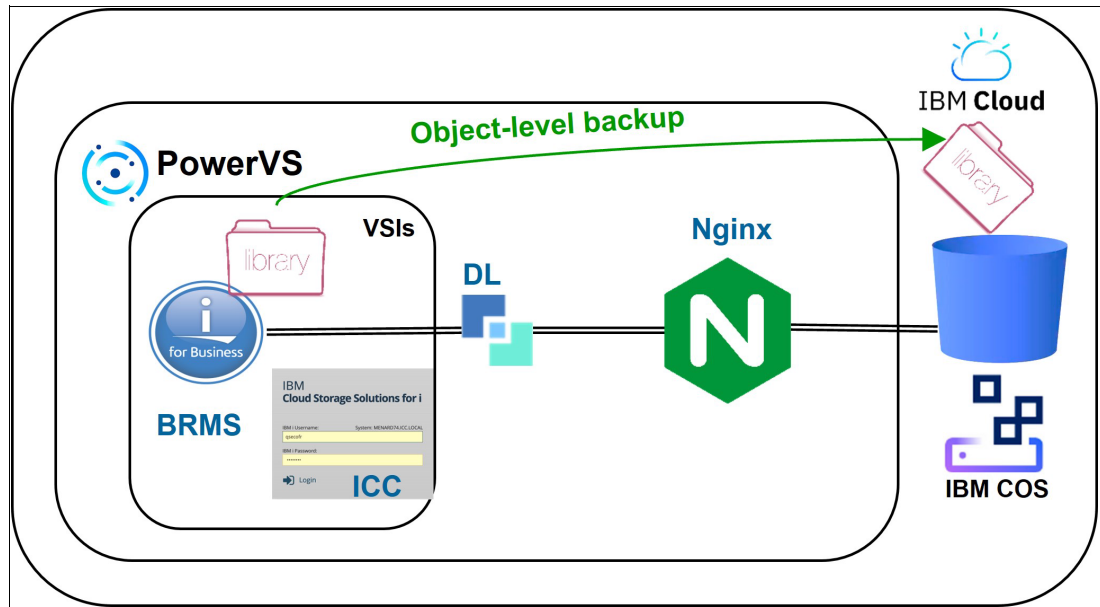


Figure 4-34 Automatic transfer of media to IBM Cloud Object Storage

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 then, 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 4-35.

```

Edit Backup Control Group Entries                                I750CL21
Group . . . . . : CLOUDBKU
Default activity . . . . . : *BKUPCY
Text . . . . . : Entry created by BRM configuration
Type information, press Enter.

Seq  Backup  List ASP  Weekly  Retain Save  SWA  Sync
     Items   Type Device Activity Object While  Message  Queue  ID
-----
10  *EXIT
20  CLDLIB   *SYSBAS  *DFTACT  I I I I I I I I  *ERR  *LIB  *LIB  *NONE
30  *EXIT
  
```

Figure 4-35 Using default values in remaining fields

8. Issue the control group, as shown in Figure 4-36.

```

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 4-36 Issuing the control group

As shown in Figure 4-37, the volume status is *ACT and the location point is cloud resource.

```

Work With Media
Position to . . . . . Starting characters System: I750CL21
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 ...

Opt  Volume  Creation Expiration  Move  Media  Dup
     Serial  Status  Date  Date  Location  Date  Class  Sts
-----
  BRMCL2  *ACT  12/09/16  12/30/16  CLOUD  12/09/16  BRMSCLDTAP
  
```

Figure 4-37 Volume status and location

4.4 Dynamic Solutions International Virtual Tape Library

This section describes the Dynamic Solutions International (DSI) Virtual Tape Library (VTL).

4.4.1 Introduction

IBM i customers require backups and many use tape or virtual tape, which requires a software-based VTL. VTL is a requirement for IBM Power Systems Virtual Server to be successful with IBM i.

Today, you can use a VTL to back up IBM i virtual machines in IBM Power Systems Virtual Server by going directly to DSI. For more information, see this [Dynamic Solutions web page](#).

4.4.2 Overview

The DSI VTL solution is a backup and recovery hardware and software solution for IBM i systems. The solution implements a backup solution that uses DSI VTL iSCSI attached device to provide a NFS as a DR solution.

Virtualized data protection system

The Data Protection System includes the following characteristics:

- ▶ Virtualized VTL/Dedupe system
- ▶ Local or cloud VM
- ▶ All features of appliance VTLs:
 - Tape emulation and deduplication
 - Secure replication
 - iSCSI support, which is ideal for IBM i and PowerVS backup and recovery
- ▶ Replicate to and from hardware VTL
- ▶ Archive and export direct to cloud storage

Virtualized environments

Restore can run in the following on-premises virtualized (VM) environments:

- ▶ IBM Power Systems LPAR/VM
- ▶ Red Hat Enterprise Linux 7 or 6
- ▶ VMware:
 - ESXi 5.x, 6.x
 - OEL 7 or 6
- ▶ Microsoft Hyper-V: OEL 7 or 6

Cloud environments

Restore can run in the following cloud virtual machines:

- ▶ IBM Power Systems Virtual Server: Red Hat Enterprise Linux 7 or 6
- ▶ Amazon Web Services (AWS): OEL 7 or 6
- ▶ Microsoft Azure: OEL 7 or 6

Figure 4-38 on page 133 shows the cloud environment for an IBM i iSCSI that uses DSI VTL.




Figure 4-38 Cloud environment by Dynamic Solutions

Building a restore VM

The following VM resources are available:

- ▶ Virtual processors
- ▶ Memory
- ▶ Storage (operating system, VTL database, VTL landing space, deduplicated storage, and deduplication index)
- ▶ Network (if iSCSI is used)
- ▶ DSI for sizing help
- ▶ DSI-provided sample OVF (expected to be modified)

Note: The following supported IBM i backup tools can be used with DSI VTL:

- ▶ IBM Backup, Recovery, and Media Services for i (see [IBM BRMS](#)).
- ▶ Native save/restore IBM i commands.
- ▶ Robot Save, for more information, see [Robot/SAVE by HelpSystems](#).
- ▶ LXI/MMS, see [LXI](#).
- ▶ DSI Tracker, further information, refer to [DSI Tracker:Media Management System for IBM i](#).

IBM i does not support DSI VTL iSCSI attached devices as an installation device.

4.4.3 IBM BRMS requirements for DSI VTL iSCSI

IBM BRMS requires the licensed program products that are listed in Table 4-3.

Table 4-3 IBM BRMS products required for DSI VTL iSCSI

Product	Option	Description
5770-BR1	BASE	IBM Backup, Recovery and Media Services for i.
5770-BR1	2	BRMS - Advanced Feature.

IBM BRMS also requires the following PTFs:

- ▶ IBM i 7.2, SI75082
- ▶ IBM i 7.3, SI75083
- ▶ IBM i 7.4, SI75084

Also, more PTFs are required to support IBM i with Dynamic Solution International Virtual Tape Library iSCSI devices. For more information, see [IBM i Support for Attaching an iSCSI VTL](#).

Note: DSI VTL iSCSI configuration and support is not owned by IBM BRMS for i. IBM Lab Services implements the DSI VTL solution, but the support is provided by Dynamic Solutions.

Setting up IBM BRMS for DSI VTL iSCSI

Ensure that the IBM BRMS-Advanced Feature is installed before issuing **INZBRM OPTION(*NFSCTLG)**. Also, enable the default control groups that define how objects are saved. IBM BRMS creates the following control groups:

- ▶ QNFSIPLFUL
- ▶ QNFSIPLINC
- ▶ QNFSSYSFUL
- ▶ QNFSSYSINC

Note: For more information about the control groups, such as QNFSIPLFUL, QNFSIPLINC, QNFSSYSFUL, and QNFSSYSINC, see this [IBM IPP wiki page](#).

Control groups are the default for SYSBASE; however, this setting can be changed for an independent auxiliary storage pool (IASP).

4.4.4 Backing up IBM BRMS by using DSI VTL iSCSI

Complete the following steps to fully back up IBM BRMS by using DSI VTL iSCSI that uses the control groups:

1. Issue the IBM BRMS QNFSSYSFUL control group to save the IBM i system user data to the DSI VTL iSCSI device.
2. Run the QNFSIPLFUL control group to a virtual optical tape device on an image catalog on IBM i, and back up the IBM i system data.
3. After the backups are taken, the QNFSIPLFUL control group automatically saves virtual optical image catalog to DSI VTL iSCSI.
4. After the backup is transferred successfully from the image catalog to the DSI VTL iSCSI, remove the image catalog and its entries.
5. Run the **STRRCYBRM** command to generate a report in case of a DR.

Figure 4-39 shows the three first steps to back up on IBM BRMS by using DSI VTL.

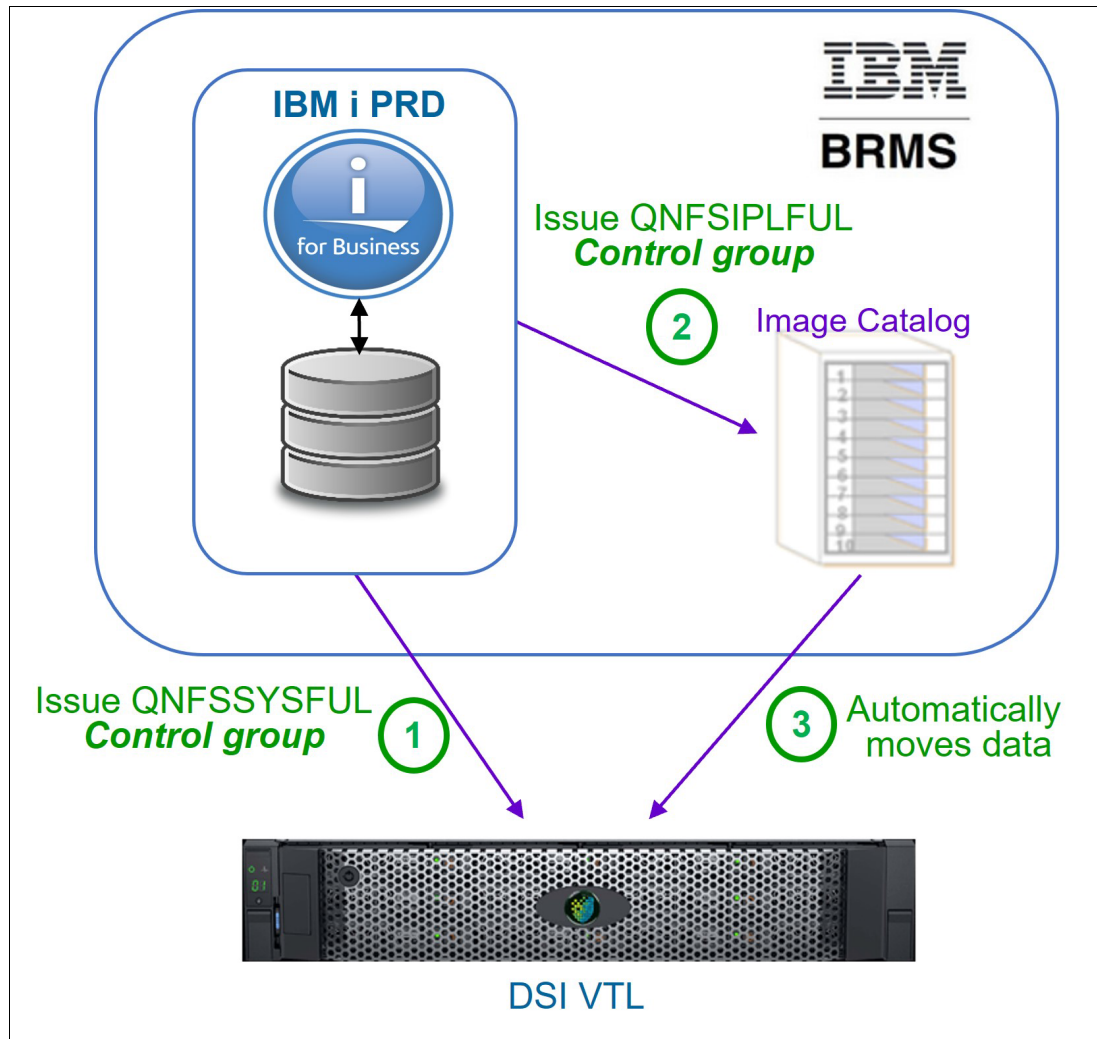


Figure 4-39 Backing up by using IBM BRMS on DSI VTL iSCSI

4.4.5 Recovering IBM BRMS by using DSI VTL iSCSI

To restore data on an IBM i NFS client LPAR, an IBM i NFS server must exist at the same release or higher than the IBM i NFS client LPAR. It also must be configured with the DSI VTL iSCSI device.

Complete the following steps to recover an IBM i NFS client LPAR:

Note: The recovery report that is cited in the volume sequence is used for the correct image files from the DSI VTL, and is restored to the IFS on IBM i NFS Server. Refer to Figure 4-40.

1. Create the image catalog and configure it as an IBM i NFS installed.
2. Verify that the remote network is configured.
3. Restore the system data and QUSRBRM information by using a virtual optical network device.
4. Use the Recovery report as a guide to restore the remainder of the user data from DSI VTL iSCSI device.

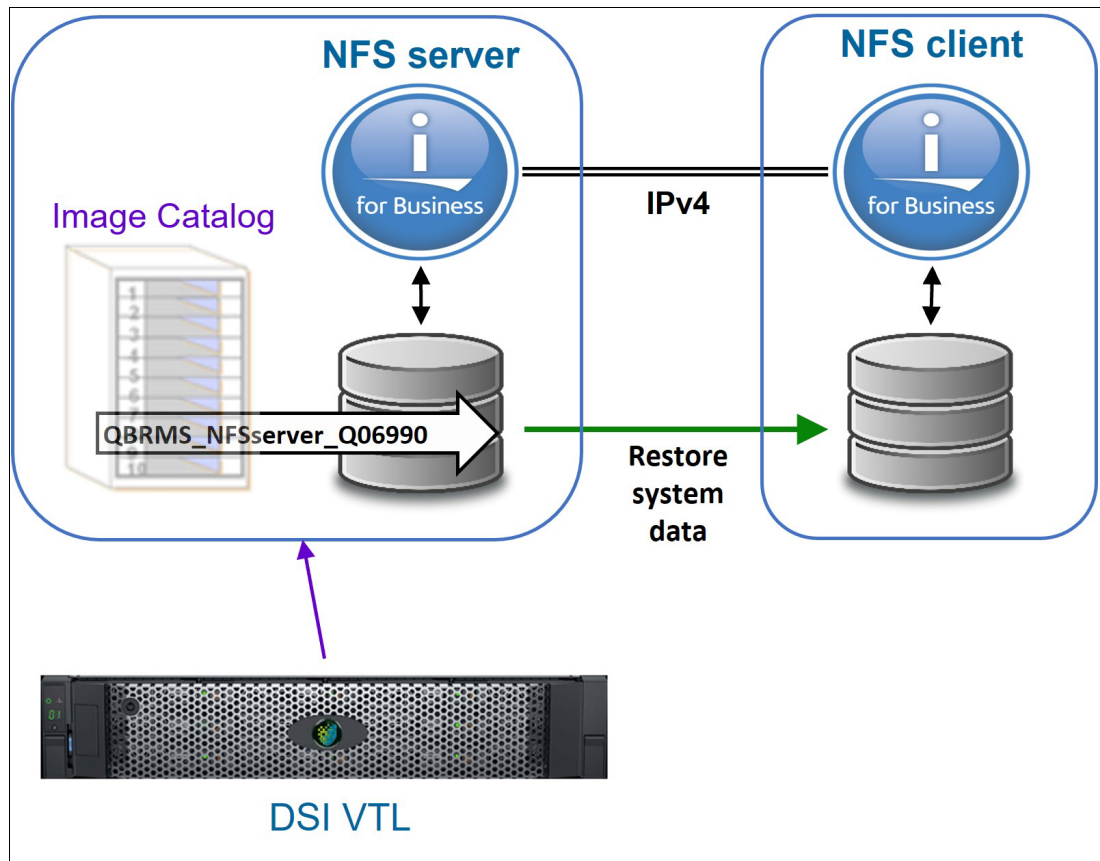


Figure 4-40 Restore using IBM BRMS and NFS server to the NFS client

4.4.6 DSI VTL restore on IBM PowerVS

DSI Restore deduplicates backup and recovery virtual software. Based on DSI's world class enterprise VTL, a restore runs on your virtual infrastructure and scales from TBs to PBs of backup data.

Licensed per TB, Restore is a fit for small businesses and global enterprises. Bandwidth efficient, encrypted data replication to a second site or cloud instance meets your compliance and DR requirements. For more information, see [this website](#) (log in required).

4.4.7 Sizing and ordering DSI VTL

For more information about sizing and ordering, see this [IBM Support web page](#).

If you need to display a TCO Calculator -VTL, see [this website](#) (log in required).

4.5 Full-system snapshot, cloning, and restore

This section describes how to obtain a full system snapshot, clone, and restore.

4.5.1 Overview

IBM Power Systems Virtual Server introduced the capability to perform snapshots, restores, and clones of Power Virtual Server Instances (VSIs).

4.5.2 Snapshot

By using the snapshot interface, a relationship can be created between source disks and a target disks (target disks are created as part of the snapshot API) at time T1. 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 want to revert the middleware to its original state before proceeding with an upgrade.

If the middleware 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 in the previous step.
4. If the upgrade succeeds, delete the snapshot that was created in the first step.

Note: Multiple snapshot operations can be conducted simultaneously. However, these concurrent snapshot operations 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 that all of the applications that are on the snapshot volume be quiesced.

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 cloned (full-copy) IBM FlashCopy® operations are running in the background. The FlashCopy operations must complete 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.

4.5.3 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.

4.5.4 Restore a snapshot

The restore operation restores all of the volumes that are part of a VM snapshot back to the source disks. Although 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, pass in 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, you cannot perform the snapshot, restore, clone, or capture operations on the other VMs that use the shared volume (while the restore operation is running).

4.5.5 IBM PowerVS API use

Important: The following samples are as-is and no support is provided.

Example 4-24 shows how to use Cloud API. The example focuses on use of VM shutdown and such operations. This sample can be modified to perform any snapshot or restore such operations.

Example 4-24 Sample how to use an API from PowerVS

```
#!/bin/bash
### START OF VARIABLES
API_KEY="ENTER YOUR API KEY HERE"
CLOUD_CRN="ENTER YOUR CLOUD CRN"
INSTANCE_NAME="ENTER YOUR INSTANCE NAME"
## Acceptable values are stop, start ,hard-reboot, soft-reboot
OPERATION="stop"
#####
IFS=":" read -ra ADDR <<< "${CLOUD_CRN}"
CLOUD_INSTANCE_ID=${ADDR[7]}
CLOUD_URL=(${ADDR[5]}.power-iaas.cloud.ibm.com)

## FIRST WE GET THE TOKEN FROM THE CLOUD IAM SERVICE USING THE API KEY

GET_TOKEN=$(curl -X POST -H "Content-Type: application/x-www-form-urlencoded" -H "Accept:
application/json" -d
"grant_type=urn%3Aibm%3Aparams%3Aoauth%3Agrant-type%3Aapikey&apikey=$API_KEY"
https://iam.bluemix.net/oidc/token | jq -r '.access_token')

## THIS IS THE POST CALL TO INVOKE the OPERATION
curl -X POST
https://$CLOUD_URL/pcloud/v1/cloud-instances/$CLOUD_INSTANCE_ID/pvm-instances/$INSTANCENAME/acti
on -H "CRN: $CLOUD_CRN " -H "Authorization: Bearer $GET_TOKEN" -H
'Content-Type:application/json' -d '{ "action":"'$OPERATION'" }'

sleep 30
```

```
## THIS IS A GET CALL
curl -X GET
https://$CLOUD_URL/pcloud/v1/cloud-instances/$CLOUD_INSTANCE_ID/pvm-instances/$INSTANCE_NAME -H
"CRN: $CLOUD_CRN " -H "Authorization: Bearer $GET_TOKEN" -H 'Content-Type:application/json'
```

Sample of volume snapshot, restore, and clone operations

In this section, we provide samples of snapshot, restore, and cloning operations.

Preconditions

The body of the Snapshot, Restore and Clone (PVM and Volume) API can be modified by using user-defined values.

Before running the Restore API, the PVM instance must be shut off.

Creating a snapshot

An example of creating a snapshot is shown in Example 4-25.

Example 4-25 Creating a snapshot

```
curl -X POST https://< Cloud IP >/pcloud/v1/cloud-instances/<Cloud Instance
ID>/pvm-instances/<PVM Instance ID>/snapshots \
-H "authorization: <AuthToken>" \
-H "content-type: application/json" \
-H "crn: <CRN>" \
-d "{\"name\": \"VM1-SS\", \"description\": \"Snapshot for VM1\", \"volumeIDs\":
[\"VM1-7397dc00-0000035b-boot-0\", \"vmdv1\"]}"
```

The expected response is shown in Example 4-26.

Example 4-26 Response after snapshot

```
{"snapshotID": "65ea39fd-cab6-46b3-b88c-3c28479ab019"}
```

Obtaining snapshot details

Figure 4-27 shows how to obtain the details of a snapshot.

Example 4-27 Obtaining snapshot details

```
curl -X GET https://< Cloud IP >/pcloud/v1/cloud-instances/<Cloud Instance
ID>/snapshots/<Snapshot ID>\
-H "authorization: <AuthToken>" \
-H "content-type: application/json" \
-H "crn: <CRN>"
```

The expected response is shown in Example 4-28.

Example 4-28 Response of snapshot details

```
{
  "action": "snapshot",
  "creationDate": "2020-04-13T08:51:21.000Z",
  "description": "Snapshotfor VM1",
  "lastUpdateDate": "2020-04-13T08:51:54.000Z",
  "name": "VM1-SS",
  "percentComplete": 100,
  "pvmInstanceID": "7397dc00-f328-4bfb-bef2-27200ca42cb9",
  "snapshotID": "65ea39fd-cab6-46b3-b88c-3c28479ab019",
  "status": "available",
  "volumeSnapshots": {
    "398344bb-a64d-4fd5-b3cd-14ddfea6dd0e":
"72f07383-ca5b-46a0-94a2-3d1e7a7faceb",
    "7a7a5b6e-1177-400a-82a4-0784957bbe75":
"33f91096-f204-4ed2-8110-c497a258c29c"
  }
}
```

Restoring to snapshot

Example 4-29 shows an example of restoring to snapshot.

Example 4-29 Restoring to snapshot

```
curl -X POST "https://< Cloud IP >/pcloud/v1/cloud-instances/<Cloud Instance
ID>/pvm-instances/<PVM Instance ID>/snapshots/<Snapshot
ID>/restore?restore_fail_action=" \
-H "authorization: <AuthToken>" \
-H "content-type: application/json" \
-H "crn: <CRN>" \
-d "{\"forceRestore\":\"false\"}"
```

The expected response is shown in Example 4-30.

Example 4-30 Expected response to restoring the snapshot

```
{
  "action": "restore",
  "creationDate": "2020-04-13T08:51:21.000Z",
  "description": "Snapshotfor VM1",
  "lastUpdateDate": "2020-04-13T08:55:28.000Z",
  "name": "VM1-SS",
  "pvmInstanceID": "7397dc00-f328-4bfb-bef2-27200ca42cb9",
  "snapshotID": "65ea39fd-cab6-46b3-b88c-3c28479ab019",
  "status": "available",
  "volumeSnapshots": {
    "398344bb-a64d-4fd5-b3cd-14ddfea6dd0e":
"72f07383-ca5b-46a0-94a2-3d1e7a7faceb",
    "7a7a5b6e-1177-400a-82a4-0784957bbe75":
"33f91096-f204-4ed2-8110-c497a258c29c"
  }
}
```

Creating a volume clone

Example 4-31 shows how to create a volume clone.

Example 4-31 Creating a volume clone

```
curl -X POST \  
  https://<Cloud IP>/pcloud/v1/cloud-instances/<Cloud Instance ID>/volumes/clone \  
  -H 'authorization: <Auth Token>' \  
  -H 'content-type: application/json' \  
  -H 'crn: <CRN>' \  
  -d '{  
    "displayName": "PerfClone",  
    "volumeIDs": ["VMT-1422dbc9-00000063-boot-0", "vmtdv1"]  
  }'
```

The expected response is shown in Example 4-32.

Example 4-32 Sample after creation of volume clone

```
{  
  "clonedVolumes": {  
    "6342e6a9-716d-4686-b644-7f089bceb332":  
    "fd99a7ae-3e15-4f7e-af79-f5637e9a27f8",  
    "8461389f-e8fb-403f-8f48-81edcc9ef46f":  
    "16ed7611-26cc-4b93-945d-760cd6a52c58"  
  }  
}
```

4.5.6 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 4-41:

- ▶ Capture the LUN's state before an upgrade; for example, upgrade of 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.

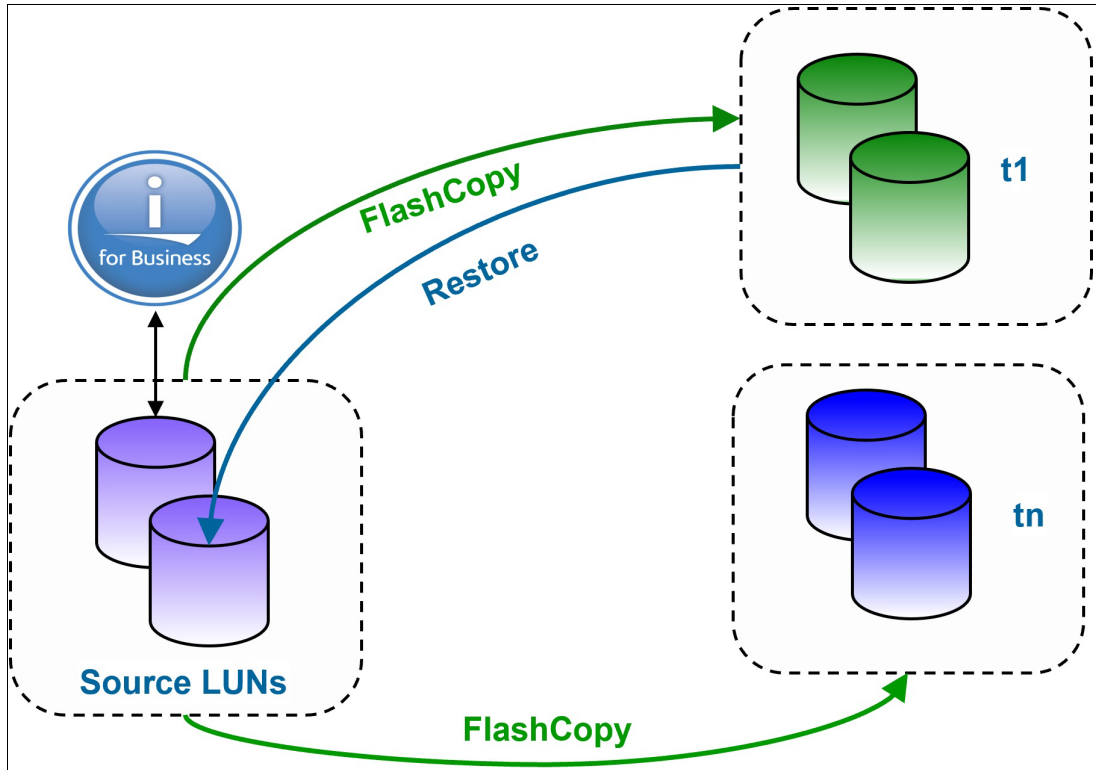


Figure 4-41 Use case - Snapshot and restore

Clone use case

A clone use case scenario is shown in Figure 4-42:

- ▶ A production environment or running environment are created and attached to a test environment to check production data and analyze data copy.
- ▶ Save the clone to IBM Cloud Object Storage. Use IBM Cloud Object Storage to regularly save as a DR strategy.

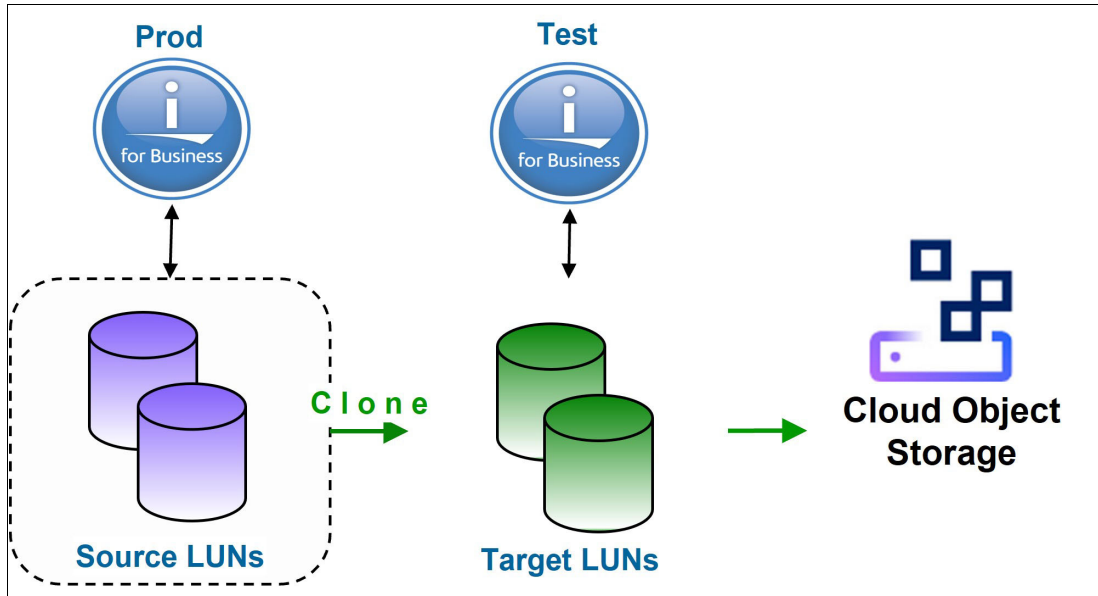


Figure 4-42 Use case - cloning

Important: For more information about taking snapshots, see Chapter 7, “Reference architectural decisions to migrate IBM i on-premises to IBM Power Systems Virtual Server” on page 243.



Hints and tips for IBM i deployments on IBM Power Systems Virtual Server

This chapter describes navigating in IBM Power Systems Virtual Server environment on IBM Cloud for IBM i.

The content in 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.

This chapter includes the following topics:

- ▶ 5.1, “Overview” on page 146.
- ▶ 5.2, “Using snapshots on IBM i instances” on page 157.

5.1 Overview

The hints and tips that are presented here are from real-world problem experiences. In this chapter, you become familiar with the process of the VNC console, SSH tunneling, the console by way of the LAN adapter, and take snapshots on IBM i instances.

5.1.1 Connecting to an IBM i virtual machine

This section describes how to connect to an IBM i virtual machine (VM) after it is deployed your system. See [this IBM Cloud Docs web page](#) for more information about configuring your IBM i VM before connecting to an IBM i VM. Then, complete the following steps:

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 then, click **Open console**. Use the VNC to access the 5250 console in the IBM Power Systems Virtual Server for IBM i.

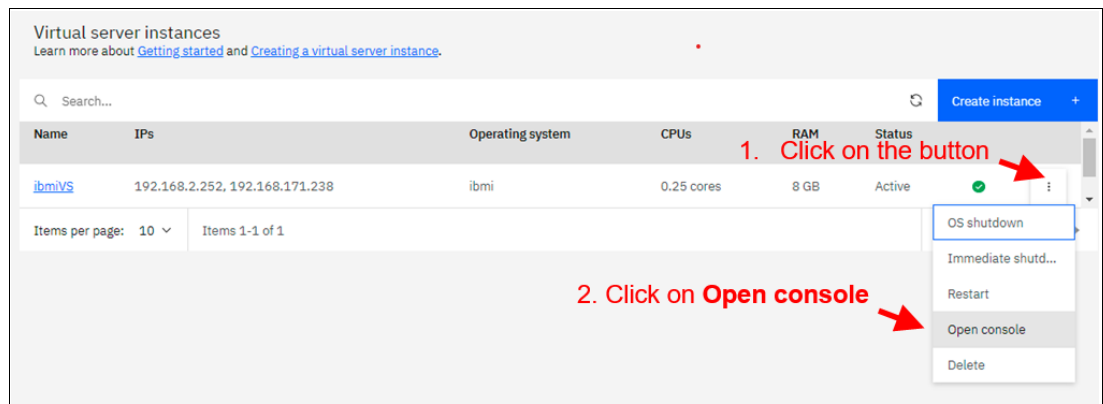


Figure 5-1 Opening console on PowerVS for IBM i

- 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 shown in Figure 5-2.

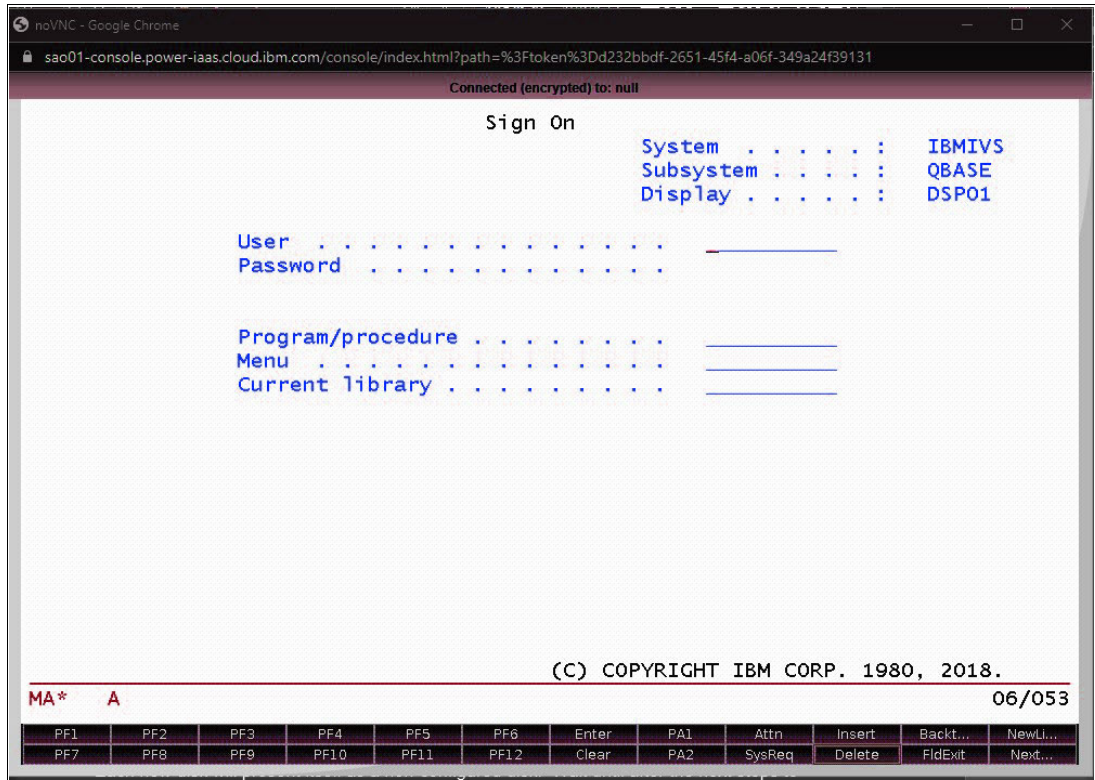


Figure 5-2 Opened console for IBM i on PowerVS

- Change the password for user profile QSECOFR, as shown in Figure 5-3.

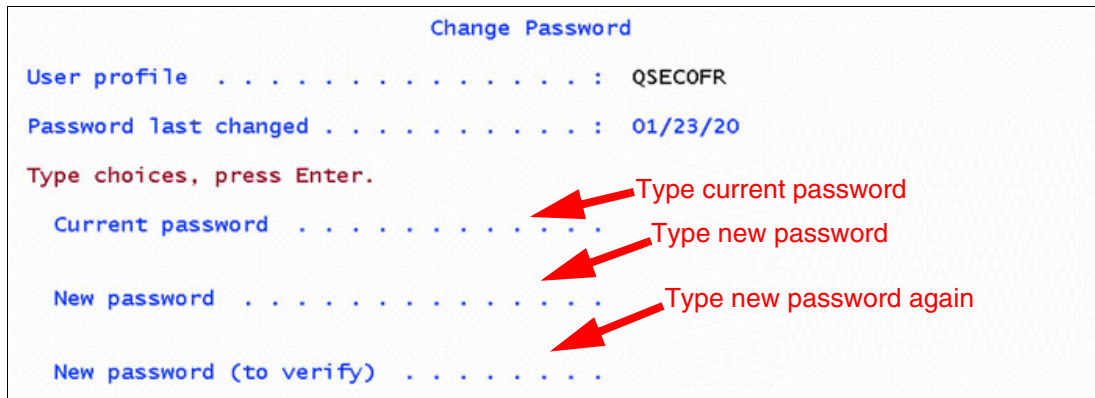


Figure 5-3 Change QSECOFR's password

- Accept the software agreements by selecting Option 5 for each. Then, press **Enter** (see Figure 5-4).

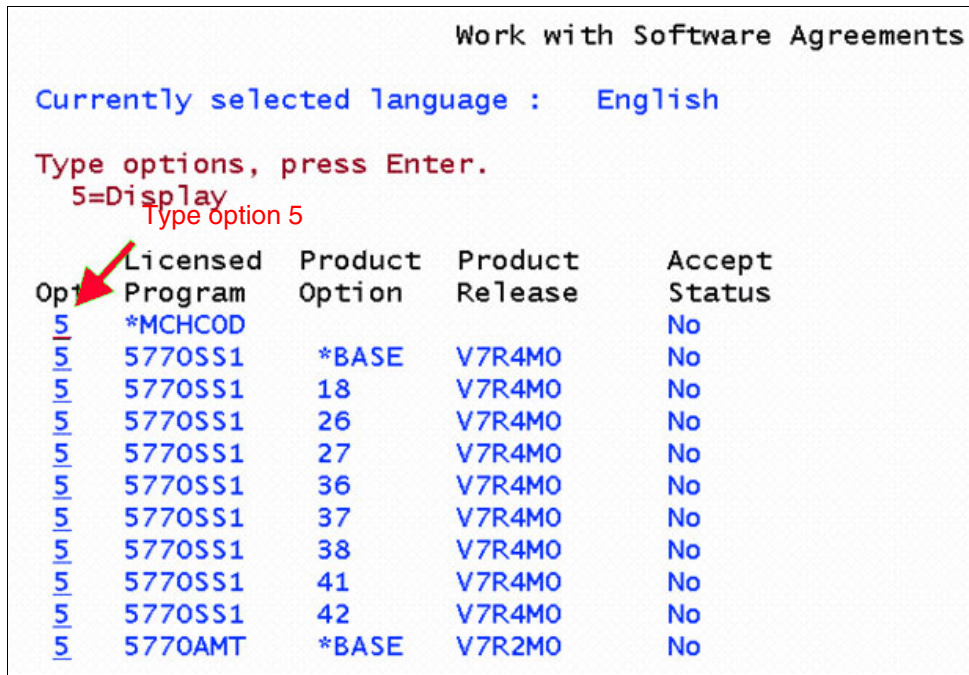


Figure 5-4 Accept software agreements

- To use the keys that are available from VNC, click **Next**. Then, select **PF15= Accept All**, as shown in Figure 5-5.

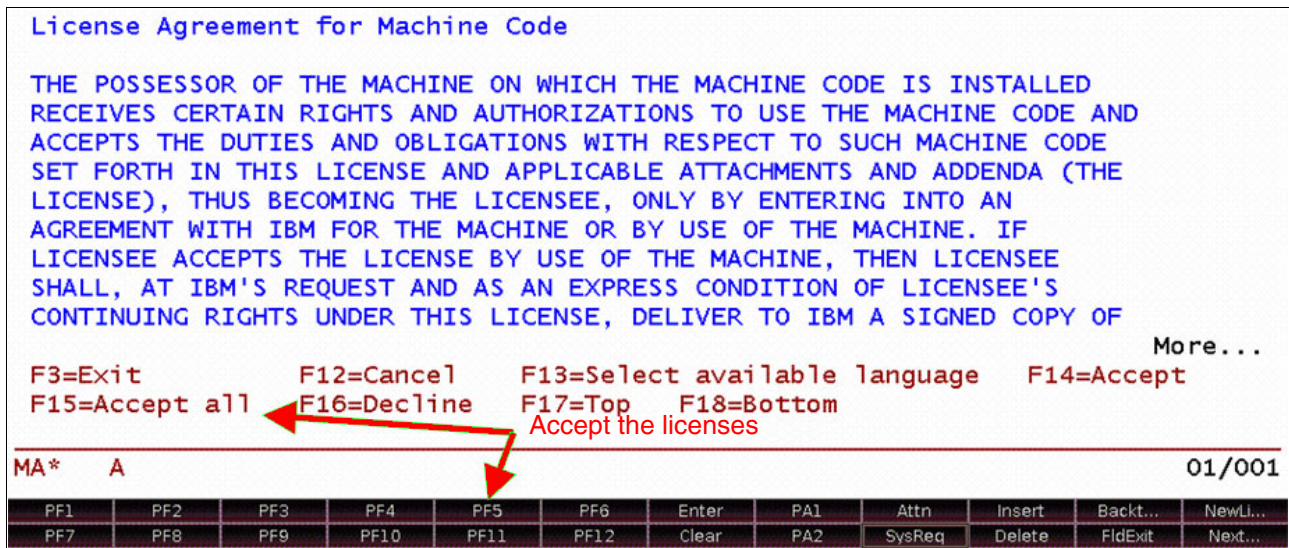


Figure 5-5 Accept the licenses agreements for Machine code

- Repeat Step 5 until the display shows all of the software agreements are in the Accept status (see Figure 5-6). Then, press **PF3**.

