Redbooks

ibm.com/redbooks

Red Hat OpenShift V4.3 on IBM Power Systems Reference Guide

Dino Quintero

Daniel Casali

Alain Fisher

Federico Fros

Miguel Gomez Gonzalez

Felix Gonzalez

Paulo Sergio Lemes Queiroz

Sudipto Pal

Bogdan Savu

Richard Wale

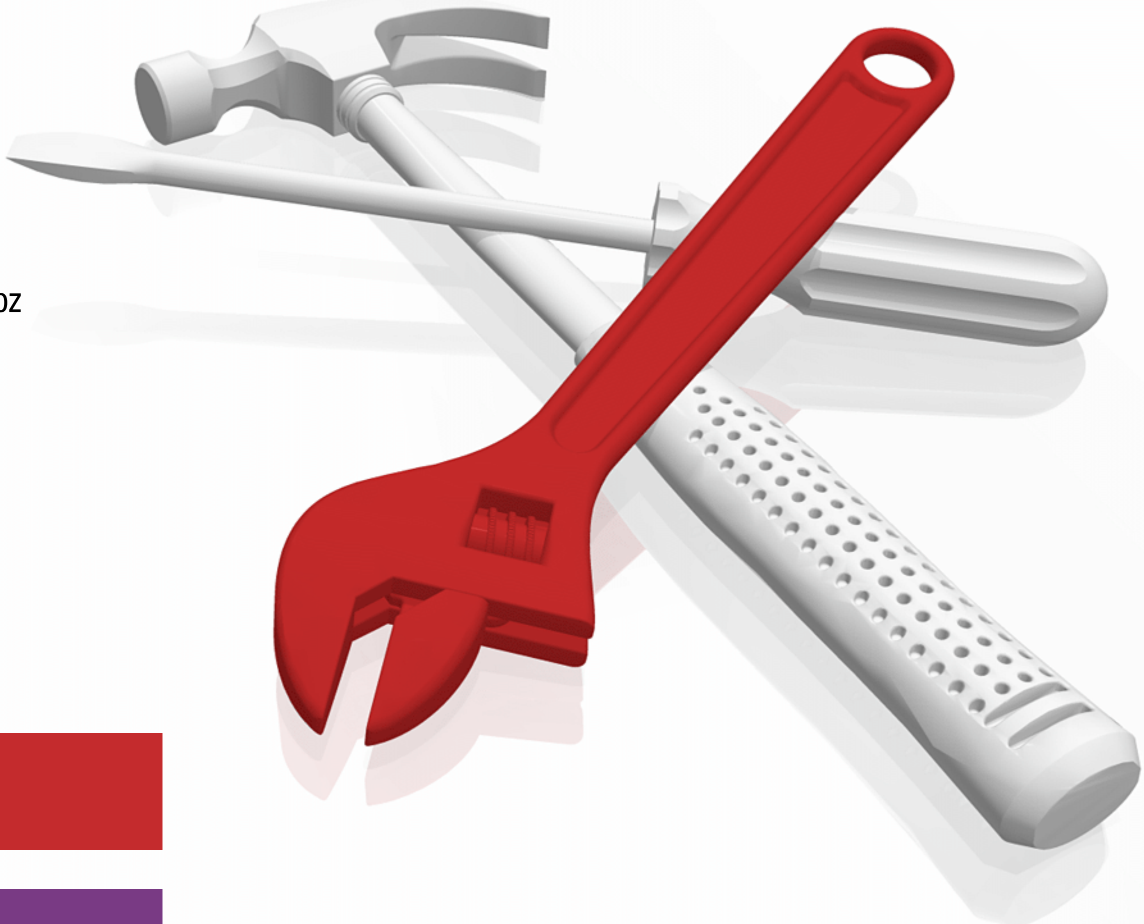

Cloud

Power Systems

IBM

Redpaper

IBM Redbooks

Red Hat OpenShift V4.3 on IBM Power Systems Reference Guide

September 2020

REDP-5599-00

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

First Edition (September 2020)

This edition applies to:

- Red Hat Enterprise Linux V7
- Red Hat OpenShift Container Platform for Power Enterprise V4.3
- Red Hat CoreOS V4.3
- IBM AIX V7.2

Contents

Notices

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

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

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

Trademarks

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

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

AIX®
Cognos®
DataStage®
DB2®
Db2®
IBM®
IBM Cloud®
IBM Cloud Pak®
IBM Spectrum®
IBM Watson®
InfoSphere®
Interconnect®
OpenCAPI™
POWER®
POWER8®
POWER9™
PowerHA®
PowerVM®
Redbooks®
Redbooks (logo) ®
SPSS®
SystemMirror®
Watson™

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.

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

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

RStudio, and the RStudio logo are registered trademarks of RStudio, Inc.

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® Redpaper publication describes how to deploy Red Hat OpenShift V4.3 on IBM Power Systems servers.

This book presents reference architectures for deployment, initial sizing guidelines for server, storage, and IBM Cloud® Paks. Moreover, this publication delivers information about initial supported Power System configurations for Red Hat OpenShift V4.3 deployment (bare metal, IBM PowerVM® LE LPARs, and others).

This book serves as a guide for how to deploy Red Hat OpenShift V4.3 and provide start guidelines and recommended practices for implementing it on Power Systems and completing it with the supported IBM Cloud Paks.

The publication addresses topics for developers, IT architects, IT specialists, sellers, and anyone who wants to implement a Red Hat OpenShift V4.3 and IBM Cloud Paks on IBM Power Systems. This book also provides technical content to transfer how-to skills to the support teams, and solution guidance to the sales team.

This book compliments the documentation that is available at IBM Knowledge Center, and also aligns with the educational offerings that are provided by the IBM Systems Technical Education (SSE).

Authors

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

Dino Quintero is an IT Management Consultant and an IBM Level 3 Senior Certified IT Specialist with IBM Redbooks® in Poughkeepsie, New York. He has 24 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, cloud computing, artificial intelligence including machine and deep learning, and cognitive solutions. He also is a Certified Open Group Distinguished IT Specialist. Dino is formerly from the province of Chiriqui in Panama. Dino holds a Master of Computing Information Systems degree and a Bachelor of Science degree in Computer Science from Marist College.