```

Work with Software Agreements
System:  IBMIVS
Currently selected language :  English
Type options, press Enter.
5=Display

```

Opt	Licensed Program	Product Option	Product Release	Accept Status
-	5733XT2	*BASE	V1R2M0	Yes
-	5770XW1	*BASE	V7R4M0	Yes

Note: A red arrow points to the 'Yes' status in both rows, with the text "Accepted state in both" next to it.

Figure 5-6 Work with Software agreements window

The Main Menu is displayed by the VNC.

- 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'
```

- Find the **#PermitRootLogin yes** on the file and remove the # symbol (see Figure 5-7). Then, press **PF3** to save and exit.

```

Edit File: /QOpenSys/QIBM/UserData/SC1/OpenSSH/etc/sshd_config
Record : 41 of 116 by 10 Column
Control : █

```

Note: A red arrow points to the '#' symbol in the line '#PermitRootLogin yes', with the text "Remove the symbol '#'" next to it.

```

CMD .....1.....2.....3.....4.....5.....6
#PermitRootLogin yes
#StrictModes yes
#MaxAuthTries 6

```

Figure 5-7 Edit file on IFS

Attention: This QSECOFR profile is for demonstrative purposes. The QSECOFR profile is used for the SSH connection. However, the preferred practices that is suggested *not* to use QSECOFR profile with SSH on IBM i. Use a different user profile with similar authority. The new user profile has a Home directory as well; this directory often is /home/<userprofile>.

- Run the `WRKSYSVAL QAUTOVRT` command. Then, press **Enter**, 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, but how many you use depends on how many you need. Press **Enter** (see Figure 5-8).

```

Change System Value

System value . . . . . : QAUTOVRT
Description . . . . . : Autoconfigure virtual devices

Type choice, press Enter.

Number of devices to
autoconfigure . . . . . 4 0-32500, *NOMAX
  
```

Figure 5-8 Change QAUTOVRT system value

Note: You can add as many devices that you require on QAUTOVRT.

- Issue the `WRKSYSVAL QLMTSECOFR` command. Then, press **Enter** and select **option 2 = Change**. Then, change the value from 1 to 0. Press **Enter**.
- Verify that `cloud-init` is configured as shown in Figure 5-9. Run the `CFGTCP` command, select **Option 1** and then, press **Enter**.

```

192.168.2.17      255.255.255.0   CLOUDINIT1     *ELAN
192.168.171.236  255.255.255.248 CLOUDINIT2     *ELAN
  
```

Figure 5-9 TCP IPs on IBM i configured by cloud-init

- Issue the following commands to start the Telnet and SSH services:

```

STRTCPSVR *TELNET
STRTCPSVR *SSH
  
```

- Verify that the ports for SSH and Telnet are listening, as shown in Figure 5-10. Issue the `NETSTAT *CNN` command. Then, press **Enter** and then, press **Shift + PF2**.

Remote Address	Remote Port	Local Port	Idle Time	State
*	*	ssh	000:02:58	Listen
*	*	telnet	000:03:04	Listen

Figure 5-10 Display NETSTAT on IBM i

5.1.2 Remote access to IBM i by 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 (*SEC0FR) 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 5-11.

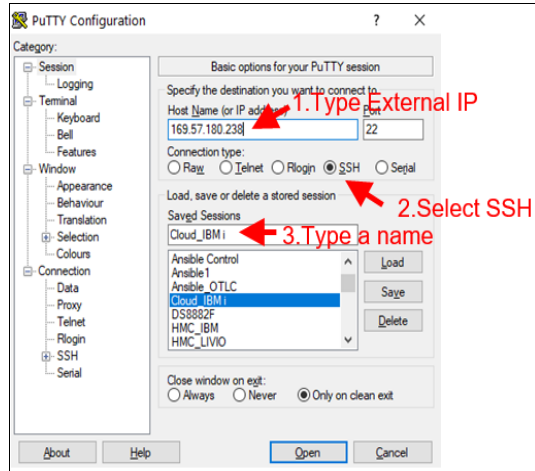


Figure 5-11 Set PuTTY terminal

The ports that are required to configure on PuTTY are listed in Table 5-1.

Table 5-1 Ports that are required to set up 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

Source port	Destination
9475	localhost:9475
9476	localhost:9476

Complete the following steps:

- a. Set up the source port and destination as shown in Figure 5-12.

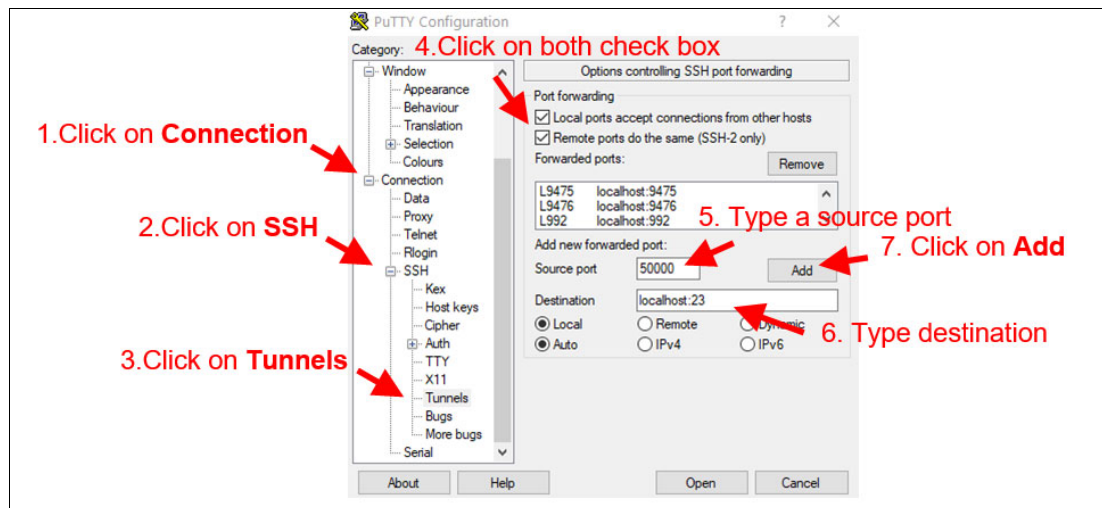


Figure 5-12 Set the source ports on PuTTY

Note: Steps 5, 6, and 7 that are shown in Figure 5-12 are repeated for each port that is in Table 5-1 on page 151.

- b. Return to Session on Category that is on the left side of PuTTY and click **Save**.
- c. Click **Open** at the session in PuTTY and click **Yes** to trust this host and connection.
- d. A new window terminal is opened. Enter QSECOFR and the password that was changed, and then, press **Enter**. The tunneling is now ready. Hold the terminal open, as shown in Figure 5-13.

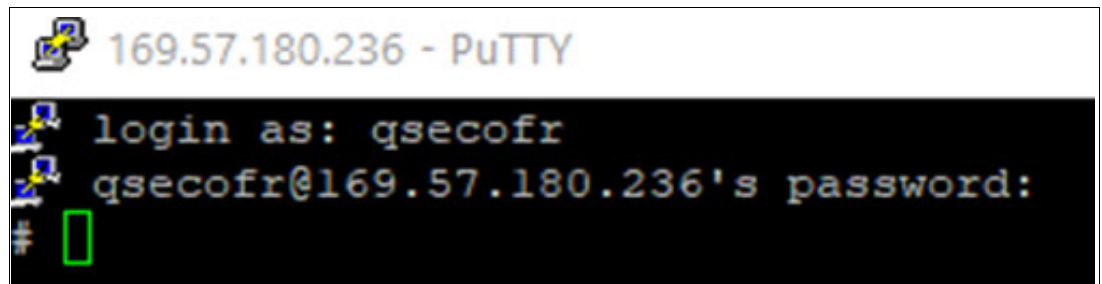


Figure 5-13 PuTTY log in for QSECOFR

Important: The tunneling that is done by using PuTTY is for a Windows system. If you use another operating system, such as Linux or Mac, the SSH tunneling to allow ACS to connect over the External IP is different. For more information, see this [IBM Cloud Docs web page](#).

2. Select **Access client solutions** → **Management** → **5250 Session Manager** → **New Display Session**. A new window opens, and the 5250 display setup is shown (see Figure 5-14).

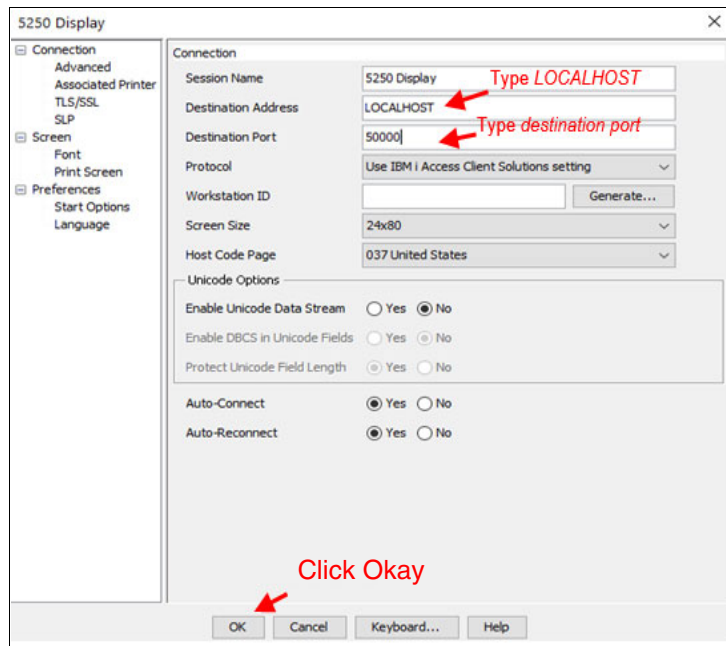


Figure 5-14 5250 display configurations

Important: For this example, 50000 was chosen as the source port number. This value also was configured in PuTTY. Do not change the source port numbers. When telnetting, avoid making the source port the same as the destination.

3. A window opens in which the connection with the PuTTY tunnel is established. Enter the credentials of the QSECOFR profile and their password. Finally, a window capture shows the Main Menu. Select **File** → **Save as** to save the 5250 display.

Note: A commonly secure Telnet session is used on-premises. However, because the use of SSH tunnel encryption is working in this example, 992 port is *not* necessary.

5.1.3 IBM i 5250 console through LAN adapter

Working with IBM i instances on IBM Power Systems Virtual Server, you need a 5250 console to perform migration, backups, and maintenance procedures.

IBM Power Systems Virtual Server instances provide a web 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 147 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 data 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. Start the Service Tools by using QSECOFR or a user with *SERVICE authority: STRSST.
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 5-15. (Use NEXT to see more Fn options.)

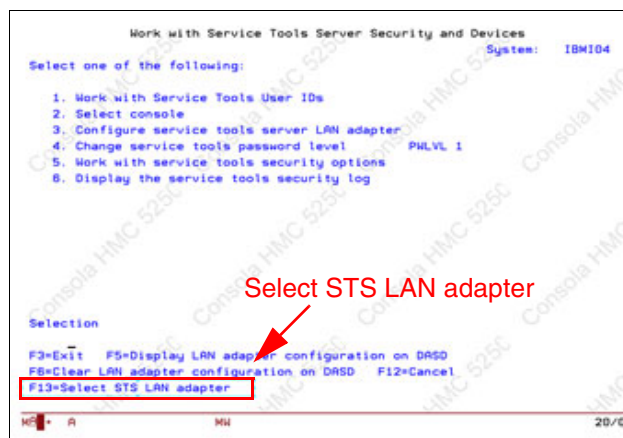


Figure 5-15 Select STS LAN adapter at SST on IBM i

- 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 5-16).



Figure 5-16 Select Service Tools Server LAN adapter

- Assign the LAN Adapter its own IP address to the LAN Adapter, as shown in Figure 5-17. Press **F7** to save, **F13** to deactivate and then, **F14** to reactivate the adapter.

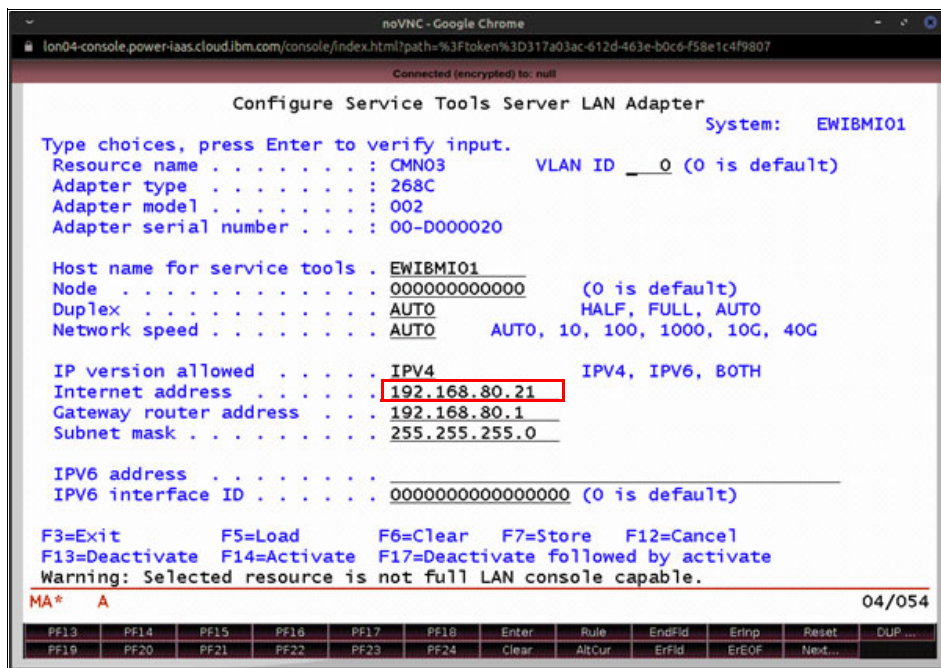


Figure 5-17 Assigning IP on LAN adapter

Note: When an adapter is in use, IPL your IBM i system.

- Configure the console at IBM Access Client Solutions, as shown in Figure 5-18. Click **System Configuration** → **New**. Then, enter the System Name, which often is the IP address. If a hostname.

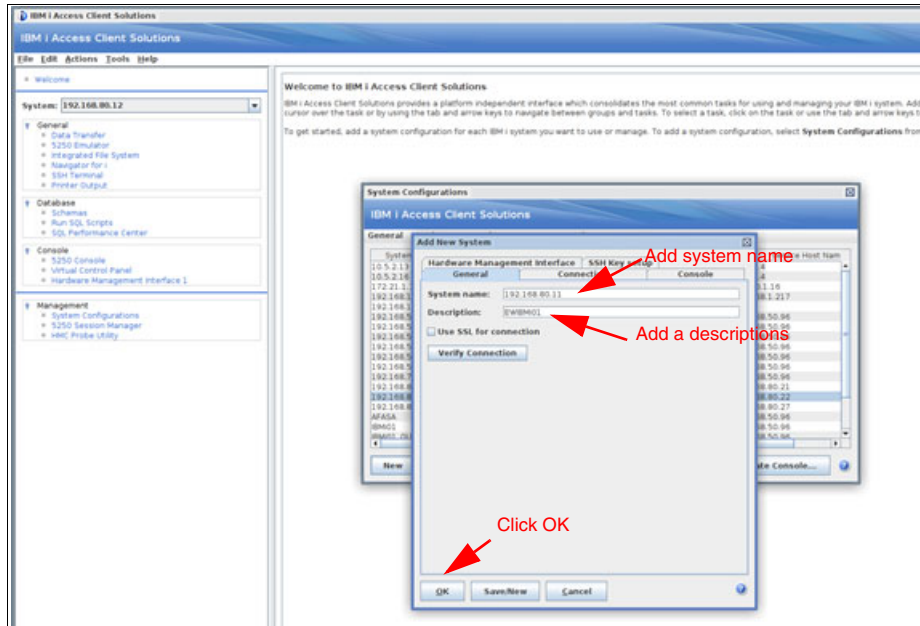


Figure 5-18 Adding a system on Access Client Solutions

- Go to the Console tab and select the Service hostname, as shown in Figure 5-19. This name is the IP address for the Service Tools Server LAN Adapter.

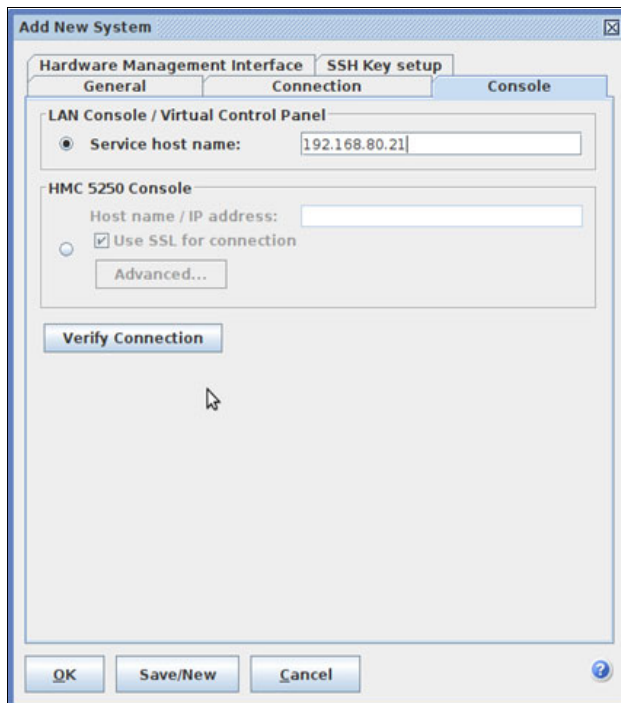


Figure 5-19 Insert the IP address

- 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

Snapshots are a resource that is 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.

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
- ▶ Making changes to system values
- ▶ Updating application programs

Remember: Snapshot is *not* a backup mechanism. It cannot use one 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.

5.2.1 Taking snapshots

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.

Installing and using IBM Cloud CLI

For more information about installing the IBM Cloud CLI on Windows, see this [IBM Cloud Docs web page](#).

Note: For Windows, some functions are not supported unless you are running Windows 10 Pro.

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-cloud-developer-tools/master/windows-installer/idt-win-installer.ps1')
```

This process can take some minutes. Wait until the end of the process and then, restart the system.

2. In a command window, run the command that is shown in Example 5-2 to check whether the setup processes succeed.

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 services are then listed (see Example 5-5).

Example 5-5 List available services

```
ibmcloud pi service-list
```

Figure 5-20 shows the list of available services.

```
crn:v1:bluemix:public:power-iaas:lon04:a/af339adbfd124f99a5cea8271bf030cc:5fa3d138-6bc4-409f-ba6d-641c071b8e51:: Power Systems Virtual Server - LON04
crn:v1:bluemix:public:power-iaas:osa21:a/af339adbfd124f99a5cea8271bf030cc:35567417-c10f-4a4a-afd3-0f57a81d6eed:: Power Systems Virtual Server-OSA21
crn:v1:bluemix:public:power-iaas:lon06:a/af339adbfd124f99a5cea8271bf030cc:0a22f7d5-6e38-4fcb-b036-9c715d3cb63a:: Power Systems Virtual Server-73 - LON06
```

Figure 5-20 List available services

6. Target the service by running the **ibmcloud pi service-target <crn>** command, as shown in Example 5-6.

Example 5-6 Target the service to use

```
ibmcloud pi service-target
crn:v1:bluemix:public:power-iaas:lon04:a/af339adbfd124f99a5cea8271bf030cc:5fa3d138-6bc4-409f-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 ins
```

8. By using this list, you can copy the instance ID that is must be 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 completes.

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, continue `ibmcloud cli` and take the snapshot. Run the `ibmcloud pi snapshot-create` command and target your VM instance. Then, choose a name to identify the snapshot (see Example 5-9).

Example 5-9 Taking the snapshot by using `ibmcloud cli`

```
ibmcloud pi snapshot-create c451ccd1-54e8-4953-9402-XXXXXXXXXXXX --name
SNP01_EW01-EWIBMi01
```

11. Run the `ibmcloud pi snaps` command to list the snapshot. You see that 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

```
CHGASPECT 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:

```
ibmcloud pi snapshot-restore <instance_id> --snapshot <snapshot_id>
```

Example 5-11 Restoring the snapshot

```
ibmcloud pi snapshot-restore daa5033a-269a-4bb4-9537-XXXXXXXXXXXX --snapshot
57122996-ac9b-4aec-a356-YYYYYYYYYYYY
```

15. Run the `ibmcloud pi snaps` command to see the restore operation status.

16. Start the VM instance after the snapshot is restored.



Disaster Recovery with IBM Power Systems 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:

- ▶ 6.1, “Introduction” on page 162.
- ▶ 6.2, “Solution components and requirements” on page 162.
- ▶ 6.3, “Operating system level replication use case with PowerHA SystemMirror for i geographic mirroring” on page 192.
- ▶ 6.4, “Logical replication use case with Bus4i” on page 213.

6.1 Introduction

Uptime is a key client expectation for critical IBM i workloads.

Whether for local solutions at data center level high availability (HA) or for multi-site solutions between different geographical locations (DR), these objectives are achieved through HA/DR solutions.

IBM Power Systems Virtual Server meet this requirement by enabling clients to take advantage of DR solutions between two IBM i Virtual Server Instances (VSIs) in separate IBM Cloud data centers.

DR solutions for IBM Power Systems Virtual Server are based on operating system level replication or logical replication.

As of this writing, DR solutions that are based in Storage replication methods are not an option available with IBM Power Systems Virtual Server.

Replication solutions between two data centers always involve prerequisite network configuration between them to allow the necessary data flow to occur securely. This requirement also applies to DR with IBM Power Systems 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:

- ▶ Performing the required network configuration
- ▶ Implementing the DR solution

Regarding the DR solutions, two use cases are included:

- ▶ Operating system level replication use case with IBM PowerHA SystemMirror for i geographic mirroring.
- ▶ Logical Replication use case with TSP Bus4i.

6.2 Solution components and requirements

This section describes the solution components and requirements for deployment.

6.2.1 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 i VSIs that is 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.

6.2.2 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. The Power Systems Virtual Server location service is a service within IBM Cloud that allows you to provision IBM i VSIs.

Only one Power Systems Virtual Server location service can be set up per data center in IBM Cloud. In this scenario, we created two Power Systems Virtual Server data center locations: one in Toronto and one in London.

Before provisioning Power Systems Virtual Server Instance in the Power Systems Virtual Server location, you must 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.

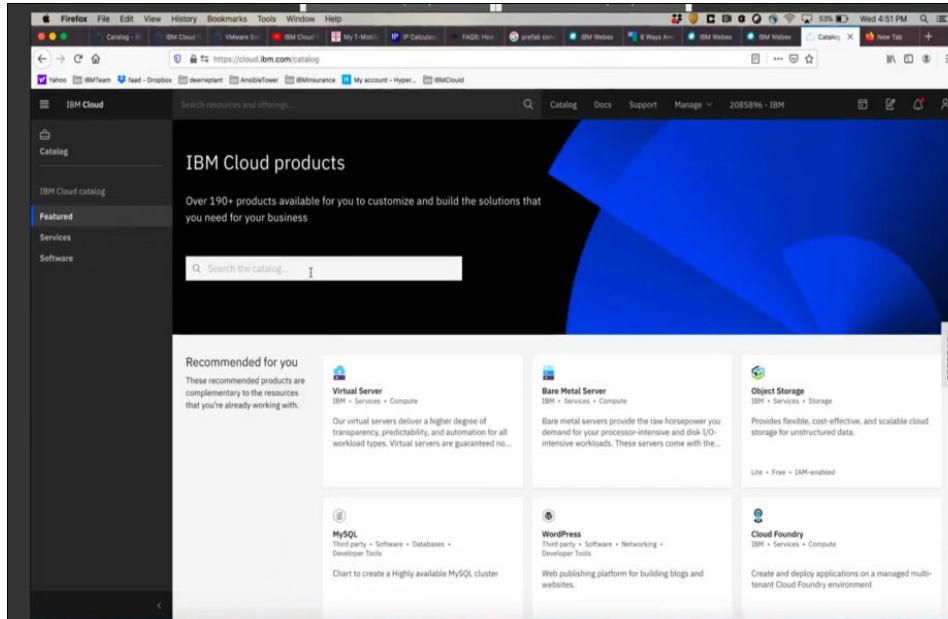


Figure 6-1 IBM Cloud products

4. Search the results for Power and select **Power Systems Virtual Servers**, as shown in Figure 6-2.

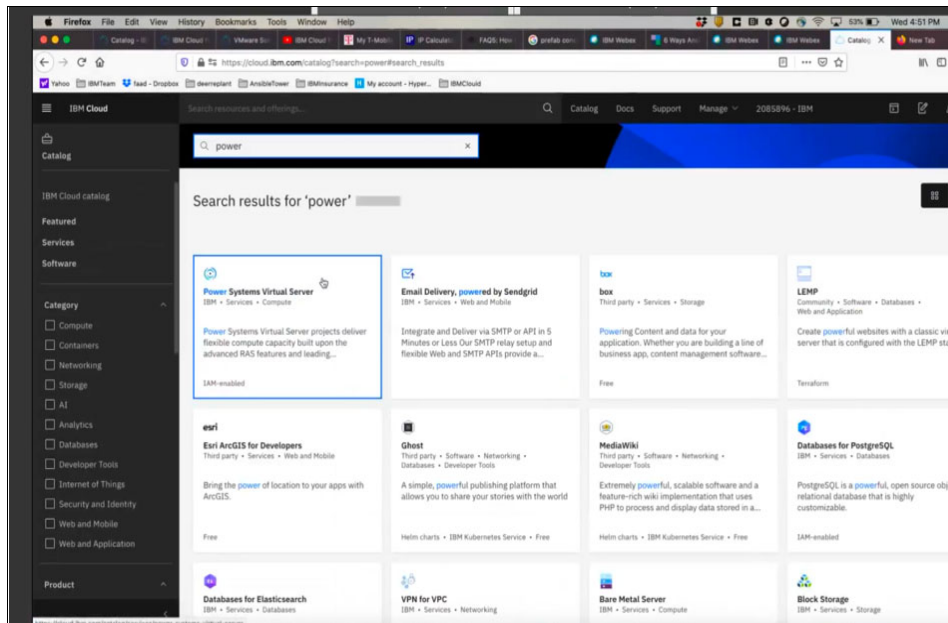


Figure 6-2 Results for Power

- Under Select Region, choose your region, as shown in Figure 6-3. You are limited to only one service per region.

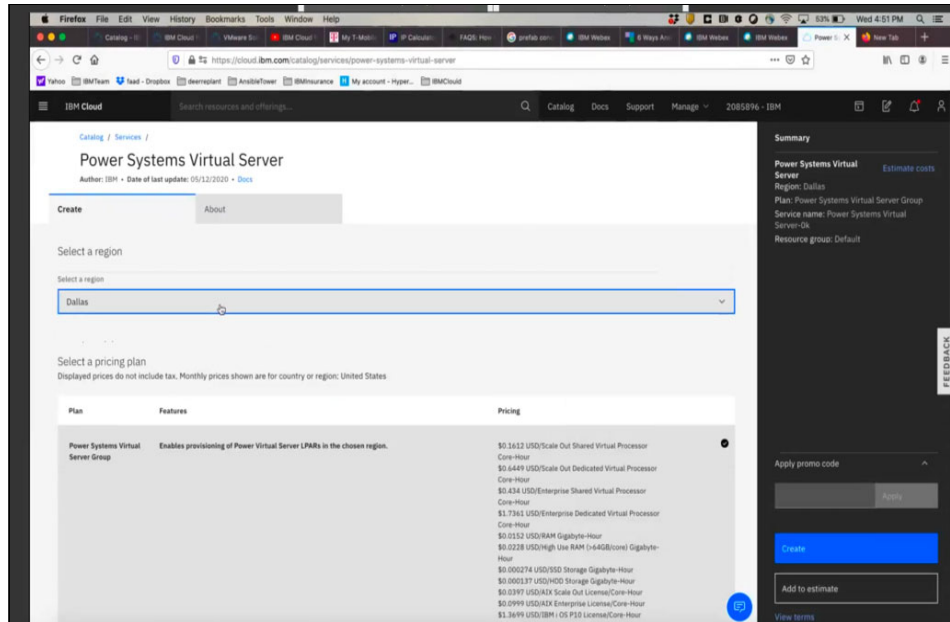


Figure 6-3 Selecting region

- Select a Service Name or chose the default name. Then, click **Create**, as shown in Figure 6-4.

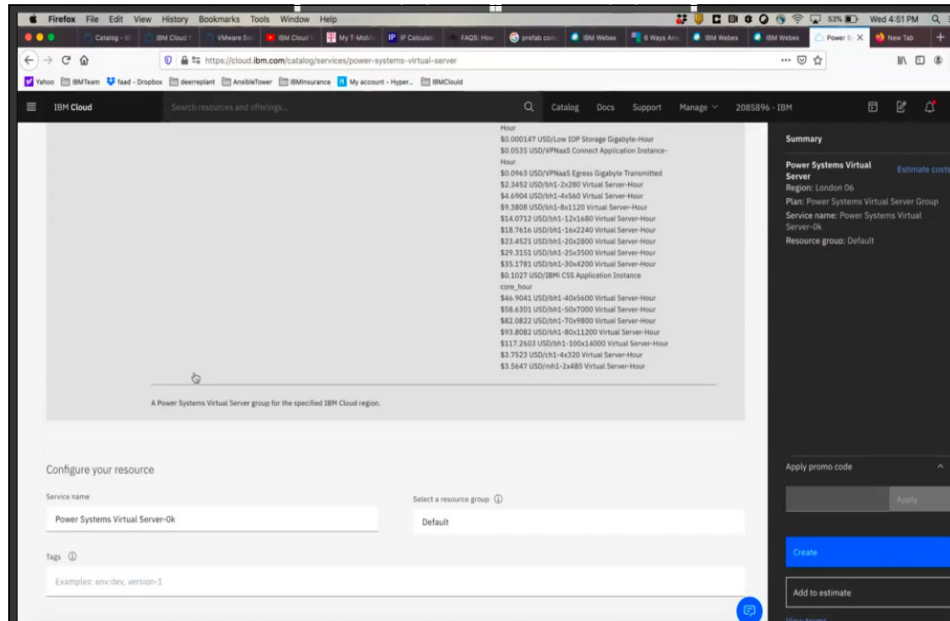


Figure 6-4 Selecting a service name

Your Power Systems Virtual Server location service now appears under the Services tab, as shown in Figure 6-5.

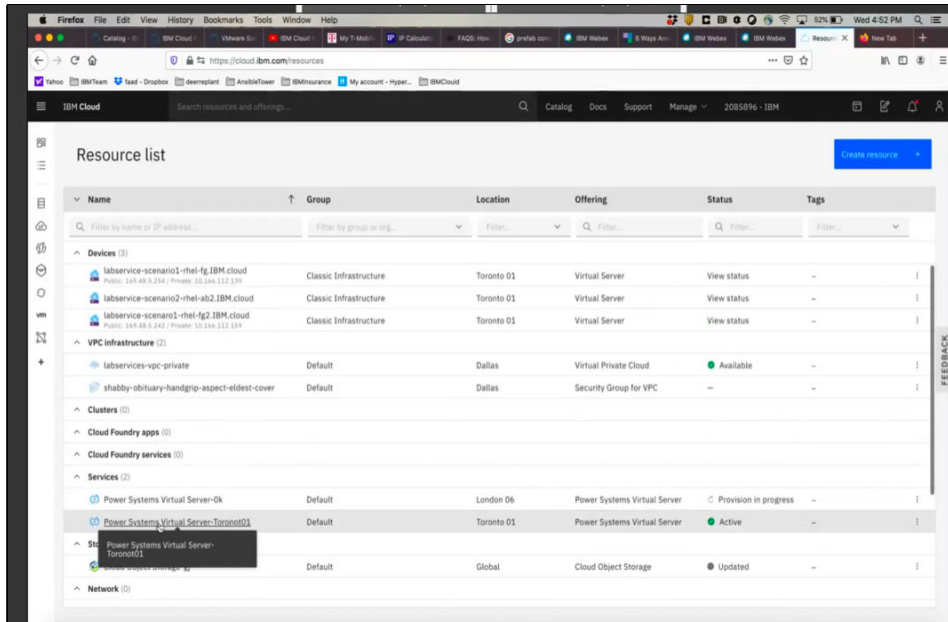


Figure 6-5 Power Systems Virtual Server instance

Repeat this process to create a second Power Systems Virtual Server location service.

7. 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).