Daniel Casali is a Thought Leader Information Technology Specialist working for 15 years at IBM with Power Systems, High Performance Computing, Big Data, and Storage. His role at IBM is to bring to reality solutions that address client's needs by exploring new technologies and for different workloads. He is also fascinated by real multicloud implementations, and always trying to abstract and simplify the new challenges of the heterogeneous architectures that are intrinsic to this new consumption model, be it on-premises or in the public cloud.

Alain Fisher is an IT specialist and DevOps Developer and supports many IBM development teams at the IBM Hursley Lab, UK. He holds a B.Sc. (Hons) degree in Computer Science from The Open University, England. He joined IBM in 2001 supporting DB2® middleware products, such as DB2, before moving to the Hursley Lab. His areas of expertise include the OpenStack cloud platform, cloud deployments, automation, and virtualization. He has participated on two previous IBM Redbooks since 2005.

Federico Fros is an IT Specialist who works for IBM Global Technologies Services leading the UNIX and Storage team for the IBM Innovation center in Uruguay. He has worked in IBM for more than 15 years, including 12 years of experience in IBM Power Systems and IBM Storage. He is an IBM Certified Systems Expert for UNIX and High Availability. His areas of expertise include IBM AIX®, Linux, PowerHA® SystemMirror®, IBM PowerVM, SAN Networks, Cloud Computing, and IBM Storage Family, including IBM Spectrum® Scale.

Miguel Gomez Gonzalez is an IBM Power Systems Integration engineer based in Mexico with over 11 years experience in Linux and IBM Power Systems technologies. He has extensive experience in implementing several IBM Linux and AIX clustering solutions. His areas of expertise are Power Systems performance boost, virtualization, high availability, system administration, and test design. Miguel holds a Master in Information Technology Management from ITESM.

Felix Gonzalez is an IT Specialist who currently works in UNIX and Storage Team for IBM Global Technologies Services in Uruguay. He is studying Electronic Engineering at Universidad de la República in Uruguay. He is a Red Hat Linux Certified Specialist with 5 years of experience in UNIX systems, including FreeBSD, Linux, and AIX. His areas of expertise include AIX, Linux, FreeBSD, Red Hat Ansible, IBM PowerVM, SAN, Cloud Computing, and IBM Storage.

Paulo Sergio Lemes Queiroz is a Systems Consultant for AI Solutions on Power Systems with IBM in Brazil. He has 19 years of experience in UNIX and Linux, ranging from systems design and development to systems support. His areas of expertise include AIX Power Systems, AIX, GPFS, RHCS, KVM, and Linux. He is a Certified Advanced Technical Expert for Power Systems with AIX (CATE) and a Red Hat Certified Engineer (RHCE).

Sudipto Pal is an Application Architect and Technical Project manager in GBS with over 20 years of leadership experience in designing innovative business solution for clients from various industry, domain, and geographical locations. He is skilled in Cloud computing, Big Data, application development, and virtualization. He successfully delivered several critical projects with IBM clients from USA and Europe. He led Cognos® administration competency and mentor several candidates. He co-authored *Exploiting IBM PowerVM Virtualization Features with IBM Cognos 8 Business Intelligence*, SG24-7842.

Bogdan Savu is a Cloud Infrastructure Architect at IBM POWER® IaaS and works for IBM Systems in Romania. 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 an IBM Certified Advanced Technical Expert for Power Systems, ITIL 4 Managing Professional Certificate, TOGAF 9 Certified, Red Hat Certified Engineer, Red Hat Certified Specialist in Ansible Automation, Red Hat Certified Specialist in Red Hat OpenShift Administration, and VMware Certified Professional. His areas of expertise include Cloud Computing, Virtualization, DevOps, and scripting.

Richard Wale is a Senior IT Specialist, supporting many IBM development teams at the IBM Hursley Lab, UK. He holds a B.Sc. (Hons) degree in Computer Science from Portsmouth University, England. He joined IBM in 1996 and has been supporting production IBM AIX systems since 1998. His areas of expertise include IBM Power Systems, PowerVM, AIX, and IBM i. He has participated in co-writing many IBM Redbooks publications since 2002.

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

Wade Wallace
IBM Redbooks, Poughkeepsie Center

John Banchy, John Engel
Sirius, an IBM Business Partner

Susan Forbes, Bryan Kelley, Proshanta Saha, Angshuman Roy
IBM USA

Theresa Xu
IBM Canada

Claus Huempel
IBM Germany

Now you can become a published author, too!

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

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

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

  ibm.com/redbooks

- Send your comments in an email to:

  redbooks@us.ibm.com

- Mail your comments to:

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

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

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

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

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

- Stay current on recent Redbooks publications with RSS Feeds:

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

1

Introduction to Red Hat OpenShift V4.3

This chapter provides an overview of the scope of this publication.

This chapter includes the following topics:

1.1 Introduction

This publication provides an overview of the 4.3 release of Red Hat OpenShift Platform on IBM Power Systems. It is interim successor to the publication *Red Hat OpenShift and IBM Cloud Paks on IBM Power Systems: Volume 1*, SG24-8459. That publication provided a summary overview of containers and Kubernetes, and introduced of Red Hat OpenShift onto the IBM Power Systems platform.

With over 25 years since the release of the original models, IBM Power Systems continues to be designed for traditional Enterprise workloads and the most demanding, data-intensive computing. The range of models offer flexibility, scalability, and innovation.

We made a statement of intent within the previously mentioned publication, that subsequent volumes be published in due course. We felt this approach better served the agile nature of the Red Hat OpenShift product. The window is always moving, the next release is already on the horizon. At the time of writing, Volume 2 is in development; however, with some of the changes and improvements that are provided by Red Hat OpenShift V4.3, we felt an interim publication was needed to follow the product release.

Red Hat OpenShift is one of the most reliable enterprise-grade containers. It is designed and optimized to easily deploy web applications and services. Categorized as a cloud development Platform as a Service (PaaS), it is and based on industry standards, such as Docker and Kubernetes.