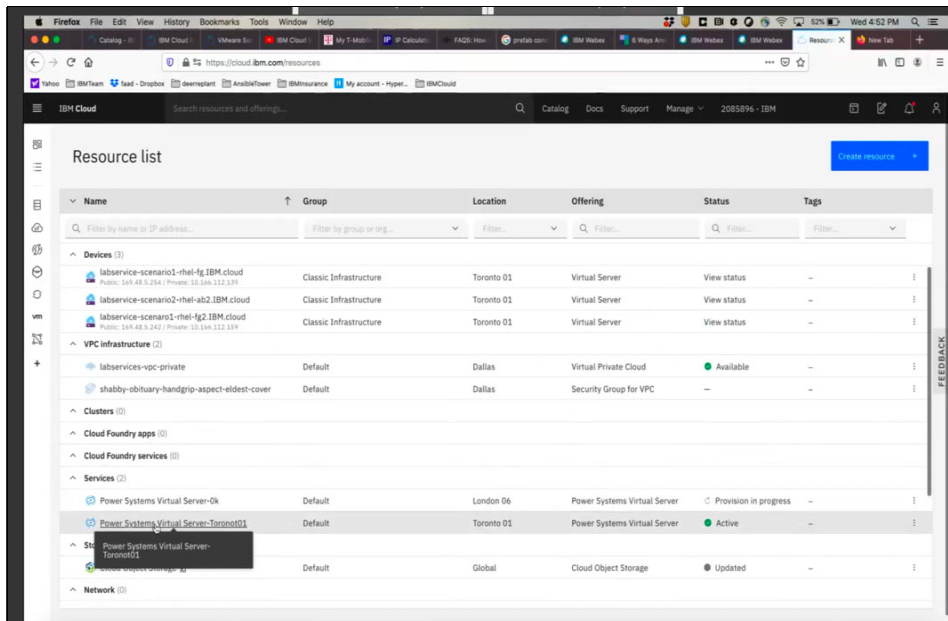


Figure 6-6 Resource list

8. Choose **Subnets** from the menu on the left, as shown in Figure 6-7.

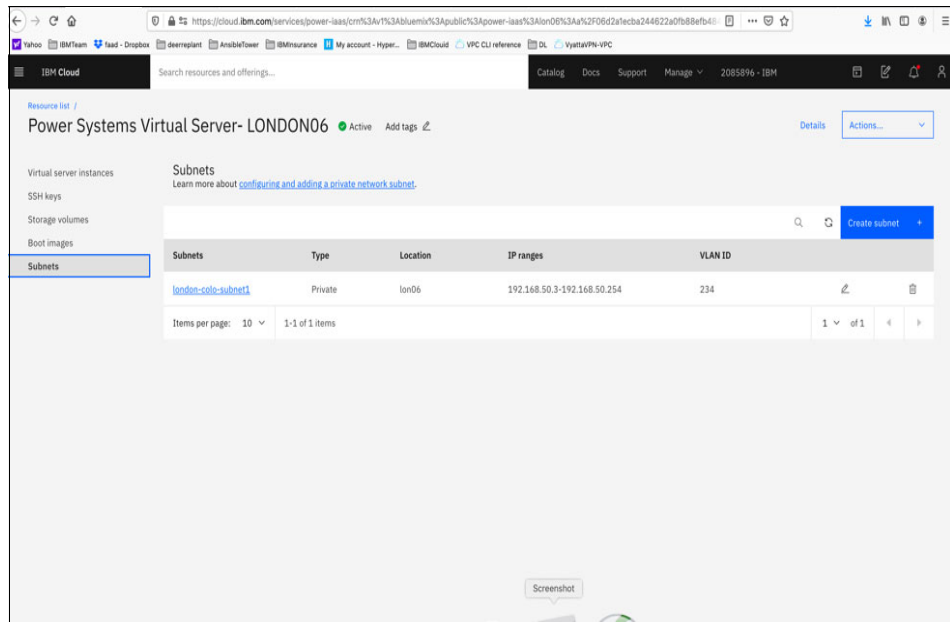


Figure 6-7 Subnets

Provide the following information:

- A name for your subnet, as shown in Figure 6-8.
- CIDR range. This range can be any private IP subnet ranges. You choose /21 - /30 based on how many IPs you require. You also can use your own private CIDR if you want.

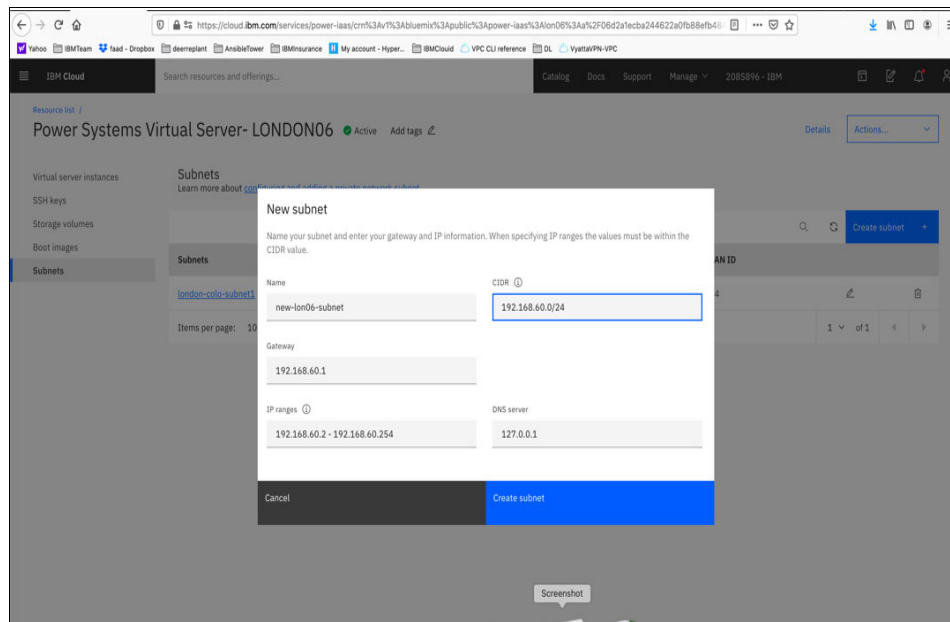


Figure 6-8 Subnet configuration

The rest of the fields are automatically populated based on the CIDR you provided.

9. Click **Create Subnet**. A VLAN ID is then associated with the subnet, as shown in Figure 6-9.

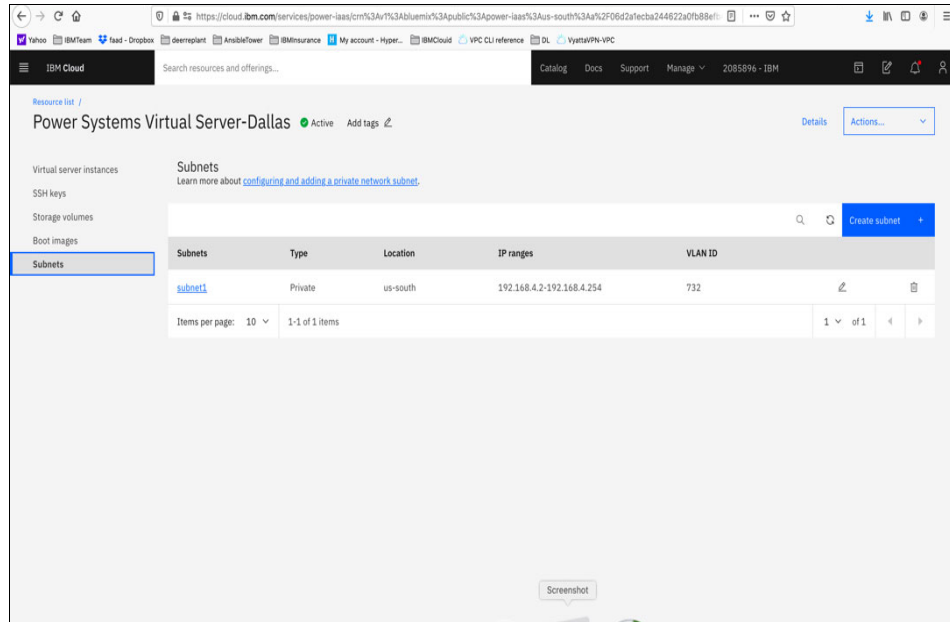


Figure 6-9 Creating a subnet

10. 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 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 upper left of the UI, as shown in Figure 6-10.

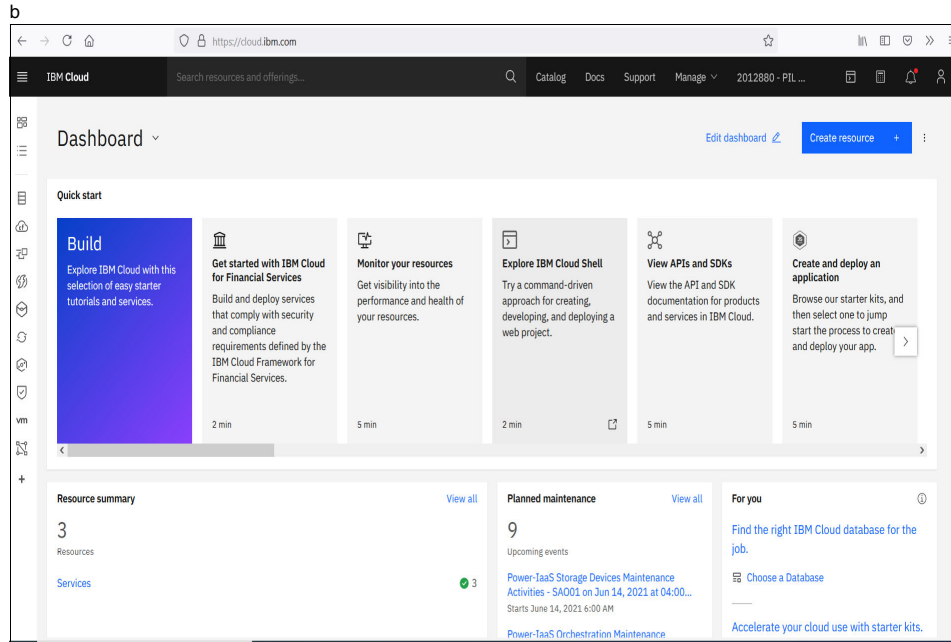


Figure 6-10 Selecting the IBM Cloud option

2. Choose **Services** from the list, as shown in Figure 6-11.

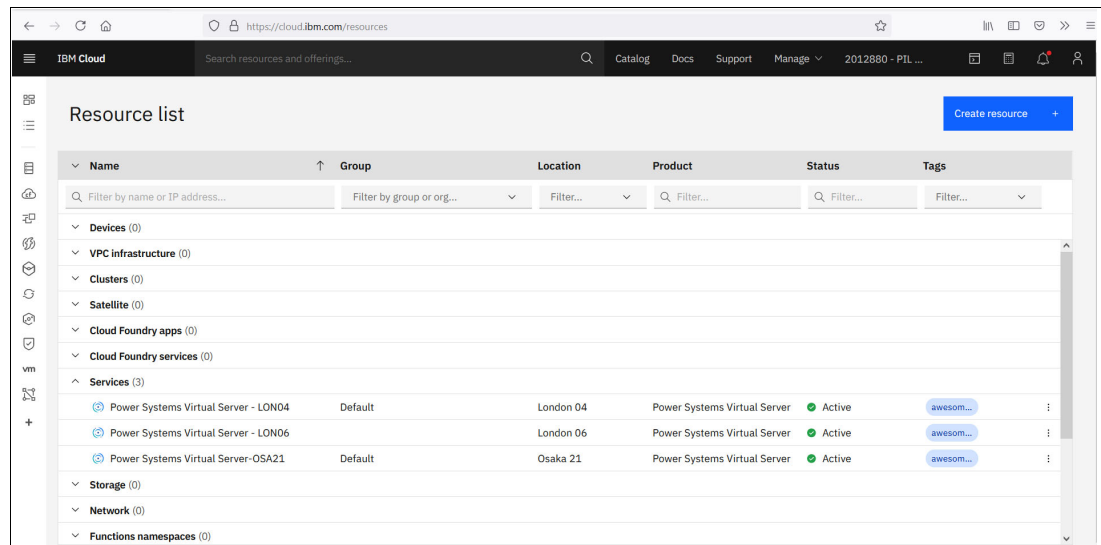


Figure 6-11 Choosing Services option

3. Click the service for data center in which you created a Power Systems Virtual Server location service. In this example, we choose **London06** Power Systems Virtual Server location, as shown in Figure 6-11.

4. Because we provisioned several VSIs, we see the list as shown in Figure 6-12. If you are creating VSIs for the first time, the list is empty. Click **Create Instance**.

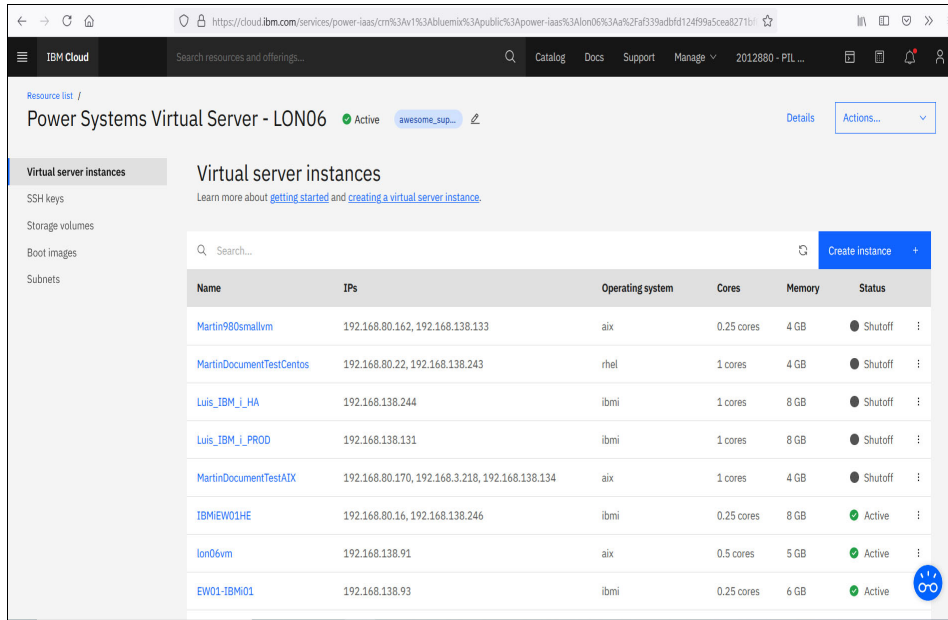


Figure 6-12 London06 Virtual Server instances

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 VSI 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.

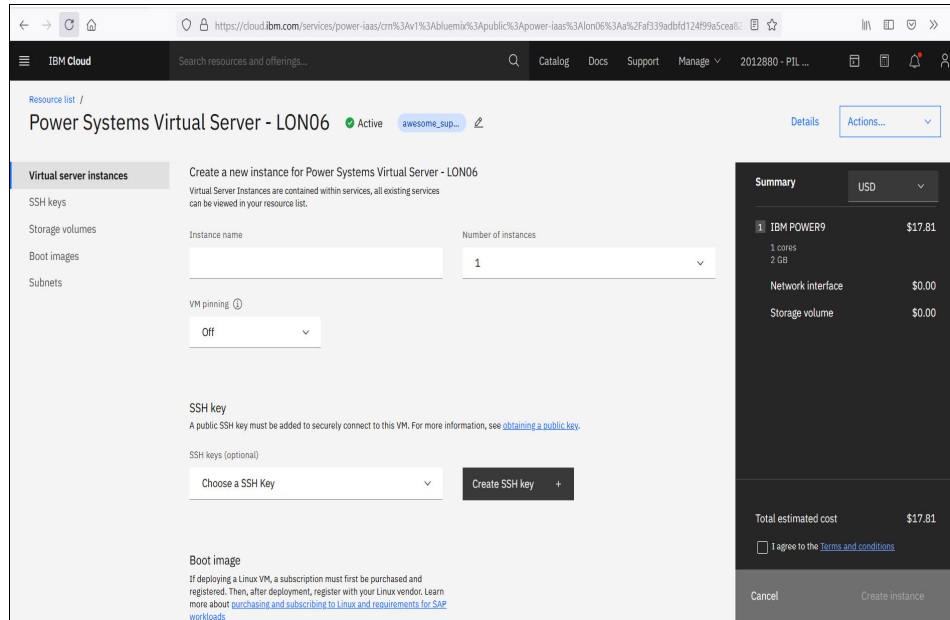


Figure 6-13 Adding SSH key

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 we selected.
7. For the Machine type, choose between S922 or E980. S922 is the cheaper of the two options, which we selected.

8. The processor resources can be assigned as Dedicated, Shared, or Shared Capped, as shown in Figure 6-14. We chose Shared because it is less expensive.

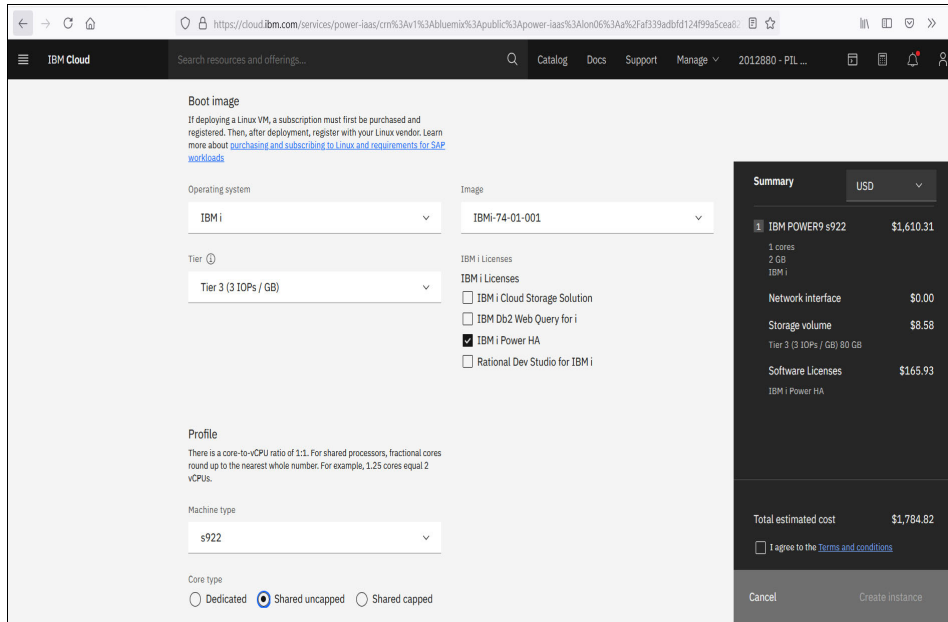


Figure 6-14 Choosing Shared process resource

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).

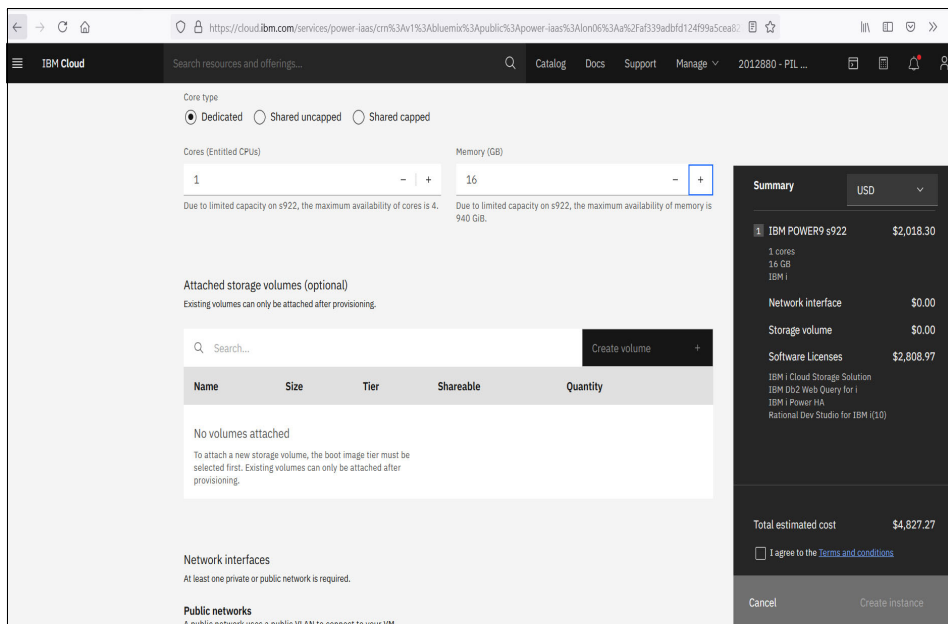


Figure 6-15 Using the minimum memory value

10. You can also attach more volumes to the VSI. In our example, we did not need to do so and used only the root volume that is included. It is highly recommended to add the disks individually after the initial Load Source (LS) volume is created. Doing so makes it easier to keep track of 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 one or more subnets were created *before* this step.
- 14.Click **Attached Existing** (see Figure 6-16).

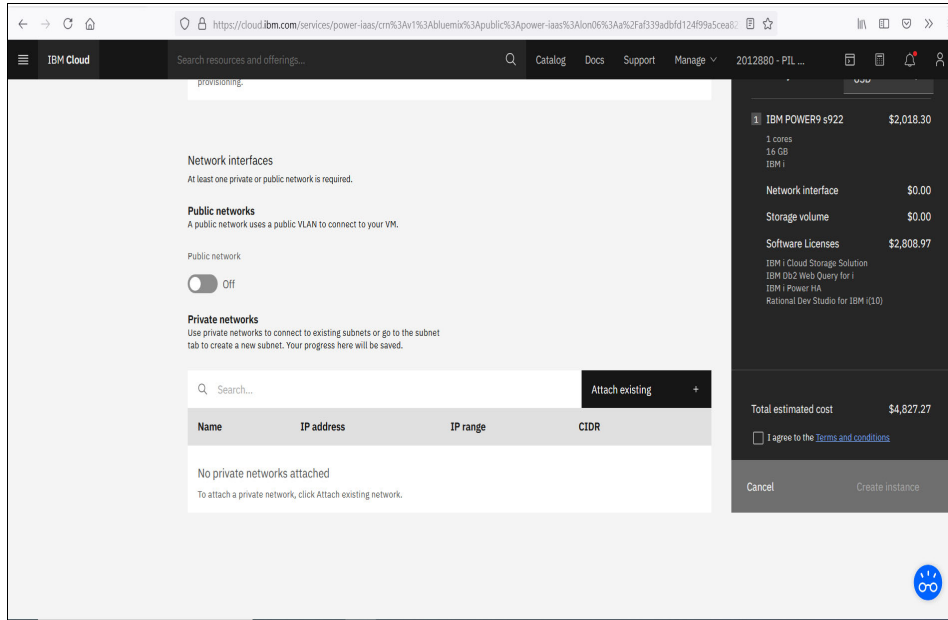


Figure 6-16 Selecting Attached Existing option

- 15.Choose the subnet that you want to attach and then, click **Attach**, as shown in Figure 6-17.

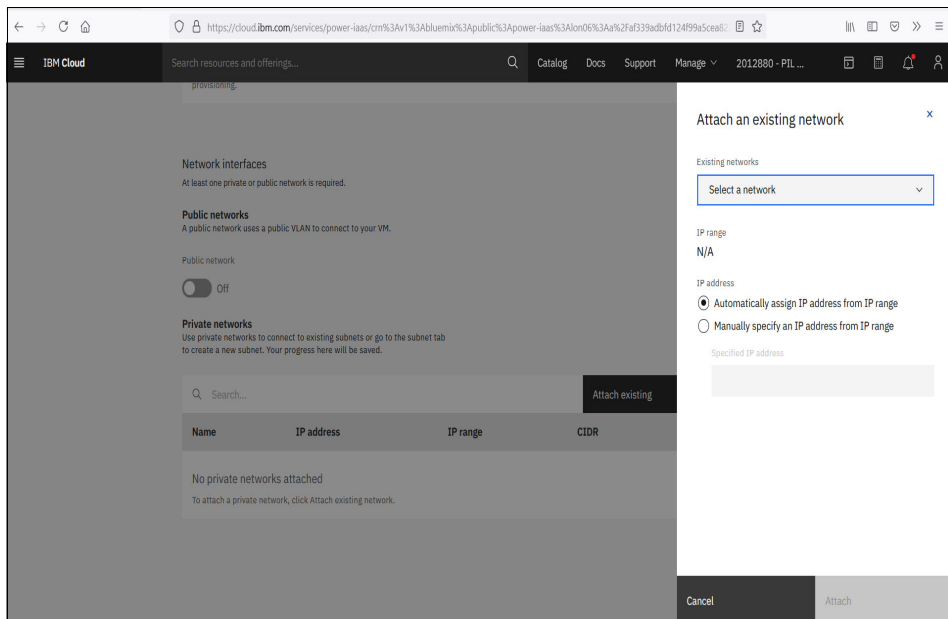


Figure 6-17 Attaching an existing network

16. Select the **I agree to the....** option and then, click **Create Instance**, as shown in Figure 6-18.

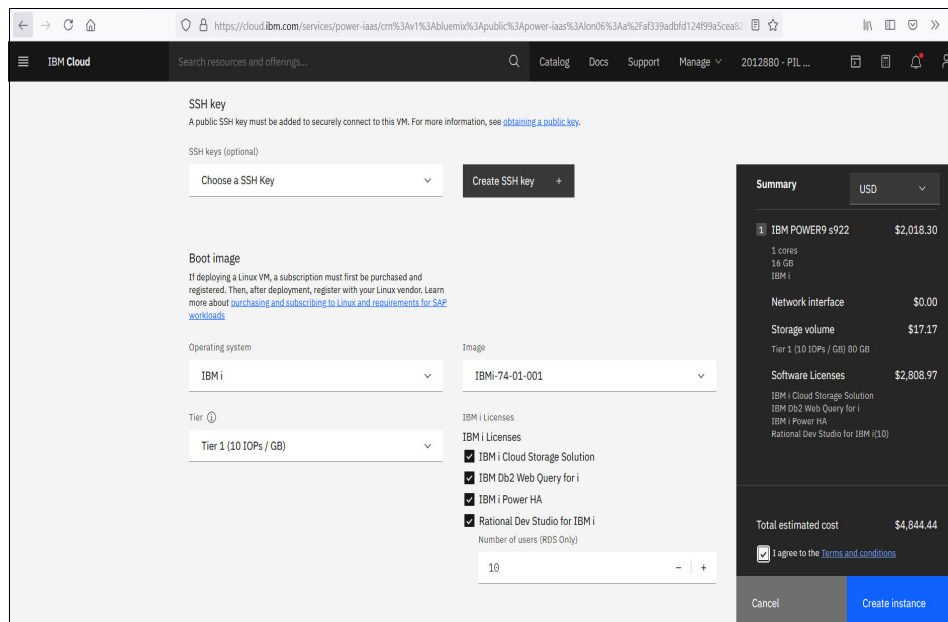


Figure 6-18 Creating an instance

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 communication with Linux or Window VSIs in IBM Cloud and with all other IBM Cloud services, such as VMware VMs, and Cloud Object Storage. Ordering a DL can take 1 - 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 then, search for **direct**, as shown in Figure 6-19.

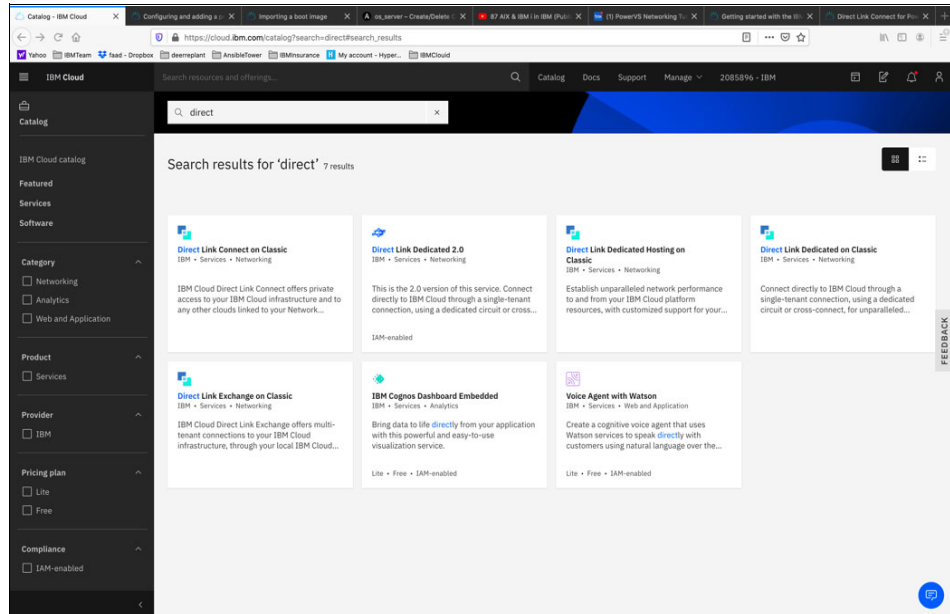


Figure 6-19 Searching for direct

3. Select **Direct Link Connect on Classic** and then, click **Create**. No options are available to select, as shown in Figure 6-20.

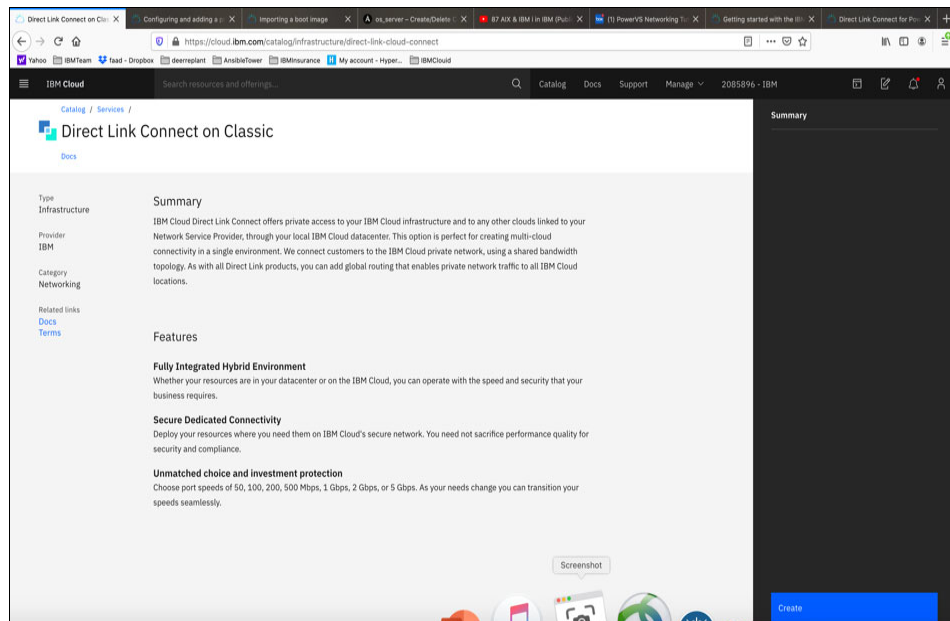


Figure 6-20 Direct Link Connect on Classic

4. Choose **Order Direct Link Connect**, as shown in Figure 6-21.

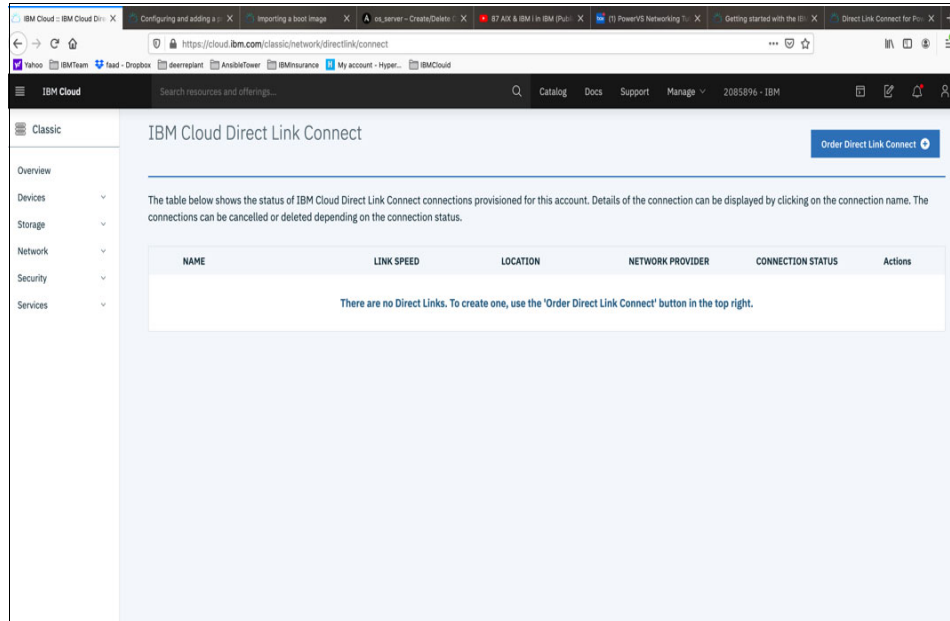


Figure 6-21 Choosing Order Direct Link connection

5. Make the following choices:

- 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)**, see Figure 6-22.

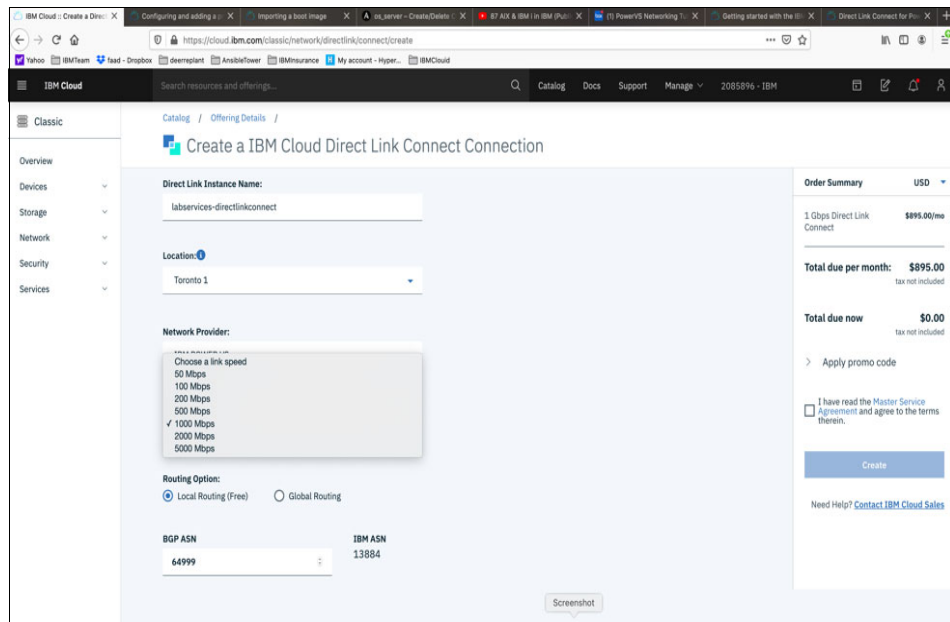


Figure 6-22 Configuring the direct link

Global routing requires extra charges and allows for easier Power Systems Virtual Server location-to-Power Systems Virtual Server location communication. You also must 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 our 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 then, click **Create**.

A support case is opened that includes the required information, as shown in Figure 6-23.

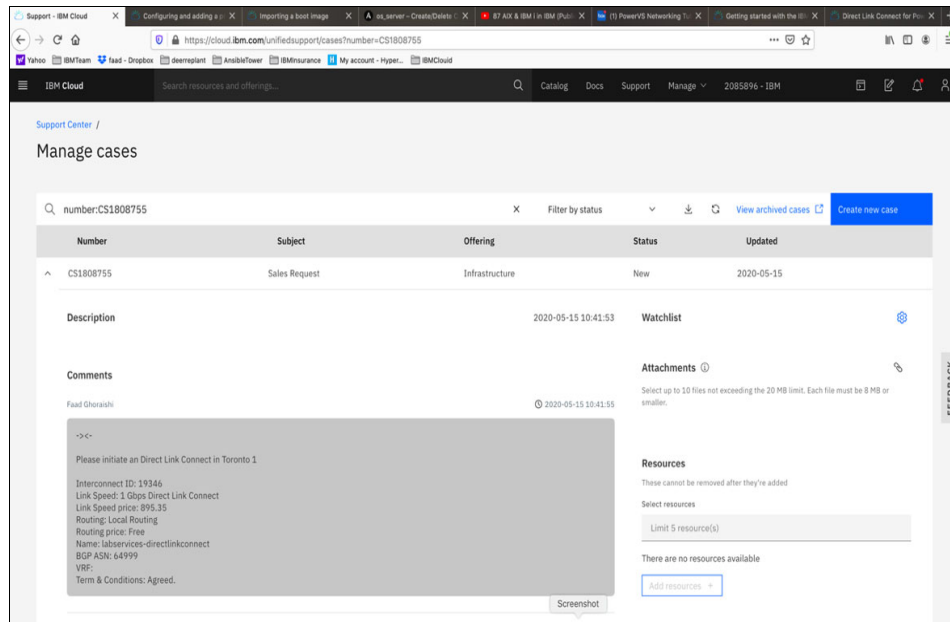


Figure 6-23 Support case window

After this process is complete, you are contacted by IBM Support and requested to complete and answer some questions in an attached document. You must 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 works. You can test your connection by pinging IBM Cloud Linux or Windows VSI from your Power Systems Virtual Server Instances and vice versa.

In our example, we used 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.

7. Log in to IBM Cloud and click the catalog and then, search for Vyatta.

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 GRE tunnels are provisioned on Vyatta Gateways, you must order one Vyatta Gateway in each Power Systems Virtual Server location.

Our example involves ordering one Vyatta in LON06 and the other in TOR01 data centers where the Power Systems Virtual Server locations exist.

In our example, we used 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**. Then, search for Vyatta, as shown in Figure 6-24.