This publication explains what is new with Red Hat OpenShift on IBM Power Systems, and provides updated installation instructions and sizing guides.

1.2 Red Hat OpenShift V4.3

The *Red Hat OpenShift and IBM Cloud Paks on IBM Power Systems: Volume 1*, SG24-8459, publication relates to Red Hat OpenShift V3.11. This book focuses on Red Hat OpenShift V4.3. Although V4.x releases exist for other hardware platforms, V4.3 is the first official V4.x release for IBM Power Systems.

> **Note:** Red Hat OpenShift V4.3 for IBM Power Systems was officially announced by IBM on 28th April 2020. It was subsequently announced two days later by Red Hat on their Red Hat OpenShift blog.

Red Hat OpenShift V4.3 provides several important changes compared to the previous V3.x release written about in Volume 1. The most significant change is the move from Red Hat Enterprise Linux to Red Hat CoreOS for the operating system that is used for the cluster nodes.

CoreOS is a lightweight Linux that is specifically designed for hosting containers across a cluster of nodes. As it is patched and configured as part of Red Hat OpenShift, it bring consistency to the deployed environment and reduces the overhead of ongoing ownership.

For more information about Red Hat OpenShift V4.3 on IBM Power Systems, see Chapter 2, “Supported configurations and sizing guidelines” on page 5.

1.3 Publication overview

Within the subsequent chapters in this publication, we discuss Red Hat OpenShift V4.3 on IBM Power Systems; describe the different installation methods (compared to the previous V3.11 release); and document supported environments, and suggested sizing considerations.

Note: Documented information regarding supported environments, configurations, and sizing guides are accurate at the time of this writing. Be aware that because of the agile nature of Red Hat OpenShift, elements and aspects can change with subsequent V4.3.x updates.

If major changes are required, a revised edition of this IBM Redpaper publication can be published. However, we always recommend to check official resources (release notes, online documentation, and so on) for any changes to what is presented here.

2

Supported configurations and sizing guidelines

This chapter provides information about the supported configurations, sizing guidelines, and recommended practices to help you size and deploy Red Hat OpenShift on Power Systems.

This chapter includes the following topics:

2.1 IBM Power Systems

Over the years, the IBM Power Systems family has grown, matured, was innovated, and pushed the boundaries of what clients expect and demand from the harmony of hardware and software.

With the advent of the POWER4 processor in 2001, IBM introduced logical partitions (LPARs) outside of their mainframe family to another audience. What was seen as radical then grew into the expected today. The term *virtualization* is now commonplace across most platforms and operating systems. These days, virtualization is the core foundation for Cloud Computing.

IBM Power Systems is built for the most demanding, data-intensive, computing on earth. The servers are cloud-ready and help you unleash insight from your data pipeline, from managing mission-critical data, to managing your operational data stores and data lakes, to delivering the best server for cognitive computing.

IBM POWER9, the foundation for the No. 1 and No. 2 supercomputers in the world, is the only processor with state-of-the-art I/O subsystem technology, including next generation NVIDIA NVLink, PCIe Gen4, and IBM OpenCAPI™.

POWER9 processor-based servers can be found in three product families: Enterprise servers, scale-out servers, and accelerated servers. Each of these three families is positioned for different types of client requirements and expectations.

IBM Power Systems servers based on POWER9 processors are built for today's most advanced applications from mission-critical enterprise workloads to big data and AI, as shown in Figure 2-1.

Figure 2-1 IBM Power Systems

Robust scale-out and enterprise servers can support a wide range of mission-critical applications running on IBM AIX, IBM i, and Linux operating systems. They also can provide building blocks for private and hybrid cloud environments.

Scale-out servers deliver the performance and capacity that is required for big data and analytics workloads.

Servers provide intelligent, accelerated infrastructure for modern analytics, high-performance computing (HPC), and AI workloads. They are advanced enterprise servers that deliver fast machine learning performance.

2.1.1 Mission-critical workloads

To handle the most data-intensive mission-critical workloads, organizations need servers that can deliver outstanding performance and scalability. Whether they are supporting small business groups or building large private and hybrid cloud environments, they need no-compromise infrastructure.

Enterprise (scale-up) servers

IBM Power Systems Enterprise (scale-up) servers offer the highest levels of performance and scale for the most data- intensive, mission-critical workloads, as shown in Table 2-1. They can also serve as building blocks for growing private and hybrid cloud environments. Support for AIX, Linux, and IBM i (for the IBM Power Systems E980 server) gives organizations the flexibility to run a wide range of applications.

Table 2-1 IBM Power Systems: Enterprise servers

	E950	**E980**
Key features	► Enterprise-class capabilities in a reliable, space-efficient form factor. ► Exceptional performance at an affordable price.	► Ideal foundation for world-class private or hybrid cloud. ► Can power large-scale, mission-critical applications. ► Flagship, high-end server.
Machine type and model (MTM)	9040-MR9	9080-M9S
Form factors	4U	5U system node and 2U system controller unit.
Sockets	2 - 4	4 per node
Processor cores	► Up to 48 cores – 12 core processor sockets at 3.15 to 3.80 GHz (max). ► Up to 44 cores – 11 core processor sockets at 3.2 to 3.80 GHz (max). ► Up to 40 cores – 10 core processor sockets at 3.40 to 3.80 GHz (max). ► Up to 32 cores – 8 core processor sockets at 3.60 to 3.80 GHz (max).	One node: 4x POWER9 CPUs; 8, 10, 11 or 12 cores each System maximum: 16x POWER9 CPUs; 8, 10, 11 or 12 cores each
Memory slots	128	128 per node
Memory max.	16 TB	64 TB per node
PCIe G4 slots	10	8 per node; max. 32
Supported operating systems	AIX and Linux	AIX, IBM i, and Linux

The IBM Power Systems E950 server is the correct fit for growing midsize businesses, departments, and large enterprises that need a building-block platform for their data center. The IBM Power Systems E980 server is designed for large enterprises that need flexible, reliable servers for a private or hybrid cloud infrastructure.