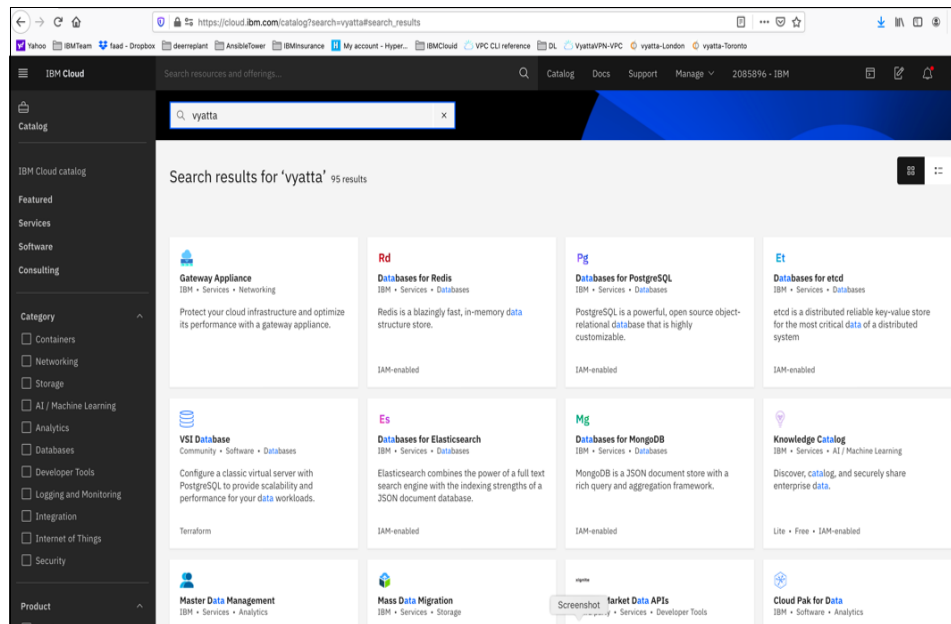


Figure 6-24 Searching for Vyatta

2. Select **Gateway Appliance** and then create it, as shown in Figure 6-25.

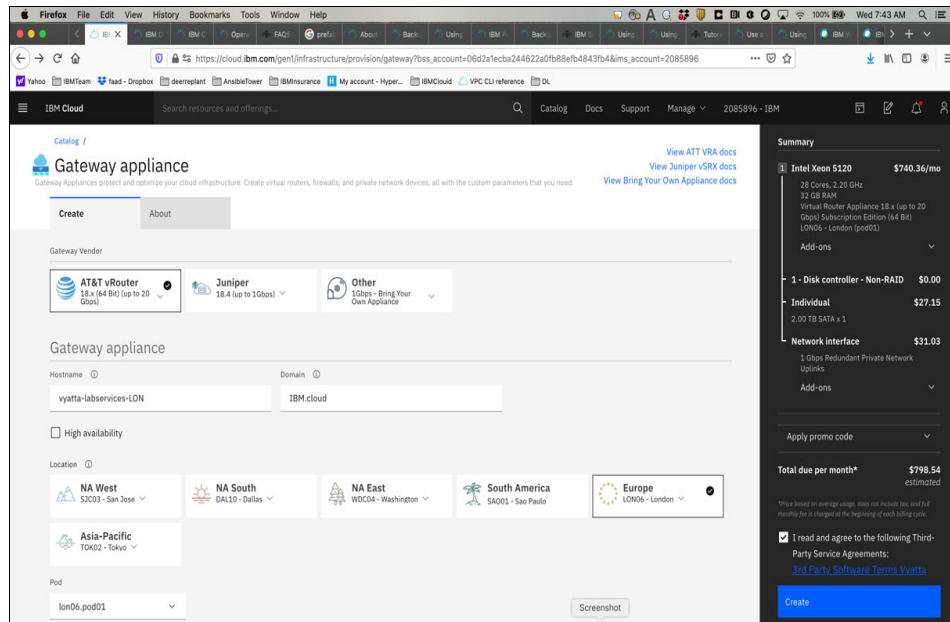


Figure 6-25 Selecting the location

3. Select **AT&T vRouter**, which is the Vyatta Gateway. Other choices of gateways are available, but we use Vyatta for our 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. We cleared this option for our example.
- Select the location by pressing the arrow key in each location to find the specific data center where you Power Systems Virtual Server location are available.
- Choose a POD if several PODs are in the selected data center location.
- Select the CPU single or dual processor. We chose **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. We chose private network interface only (see Figure 6-26).

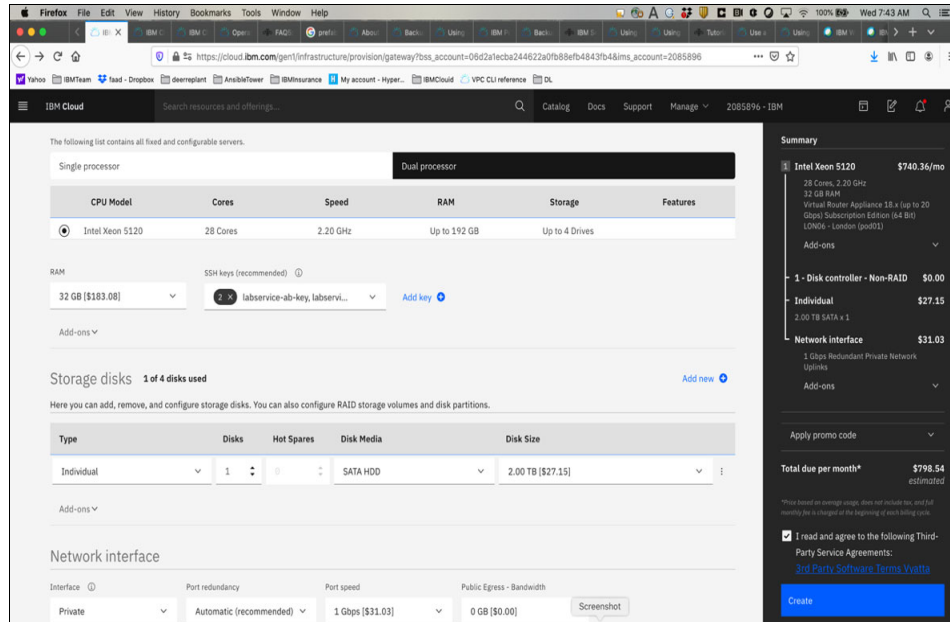


Figure 6-26 Choosing the private network interface

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).

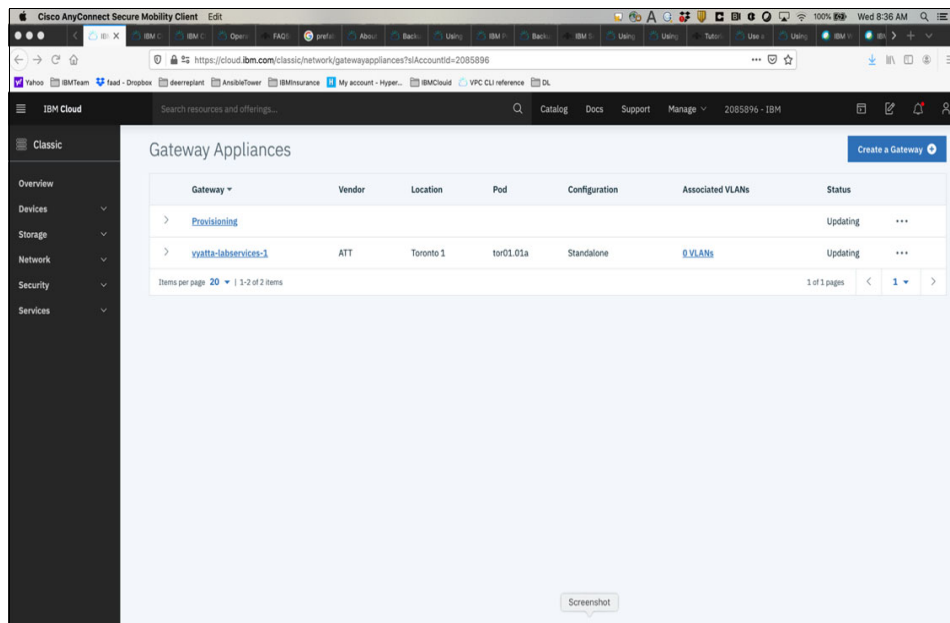


Figure 6-27 Listed gateways

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

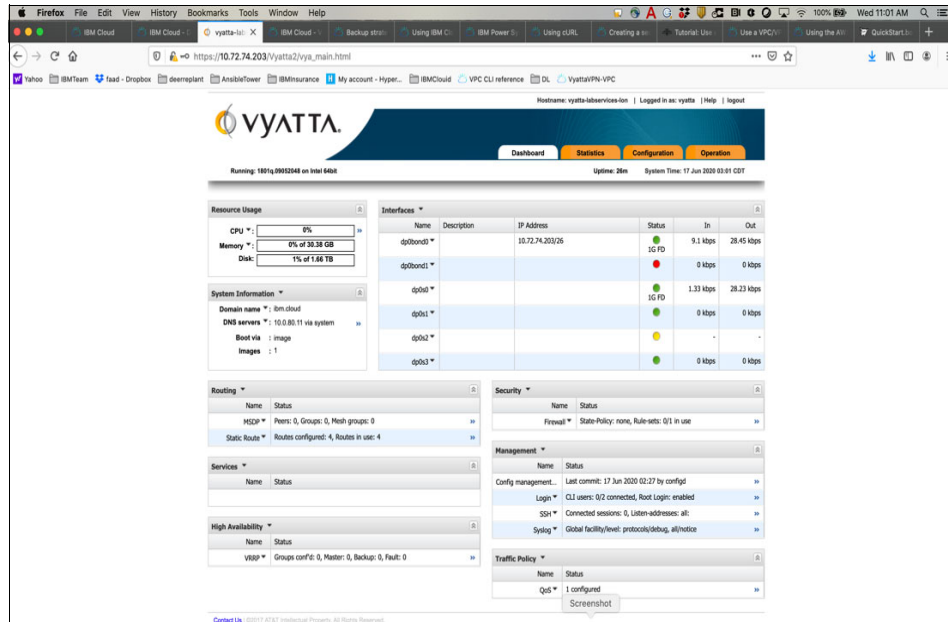


Figure 6-28 Vyatta gateway

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

You must 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.

You also must 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 complete their end of the GRE tunnel.

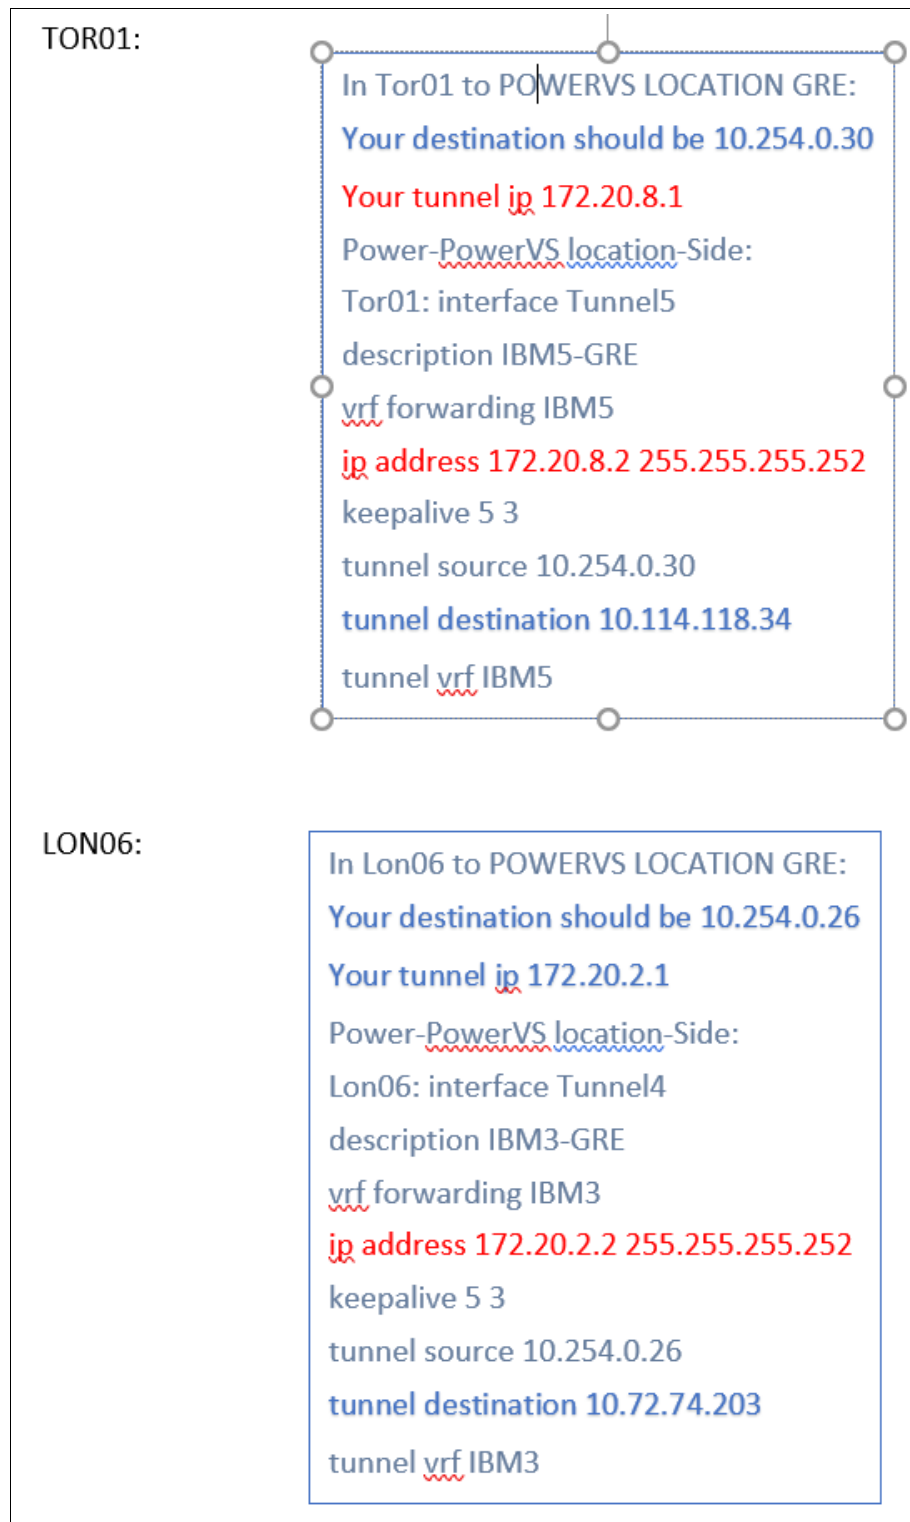


Figure 6-29 GRE information

The items that are shown in red in Figure 6-29 is 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

Verify your Vyatta Gateway access.

The Vyatta Gateway address can be find in the IBM Cloud UI under Devices.

Log in to IBM Cloud UI and click **IBM Cloud** on upper left side. Then, click **Devices** and choose the Vyatta system that you want to configure:

- ▶ `vyatta-labservices-lon.ibm.cloud`
- ▶ `vyatta-labservices-tor.ibm.cloud`

For LON06, click **London Vyatta: vyatta-labservices-lon.ibm.cloud** (see Figure 6-30).

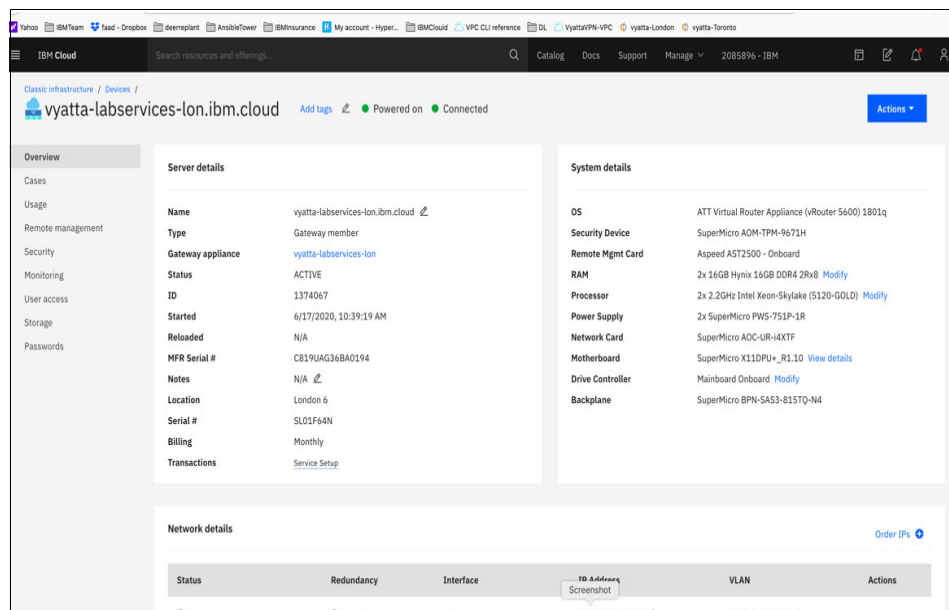


Figure 6-30 London Vyatta

Under Network Details, you see your Vyatta Gateway IP address.

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.

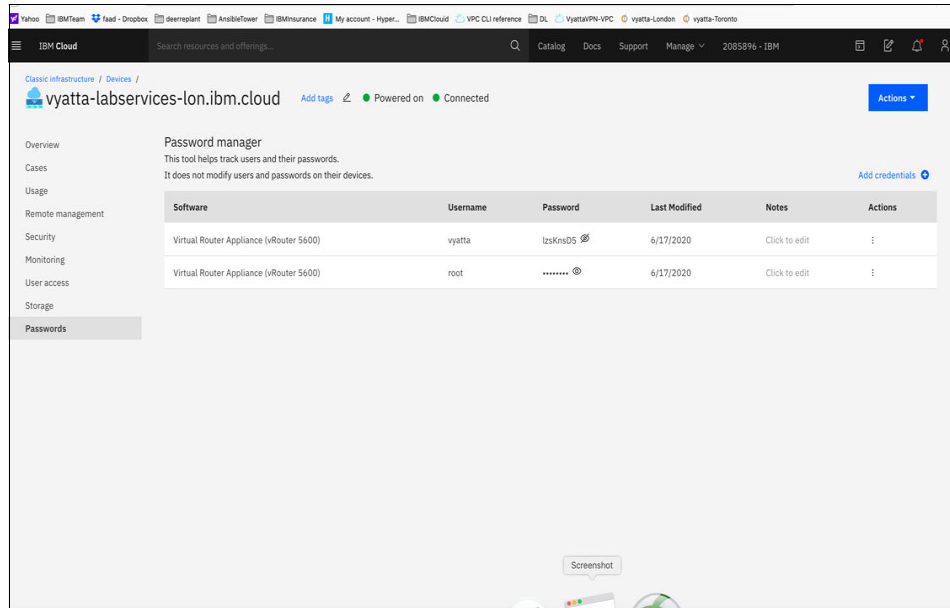


Figure 6-31 Vyatta password

Open a browser and log in to the Vyatta Gateway by using the following credentials (see Figure 6-32 on page 185):

- ▶ 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 IPs in IBM Cloud, you must start your MotionPro Plus VPN access, which gives you access to IBM Cloud private IPs.

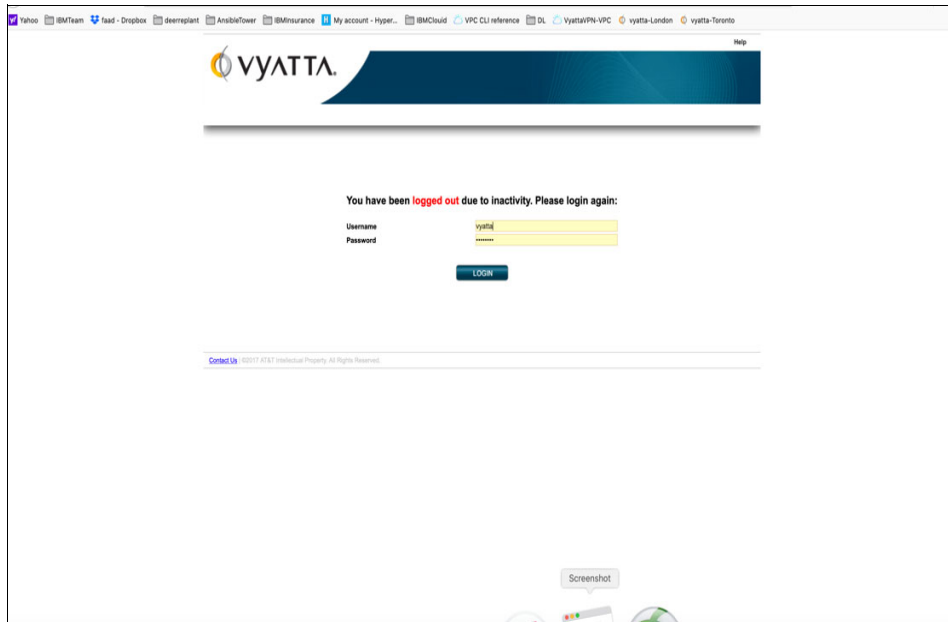


Figure 6-32 Vyatta login window

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.

Set up Power Systems Virtual Server location GRE tunnels in the Vyatta Gateways.

For more information about configuring GRE tunnels, see the following resources:

- ▶ IBM Cloud Docs: [VPN into a secure private network](#)
- ▶ Vyatta System's [Tunnels Reference 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.

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 L0N06 (or “data center 1”).
- ▶ GRE tunnel on Vyatta to end the Power Systems Virtual Server location GRE in T0R01 (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 Reference 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

```
The default interactive shell is now zsh.
To update your account to use zsh, please run `chsh -s /bin/zsh`.
For more details, please visit https://support.apple.com/kb/HT208050.
Faads-MacBook-Pro:~ faadghoraishi$
Faads-MacBook-Pro:~ faadghoraishi$ ssh vyatta@10.72.74.203
Welcome to AT&T vRouter 5600

Welcome to AT&T vRouter
Version:      1801q
Description: AT&T vRouter 5600 1801q
Linux vyatta-labservices-lon 4.9.0-trunk-vyatta-amd64 #1 SMP Debian 4.9.124-0vyatta2+2.1 (2018-09-05) x86_64
Last login: Tue Jun 30 00:58:24 2020 from 10.1.232.20

vyatta@vyatta-labservices-lon:~$
```

Figure 6-34 Logging in to GRE Tunnel in London06

For this example, we use the information that is provided by IBM Support for LON06 GRE, as shown in Figure 6-35.

```
In Lon06 to POWERVS LOCATION GRE:
Your destination should be 10.254.0.26
Your tunnel ip 172.20.2.1
Power-PowerVS location-Side:
Lon06: interface Tunnel4
description IBM3-GRE
vrf forwarding IBM3
ip address 172.20.2.2 255.255.255.252 (172.20.2.2/30)
keepalive 5 3
tunnel source 10.254.0.26
tunnel destination 10.72.74.203
tunnel vrf IBM3
```

Figure 6-35 London06 GRE information

Run the following commands (in our example, our tunnel is named tun0 on the Vyatta Gateway:

- ▶ 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 show in the GUI
- ▶ ssh vyatta@10.114.118.34
- ▶ ssh to Tor01 Vyatta Gateway

For this example, we use the information that was provided by IBM Support for TOR01 GRE, as shown in Figure 6-36.

```
In Tor01 to POWERVS LOCATION GRE:
Your destination should be 10.254.0.30
Your tunnel ip 172.20.8.1
Power-PowerVS location-Side:
Tor01: interface Tunnel5
description IBM5-GRE
vrf forwarding IBM5
ip address 172.20.8.2 255.255.255.252
keepalive 5 3
tunnel source 10.254.0.30
tunnel destination 10.114.118.34
tunnel vrf IBM5
```

Figure 6-36 GRE Tunnel in Toronto01

Run the following commands (for this example, we named our tunnel tun0 in the Vyatta Gateway, which is the same name as the other Vyatta Gateway):

- ▶ 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

In this section, we describe setting up a new tunnel in each of the two Vyatta gateways to allow for cross-Vyatta connection by way of a GRE tunnel.

In this example, we chose the tunnel address and tunnel source and destination IPs. The tunnel address can be any IP subnet that you choose (we named our tunnel tun1 in both Vyatta gateways).

We also selected a similar IP as the IPs that are used in the Power Systems Virtual Server location GRE tunnels. We choose a CIDR of /30 because we need only two IP address: 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 following command to configure 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 following commands to configure LON06 GRE:

- ▶ 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 set up static routes in each Vyatta to point the subnets for our 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 LON06 must to point to the subnets in TOR01 and vice versa.

We configure both GREs to the Power Systems Virtual Server location and between Vyattas.

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 not exists and you can ping between your Power Systems Virtual Server Instances in each Power Systems Virtual Server location. You also can 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 to address this issue. IBM Support must address this issue from their Cisco Router side.

Configuring a reverse-proxy Centos VSI

In this section, we discuss the procedure to configure a reverse proxy to allow access to 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 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 our deployment is shown in Figure 6-37.

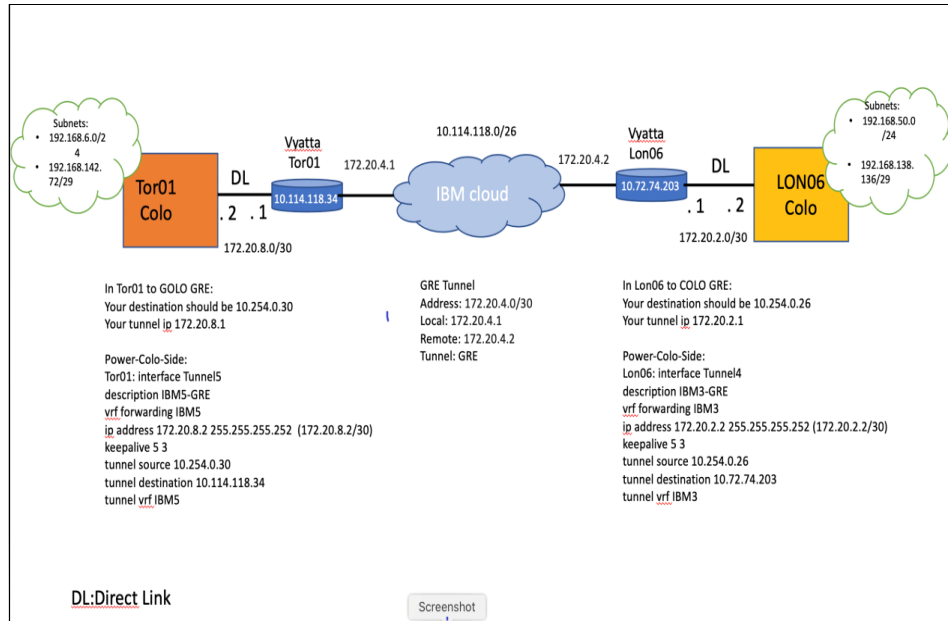


Figure 6-37 Network architecture

6.3 Operating system level replication use case with PowerHA SystemMirror for i geographic mirroring

This section provides a use case for system level replication.

6.3.1 Introduction

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).

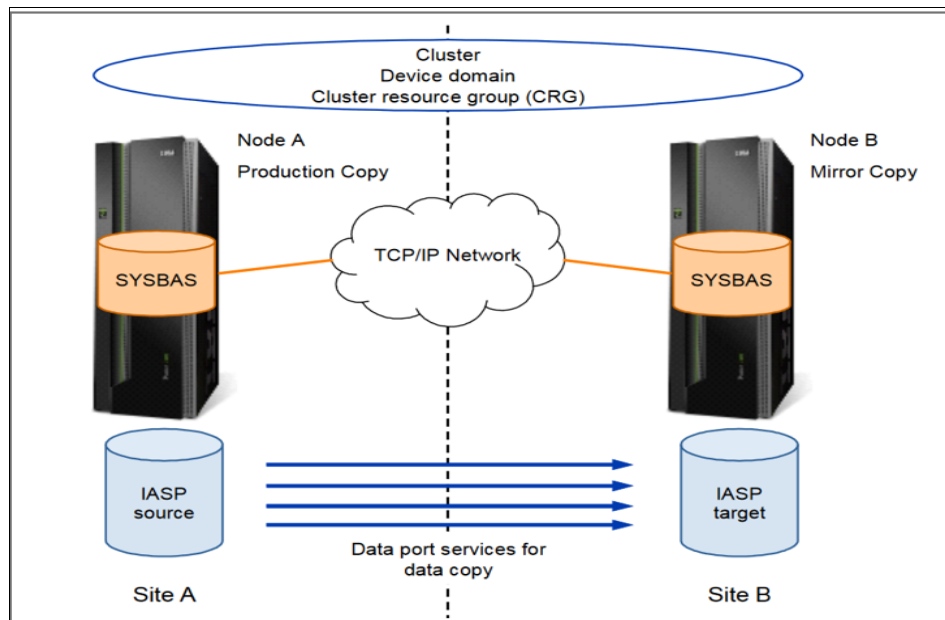


Figure 6-38 PowerHA SystemMirror for i geographic mirroring

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.

6.3.2 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, include PowerHA for i Enterprise Edition, 5770-HAS option 1.

6.3.3 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 also can 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.

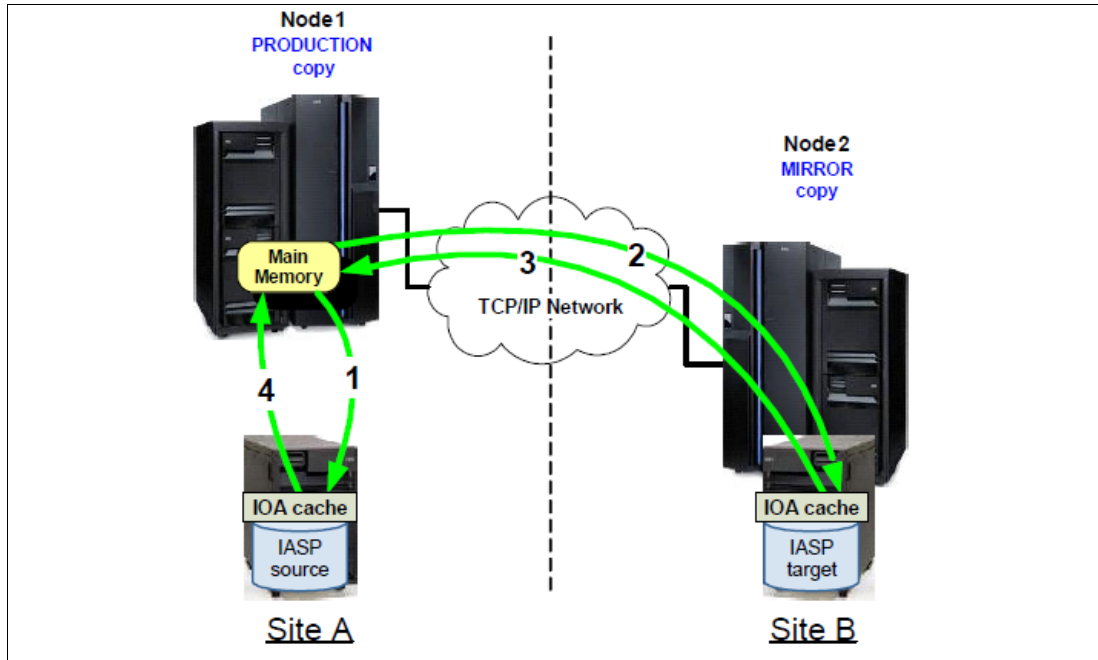


Figure 6-39 Synchronous transmission mode with synchronous mirroring mode

These writes must complete to the disk (often, disk cache) on the preferred target system (acknowledgment operation [3]) and the production copy system (acknowledgment operation [4]) before sending the acknowledgment 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).

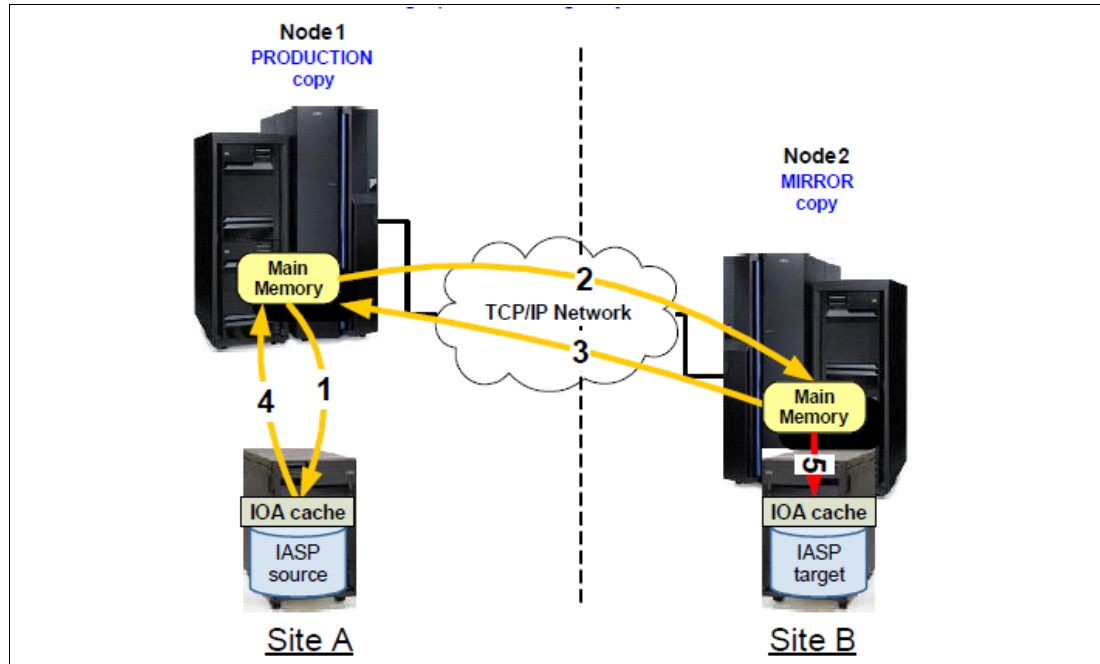


Figure 6-40 Synchronous transmission mode with asynchronous mirroring mode

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 complete 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, switchover or failover 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 on 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 tie up 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. With tracked changes, geographic mirroring can avoid full resynchronization after it resumes in many cases, which minimizes the exposure of timeframes where no valid mirror copy is available.

Tracking space is configured when geographic mirroring is configured or later by using the Change Auxiliary Storage Pool Session (**CHGASPSN**) command.

Tracking space is allocated inside of the IASPs. The more tracking space that is specified, the more changes the system can track. The amount of space for tracking can be defined by the user up to 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 Change Auxiliary Storage Pool Session (**CHGASPSSN**) command. The default value of the Suspend timeout (**SSPTIMO**) parameter is 120 seconds.

Network topology

Geographic mirroring can be used in environments over any distance; however, only 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 up to 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 causes 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).

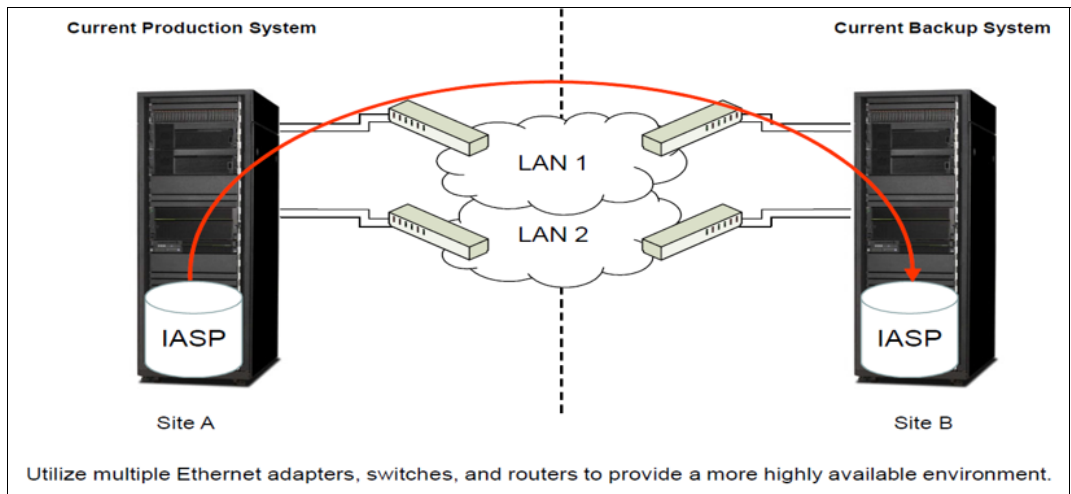


Figure 6-42 Preferred network configuration for Geographic Mirroring

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 more flexibility for the distance between systems, implications still result from undersizing the source, target, or the network infrastructure between the sites.

Minimizing the latency (that is, the time that the production system waits for the acknowledgment that the information was received on the target system) is key to good application performance.

Back up planning for geographic mirroring

Before you implement HA that is based on geographic mirroring, you must follow a backup strategy that adheres to your business needs.

If the current strategy is to end production applications and perform backups on the production system and your business has no requirement to change this strategy, you can continue to run backups as usual. 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.

6.3.4 Implementation of PowerHA SystemMirror for i geographic mirroring in the Power Systems 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 creating the IBM i instances by way of IBM Cloud Services, take note of 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, and then, add each new volume (SYSBAS or IASP) individually (later on in the process) to keep track of each disk unit ID as 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, they are visible from the Cloud interface only. Therefore, it is useful to individually keep track of the disk unit ID (from IBM i interface) and disk name (from Cloud interface) so you can assign your disks to the suitable ASP. These disks are best created at a later step in the process.

When choosing a name for the instance, consider that the 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.

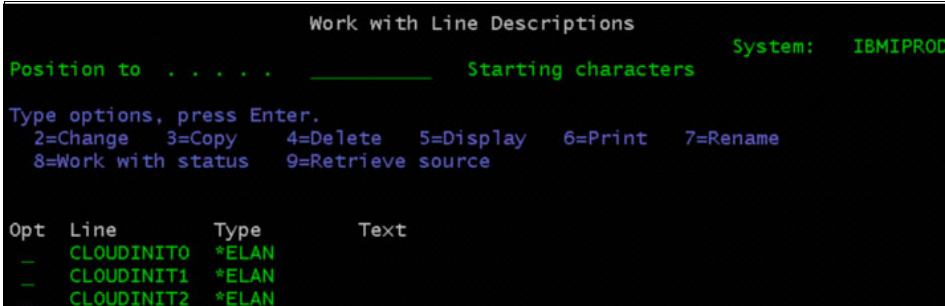
Complete the following steps:

1. In the cloud GUI, choose IBM PowerHA under Software Licenses, which loads and enable 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**. Then, press **F15 (Accept ALL)** on each license. After all licenses are accepted, press **F3 (Exit)** to main menu.

During this time, the IP interfaces and line descriptions are still being configured. This process can take up to 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
--  -
 0   CLOUDINIT0 *ELAN
 1   CLOUDINIT1 *ELAN
 2   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).

```

Work with TCP/IP Interfaces
System:  IBMIPROD
Type options, press Enter.
1=Add  2=Change  4=Remove  5=Display  9=Start  10=End

  Opt  Internet      Subnet      Line      Line
      Address      Mask       Description  Type
-----
  ___  127.0.0.1      255.0.0.0   *LOOPBACK  *NONE
  ___  192.168.6.118  255.255.255.0 CLOUDINIT0 *ELAN
  ___  192.168.142.74 255.255.255.248 CLOUDINIT1 *ELAN

```

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 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 SVRSPCVL(*TELNET) AUTOSTART(*YES)
- CHGTCPSVR SVRSPCVL(*FTP) AUTOSTART(*YES)
- CHGTCPSVR SVRSPCVL(*SSHD) AUTOSTART(*YES)
- CHGTCPSVR SVRSPCVL(*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 a **PWRDWN SYS 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 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 way of the internal IP. Optionally, an external IP interface can be assigned to the instance by way of 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)
```

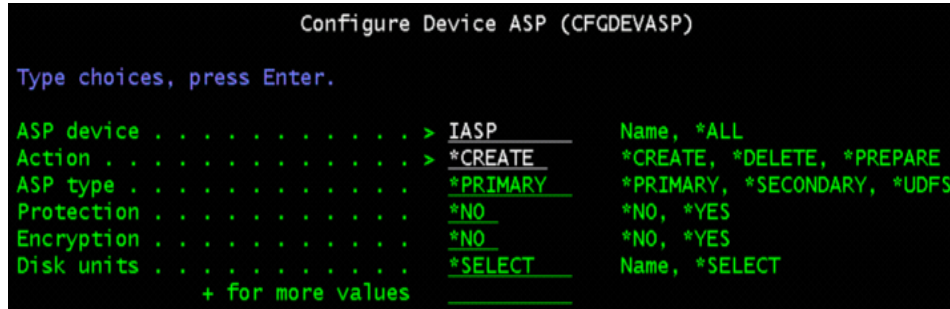


Figure 6-45 Creating an IASP

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.

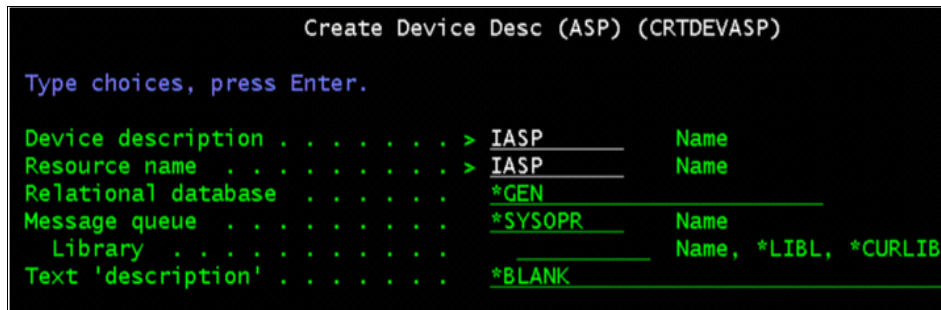


Figure 6-46 Creating an ASP device

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>) RSRNAME(<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 is correctly IASP-enabled, a DR solution with PowerHA Geographic Mirroring can be implemented.

On both instances, run the following commands:

- ▶ CHGTCPSPVR SVRSPCVL(*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 then, 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 now 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.

```

PowerHA                               Work with Cluster Nodes
Local node . . . . . : CLOUDPRD
Consistent information in cluster . . . : Yes

Type options, press Enter.
 1=Add  2=Change  4=Remove  5=Display more details  6=Work with monitors
 8=Start 9=End

Opt  Node      Status      Device Domain
--  -
 6   CLOUDDR   Active     CLDDEVDMN
 6   CLOUDPRD  Active     CLDDEVDMN

```

Figure 6-47 Working with cluster nodes

Adding cluster nodes to the device domain

From the Production IBM i instance, complete the following steps to add the PowerHA cluster nodes to the same device domain (from WRKCLU, Option 7 [Work with device domains], does not show a device domain listed):

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 now shows the new device domain with one node (the production node) listed to the right (see Figure 6-48).

```

POWERHA                               WORK WITH DEVICE DOMAINS
Consistent information in cluster . . . : Yes

Type options, press Enter.
 1=Add  6=Work with nodes  7=Work with switchable hardware

Opt  Device Domain      Number of Nodes  -----Nodes-----
 6   CLDDEVDMN          2   CLOUDPRD CLOUDDR

```

Figure 6-48 Working with device domain

3. Select **Option 6 (Work with nodes)** and then, 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 now shows the device domain with both nodes listed, as shown in Figure 6-49.

```

PowerHA                               Work with Device Domain Nodes
Device domain . . . . . : CLDDEVDMN
Consistent information in cluster . . . : Yes

Type options, press Enter.
  1=Add    4=Remove

Opt      Node          Status
--      -
  _      CLOUDDR       Active
  _      CLOUDPRD      Active
  
```

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 then, press **Enter** (see Figure 6-50).

```

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

Opt      Container/CRG      Type      Status      Primary Site/Node
--      -
  _6     CLDGEOMIR          *DEV     Active     CLOUDPRD
  
```

Figure 6-50 Working with cluster resource groups

2. Select **Option 1 (Cluster resource group)** from the pop-up menu (at V7R4) and press **Enter**.
3. Specify Type=*DEV, Exit Program=*NONE, and User Profile=*NONE.
4. Enter a "+" next to Recovery domain node list and then, press **Enter**.

- 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). Then, page down.

```
PowerHA                               Work with Recovery Domain

Cluster resource group . . . . . : CLDGEOMIR
Consistent information in cluster . . . : Yes

Type options, press Enter.
1=Add node   4=Remove node   5=Display more details

Opt   Node           Status      Current Node Role   Preferred Node Role   Site Name
---   ---           ---        ---                ---                  ---
_     CLOUDDR         Active     *BACKUP 1          *BACKUP 1             TORONTO2
_     CLOUDPRD        Active     *PRIMARY           *PRIMARY              TORONTO1
```

Figure 6-51 Work with recovery domains

- 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 then, page down.
- (Optional) Enter a text description for the CRG.
- 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):

- 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

Opt   Container/CRG   Type   Status   Primary Site/Node
---   ---           ---   ---     ---
_     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 *DEV D and then, 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

Opt      Configuration Object Name      Object Type      Device Type      Device Subtype      Vary Online
-----
 1      IASP01                          *DEV D          *ASP             Primary             *ONLINE

```

Figure 6-53 Working with configuration objects

3. If automatic vary-on of the IASP is wanted 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(*ASYN C) 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
Action . . . . . > *CREATE
Source site . . . . . TORONTO1
Target site . . . . . TORONTO2
Session . . . . . CLDGEOMIR
Source ASP copy description . . . . . CLOUDPRD
Target ASP copy description . . . . . CLOUDDR
Transmission delivery . . . . . *ASYN C
Disk units . . . . . *SELECT
+ for more values
Name
*CREATE, *DELETE
Name, *
Name, *
Name, *NONE
Name
Name
*SYN C, *ASYN C
Name, *SELECT

```

Figure 6-54 Configuring 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.

- 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.
- 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

Opt   ASP Device      ASP Copy      ASP Session      Session
-----
25    IASP01          CLOUDPRD     CLDGEOMIR       *GEOMIR
      IASP01          CLOUDDR      CLDGEOMIR       *GEOMIR

```

Figure 6-55 Working with ASP copy descriptions

- After the configuration process is complete, DSPASPSSN SSN(<ASP Session Name>) shows the status as *RESUMPND.
- Vary on the IASP device, and recheck the ASP session status. It shows as *SUSPENDED.
- 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).

```

PowerHA                               Display ASP Session                               IBMIPROD
                                         07/16/20 09:38:08
Session . . . . . : CLDGEOMIR
Type . . . . . : *GEOMIR

Source node . . . . . : CLOUDPRD
Target node . . . . . : CLOUDDR
Transmission Delivery . . . . . : *ASYNC
                                         More...

Copy Descriptions

ASP Device      ASP Copy      Role      Vary State      Replication State      Data State
IASP01          CLOUDPRD     PRODUCTION AVAILABLE     ACTIVE     USABLE
                CLOUDDR      MIRROR     VARYON      ACTIVE     UNUSABLE

Press Enter to continue
                                         Bottom

F3=Exit  F5=Refresh  F11=View 2  F12=Cancel  F19=Automatic refresh

```

Figure 6-56 Displaying ASP session

CHGASPSSN SSN(<ASP Session Name>) OPTION(*RESUME) restarts synchronization of geographic mirroring.

6.3.5 Switching a geographic mirroring environment

In this section, we describe the geographic mirroring synchronization process and show 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 also are 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. You must 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 Display ASP Session DSPASPSSN SSN (**CLDGEOMIR**) command to confirm that the nodes and ASP copies were reversed.

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.

We use the following failure scenarios 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

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, a CPABB02 Cluster resource groups are failing over to node backup-node. (C G) inquiry message 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 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. The backup node now has the production copy:

```
DSPASPSSN SSN(CLDGEOMIR)
```

Important: The terms can be confusing here. After a switch or failover, the “backup” node (which refers to a physical system or partition) is the “source” node and contain 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 the **STRCRG CLUSTER(CLOUDCLU) CRG(CLDGEOMIR)** command.

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 then, press **F4**.
10. On the Change ASP Session display, enter the *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 often is required. The full resynchronization can take an extended time.

Consider the backup copy is 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".

6.3.6 Troubleshooting

With any new set-up, 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 noncloud 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 (that is, "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, or 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.

6.4 Logical replication use case with Bus4i

This section provides a logical replication use case with Bus4i.

6.4.1 Introduction

For many years, logical replication solutions were the only choice in the high availability (HA) and Disaster Recovery (DR) product market for IBM i and its predecessors.

IBM's introduction of PowerHA and its hardware-based replication solutions (at the operating system level or External Storage-based) provided a different approach for HA and DR in the IBM i platform.

However, even 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, allowing 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 and enables the customer to 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 type 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
- ▶ [RoboHA](#) 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, we use the Bus4i product.

6.4.2 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

6.4.3 Planning

Unscheduled computer downtimes are any IT user's ultimate nightmare. This issue is 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 assures in this process that operations can be carried out immediately with the secondary system when scheduled or unscheduled downtimes occur (see Figure 6-57).

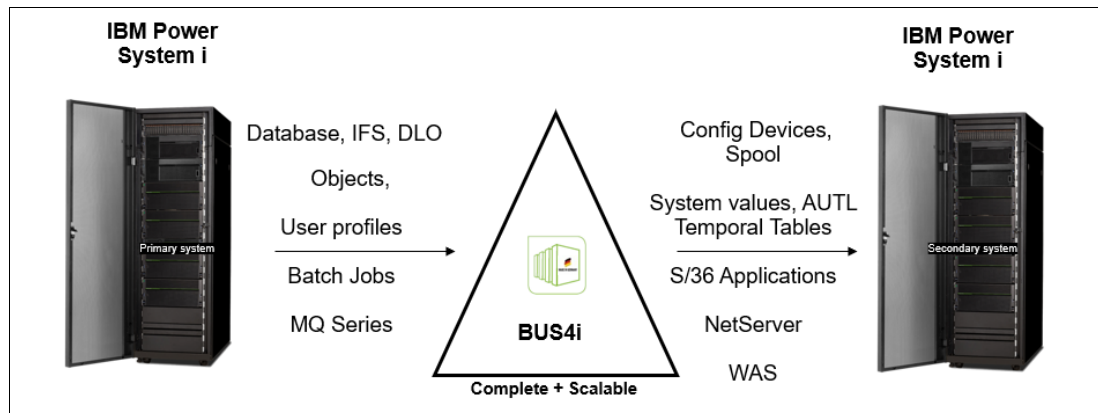


Figure 6-57 BUS4i replicates everything

When the nonfunctional server unit is restored, the systems are resynchronized.

To keep users consistently updated about the mirroring process status, they are notified about malfunctions by using SMS or email.

To replicate a system that uses BUS4i, some products 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 here because it also is 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 assures that the data is applied in an identical order as in the primary system. Administrative interventions into the primary system automatically are duplicated to the secondary system (see Figure 6-58).

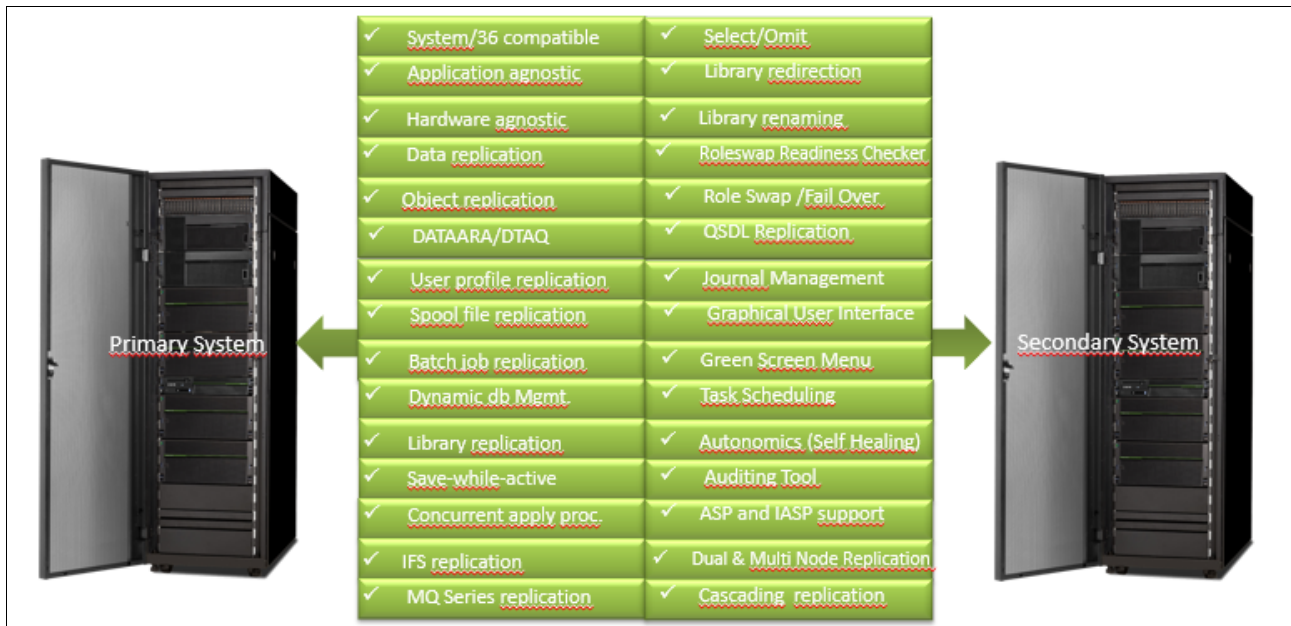


Figure 6-58 Entry to professional and enterprise functions in one solution

Intelligent sequence control restricts 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 then, 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 assures 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.

Mirroring in a test system enables you to verify alterations to the test system, which allows precise definition and minimization of downtimes in the real system.

Autonomy from user applications

User applications must not be adapted 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.

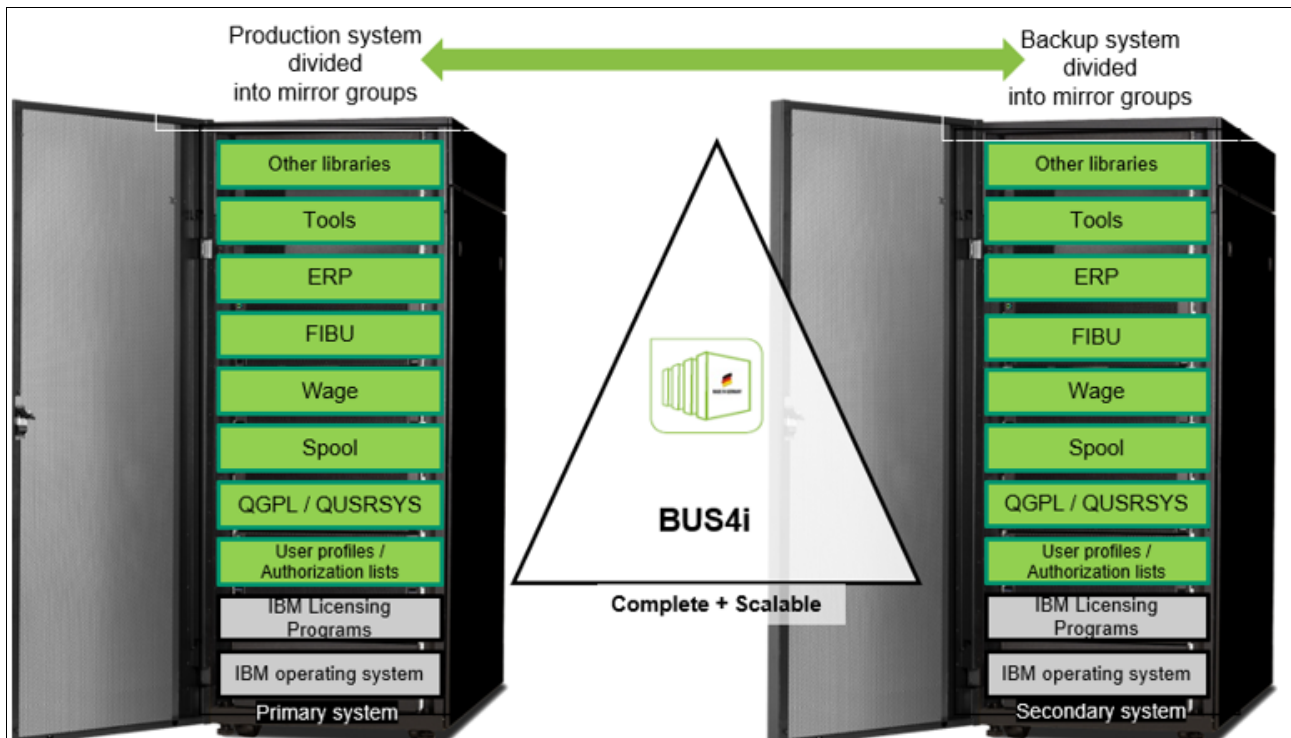


Figure 6-59 Most flexible configuration options

Secondary system with differing hardware models

One or multiple secondary systems must not be identical to the original system. The operating system release also must be identical only for 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 password) also are 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, and it is here that the synchronized process is suited. The transmission system waits with further processing until the receiving function acknowledges receipt and processing.

Processing routines that are independent from transmitting and receiving processes

Owing to autonomous processes, users have more flexibility in designing their systems and routines (see Figure 6-60).




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 (print-out) cannot be repeated; for example, 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 still must 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 endangered 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: locally with only one Power Systems, locally with two Power Systems, and remotely with on-premises sites, all in a hybrid architecture or in the cloud (see Figure 6-61).

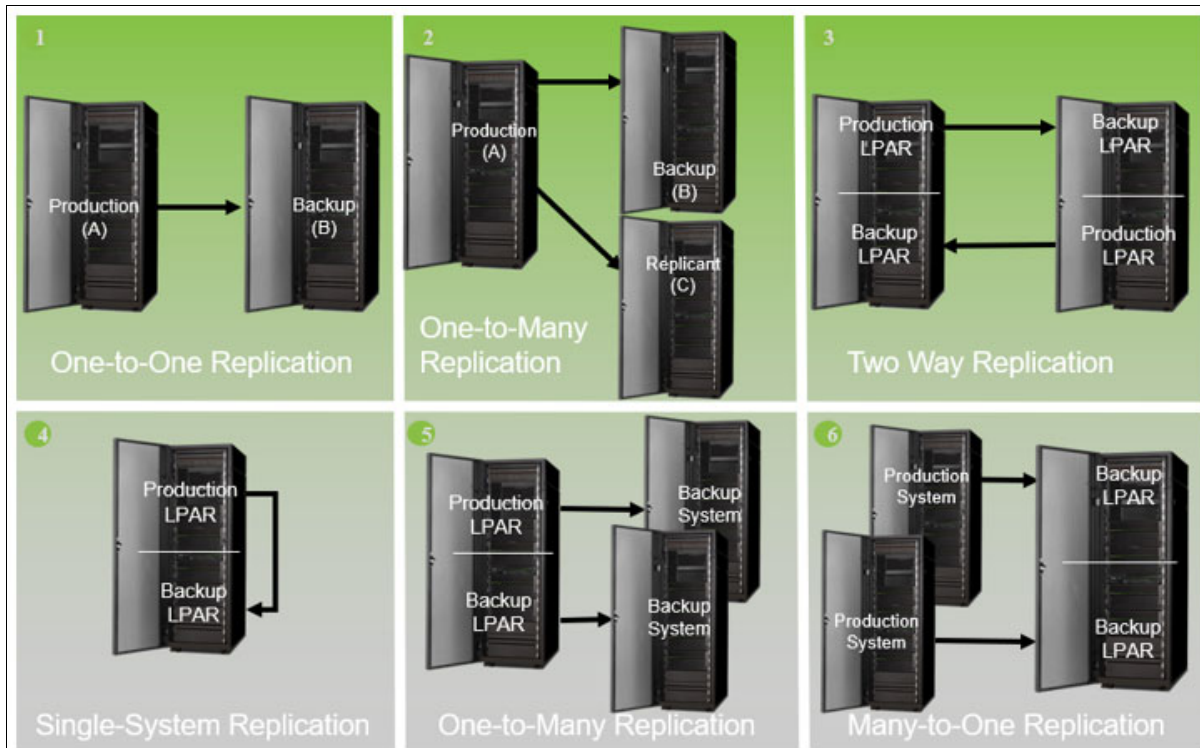


Figure 6-61 Multiple replication scenarios

The T.S.P. system BUS4i is based on the operational principle of journaling. *Journals* 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 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-tier 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 (see Figure 6-62):



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/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 a quick overview of the available information.

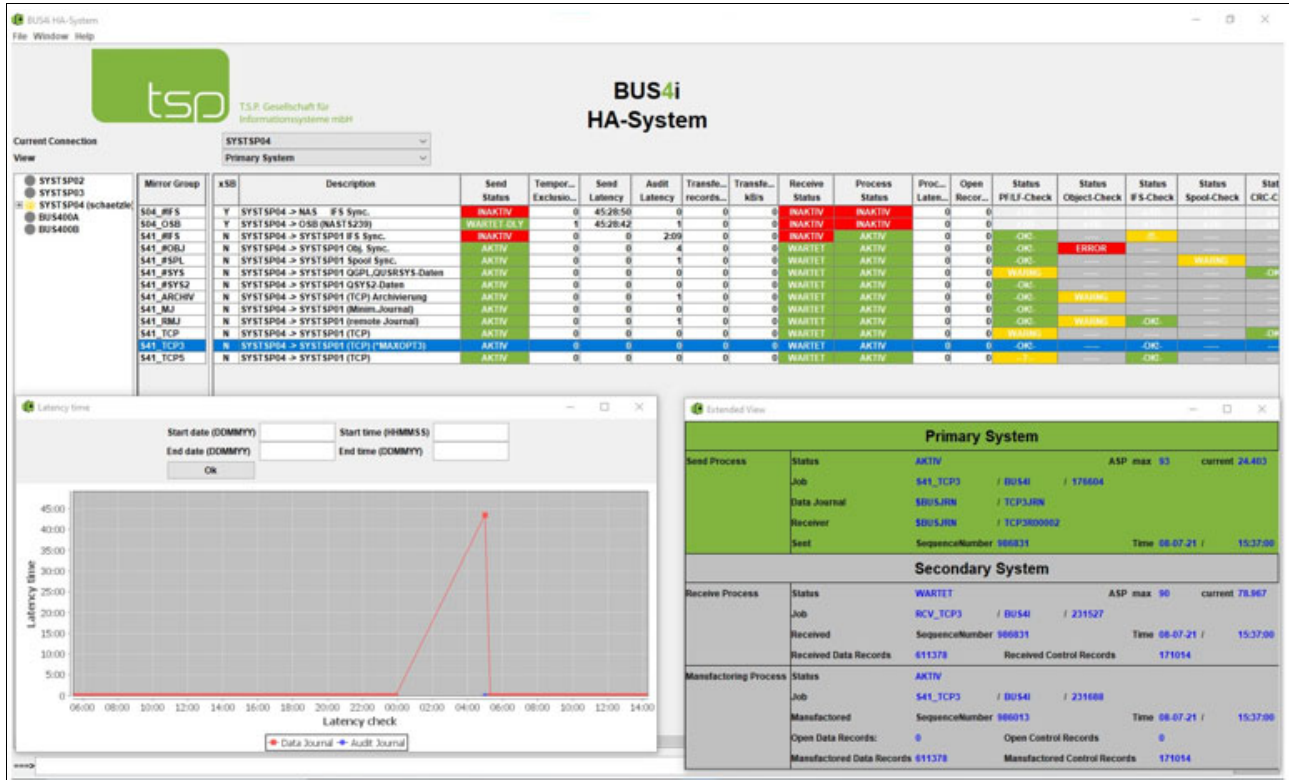


Figure 6-64 Graphical user interface

6.4.4 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).

```
Sitzung A - [24 x 80]
Datei Bearbeiten Ansicht Kommunikation Aktionen Fenster Hilfe
SRC000 BUS4i High Availability System R7.1 System: IBMI04
----- Main Menu Primarysystem -----
Options:
  1. Work with transmission processes
  2. Administer mirror groups
  3. Service functions
  4. Administer filters and rules
  5. Process control
  6. Backup tasks
  9. Switch System (Primary to secondary System)
  20. Administer journals
  21. Maintenance and analyses
  90. Signoff
Option or Command
==>
F1=Help F3=Exit F4=Prompt F8=Spool-Files F9=Retrieve
F10=Submitted Jobs F11=Active Jobs F12=Cancel F19=Active BUS4i Jobs
© Copyright T.S.P GmbH - Your Partner for High Availability -
IA A 20/007
|902 - Sitzung wurde erfolgreich gestartet. |Microsoft Print to PDF EIN PORTPROMPT:
```

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 (see Figure 6-66).

```
Sitzung A - [24 x 80]
Datei Bearbeiten Ansicht Kommunikation Aktionen Fenster Hilfe
TGT000 BUS4i High Availability System R7.1 System: IBMI04
----- Main Menu Secondarysystem -----
Options:
  1. Processing procedures
  2. Administer mirror groups
  3. Service functions
  5. Administer receiving processes
  6. Administer Remote Journal reading processes
  7. Menu Process Control
  8. Data Recovery for xSB-Versions
  9. Switch System (Secondary to Primary System)
  20. Administer journals
  21. Maintenance and analyses
  Weitere ...
Option and Command
==>
F1=Help F3=Exit F4=Prompt F8=Spool-Files F9=Retrieve
F10=Submitted Jobs F11=Active Jobs F12=Cancel F19=Active BUS4i Jobs
© Copyright T.S.P GmbH - Your Partner for High Availability -
MA A 09/044
|902 - Sitzung wurde erfolgreich gestartet. |Microsoft Print to PDF EIN PORTPROMPT:
```

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 in which, 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 for Payroll and Human Resources applications.

A mirror group P2S_#NEU also is available that can be used to include in it 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.

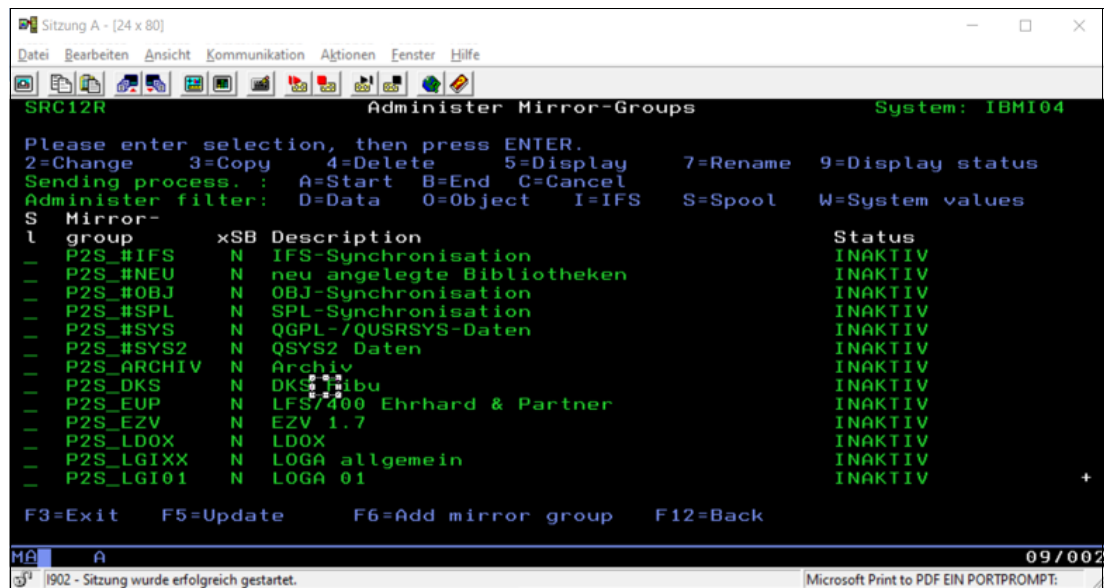


Figure 6-67 BUS4i mirror groups administration

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.

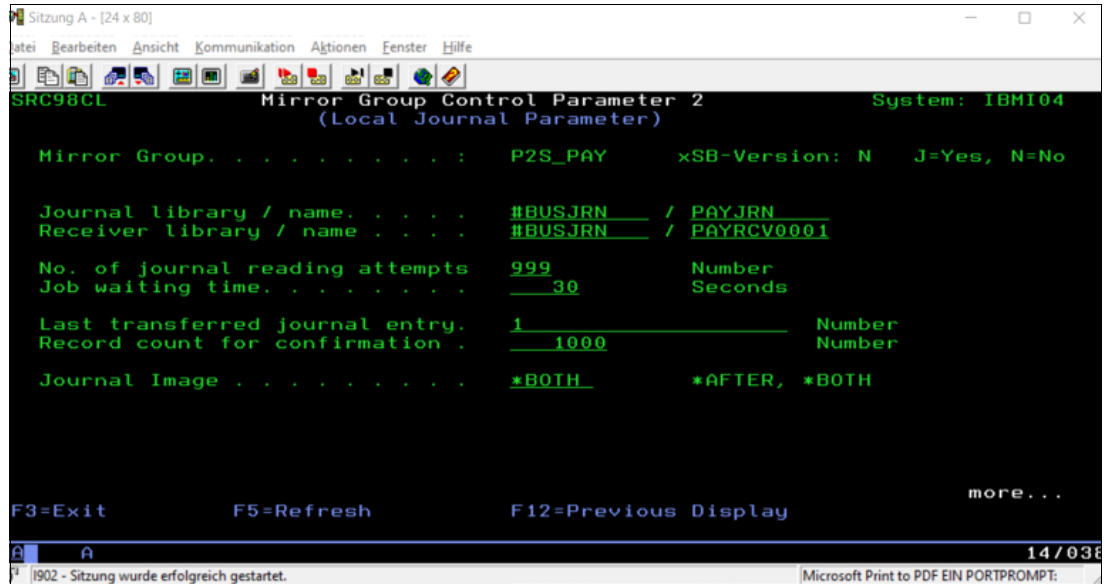


Figure 6-68 BUS4i Mirror Group Control Parameter 2 window

Figure 6-69 shows the following information:

- ▶ 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

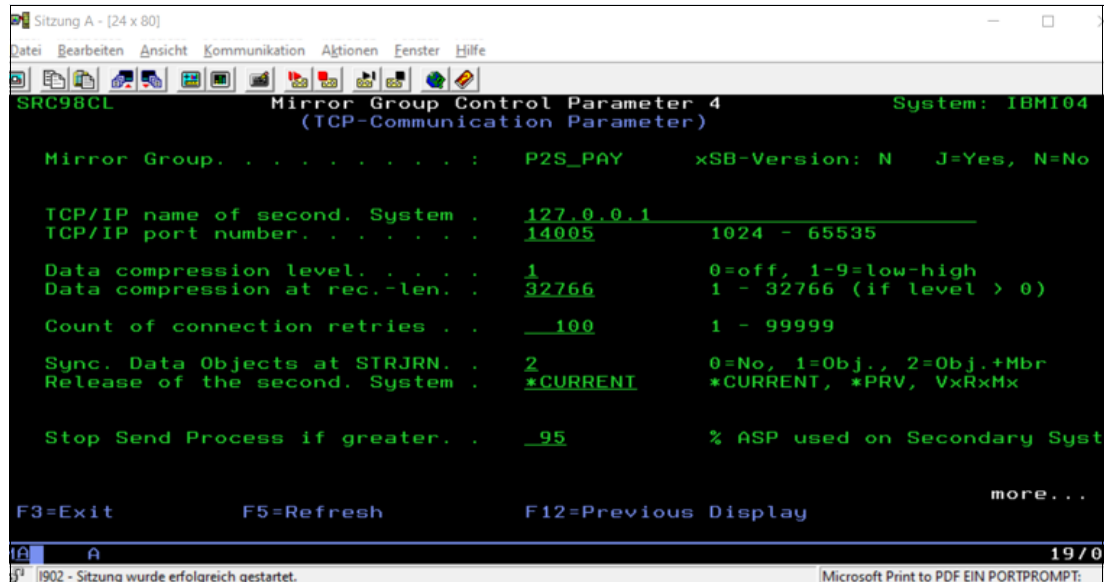


Figure 6-69 BUS4i Mirror Group Control Parameter 4 window

For mirror groups #IFS, #SYS, and QSYS2, the QAUDJRN system audit journal from the QSYS library, the QAOSDIAJRN from the QUSRSYS, and the QSQJRN from QSYS2 are used (see Figure 6-70 - Figure 6-72 on page 226).

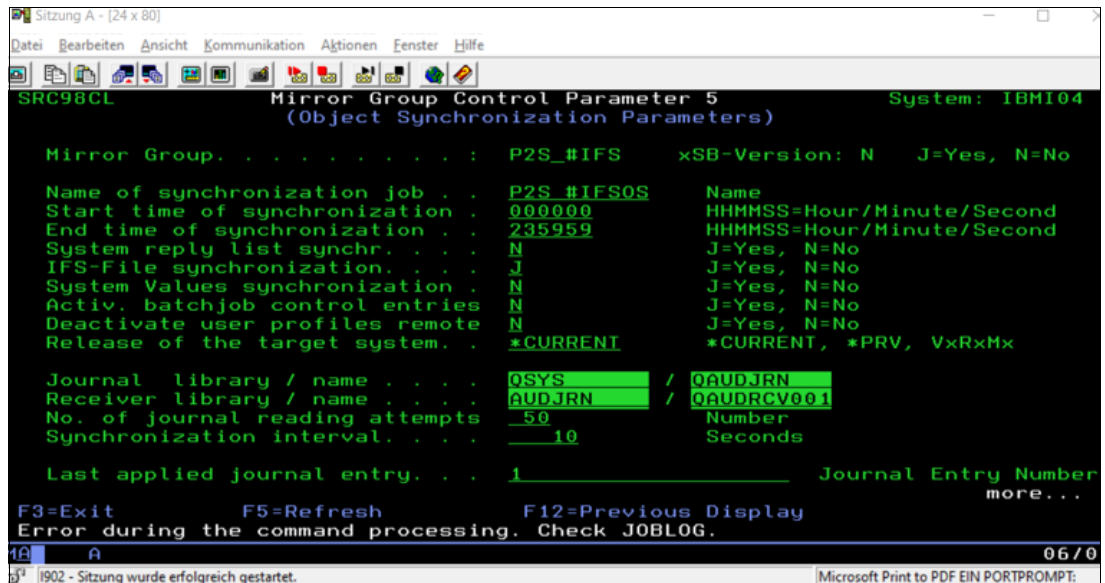


Figure 6-70 BUS4i Mirror Group Control Parameter 5 window

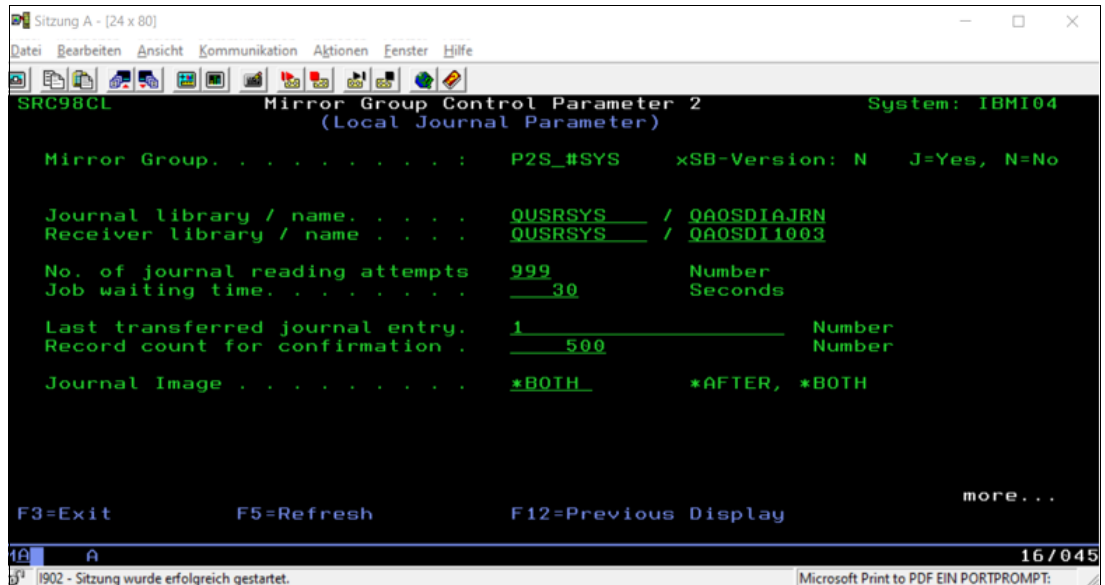


Figure 6-71 BUS4i Mirror Group Control Parameter 7 window

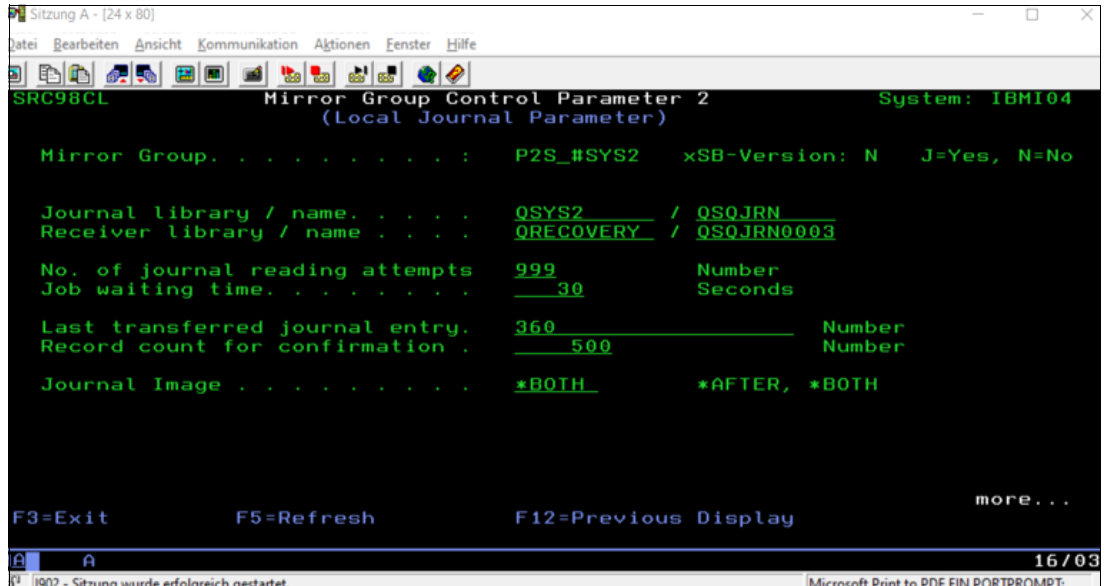


Figure 6-72 BUS4i Mirror Group Control Parameter 2 window

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

For example, 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 and Figure 6-75.