Scale-out servers

IBM Power Systems scale-out servers for mission-critical workloads offer a strong alternative to commodity x86 servers, as shown in Table 2-2. They provide a robust, reliable platform to help maximize performance and help ensure availability.

Scale-out AIX, IBM i, and Linux servers are designed to scale out and integrate into an organization's cloud and AI strategy, delivering exceptional performance and reliability.

Table 2-2 IBM Power Systems: Scale-out servers

	S914	**S922**	**S924**	**L922**
Key features	► Entry-level offering. ► Industry-leading integrated security and reliability. ► Cloud-enabled.	► Strong price-performance for mission-critical workloads. ► Dense form factor with large memory footprint. ► Cloud-enabled with integrated virtualization.	► Industry-leading price-performance for mission-critical workloads. ► Large memory footprint. ► Strong security and reliability. ► Cloud-enabled with integrated virtualization.	► Industry-leading price performance for mission-critical Linux workloads. ► Dense form factor with large memory footprint.
Machine type and model (MTM)	9009-41A	9009-22A	9009-42A	9008-22L
Form factors	4U and tower	2U	4U	2U
Sockets	1	2	2	1 or 2
Microprocessors	1x POWER9 CPUs; 4, 6 or 8 cores.	Up to 2x POWER9 CPUs; 4, 8 or 10 cores.	2x POWER9 CPUs; 8, 10 or 12 cores.	Up to 2x POWER9 CPUs; 8, 10 or 12 cores.
Memory slots	16	32	32	32
Memory max.	1 TB	4 TB	4 TB	4 TB
PCIe G4 slots	2	4	4	4
Supported operating systems	AIX, IBM i, and Linux	AIX, IBM i, and Linux	AIX, IBM i, and Linux	Linux

Scale-out servers for SAP HANA servers are designed to deliver outstanding performance and a large memory footprint of up to 4 TB in a dense form factor, as shown in Table 2-3 on page 9. These servers help deliver insights fast at the same time maintaining high reliability. They are also scalable: When it is time to grow, organizations can expand database capacity and the size of their SAP HANA environment without having to provision a new server.

Table 2-3 IBM Power Systems: Scale-out servers for SAP HANA

	H922	H924
Key features	► Optimized for SAP HANA. ► High performance, tight security. ► Dense form factor with large memory footprint ► For Linux-focused customers.	► High performance for SAP HANA. ► Strong security with large memory footprint ► For Linux-focused customers.
Machine type and model (MTM)	9223-22H	9223-42H
Form factors	2U	4U
Sockets	1 upgradeable or 2	2
Cores per socket	4, 8 or 10	8, 10 or 12
Memory slots	32	32
Memory max.	4 TB	4 TB
PCIe G4 slots	4	4
Supported operating systems	AIX, IBM i, and Linux	AIX, IBM i, and Linux

2.1.2 Big data workloads

Across industries, organizations are poised to capitalize on big data to generate new business insights, improve the customer experience, enhance efficiencies, and gain competitive advantage. However, to make the most of growing data volumes, they need servers with the performance and capacity for big data and AI workloads.

IBM Power Systems Scale-out servers for big data deliver the outstanding performance and scalable capacity for intensive big data and AI workloads, as shown in Table 2-4. Purpose-built with a storage-rich server design and industry-leading compute capabilities, these servers are made to explore and analyze a tremendous amount of data, all at a lower cost than equivalent x86 alternatives.

Table 2-4 IBM Power Systems: Scale-out servers for big data

	LC921	LC922
Key features	► High performance in a space-saving design. ► Industry-leading compute in a dense form factor.	► Highest storage capacity in the IBM Power Systems portfolio. ► Up to 44 cores and 2 TB of memory. ► High performance at lower cost than comparable x86 systems.
Machine type and model (MTM)	9006-12P	9006-22P
Form factors	1U	2U
Sockets	1 upgradeable or 2	2

	LC921	LC922
Microprocessors	1x or 2x POWER9 CPUs; 16 or 20 cores	1x or 2x POWER9 CPUs; 16, 20 or 22 cores
Memory slots	32	16
Memory max.	2 TB	2 TB
PCIe G4 slots	4	6
Supported operating system	Linux	Linux
Max. storage	40 TB	120 TB

2.1.3 Enterprise AI workloads

AI holds tremendous promise for facilitating digital transformations, accelerating innovation, enhancing the efficiency of internal processes, identifying new marketplace opportunities, and more. For organizations to take advantage of AI and cognitive technologies, such as machine learning and deep learning, they need powerful, accelerated servers that can handle these data-intensive workloads.

Accelerated servers can also play a vital role in HPC and supercomputing. With the correct accelerated servers, researchers and scientists can explore more complex, data- intensive problems and deliver results faster than before.

The IBM Power Systems Accelerated Compute Server helps reduce the time to value for enterprise AI initiatives. The IBM PowerAI Enterprise platform combines this server with popular open source deep learning frameworks and efficient AI development tools to accelerate the processes of building, training, and inferring deep learning neural networks. Using PowerAI Enterprise, organizations can deploy a fully optimized and supported AI platform with blazing performance, proven dependability, and resilience.

The new IBM Power System IC922 server is built to delivers powerful computing, scaling efficiency, and storage capacity in a cost-optimized design to meet the evolving data challenges of the artificial intelligence (AI) era (see Table 2-5).

The IC in IC922 stands for inference and cloud. The I can also stand for I/O.