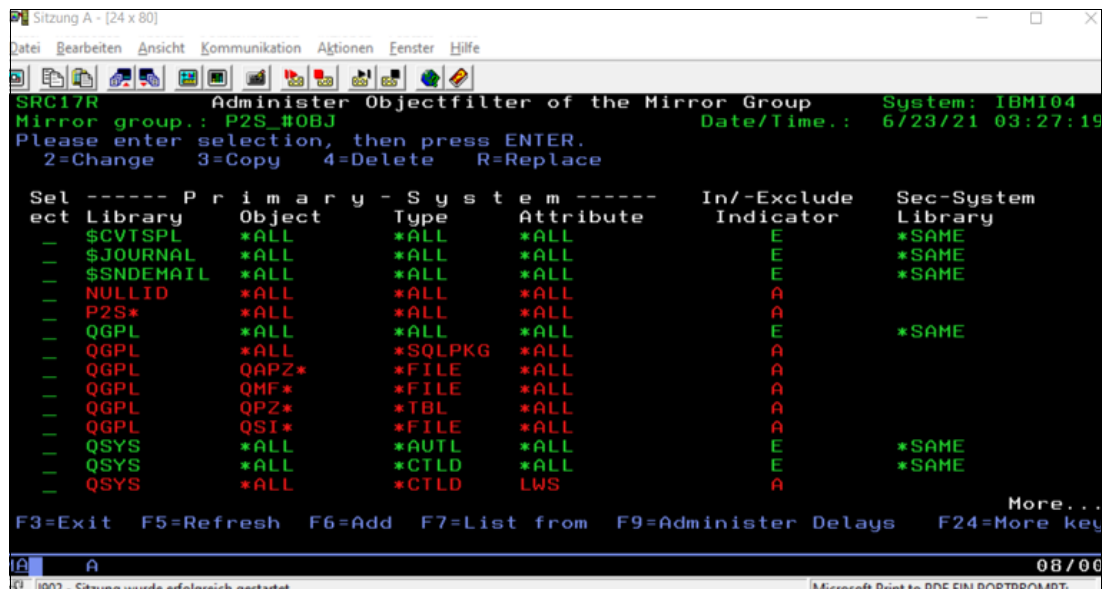


Figure 6-74 BUS4i objects filter entries

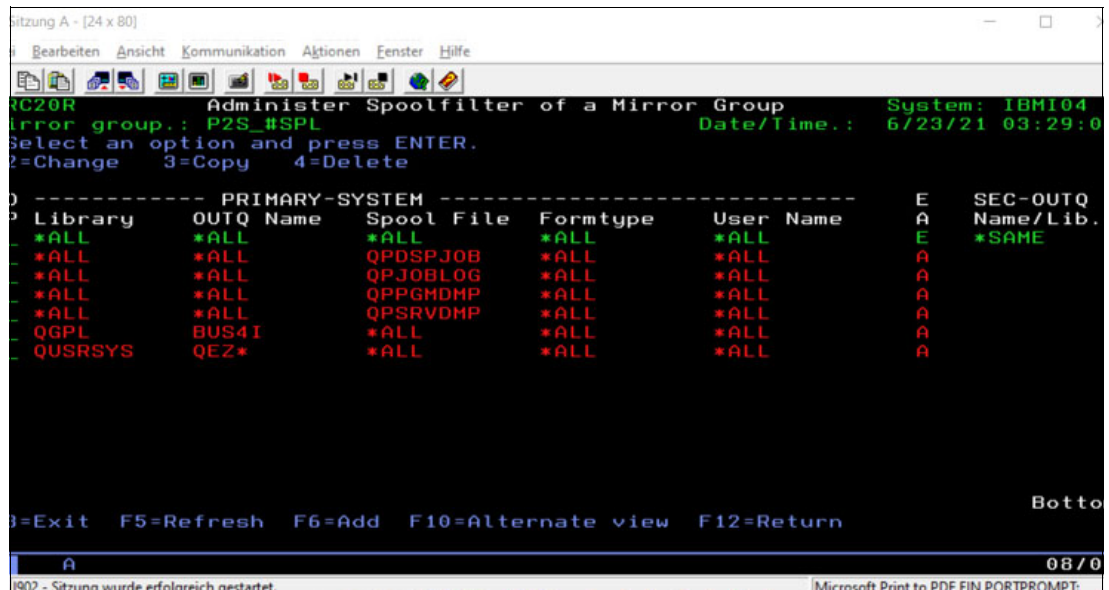


Figure 6-75 BUS4i spool files filter entries

Additionally, by using the mirror groups P2S_#SYS and P2S_#OBJ, you also can control which QGPL and QUSRSYS library objects and system values are included in or excluded from replication (see Figure 6-76 and Figure 6-77).

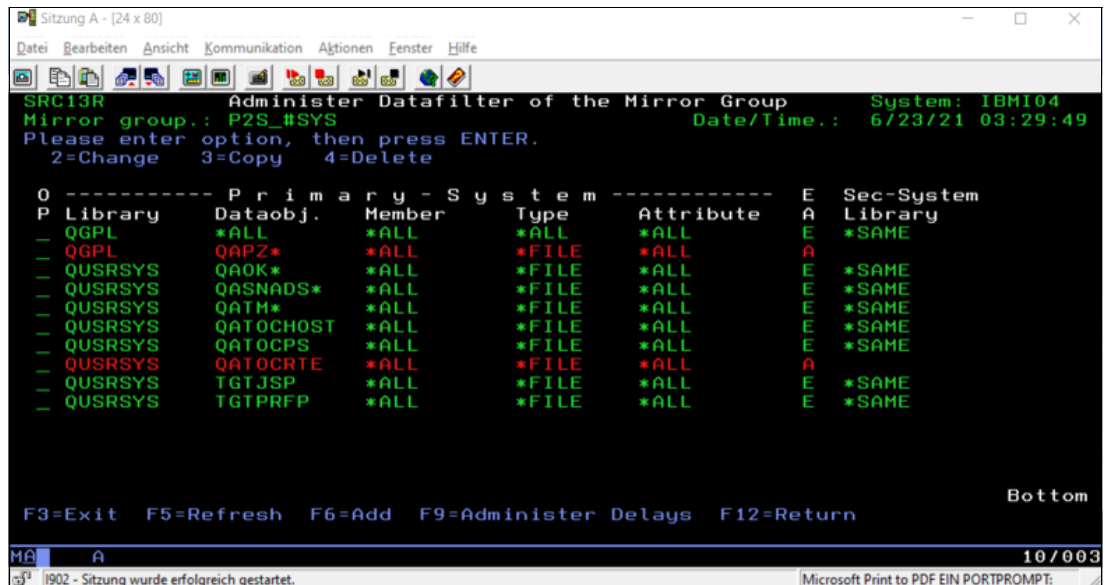


Figure 6-76 BUS4i QGPL and QUSRSYS filter entries

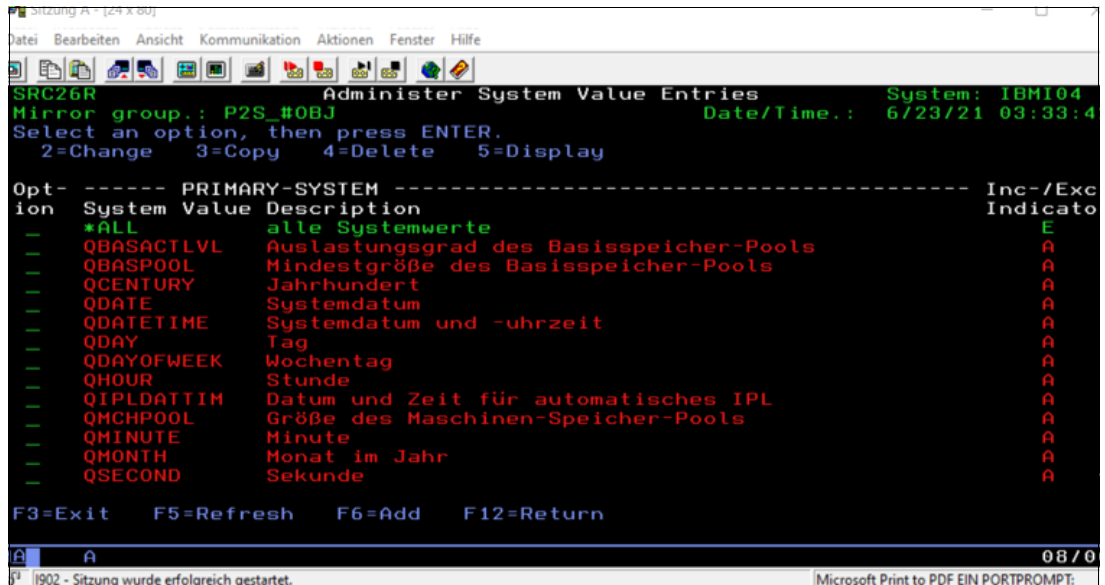


Figure 6-77 BUS4i system value entries

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 it is not interesting enough to include some of these libraries or objects to the replication (for example, in the case of new libraries with large volumes of redundant data for testing), these libraries or objects can be excluded by using an entry in red in this same mirror group (see Figure 6-78).

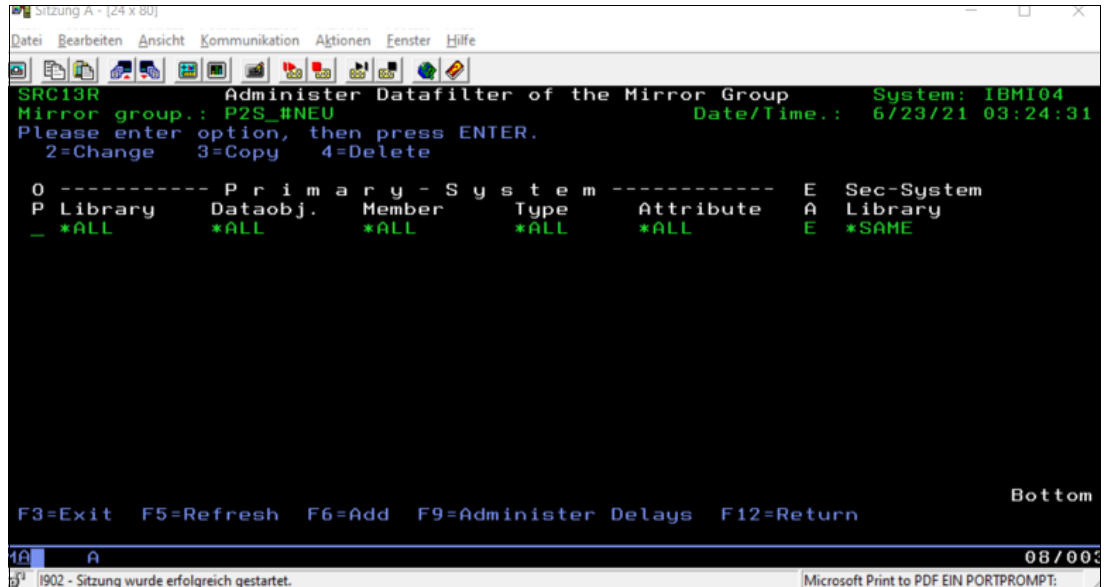


Figure 6-78 BUS4i data filter entries for new libraries

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.

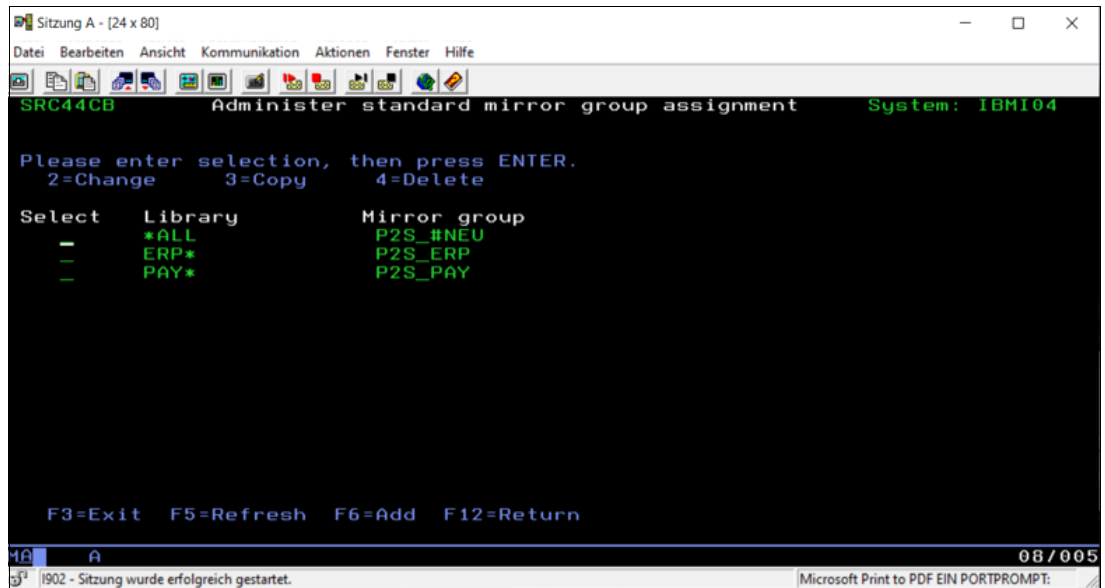


Figure 6-79 BUS4i standard mirror group assignment

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.

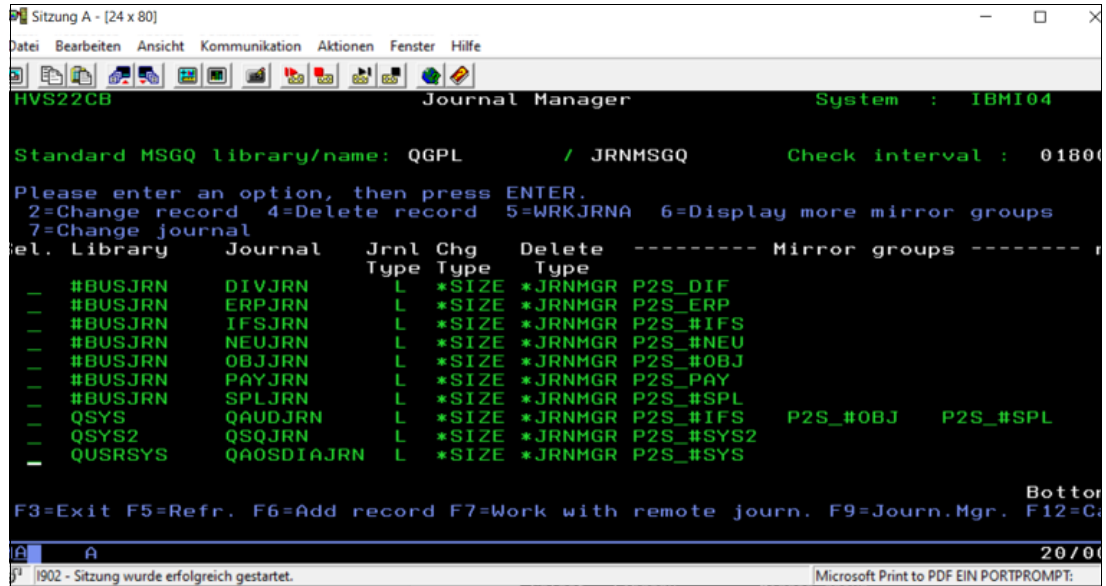


Figure 6-80 BUS4i Journal Manager window

6.4.5 Switching a Bus4i environment between two IBM i virtual machines in the IBM Power Systems Virtual Server Cloud

To document this use case, we created two IBM i 7.4 virtual machines (VMs) in the IBM Power Systems Virtual Server location London 04 called EWIBMI02. This location acts as the primary system and EWIBMI04 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 (see Figure 6-81).

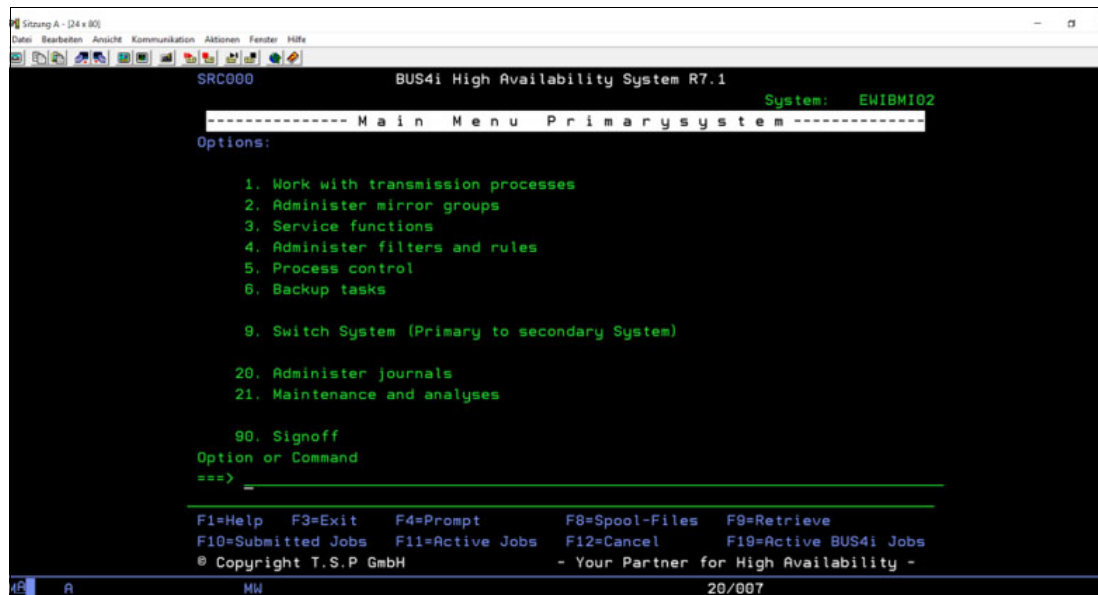


Figure 6-81 BUS4i main menu of the primary system

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.

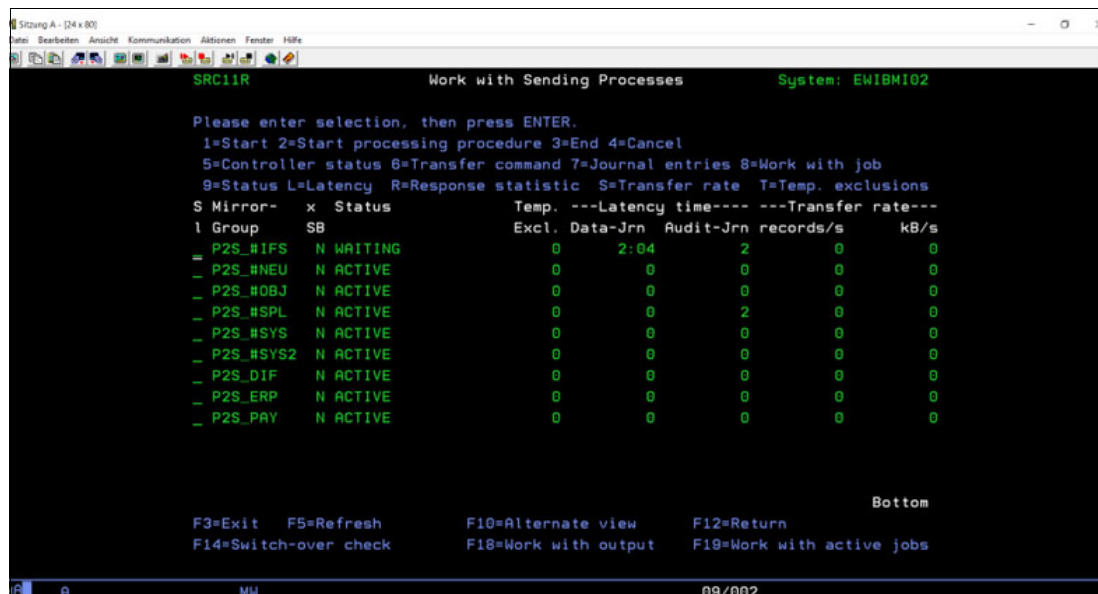


Figure 6-82 BUS4i Working with sending process

In the main menu of the primary system, Option 2. Administer Mirror Groups also can be used.

3. Check the transfer rates for each mirror group by using **Option S. Transfer Rate** (see Figure 6-83).

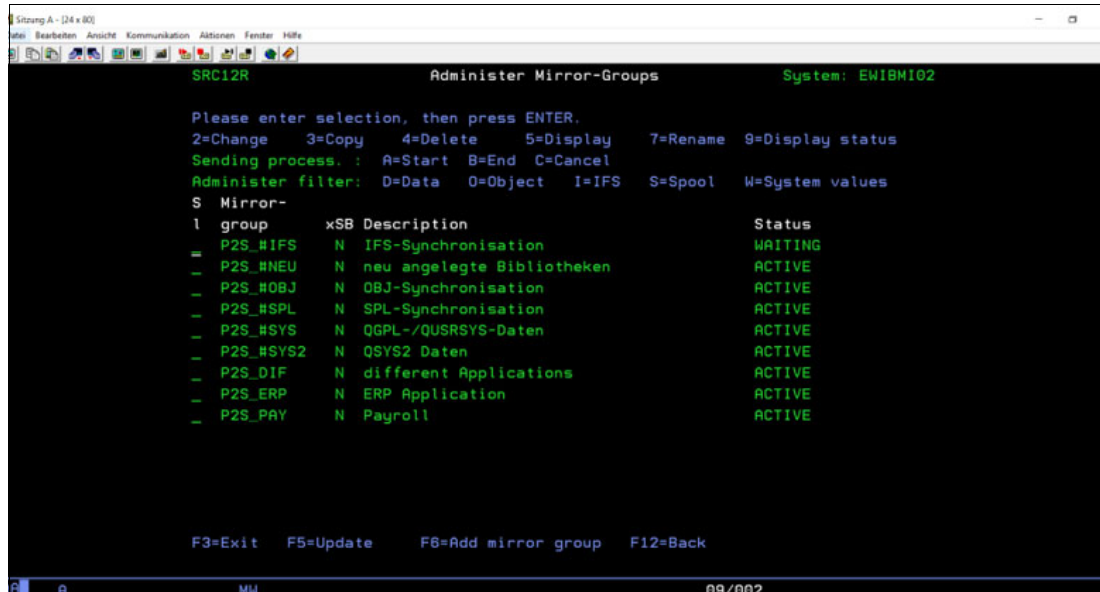


Figure 6-83 BUS4i Administering mirror groups

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.

All of the mirror groups stay in ACTIVE status and without any errors or warnings reported, as shown in Figure 6-84.

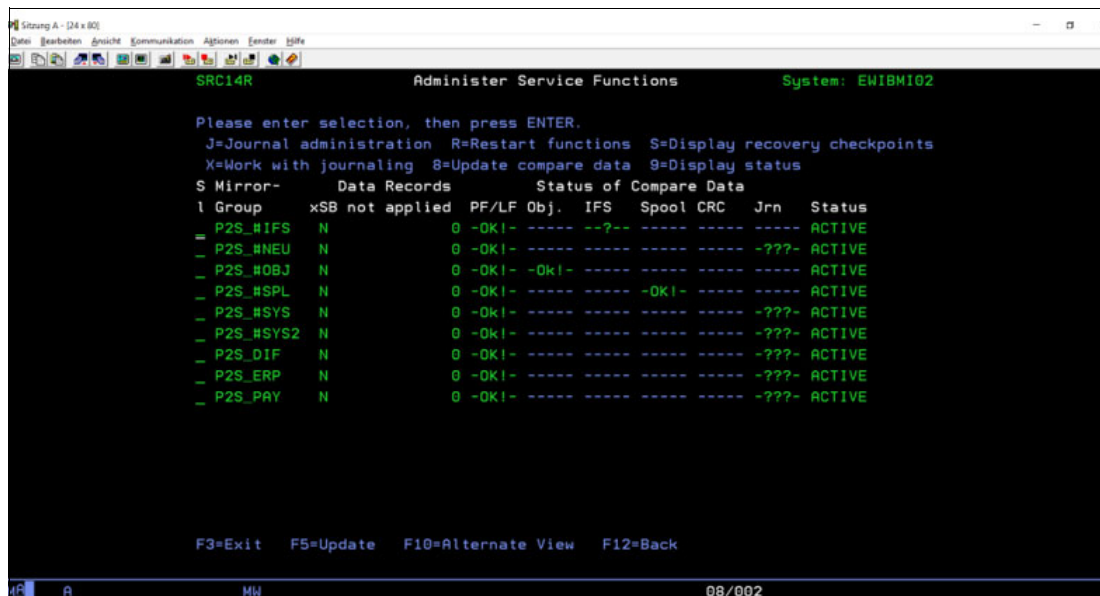


Figure 6-84 BUS4i administering service functions

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.

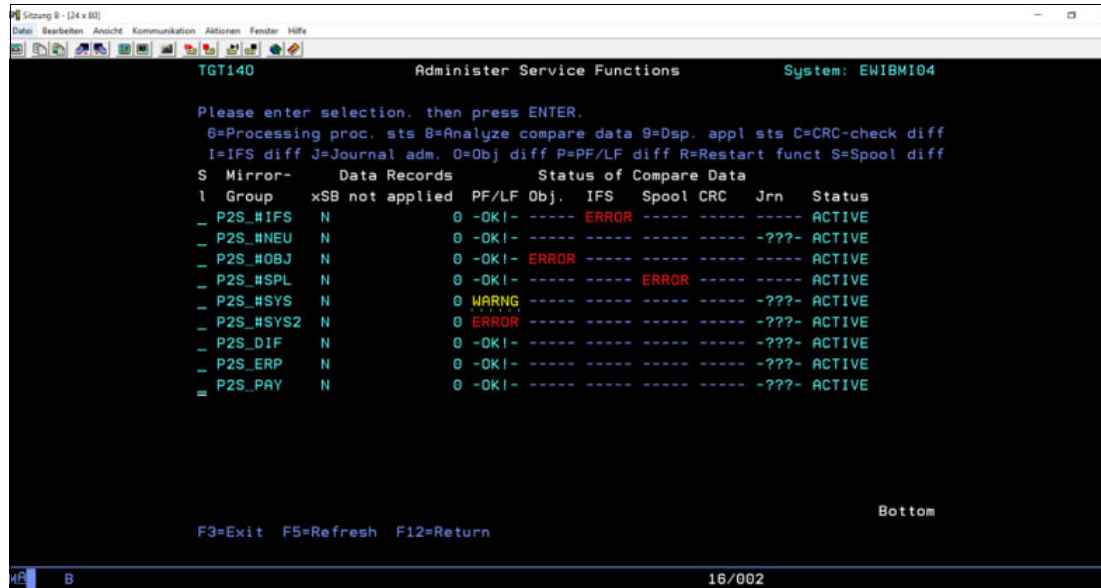


Figure 6-85 BUS4i administering service functions with examples of errors and warnings

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**. If some adjustments to the switchover process must be done, they can be performed by using this interface (see Figure 6-86).

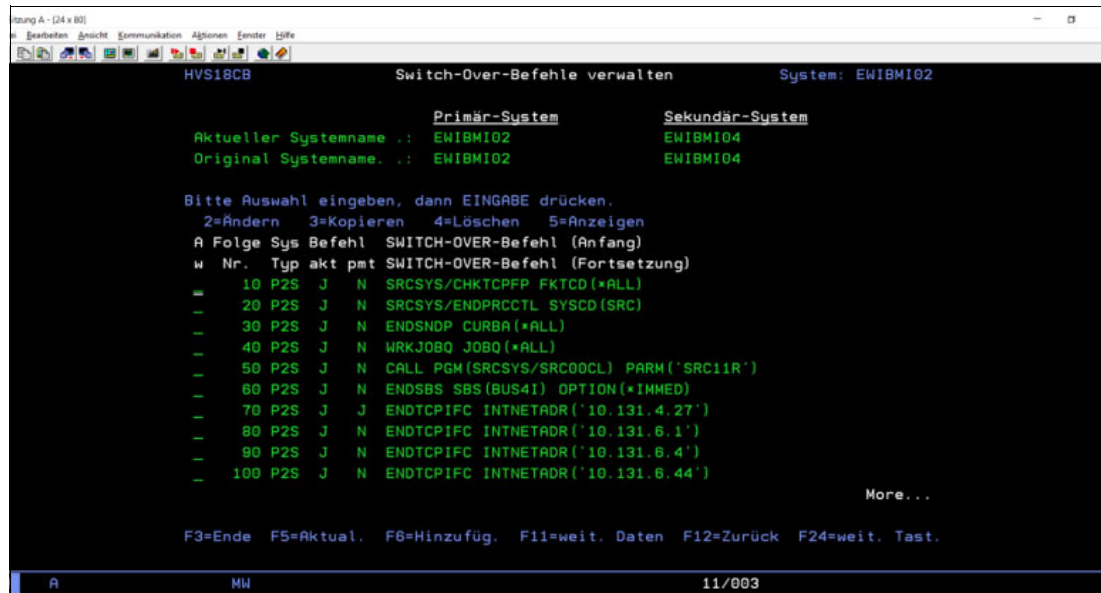


Figure 6-86 BUS4i switchover command

6. If everything is ready, start the planned switchover process from the main menu of the primary system (the system with 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.

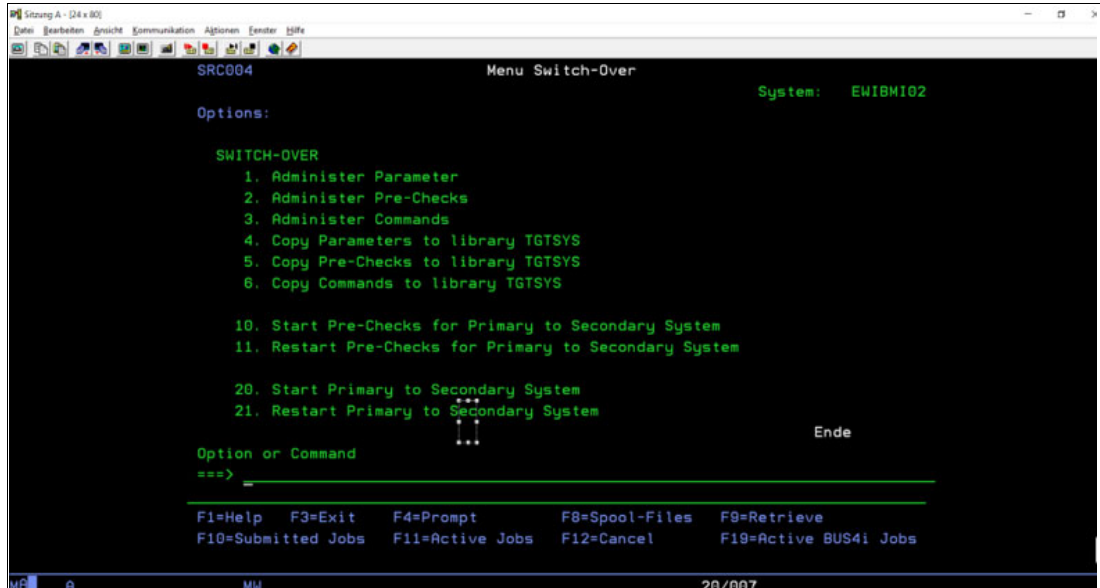


Figure 6-87 Switchover menu

A switchover message appears that indicates you are switching the PRIMARY system to the SECONDARY system. Answer with **Y=Yes** and then, press **Enter** to continue, as shown in Figure 6-88.

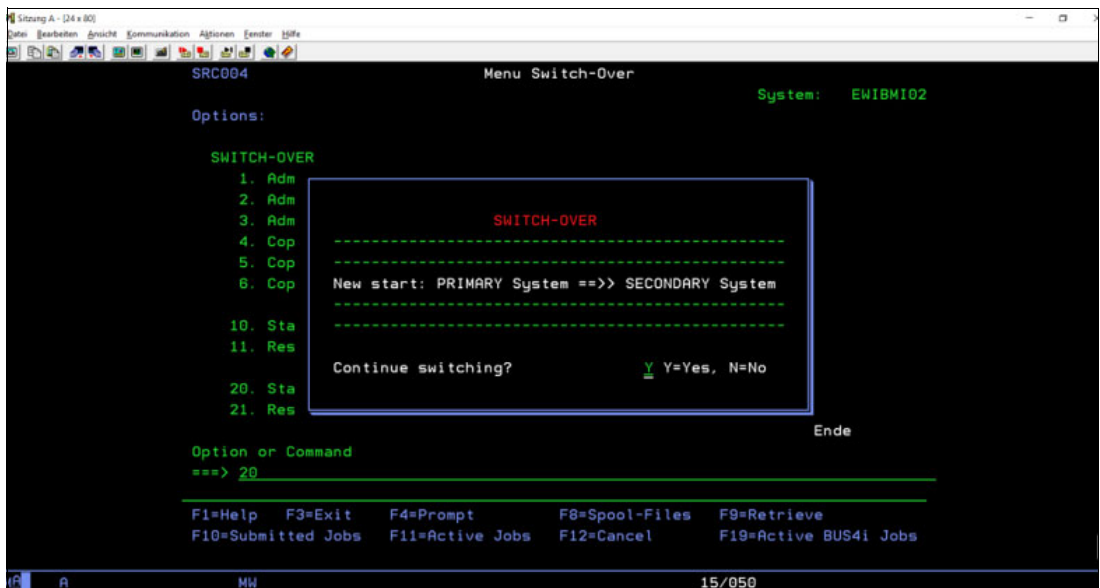


Figure 6-88 BUS4i switchover confirmation message in primary system

- 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.

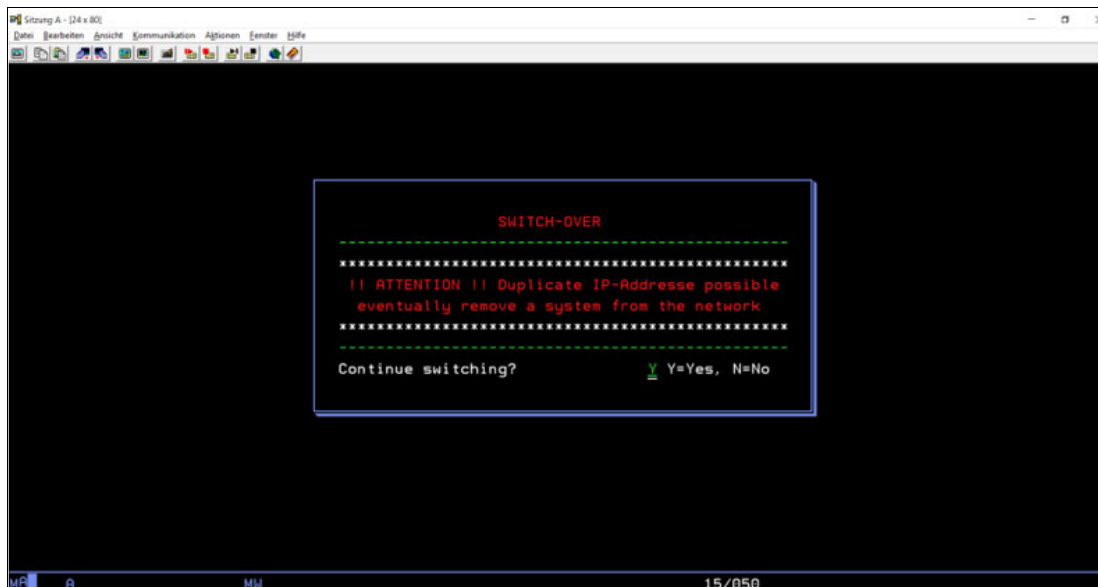


Figure 6-89 BUS4i switchover confirmation message in primary system

- 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.