Table 2-5 IBM Power Systems: Accelerated compute servers

	AC922	IC922
Key features	► Unprecedented performance for modern AI, analytics, and HPC workloads. ► Proven deployments from small clusters to the world’s largest supercomputers, with near-linear scaling. ► Simple GPU acceleration.	The Power IC922 server is engineered to put your AI models to work and unlock business insights. It uses cooptimized hardware and software to deliver the necessary components for AI inference.
Machine type and model (MTM)	8335-GTH I 8335-GTX	9183-22X
Form factors	2U	2U
Sockets	2	2

	AC922	IC922
Microprocessors	2x POWER9 with NVLink CPUs: 16 or 20 cores; or 18 or 22 cores with liquid cooling.	12-core (2.8 - 3.8 GHz), 16-core (3.35 - 4.0 GHz), and 20-core (2.9 - 3.8 GHz) POWER9.
GPUs	4 or 6 NVIDIA Tesla GPU processors.	Up to six NVIDIA T4 GPU accelerator.
Memory slots.	16	32
Memory max.	1 TB	2 TB
PCIe G4 slots	4	Multiple I/O options in the system with the standard Peripheral Component IBM Interconnect® Express (PCIe) Riser.
Supported operating systems	Linux	Linux

2.2 Red Hat OpenShift Container Platform V4.3 on IBM Power Systems

Red Hat OpenShift V4.3 for Power Systems was announced by IBM on April 28, 2020. For more information, see the announcement letter for Red Hat OpenShift V4.3 on Power Systems.

Red Hat OpenShift V4.3 for Power Systems is an enterprise-grade platform that provides a secure, private platform-as-a-service cloud on IBM Power Systems servers.

The Red Hat OpenShift V4.3 includes the following key features:

- Red Hat Enterprise Linux CoreOS, which offer a fully immutable, lightweight, and container-optimized Linux OS distribution.
- Cluster upgrades and cloud automation.
- IBM Cloud Paks support on Power Systems platforms. IBM Cloud Paks are a containerized bundling of IBM middleware and open source content.

The Red Hat OpenShift architecture builds on top of Kubernetes and is consists of the following types of roles for the nodes:

- Bootstrap

  Red Hat OpenShift Container Platform uses a temporary bootstrap node during initial configuration to provide the required information to the master node (control plane). It boots by using an Ignition configuration file that describes how to create the cluster. The bootstrap node creates the master nodes, and master nodes creates the worker nodes. The master nodes install more services in the form of a set of Operators. The Red Hat OpenShift bootstrap node runs CoreOS V4.3.

- Master

  Red Hat OpenShift Container Platform master is a server that performs control functions for the entire cluster environment. The master machines are the control plane. It is responsible for the creation, scheduling, and management of all objects that are specific to Red Hat OpenShift. It includes API, controller manager, and scheduler capabilities in one Red Hat OpenShift binary. It is also a common practice to install an `etcd` key-value store on Red Hat OpenShift masters to achieve a low-latency link between `etcd` and Red Hat OpenShift masters. The Red Hat OpenShift master node runs CoreOS V4.3 or Red Hat Enterprise Linux V7.x.

  > **Important:** Because of the consensus that is required by the RAFT[a] algorithm, the `etcd` service must be deployed in odd numbers to maintain quorum. For this reason, the minimum number of `etcd` instances for production environments is three.

  a. The Raft Consensus Algorithm: https://raft.github.io/

- Worker

  Red Hat OpenShift worker nodes run containerized applications that are created and deployed by developers. A Red Hat OpenShift worker node contains the Red Hat OpenShift node components, including the container engine CRI-O, container workloads running and stopping executor Kubelet, and a service proxy managing across worker nodes communication for Pods. A Red Hat OpenShift application node runs CoreOS V4.3 or Red Hat Enterprise Linux V7.x.

A *deployment host* is any virtual or physical host that is typically required for the installation of Red Hat OpenShift. The Red Hat OpenShift installation assumes that many, if not all the external services, such as DNS, load balancing, HTTP server, and DHCP are available in a data center and therefore they do not need to be duplicated on a node in the Red Hat OpenShift cluster.

However, experience shows that creating a deployment host node and consolidating the Red Hat OpenShift required external services on it greatly simplifies installation. After installation is complete, the deployment host node can continue to serve as a load balancer for the Red Hat OpenShift API service (running on each of the master nodes) and the application ingress controller (also running on the three master nodes). As part of providing this single front door to the Red Hat OpenShift cluster, it can serve as a jump server that controls access between some external network and the Red Hat OpenShift cluster network.

Figure 2-2 shows a high-level view of the Red Hat OpenShift Container Platform components for the various IBM Power Systems hardware platforms.

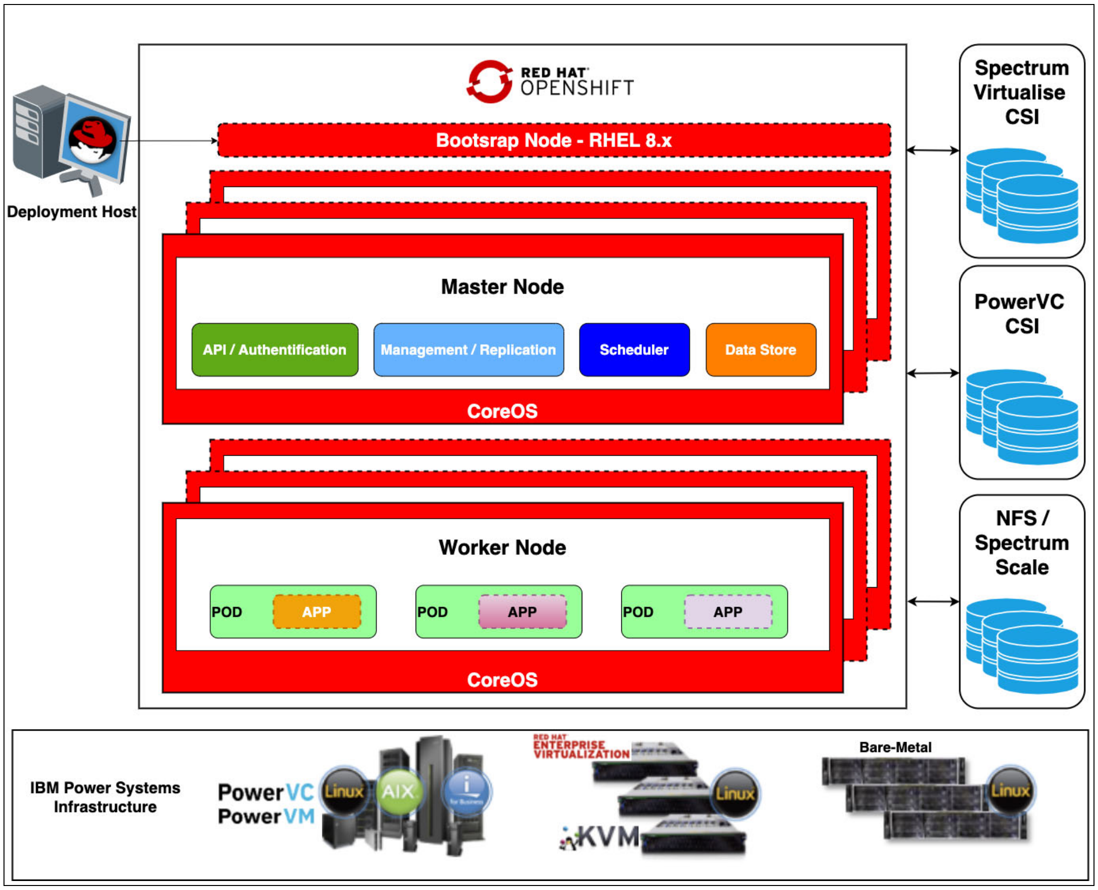

Figure 2-2 Red Hat OpenShift Container Platform for IBM Power Systems

Note: Standalone KVM is for development purposes only.

Nodes can run on top of PowerVC, PowerVM, Red Hat Virtualization, KVM, or run bare metal environment. Table 2-6 shows the IBM Power Systems infrastructure landscape for Red Hat OpenShift Container Platform V4.3.

Table 2-6 IBM Power Systems infrastructure landscape for OpenShift V4.3

IaaS	PowerVC*	N/A	RHV-M	N/A
Hypervisor	PowerVM	PowerVM	KVM/RHV	Bare metal
Guest operating system	CoreOS V4.3 or later	CoreOS V4.3 or later	CoreOS V4.3 or later	CoreOS V4.3 or later
Systems	E980, E950, S924, S922, S914, L922	E980, E950, S924, S922, S914, L922	LC922, AC922, IC922	LC922, AC922, IC922
File storage	NFS, Spectrum Scale**	NFS, Spectrum Scale**	NFS, Spectrum Scale**	NFS, Spectrum Scale**

IaaS	PowerVC*	N/A	RHV-M	N/A
Block storage	PowerVC CSI, Spectrum Virtualize CSI	Spectrum Virtualize CSI	Spectrum Virtualize CSI,	Spectrum Virtualize CSI

Consider the following points:

- PowerVC V1.4.4.1 added support for CoreOS and Red Hat Enterprise Linux V8 VMs and guests.
- Spectrum Scale on Red Hat OpenShift V4.3 requires Red Hat Enterprise Linux V7 worker nodes.

2.2.1 Differences between Red Hat OpenShift Container Platform V4.3 and V3.11

This section highlights some of the following high-level differences between Red Hat OpenShift V4.3 and Red Hat OpenShift V3.11 on IBM Power Systems:

- One key difference is the change in the base OS, as it transitions from Red Hat Enterprise Linux V7 to CoreOS for the master nodes. CoreOS is a stripped-down version of Red Hat Enterprise Linux that is optimized for container orchestration. CoreOS is bundled with Red Hat OpenShift V4.x and is not separately charged.
- The container runtime moves from Docker to Cri-O. The difference in the base operating system is that Cri-O is a lightweight alternative to the use of Docker as the run time for Kubernetes.
- The container CLI also transitions into Podman. The key difference between Podman and Docker for CLI is that Podman does not require a daemon to be running. It also shares many of the underlying components with other container engines, including CRI-O.
- Installation and configuration are done by using openshift-install (ignition-based) deployment for Red Hat OpenShift V4.3 and replacing Red Hat OpenShift-Ansible based on Ansible tool.

The two stacks are compared in Table 2-7.

Table 2-7 Red Hat OpenShift Container Platform V4.3 versus V3.11 stack differences

	Red Hat OpenShift V4.3	**Red Hat OpenShift V3.11**
Base operating system	CoreOS	Red Hat Enterprise Linux V7
Container Run-time	CRI-O	Docker
Container CLI	Podman, Buildah, Skopeo	Docker
Installation	openshift-Install	OpenShift-Ansible
Operational tools	Prometheus	Hawkular, Cassandra
Z-stream update	Automated (every week)	Manual (six weeks)
Content Update	Automated with operators	Manual

A typical Red Hat OpenShift V3.11 deployment is shown in Figure 2-7.