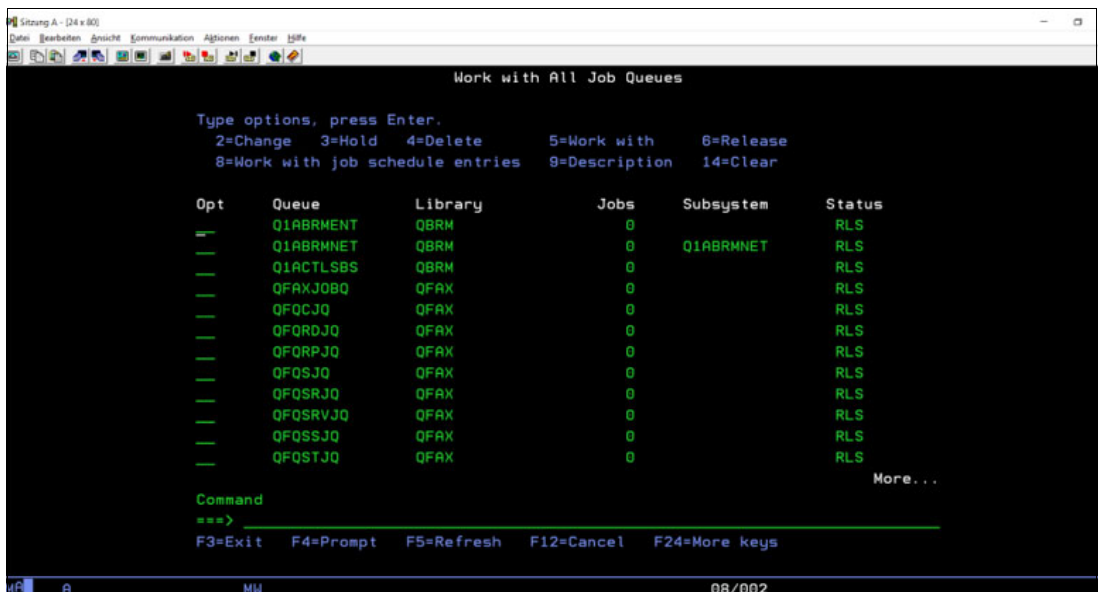


Figure 6-90 BUS4i switchover Work with All Job Queues window

10. The Work with Sending Processes window is displayed, as shown in Figure 6-91. Check that no latencies or temporary excludes appear.

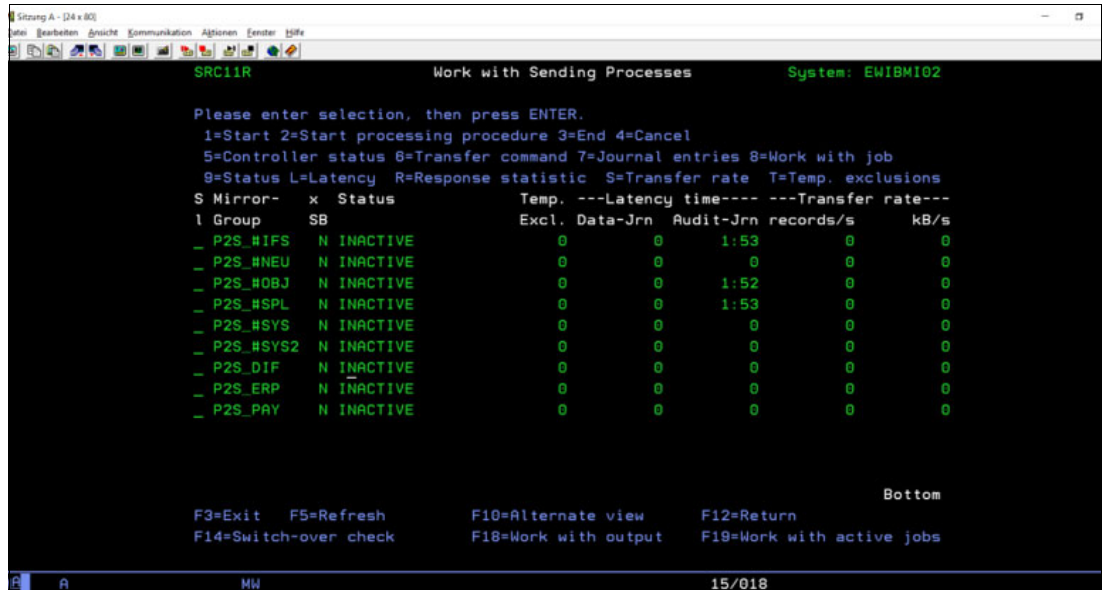


Figure 6-91 BUS4i switchover Work with Sending Processes window

After the switchover process is completed in the primary system, complete the following steps to perform the same procedure on the secondary 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 title in red).
2. Select **Option 9. Switch System (Secondary to Primary System)**.
3. In switchover menu, select **Option 20, Start Secondary to Primary System** and then, 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 237).

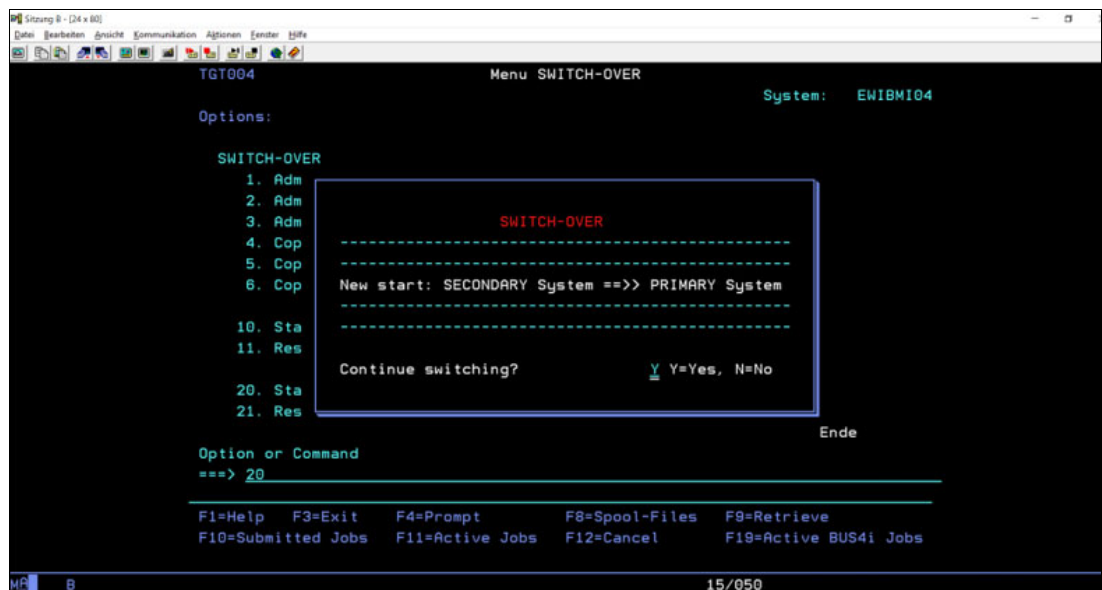


Figure 6-92 BUS4i switchover confirmation message in secondary system

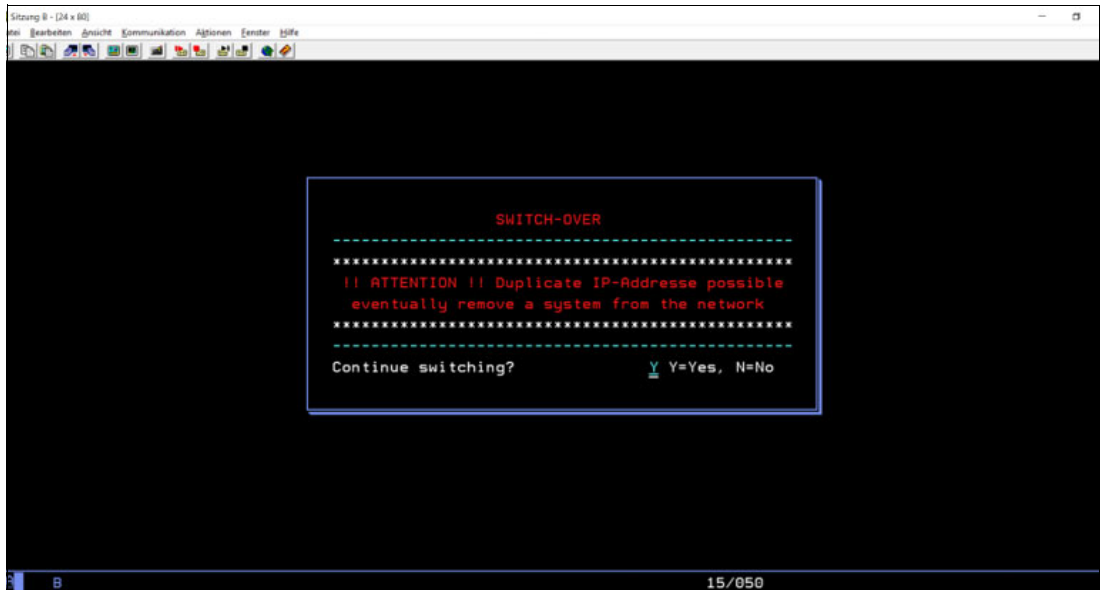


Figure 6-93 BUS4i switchover confirmation message in secondary system

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 EWIBMI04 shows the main menu for primary system (with title in white).

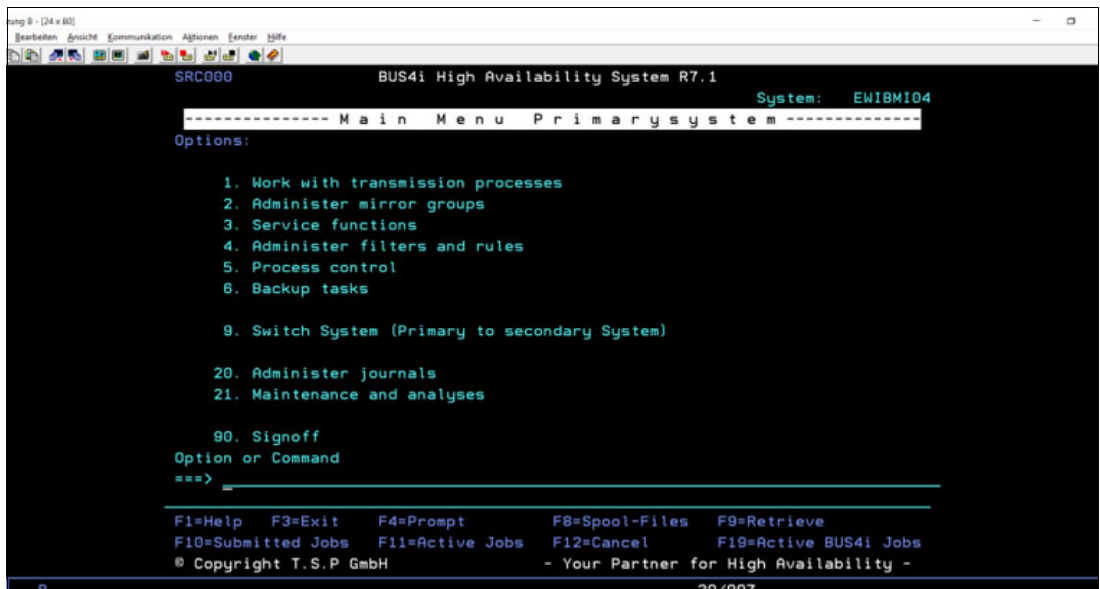


Figure 6-94 BUS4i new primary system after switchover

The preferred source system EWIBMI02 shows the main menu for secondary system (with title in red), as shown in Figure 6-95.

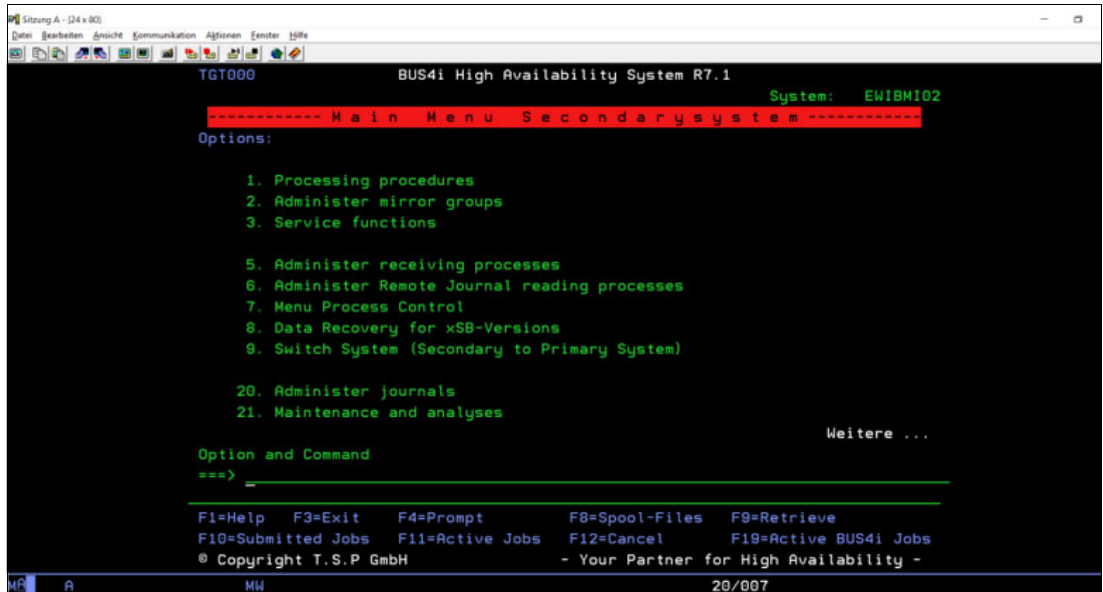


Figure 6-95 BUS4i new secondary system after switchover

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 239.

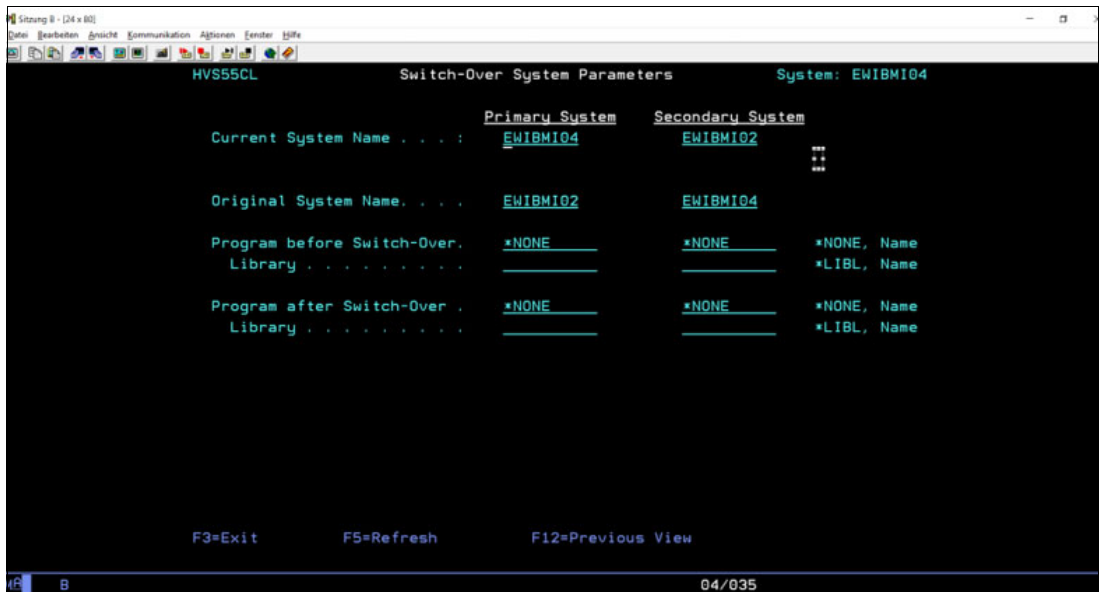


Figure 6-96 BUS4i switchover system parameters on new primary system after switchover

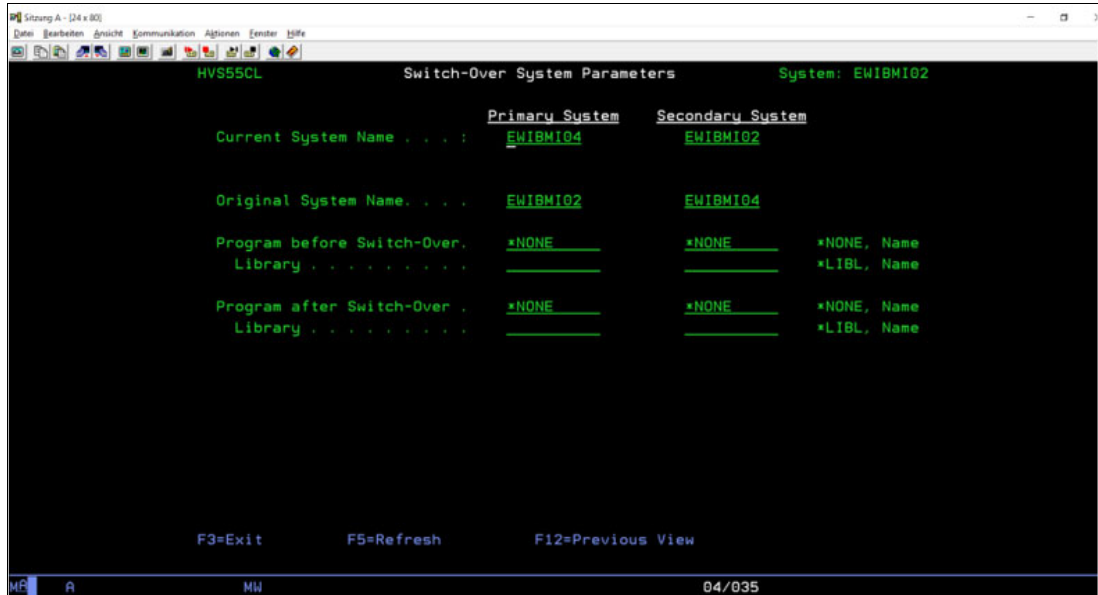


Figure 6-97 BUS4i switchover system parameters on new secondary system after switchover

- After verifying 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), start the replication process from the new primary system to the new secondary system by running the bus4i *yes command (see Figure 6-98).

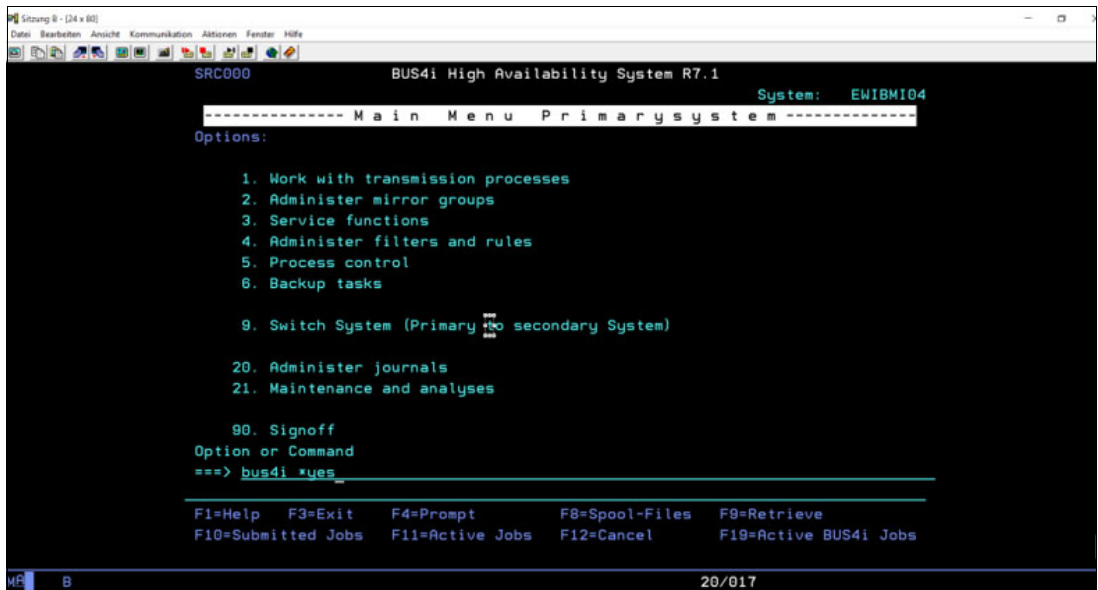


Figure 6-98 BUS4i starting the replication processes

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.

```

Work with Active Jobs                                     EWIBMI04
                                                         25/06/21 04:41:14 CDT
CPU %: 35.7   Elapsed time: 01:34:10   Active jobs: 213

Type options, press Enter.
 2=Change  3=Hold  4=End  5=Work with  6=Release  7=Display message
 8=Work with spooled files  13=Disconnect ...

Current
Opt Subsystem/Job  User      Type CPU % Function  Status
---
  -- BUS4I         QSYS      SBS    .0      DEQW     DEQW
  -- CHKSRCP      BUS4I     BCH    .0      DLY-7   DLYW
  -- JRNMGR       BUS4I     BCH    .0      PGM-HVS39CL DEQW
  -- S2P_#IFS      BUS4I     BCH    .0      DLY-11  DLYW
  -- S2P_#IFS0S   BUS4I     BCH    .0      PGM-SRC41CL DEQA
  -- S2P_#NEU     BUS4I     BCH    .0      IDX-HVS21L1 RUN

Bottom

Parameters or command
===>
F3=Exit   F5=Refresh   F7=Find   F10=Restart statistics
F11=Display elapsed data  F12=Cancel  F23=More options  F24=More keys
    
```

Figure 6-99 BUS4i subsystem in the new primary system

```

Work with Active Jobs                                     EWIBMI02
                                                         06/25/21 09:42:41 UTC
CPU %: 11.0   Elapsed time: 01:43:44   Active jobs: 245

Type options, press Enter.
 2=Change  3=Hold  4=End  5=Work with  6=Release  7=Display message
 8=Work with spooled files  13=Disconnect ...

Current
Opt Subsystem/Job  User      Type CPU % Function  Status
---
  -- BUS4I         QSYS      SBS    .0      DEQW     DEQW
  -- CHKTGTP      BUS4I     BCH    .0      DLY-7   DLYW
  -- JRNMGR       BUS4I     BCH    .0      PGM-HVS39CL DEQW
  -- RCV_#IFS      BUS4I     BCH    .0      PGM-TGT10CL SELW
  -- RCV_#NEU     BUS4I     BCH    .0      PGM-TGT10CL SELW
  -- RCV_#OBJ     BUS4I     BCH    .0      PGM-TGT10CL SELW
  -- RCV_#SPL     BUS4I     BCH    .0      PGM-TGT10CL SELW
  -- RCV_#SYS     BUS4I     BCH    .0      PGM-TGT10CL SELW
  -- RCV_#SYS2    BUS4I     BCH    .0      PGM-TGT10CL SELW

More...

Parameters or command
===>
F3=Exit   F5=Refresh   F7=Find   F10=Restart statistics
F11=Display elapsed data  F12=Cancel  F23=More options  F24=More keys
    
```

Figure 6-100 BUS4i subsystem in the secondary system

6. 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.

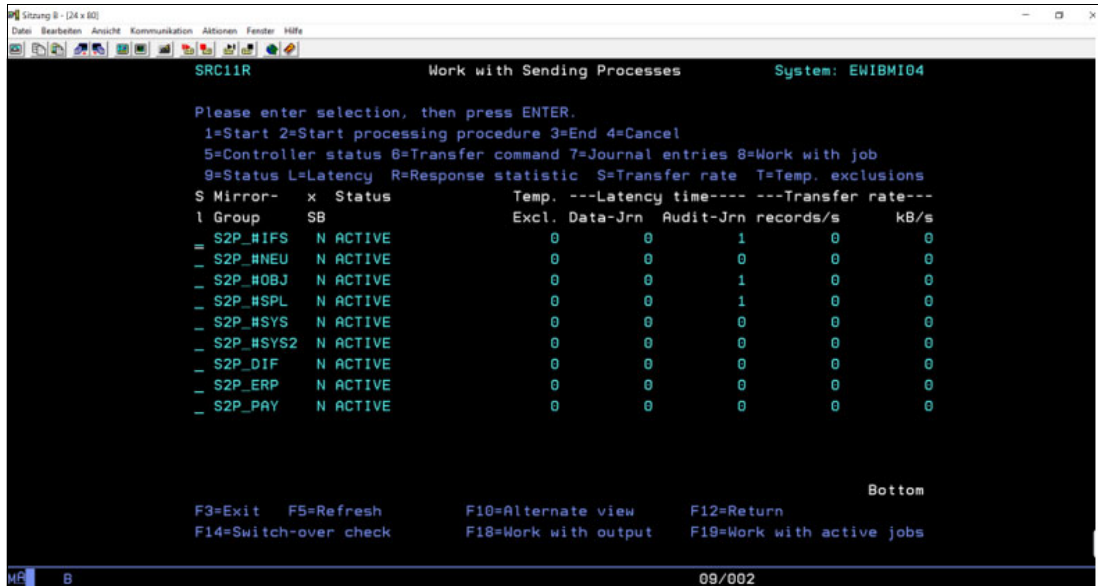


Figure 6-101 BUS4i Work with Sending Process

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.

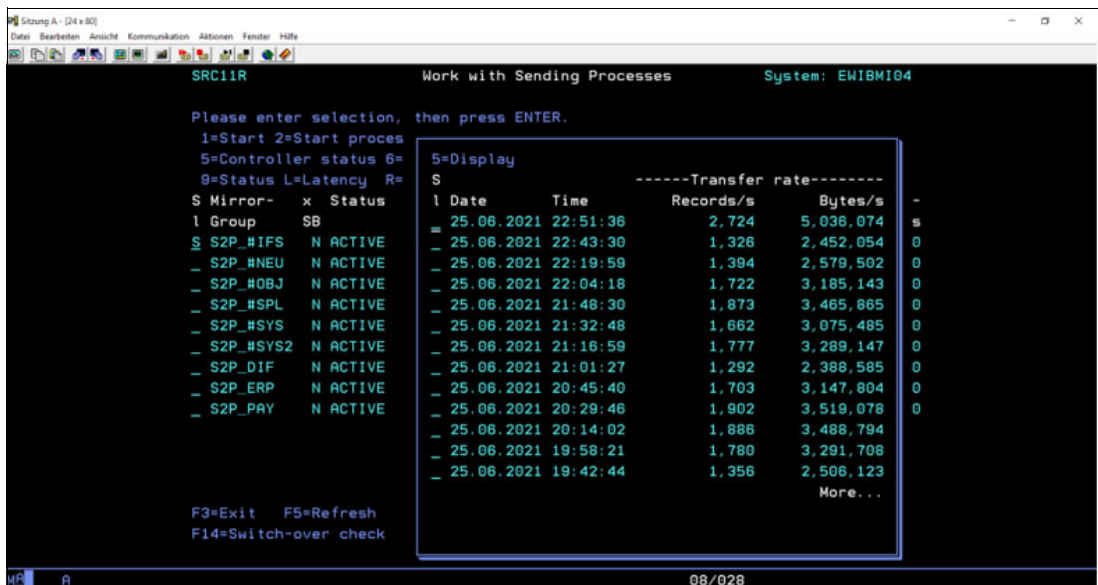


Figure 6-102 BUS4i mirror groups Transfer Rate

Important: If an unplanned switchover (failover) is required, the same procedure is applied, but only the steps that correspond to the secondary system are run.



Reference architectural decisions to migrate IBM i on-premises to IBM Power Systems Virtual Server

This chapter provides an example of migrating IBM i workloads from on-premises to off-premises IBM Power Systems Virtual Server on IBM Cloud.

We describe the considerations, architectural decisions, issues or problem statements, assumptions, motivations, alternatives, and implications that are part of such a process.

This chapter includes the following topics:

- ▶ 7.1, “Introduction” on page 244
- ▶ 7.2, “Use case overview” on page 244

7.1 Introduction

In this section, we describe important architectural considerations before migrating IBM i on-premises workloads to off-premises. We discuss one of the most common methods that is used to migrate IBM i environments to the cloud.

We also include a use case that is based on real customers experiences, which includes functional and non-functional requirements. Then, an overview of an architecture diagram is presented.

Note: For more information about customer's current IBM i systems, see this [Seismic web page](#) (log in required).

7.2 Use case overview

Suppose that a customer signed an agreement with 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 SA001 and SA004, both in Sao Paulo, Brazil.

Per the agreement signed with the customer, all IBM i VMs are to 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 SA004:

- ▶ 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 7-1).

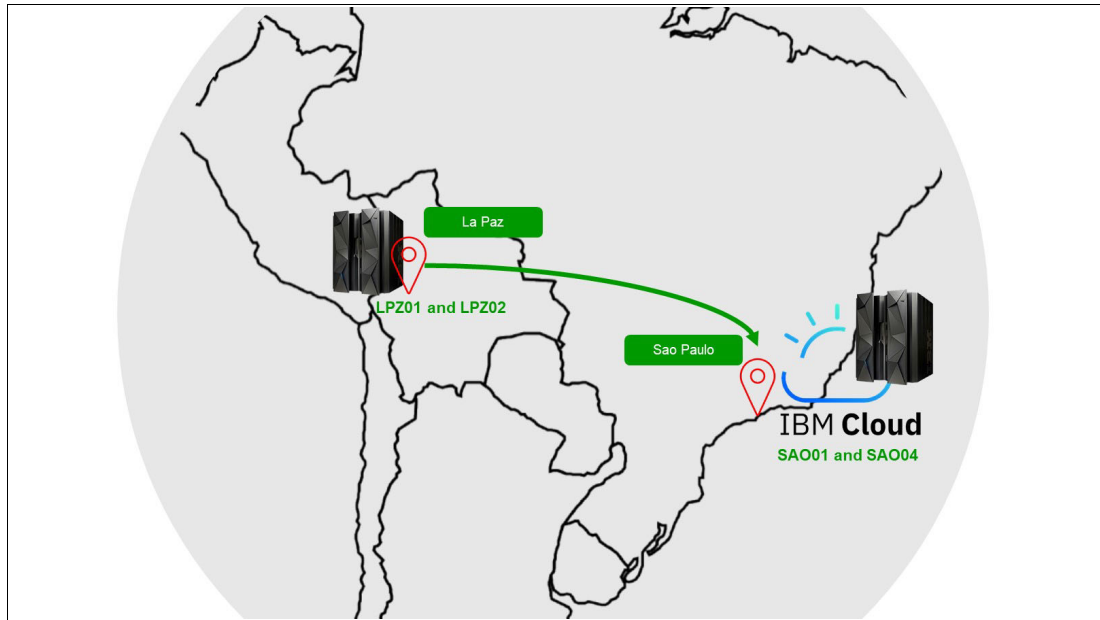


Figure 7-1 Migrating IBM i VMs from La Paz (on-premises) to Sao Paulo (off-premises)

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 cities are separated by 3000 kilometers (1684 miles). For more information about multi-zone regions, see this [IBM Cloud Docs web page](#).

7.2.1 Scope

In this use case, all 10 IBM i VMs must be 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 to be migrated to IBM Cloud by using IBM Cloud Object Storage backup. Then, they are to be restored to the target VM in IBM PowerVS on IBM Cloud. The IBM i production VMs that are on LPZ01 are to be 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 that done by the third-party vendor often are 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 customer the temporary technical information. such as: IP addresses, DNS names, and VM temporary name. This information 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 set up 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 PowerVS
- ▶ 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 us 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 PowerVS 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 Collocation), 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 SA001 site. A stand-alone firewall is deployed in SA004.

Functional requirements

The solution must satisfy functional and nonfunctional requirements in a way that best balances competing stakeholders' concerns and that considers any relevant constraints (see Table 7-1).

Table 7-1 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 establishing DR between SA001 and SA004.
Provide fault tolerant LAN infrastructure in IBM Cloud.	Provide network connectivity for application 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.
Provide WAN connectivity.	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 to be provided in primary IBM Cloud site (SA001); firewall cluster in High Availability, dual ports connectivity.
- ▶ Manageability access for customer's users to be provided by using jump servers.

Note: For more information about Worldwide Tools Solutions (WWTS), see this [IBM Support web page](#).

7.2.2 Architectural decisions

Architectural decisions help to communicate 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.
- ▶ Offer assistance to ensure that the arrangement fulfills functional and nonfunctional prerequisites. If it does not fulfill those requirements, offer assistance to make it unequivocal to the partners.
- ▶ 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 perspective:

- ▶ Table 7-2: Infrastructure
- ▶ Table 7-3 on page 250: Migration
- ▶ Table 7-4 on page 251: Servers
- ▶ Table 7-5 on page 252: Networking

Table 7-2 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 VS, which is moved from the customer's data centers to IBM Cloud in Sao Paulo in a dual site configuration.
Assumptions	Customer provides the WAN connectivity up to the network PoP Equinix ¹ 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 PowerVS, 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	<ul style="list-style-type: none"> ▶ 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.

¹For more information about Equinix, see [America Data Centers](#).

Important: 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 also can get a cross connection to Megaport and connect to IBM Power Systems Virtual Server directly.
- ▶ Before you begin, determine the location connection to IBM Cloud by verifying your colocation provider's or service provider's capabilities to reach the Meet-Me Room and cross-connect into IBM Cloud. For more information, see [this web page](#).

On SA001, the location type and Meet-Me Room (MMR) Operator, is DC and Digital Realty (Ascenty), and SA004 DC(AZ1) and ODATA.

In this example, to lift and shift 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
- NOVACORE 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 7-3 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 back up site (use of DNS for servers translation; secondary site has a different TCP IP address).

Architectural decision	Dual site infrastructure is required for HA purposes.
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 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	<ul style="list-style-type: none"> ▶ Provide WAN network connectivity to secondary site for customer's users and application connectivity. ▶ Duplicate part of the primary infrastructure in backup site.

Table 7-4 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.
Assumptions	The use of IBM Cloud Object Storage for the migration is one of the available methods for moving 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 Master Data Management (MDS) 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	<ul style="list-style-type: none"> ▶ Master Data Management is excluded because of the delay for taking backups data on 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	<ul style="list-style-type: none"> ▶ 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 7-5 Delta data synchronization

Architectural decision	A third-party vendor replication solution is used.
Problem statement	Customer has IBM i VMs with storage 10 TB - 70 TB. Some of the IBM i VMs include a journal of 1 TB that generates 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 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 PowerVS. 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 7-6 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 primary link outage.
Assumptions	WAN part is 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 SAO01 and a redundant connectivity in SAO04.
Decision	Provide redundant connectivity in SAO01 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 (see Table 7-6 on page 252). 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 7-2 on page 249 - Table 7-5 on page 252 list decisions only as an example. Real world decisions can vary according to the customer, scenario, third-party vendor applications, in-house applications, region, networking, and so on.

A certified IBM i architect can help make decisions about your scenario.

7.2.3 Architecture overview

Figure 7-2 shows an overview for IBM Cloud in Sao Paulo. The division of the different zones is delineated.

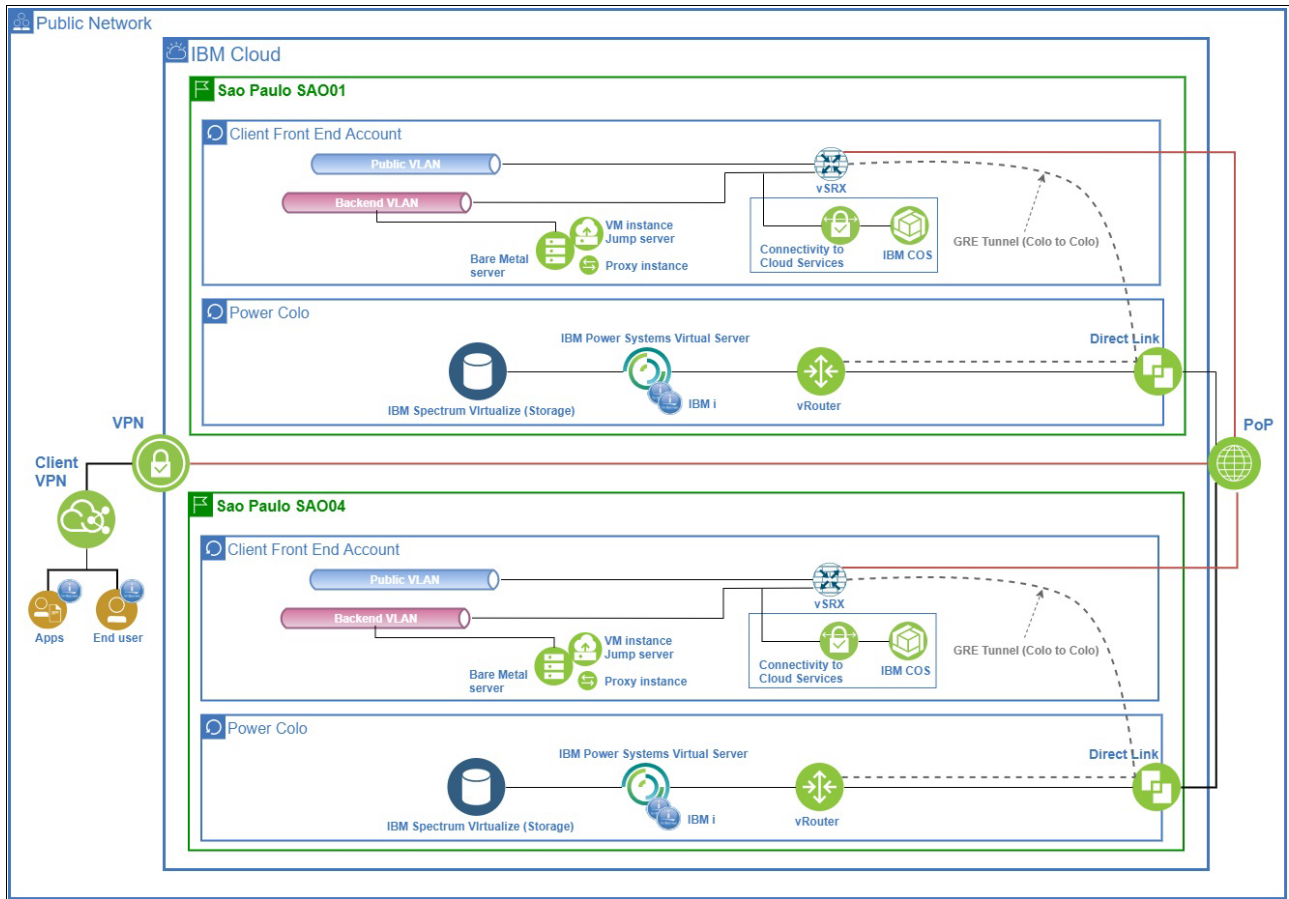


Figure 7-2 IBM i on IBM PowerVS architecture

These zones are separated from each other by physical switches, which is an activity that is controlled by firewalls. A dual-site infrastructure is used. Consider the following points:

- ▶ SA001 is where the IBM i production runs and it comprised in 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 PowerVS is installed.
- SA004 is used for DR purposes if an outage occurs at SA001.