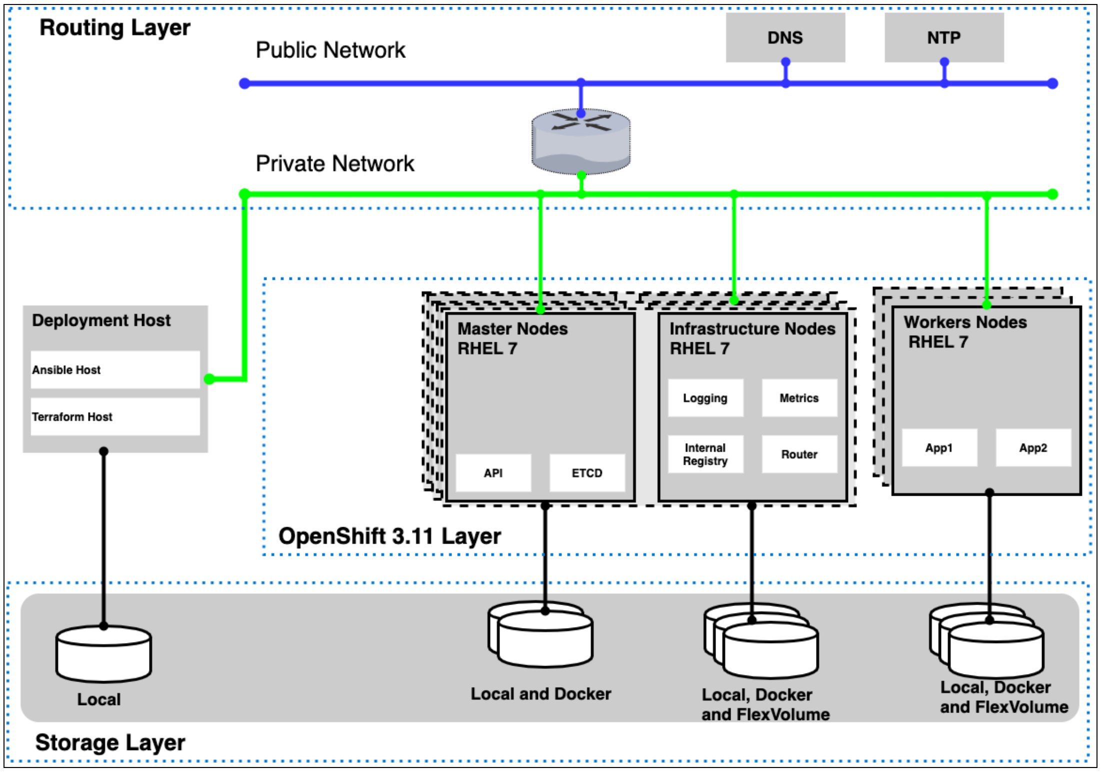

Figure 2-3 Typical Red Hat OpenShift V3.11 deployment

In Figure 2-7, the solid boxes represent nodes that are required at a minimum to run Red Hat OpenShift Container Platform v3.11. The dashed lines represent recommended configuration for production.

When deploying Red Hat OpenShift V3.11 on Power Systems, it is required to have only one master, one infrastructure, and one worker node. It is also common practice to consolidate the master and infrastructure nodes on a single VM.

Although the worker node can also be consolidated onto a single VM, it is not recommended because Red Hat OpenShift core licenses are determined by the number of cores on the worker nodes and not the control nodes. Although this installation is the minimal installation, it is recommended that at least three copies of this worker node on three separate systems for high availability and fault tolerance when moving into a production level deployment.

In Red Hat OpenShift V4.3, the three master nodes become a requirement, and at least two worker nodes must be present in the cluster. Red Hat OpenShift environments with multiple workers often require a distributed shared file system.

In most cases, the application developer does not have any control over which worker in the Red Hat OpenShift cluster the application Pod is dispatched. Therefore, regardless of which worker the application can be deployed to, the persistent storage that is needed by the application must be available.

One supported storage provider for Red Hat OpenShift V4.3 on POWER architecture is an NFS server. PowerVC CSI driver is also supported and can be used for block storage. To use NFS, you must create the NFS server that in this book was done by using Spectrum Scale CES.

For a provisioner, export the NFS to each of the Red Hat OpenShift worker nodes. You can create static persistent volumes and use a different export point for each wanted Red Hat OpenShift persistent volume. However, you do not explicitly mount the exports on the Red Hat OpenShift workers. The mounts are done by Red Hat OpenShift as needed by the containers when they request an associated PVC for the predefined persistent volumes. Also, a persistent volume provisioner for NFS is available and used in our IBM Cloud Pak® for Data scenario.

In Figure 2-4, the solid boxes represent nodes that are required at a minimum to run Red Hat OpenShift Container Platform V4.3 with NFS storage.

> **Note:** The Spectrum Scale that is shown on Figure 2-4 is a representation of a cluster with a minimum of two nodes sharing the volumes. More nodes can be required, depending on the size of the cluster.

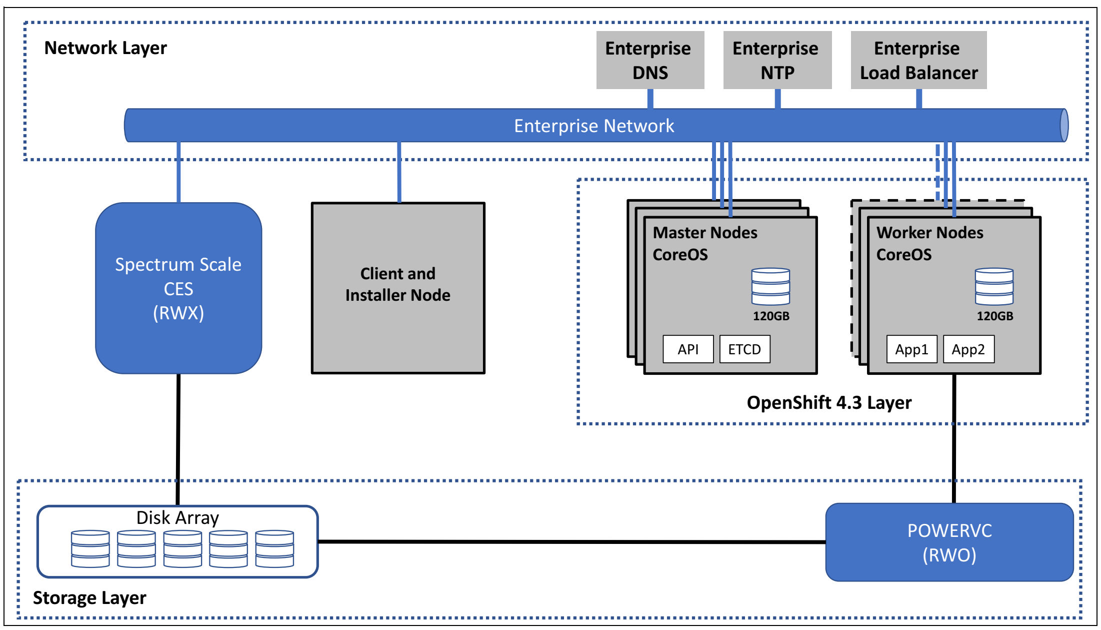

Figure 2-4 Typical Red Hat OpenShift V4.3 deployment

2.3 Supported configurations

This section provides guidance on the supported configurations for installing Red Hat OpenShift V4.3 on Power Systems.