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 regarding that the backups are saved in IBM Cloud Object Storage if backups are less than 2 TB to be a viable solution the use of BRMS/ICC to save to IBM Cloud Object Storage. For more information, see this [IBM IPP Wiki page](#).
- ▶ 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 this [IBM Training web page](#).



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 is changing the computing landscape, cloud technologies allowed customers 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 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 Licensed Program Products and operating systems” on page 258.
- ▶ “Passport Advantages software” on page 265.
- ▶ “Movable IBM i” on page 266.

IBM i Licensed Program Products and operating systems

IBM Power Systems Virtual Server presented three IBM Power Systems, such as: E880 (9119-MHE), E980 (9080-M9S), and S922 (9008-22A). Therefore, the following software is used:

- ▶ IBM i OS and base Licensed Program Products, P10
- ▶ IBM i OS and base Licensed Program Products, P30

Note: IBM i OS and Base Licensed Program Products are always ordered together; therefore, pricing of both products is always included.

For more information about CPW Benchmark performance regarding the IBM Power Systems available on IBM Cloud, see [IBM Power Systems Performance Report](#).

On IBM Power Systems Virtual Server, IBM i images are on different releases to deploy a new virtual machine, each IBM i release has a *golden image* that in addition to the IBM i operating system, other Licensed Program Products and IBM i O.S features are included in the IBM Power Systems Virtual Server on IBM Cloud.

Table A-1 lists the software that is included in the IBM i and Licensed Program Product package that is available in IBM Power Systems Virtual Server.

Table A-1 Licensed Program Product package that is available in PowerVS

Product	Option	Description
5770-SS1		IBM i processor and unlimited users.
5770-SS1	23	OptiConnect.
5770-SS1	44	Encrypted Backup Enablement.
5770-SS1	45	Encrypted ASP Enablement.
5770-SS1	18	Media and Storage Extensions.
5770-SS1	26	IBM DB2 Symmetric Multiprocessing.
5770-SS1	27	DB2 Multi-system.
5770-SS1	38	PSF for IBM i Any Speed Printer Support.
5770-SS1	41	HA Switchable Resources.
5770-SS1	42	HA Journal Performance.
5770-SS1	All other IBM i options are included with IBM i at no extra charge.	For example: <ul style="list-style-type: none"> ▶ Option 8, AFP Compatibility Fonts ▶ Option 12, Host Servers ▶ Option 13, System Openness Includes
5770-DG1	Base	HTTP Server for i.
5770-JV1	Base	Developer Kit for Java.
5770-NAE	Base	Network Authentication Enablement for i.
5733-SC1	Base	Portable Utilities for i.
5770-TC1	Base	TCP/IP.
5770-TS1	Base	Transform Services for i.

Product	Option	Description
5770-UME	Base	Universal Manageability Enablement for i.
5770-XE1		IBM i Access for Windows.
5733-ACS		IBM i Access Client Solutions.
5733-ARE	Base	IBM Administration Runtime Expert.
5798-FAX	Base	IBM Facsimile Support for i.
5770-SM1	Base	IBM System Manager for i.
5770-DFH	Base	IBM CICS® Transaction Server for i.
5770-MG1	Base	IBM Managed System Services for i.
5770-AF1		Advanced Function Printing for i.
5761-AMT		IBM Rational Application Management Toolset.
5761-AP1	Base	Advanced DBCS Printer Support.
5733-B45		AFP Font Collection for i.
5770-BR1	Base	Backup, Recovery and Media Services.
5761-DB1		System/38 Utilities.
5761-CM1	Base	Communications Utilities.
5648-E77		InfoPrint Fonts.
5769-FN1		AFP DBCS Fonts.
5769-FNT		AFP Fonts.
5722-IP1	Base	InfoPrint Server for i.
5770-JS1	Base	Advanced Job Scheduler for i.
5770-PT1	Base	Performance Tools.
5770-QU1	Base	Query for i.
5770-ST1	Base	Db2 Query Manager and SQL Dev Kit for i.
5733-XT2		XML Toolkit.
5770-XH2	Base	IBM i Access for Web.
5770-XW1	Base	IBM i Access Family.
5761-DS2		Business Graphics Utility.

Note: For all Licensed Program Products (LPP) that are listed in Table A-1, all features of the LPPs are included. Therefore, each LPP in the package contains optional features and not just the Base option.

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 can be included 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-1. For more information, see [IBM Cloud Storage Solutions for i](#).




Figure A-1 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; that is, 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 (see Figure A-2), a customer can use IBM i operating system capabilities to set up a DR backup site in PowerVS. 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).

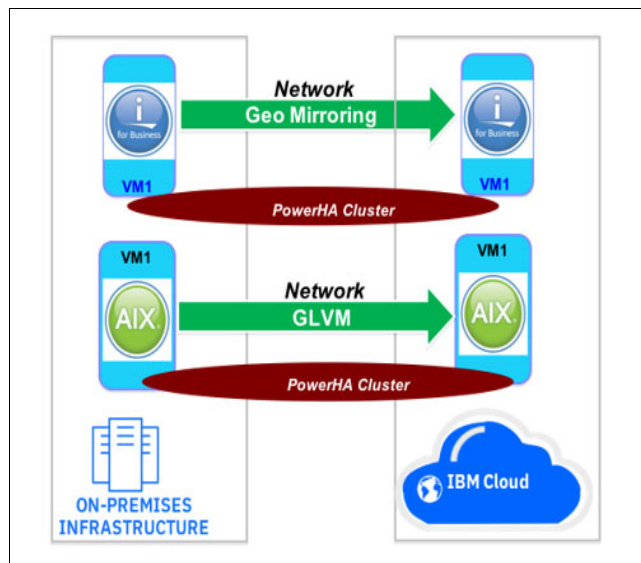


Figure A-2 Private-to-public cloud DR solution

In scenario 2 (see Figure A-3), a customer can deploy across IBM Cloud regions. For example, PowerHA Enterprise Edition can be used to enable DR across Washington and Dallas.

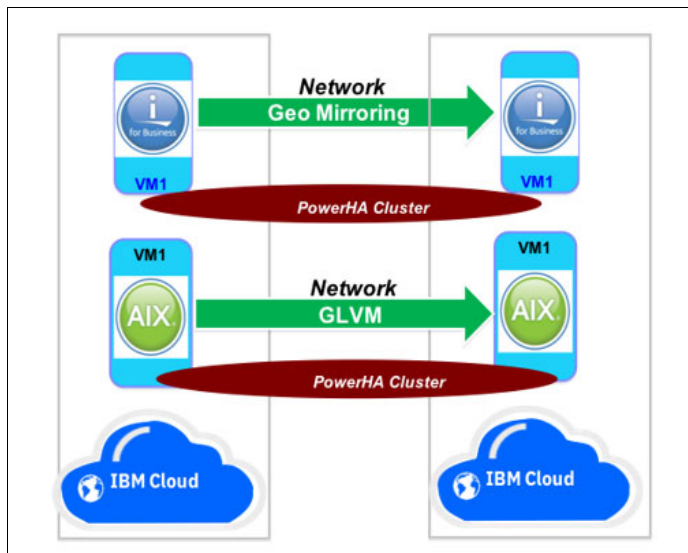


Figure A-3 Public-to-public cloud DR Solution

Important: IBM PowerHA for i (5770-HAS) Enterprise Edition is available in only the IBM Power Systems Virtual Server on IBM Cloud.

Table A-2 lists the IBM PowerHA for i Enterprise Edition license pricing.

Table A-2 Prices for IBM PowerHA for i (5770-HAS)

PowerHA for i Enterprise Edition	Monthly license list price	Perpetual license list price
P10 – s922 (Small)	\$155 per core, per month	\$3740 per core
P30 – e880/E980 (Medium)	\$220 per core, per month	\$5250 per core

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 [DB2 Web Query for i](#).

Figure A-4 shows the pricing for software license IBM DB2 Web Query for i while creating an IBM virtual machine in IBM Power Systems Virtual Server on IBM Cloud.

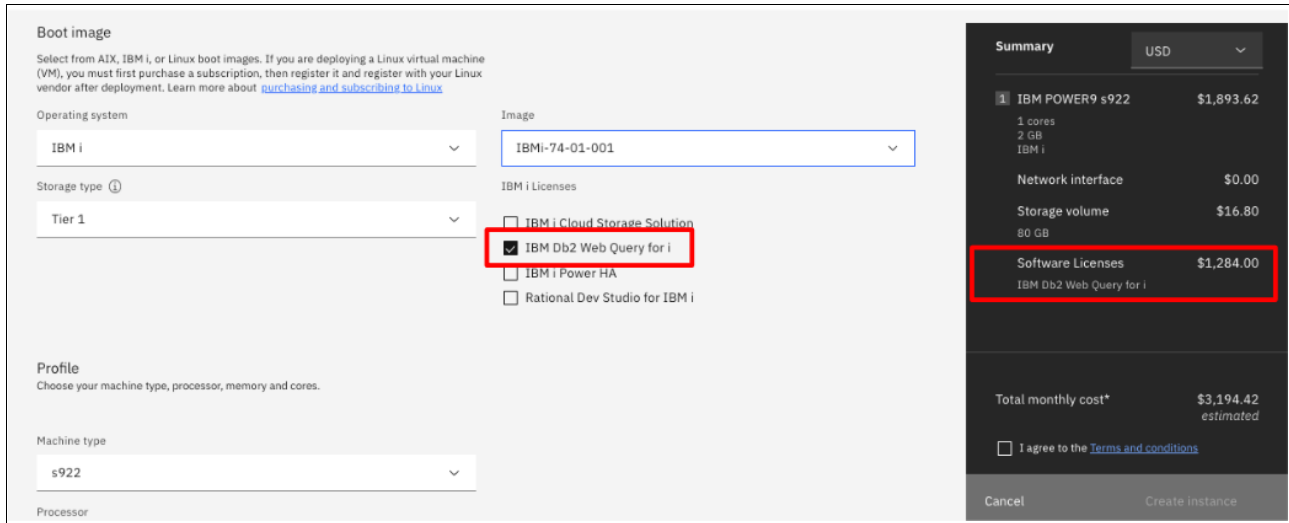


Figure A-4 Example: DB2 Web Query for i in the PowerVS offering

The price is the same for all system sizes. Price is per core per month and includes (see Table A-3):

- ▶ Standard Edition function
- ▶ Unlimited Developer Users, Developer Workbench Users, and Runtime Enablement Groups

Table A-3 Prices for DB2 Web Query for i Standard Edition (5733-WQS)

DB2 Web Query for i Standard Edition	PowerVS Monthly price
Any size	\$1063 per core, per month

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 the IBM i. This package consolidates all of the key IBM i development tools into one offering.

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 any number of 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 directly or indirectly (for example, by way of a multiplexing program, device, or application server) through any means. 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.

Rational Development Studio for i (5770-WDS), the per-user price includes all of 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-4.

Table A-4 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	<ul style="list-style-type: none"> ▶ \$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 then, click **Save edits and order** to complete the instance modification process and accept the price.

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 this [IBM Passport Advantage web page](#).

Movable IBM i

When matched with IBM Power Systems Virtual Server, IBM i systems perform at the most elevated of speed and agility while providing infrastructure flexibility and cutting edge capabilities with the opportunity to lower operations costs. These key benefits (reliability, availability, and flexibility) allow many customers to gain huge advantages in running business-critical applications, incorporating cognitive technologies, such as AI, IoT, and Machine Learning.

Clients cite the low total cost of ownership (TCO) as a significant factor in their choice to run their businesses on IBM i and Power Systems. IBM creates, tests, and preinstalls the core middleware components of IBM i up front. On other platforms, operating systems databases and middleware integration is done by customers in their own data center.

Moving IBM i workloads to Power Systems Virtual Server can help to improve TCO calculations.

Understanding Movable IBM i

IBM i clients around the world that must migrate IBM i VMs to IBM Power Systems Virtual Server can qualify to get credit for your IBM Power Systems on-premises investments in IBM i entitlements. The advantage is that this credit reduces pricing for an IBM i subscription when enrolling in capacity on the IBM Power Systems Virtual Server on IBM Cloud.

IBM i customers can take advantage of digital transformation by adopting hybrid cloud environments that use IBM Power Systems Virtual Server. Lower subscription fees and increased return on investment (ROI) are realized by Movable IBM.

The use of cloud for data center divestiture of production, development, test, and DR environments are some of the more common use cases for which customers use IBM Power Systems Virtual Server.

From the customer's perspective, Movable IBM i protects their investment in on-premises entitlements on IBM Power Systems while migrating to a hybrid cloud environment with IBM Power Systems Virtual Server. This protection helps to reduce costs instead of adding costs by purchasing new licenses.

Movable IBM i process

The Movable IBM i for IBM Power Systems Virtual Server offering requires an agreement to be signed with IBM. IBM checks your current on-premises IBM Power Systems with IBM i workloads entitlements and documents the quantity to be used in IBM Power Systems Virtual Server.

Note: For more information about the license agreement, see this [Seismic web page](#) (log in required).

Figure A-5 shows the Movable IBM i process from on-premises to off-premises.

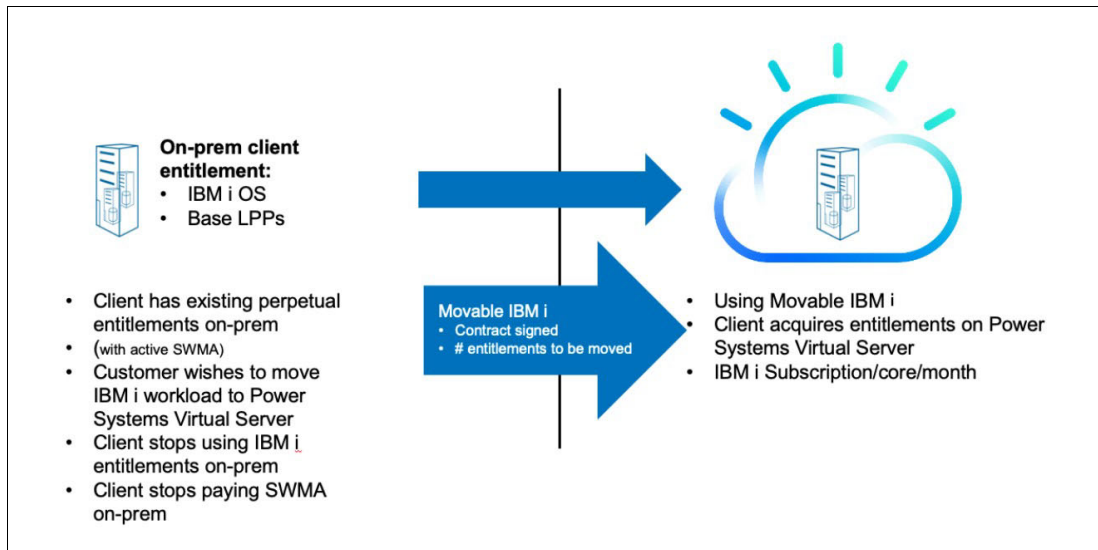


Figure A-5 Movable IBM i process

For more information about Movable IBM i, see this [Seismic web page](#) (login required).

Important: Movable IBM i pricing is subject to change without notice.



B

Frequently asked questions about IBM i on IBM Power Systems Virtual Server

This appendix provides answers to the most commonly asked questions about IBM i on IBM Power Systems Virtual Server on IBM Cloud.

This appendix includes the following topics:

- ▶ “FAQs about IBM i on IBM Power Systems Virtual Server” on page 270.
- ▶ “Miscellaneous FAQs about IBM i on PowerVS” on page 270.

FAQs about IBM i on IBM Power Systems Virtual Server

In this section, we answer the most commonly asked questions about IBM i on PowerVS.

Miscellaneous FAQs about IBM i on PowerVS

1. Estimating the amount of IBM Cloud Object Storage that is needed for IBM BRMS backups on IBM i PowerVS, what compression ratio can be expected?

There is almost no compression by default. It can be enabled 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. Does IBM i 7.1 support on top of S922?

IBM 7.1 is supported on-premises or off-premises (PowerVS) on the S922 or E980 under VIOS. Recall that IBM i on the S922S922 is limited to four core maximum size partitions.

3. In which data centers can it be IBM i release 7.1?

Currently, DAL12 and FRA04. FRA05 is next, followed by TOK04. But, there exists a chargeable Service Extension charge on IBM i 7.2 as is today on 7.1.

IBM i 7.1 is not available everywhere because the firmware 950 is the minimum that is required on S922 and E980. For more information, see [this web page](#).

4. Is IBM i 7.1 eligible for Power Virtual Server's Movable IBM i (also referred to as MOL or IBM i mobility licensing)?

Yes. For Movable IBM i, refer to the Seismic materials, including the validation and approval process that are available at [this web page](#).

5. What is the price for IBM i 7.1 in Power Virtual Server?

For IBM i 7.1, a Service Extension price is automatically added to the IBM i price. Use the latest version of [the pricing estimator](#).

Is the price for Service Extension built into the IBM i price in the Price Estimator V9i?

On the **Multiple LPAR Price Estimate** tab of Price Estimator V9i, select **Yes** for IBM i Service Extension for IBM i 7.1. (Refer to the figures in the instructions on chart 5.) Also, select **Yes** for IBM i Service Extension when estimating the pricing for IBM i 7.2 for customers who run IBM i 7.2 after 1 May 2021 when Service Extension starts for 7.2.

6. What is the duration of Service Extension for IBM i 7.1 in Power Virtual Server?

Because Power Virtual Server is running on IBM POWER9 servers, IBM i 7.1 Service Extension is available through 30 April 2024.

7. What type of support coverage is provided with IBM i 7.1 Service Extension? Which products are covered?

- From as of this writing to 30 April 2021, some of the IBM i 7.1 products and features are still available under new defect, plus usage and known defect support coverage.
- Starting 1 May 2021, IBM i 7.1 Service Extension is usage and known-defect support.

Note: For more information, see [Service Extension for IBM i 7.1 and IBM i 7.2](#).

8. 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. We provide stock images to help customers move their image easily to Power Virtual Server.

9. In the IBM Power Systems Virtual Server Price Estimator, there is an option for IBM i Service Extension. Does that mean that instead of going through a separate contract with IBM to get a Service Extension, we can streamline it all under Power Virtual Server pricing?

Yes. In Power Virtual Server, there is no need for a separate contract for Service Extension. Service Extension is automatically added for both IBM i 7.1 and IBM i 7.2 (after the Service Extension starts for 7.2 on 2021 May 1).

10. Does it need to install cloud-init on the IBM i LPAR -before creating the SAVSYS backup? Or, is this only necessary if we setup and use IBM Cloud Storage Solution?

Refer to slide 21 of [IBM i Considerations](#).

[IBM i Migration with IBM Power Systems Virtual Server](#) also has information (for example, see page 5), and [IBM i Backups with IBM Power Systems Virtual Server](#).

11. Can a QSECOFR user access PowerVS by using SSH?

The QSECOFR user cannot access PowerVS by using SSH. Therefore, a new user is needed. However, PowerVS is available if you modify the `sshd_config` with `PermitRootLogin yes` (this practice is not recommended; therefore, you create another user with similar authority than QSECOFR).

12. Does PowerVS support virtual serial numbers?

There is no support for virtual serial numbers in PowerVS. However, that feature is on the roadmap.

13. What version of PowerVC is running on IBM PowerVS?

PowerVS is based on PowerVC that uses 1.4.4.2 version and manages the servers, such as S922 and e980 by Novalink.

14. 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 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 SWMA lapsed on the donor machine, and the customer needs support on-premises, they must reactivate SWMA. An after license fee applies.

15. What is VTL and how is it applicable to PowerVS?

DSI Restore is deduplicating backup and recovery virtual software. Based on DSI's world-class enterprise Virtual Tape Library (VTL), you can restore runs on your virtual infrastructure and scale from TBs to PBs of backup data. Restore is a fit for small businesses and global enterprises because it is licensed per TB. Bandwidth efficient, encrypted data replication to a second site or cloud instance meets your compliance and disaster recovery requirement.

16. What is driving the business need for a software-based VTL?

IBM i customers require backup, and many use tape or virtual tape, meaning that a software-based VTL is needed.

17.What is available today with VTL?

For more information about VTL DSI for IBM i, see [this web page](#).

18.How does VTL affect current customers?

Customers that are running applications in PowerVS can use DSI VTL software to back up in the cloud. There are no changes to applications or process. This is different than previous solution, which required IBM Cloud Storage and IBM Cloud Object Storage. Now, faster and less data space is required.

19.What is coming next with VTL on PowerVS roadmap?

Coming soon is the ability of a customer to order on-premises to use migration from Lab services. Software VTL for IBM i PowerVS will be available following by integration that allows VTL to be ordered directly from our GUI.

20.What about those customers that want to move in the future?

Customers intending to move to the Cloud use DSI VTL Software technology to migration of on-premises data to PowerVS. Customers with large amounts of data need to move it efficiently.

21.How does a restore is done from backup 21 when using the DSI VTL?

PowerVS supports iSCSI; therefore, it is not possible on IBM i to perform a D-mode IPL. Special steps must be taken. For more information, see the following web pages:

- [Support for iSCSI VTL](#)
- [BRMS for IBM Cloud using DSI VTL iSCSI Devices](#)

Hence, NFS servers are required for D-mode IPLs.

22.What is Movable IBM i (IBM i MOL)?

To protect customer investments in IBM i entitlements, IBM i MOL allows customers to move on-premises entitlements to PowerVS instead of purchasing new entitlements.

It is applicable to customers who have valid IBM i entitlements with active SWMA.

P10 and P30 IBM i OS and base LPP are available as target environments for IBM i MOL (add-on LPPs are excluded).

23.What is the process for customers that want to move their on-premises DR systems (“Donor”) into PowerVS?

Move entitled IBM i license from on-premises DR systems to PowerVS. For example, if a customer has 10 IBM i processor entitlement on the donor, they can move up to 10 IBM i processor entitles to PowerVS. These 10 entitlements are tracked and charged at the IBM i MOL price. The customer agrees to stop using those 10 IBM i entitlements on the donor.

The process includes the following steps:

- a. Sellers, Business Partners, GSIs complete the Movable IBM i validation worksheet (Seismic Link) by using the following information:
 - Customer name
 - Machine Type model; for example, 9009-22A
 - Seven-digit Machine Serial Number
 - Number of IBM i processor entitlements to move to PowerVS
 - IBM i software tier; for example, P10 or P30
- b. Submit to TSS (cmahadev@in.ibm.com).
- c. TSS to validate customer’s IBM entitlement and respond to sells, BPs, GSIs:
 - Validate the number of IBM processor entitlements
 - Validate IBM i software tier
 - Validate active SWMA

- Communicate back to sellers, BPs, GSIs
 - Save the approval
- d. Sellers, BPs, GSIs to close pricing proposal with IBM i MOL price, if approved.
 - e. Sellers, BPs, GSIs to close contract and Movable IBM i Agreement.
 - f. Sellers, BPs, GSIs to send the signed Movable IBM i Agreement to OM (mingzhi@us.ibm.com) and OM to save the Agreement.
 - g. Technical OM or Development team to implement it in tracking database.

24. How many entitlements will the customer be eligible to move to PowerVS?

A customer can move their entitled IBM i license from their on-premises machine (“donor”) to PowerVS. For example, if a customer has 20 IBM i processor entitlements on the donor, they can move up to 20 IBM i processor entitlements to PowerVS. These 20 entitlements are tracked and charged at IBM i MOL price. The customer agrees to stop using those 20 IBM i entitlements on the donor.

If in PowerVS, a customer expands the IBM i usage to 21 cores in PowerVS. Their first 20 cores are charged at IBM i MOL price and the 21 cores are charged at the standard IBM i subscription price.

25. Can a customer’s on-premises donor machine be a POWER9™, IBM POWER8, POWER7/POWER7+, or earlier machine?

Yes, the type of on-premises server is not relevant for Movable IBM i. A customer’s number of IBM i processor entitlements and active SWMA are validated for the customer’s donor machine.

26. Can a customer with a donor IBM i P10 software tier use Movable IBM i for a PowerVS machine with P30 software tier? Can a customer with a donor IBM i P30 software tier use Movable IBM i for a PowerVS machine with P10 software tier?

- P10 tier donor can use Movable IBM i on P10 tier in PowerVS.
- P30 tier donor can use Movable IBM i on P10 tier or P30 tier in PowerVS. Specifically, a customer with P30 on-premises entitlement who is moving to the PowerVS P10 tier receives the Movable IBM i pricing for P10. This is a one-for-one mapping from donor to PowerVS entitlement.

The Movable IBM i offering does not support upgrading to a larger size than the donor server; for example, P10 tier on-premises cannot use Movable IBM i on the P30 tier in PowerVS.

27. Can a customer with on-premises IBM i P05 software use Movable IBM i in PowerVS?

The Movable IBM i offering does not support upgrading to a larger tier size than the on-premises server. A P05 on-premises customer can run workload in the PowerVS cloud. The customer acquires new PowerVS subscription at standard price.

28. Can a customer with an IBM i P20 software tier donor use Movable IBM i in PowerVS?

Regarding moving from a P20 tier donor to the P10 tier in PowerVS, PowerVS intends to support the ability for P20 tier donors to use the Movable IBM i on P10 tier in PowerVS. This support is on the PowerVS roadmap and is not available as of 12 November 2020.

Regarding moving from a P20 tier donor to the P30 tier in PowerVS, the Movable IBM i offering does not support upgrading to a larger tier size than the on-premises server. Therefore, the customer with P20 donor can acquire the P30 entitlements and pay the standard IBM i P30 subscription price.

29.If a customer has remaining IBM i Software Maintenance on-premises, they receive credit with the Movable IBM i offering?

Consider the following scenario: A customer paid for three years of IBM i SWMA on their donor, one year has passed, and now the customer wants to use Movable IBM i.

Two methods are available by which customers renew SWMA and the methods include different termination provisions. Therefore, check which method under which the customer acquired SWMA:

- SWMA that was acquired by way of the Power Systems channel (for example, AAS, iERP by using e-config) cannot be canceled. There are no credits or refund for SWMA that was acquired by this method. Therefore, if a customer acquired SWMA by using this method, there are no credits issued when Movable IBM i is used for the remaining on-premises SWMA.
- SWMA that was acquired by way of Service Suite or Service Elite, which are acquired by way of IBM Technical Support Services (TSS) sales, can be terminated, subject to the termination provisions of the contract. Credit is issued to the customer for the remaining SWMA by using the standard methods.

30.What if a customer does not have active SWMA on their donor and wants to use Movable IBM i?

Active IBM i group SWMA is required on the donor. Thus, activate SWMA by using standard steps. If the IBM i group SWMA lapsed, an after-license fee applies to reinstate SWMA.

31.Is the customer required to keep SWMA active on-premises?

If the customer is moving all of their IBM i workload to PowerVS, they do not need to keep IBM i group SWMA active on their donor. However, the customer must keep SWMA active on the donor for specific situations, for example: the customer moves part of their IBM i workload to the cloud and continues to use some IBM i entitlements on the donor.

Note: The SWMA renewal for Power Systems Software that SWMA is renewed on all Program entitlements on the donor machine. For example, if a customer has 10 IBM i entitlements on the donor, moves five to PowerVS, and continues to use five on the donor, the customer keeps active IBM i group SWMA on all 10 entitlements on the donor.

32.Can a customer use Movable IBM i for IBM i 7.1?

The PowerVS intends to support IBM i 7.1 in the future. After it is available, the Movable IBM i offering can be used with IBM i 7.1.

33.How many days of migration are allowed for customers moving to PowerVS? What if the customer needs an extended migration period?

For IBM Power Systems Software, 30 days are allowed for concurrent usage of software for migrations. If more than 30 days is required, IBM i customers can acquire IBM i Temporary Licensing (5733-ITL) on the donor machine.

34.Are the Licensed Program Products (LPPs) for IBM i and IBM i optional features available with the Movable IBM i pricing?

The Movable IBM i pricing includes the IBM i operating system and the base set of LPPs for i and IBM i optional features, which are always included with the IBM i subscription pricing. The separately ordered LPPs-PowerHA for i, Cloud Storage Solutions for i, Rational Development Studio for i, and DB2 Web Query for i are not available with Movable IBM i pricing.

35. What if a Managed Service Provider (MSP) has a customer with a donor machine and IBM i entitlement, which can be eligible for Movable IBM i? Can the MSP acquire Movable IBM i pricing for their customer? What if the entitlement is owned by the MSP?

Movable IBM i can be used by the party who holds the IBM i entitlement on the donor.
Examples:

- If a customer is the IBM i entitlement holder on the donor, the customer can use Movable IBM i in PowerVS. The MSP cannot get Movable IBM i pricing in PowerVS for this customer; however, the customer can get such pricing themselves.
- If the MSP is the IBM i entitlement holder on the MSP's donor, the MSP can use Movable IBM i in PowerVS.

36. What is the expected time to deployment of IBM i VM on IBM Power Systems Virtual Server?

A few minutes. Most of the time is spent in the IBM i OS booting.

37. If a customer is looking for LPP 5733 -I D1 InfoPrint Designer/400, is this included in base license on IBM i running on PowerVS?

No. It is a withdrawn product; therefore, it is not part of the base licenses.

38. Power VS encryption (IBM i): For encryption at rest, what are the algorithms that are used and how are keys for the algorithms managed, including rotation frequency?

As on IBM Cloud, you can find Flash System 7200 and 9200. For more information, see [Spectrum Virtualize Encryption](#).

FCMs (Tier 1) have built-in encryption. Keys are not rotated. This also does not support a KMIP.

39. For IBM i BRMS/IBM Cloud Storage backup to IBM Cloud Object Storage. What is the average throughput for 1 TB backup? For how much time do you calculate?

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; it was taken from tests, and can vary on your scenario or situation.

40. 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 get less number virtual tapes but throughput is less, but time it takes to transfer is same such as without compression.
- With IBM Cloud Storage compression, it takes same number of virtual tapes such as without compression (option1) and you do not get any extra space benefit such as you get with BRMS compression. Time it takes almost 5 - 6 times to transfer to IBM Cloud Object Storage so It is not recommend IBM Cloud Storage compression. It is suggested BRMS Compression because you save extra space of your IBM i LPAR that you can use for retaining some days of virtual tapes on your IBM i LPAR as well. With BRMS Compression, you got around 59% compression ratio and with IBM Cloud Storage you got around 80% compression ratio.

41. Can IBM Cloud Object Storage be used as backup for IBM i? Does it function with on-premises solutions instead of the daily physical tape replacement?

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 must be transferred.
- Full System Save is possible; therefore, a Full System restore requires more work with NFS server and Network Install.
- 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 out to IBM Cloud Object Storage from IBM Cloud Private network.

It also is possible to perform backups from on-premises directly to IBM Cloud Object Storage by using public interface whether all requirements are met in on-premises systems. Is for this reason that you review data transfer performance.

42. If a customer takes their IBM i image from on-premises to off-premises and that system includes Rational Dev Studio licensed, does it transfer to an instance that is on IBM Power Systems Virtual Server on IBM Cloud? Does Rational Directory Server on PowerVS include any extra cost?

If you are bringing your IBM i VM to PowerVS by image file or backup or restore, the license key for 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.

Glossary

Access Client Solutions (ACS)

Provides a Java-based, platform-independent interface that runs on most operating systems that support Java.

Ansible

An open-source configuration management and provisioning tool that automates complex applications and provisions resources in the cloud.

Application Programming Interface (API)

An interface that allows an application program that is written in a high-level language to use specific data or functions of the operating system or another program.

Aspera

An IBM software product that helps users move critical files and data sets of any size at a maximum speed over your infrastructure and worldwide IP networks.

Backup and Recovery Media Services (BRMS)

Helps to implement a disciplined approach to managing your backups and provides you with an orderly way to retrieve lost or damaged data.

Cloud computing

A computing platform where users can access applications or computing resources, as services, from anywhere through their connected devices.

Cloud native

Refers to how an application is built and deployed. A cloud native application consists of discrete, reusable components that are known as microservices that are designed to integrate into any cloud environment.

Cloud Object Storage (COS)

A persistent unstructured data storage service format that are spread out across multiple networked systems.

Cloud service providers (CSPs)

Third-party companies offering a cloud-based platform, including SaaS, PaaS, and IaaS services.

CloudForms

A management platform to control and manage cloud infrastructures.

Command Line Interface (CLI)

A computer interface in which the input and output are text based.

Container

A system construct that allows users to simultaneously run separate logical operating system instances. Containers use layers of file systems to minimize image sizes and promote reuse.

Dashboard

A user interface component that provides a comprehensive summary of pertinent information from various sources to the user.

DevOps

A software methodology that integrates application development and IT operations so that teams can deliver code faster to production and iterate continuously based on market feedback.

Direct Link Connect (DLC)

A protocol layer that is used by nodes on a data link to accomplish an orderly exchange of information.

Disaster recovery (DR)

IT technologies and best practices that prevent or minimize data loss and business disruption resulted from equipment failures and localized power outages to natural disasters incidents.

Dynamic Resource Optimization (DRO)

Feature that continuously monitors the resource utilization of physical hosts and virtual machines in a cloud environment.

Entitled Systems Support (ESS)

A website from IBM to view and manage IBM Power Systems and IBM Storage Systems software and hardware.

High availability (HA)

The ability of IT services to withstand all outages and continue providing processing capability according to some predefined service level.

Hybrid cloud

A cloud computing environment that consists of multiple public and private resources.

Identity and Access Management (IAM)

The process of controlling access of authorized users to data and applications, while helping companies comply with various regulatory requirements.

Image (Container Image)

A file and its execution parameters that are used within a container runtime to create a container. The file consists of a series of layers, combined at runtime, that are created as the image is built by successive updates.

Infrastructure as a Service (IaaS)

The delivery of a computer infrastructure, including server, networking, data center, and storage functions as an outsourced service.

Infrastructure-as-Code (IaC)

Manages and provides an infrastructure through code instead of through manual processes.

Licensed Program Product (LPP)

It is a licensing scheme for IBM licensed program products.

Logical Partition (LPAR)

See “Virtual machine (VM)”.

Managed Service Provider (MSP)

Third-party company that operates day-to-day management services to its customers.

Multicloud

A cloud adoption strategy that embraces a mix of cloud models (public, dedicated, private, managed) to best meet unique business, application, and workload requirements.

On-premises

Pertaining to software that is installed and run on the local computers of a user or organization.

Open Virtual Application (OVA)

An open standard for packaging and distributing virtual appliances.

Platform as a Service (PaaS)

The delivery of a computing platform, including applications, optimized middleware, development tools, and runtime environments, in a cloud-based environment.

Pod

A group of containers that are running on a Kubernetes cluster. A pod is a runnable unit of work, which can be a stand-alone application or a set of microservices.

Private cloud

A cloud computing environment on which access is limited to members of an enterprise and partner networks.

Program Temporary Fix (PTF)

A fix approved by IBM and made available to all customers.

Public cloud

A cloud computing environment on which access to standardized resources, such as infrastructure, multi-tenant hardware, and services, is available to subscribers on a pay-per-use basis.

Software as a Service (SaaS)

A model of software deployment whereby software including business processes, enterprise applications, and collaboration tools, are provided as a service to customers through the cloud.

Terraform

An open source “Infrastructure as Code” tool that was created by HashiCorp.

Virtual machine (VM)

An instance of a data-processing system that appears to be at the exclusive disposal of a single user, but whose functions are accomplished by sharing the resources of a physical data-processing system.

Virtual Server Instance (VSI)

See “Virtual machine (VM)”.

Virtual Tape Libraries (VTL)

A hardware component that can emulate a tape library while it uses a disk as the underlying storage hardware.

Related publications

The publications that are listed in this section are considered particularly suitable for a more detailed discussion of the topics that are covered in this book.

IBM Redbooks

The following IBM Redbooks publications provide more information about the topic in this document. Note that some publications that are referenced in this list might be available in softcopy only:

- ▶ *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

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>

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 ADDMLBRM 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

IBM Power Systems Virtual Server Guide for IBM i

SG24-8513-00

ISBN 0738460354



(0.5" spine)

0.475" x 0.873"

250 <-> 459 pages



SG24-8513-00

ISBN 0738460354

Printed in U.S.A.

Get connected