2.3.1 Supported configurations and recommended hardware

Red Hat OpenShift V4.3 on Power Systems can be deployed in the supported configurations as shown in Table 2-8 on page 17. Both PowerVM- and KVM-based Power Systems are supported for Red Hat OpenShift V4.3.

On KVM-based systems, such as the AC922 and IC922, the host operating system is Red Hat Enterprise Linux and CoreOS is the guest operating system. CoreOS as the host operating system is for the bare metal deployment.

Table 2-8 Supported configurations

POWER9 server	Guest OS	Host OS	Cloud mgmt.	File storage	Block storage
E980/E950	CoreOS V4.3	N/A	PowerVC V1.4.4.1	NFS Spectrum Scale	Spectrum Virtualize CSI PowerVC CSI
S9xx/L9xx	CoreOS V4.3	N/A	PowerVC V1.4.4.1	NFS Spectrum Scale	Spectrum Virtualize CSI PowerVC CSI
AC922/IC922	CoreOS V4.3	Red Hat Enterprise Linux V7.x Red Hat Enterprise Linux V8.1 CoreOS V4.3	Bare metal, RHV, OSP 16	NFS Spectrum Scale	Spectrum Virtualize CSI PowerVC CSI

Figure 2-5 on page 18 clarifies the supported host operating system and guest operating system for KVM-based Power Systems. For PowerVM-based systems, such as the enterprise E/S/L systems, no host operating system is used, only a guest operating system. Red Hat Enterprise Linux V7 as a guest operating system is supported in Red Hat OpenShift V3.11 only. Red Hat OpenShift V4.3 requires CoreOS.

Note: At the time of this writing, Spectrum Scale support is limited to deployment as an NFS server. IBM intends to support Spectrum Scale on CoreOS in future releases.

2.3.2 Getting started with Red Hat OpenShift V4.3 on Power Systems with a minimal configuration

This section describes the minimal configuration of Red Hat OpenShift V4.3 on Power Systems to get you started using the container platform. This initial sizing helps you to use a Red Hat OpenShift V4.3 instance (stand-alone) without a large capital investment. This option also helps you to scale to an HA production level deployment in the future.

To deploy in a minimal configuration of Red Hat OpenShift V4.3, you need the following machines:

- A temporary bootstrap machine.
- Three masters or control plane machines.
- Two workers or compute machines.

The relative sizing for these nodes is shown in Table 2-9.

Table 2-9 Minimum configuration

Machine	Operating system	vCPU	RAM	Storage
Bootstrap	Red Hat CoreOS	4	16 GB	120 GB
Control Plane	Red Hat CoreOS	4	16 GB	120 GB
Compute	Red Hat CoreOS	2	8 GB	120 GB

The bootstrap node is required only during the installation step; it is not needed after the cluster is up and running.

Note: Because of this sizing (see Table 2-9 on page 17), the number of required cores depends on the configured SMT level. Remember to account for that number to stand up a minimal instance of Red Hat OpenShift V4.3.

PowerVM, bare metal, and KVM-based Power Systems are supported by Red Hat OpenShift V4.3. Consider the following points:

- On PowerVM-based systems, such as the enterprise E/S/L systems, no host operating system is used and CoreOS is the guest operating system.
- On KVM-based systems, such as the AC922 and IC922, the host operating system is Red Hat Enterprise Linux and CoreOS is the guest operating system.
- CoreOS as the host operating system is for the bare metal deployments.

Figure 2-5 shows an example production level deployment on scale-out (IC922) systems. This configuration attempts to be as cost efficient as possible. The 22 cores 240 GB represents the total amount of compute resources that are available for worker nodes on that system. It must not be interpreted as a single VM with 20 cores and 240 GB allocated; instead, it is to what it can be scaled up.

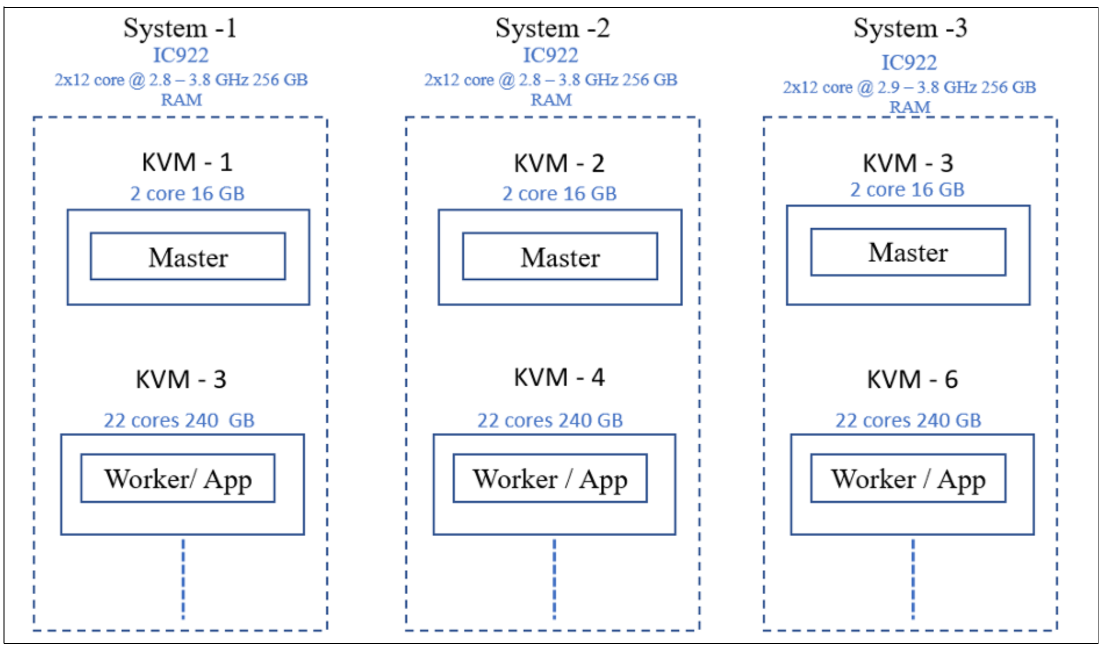

Figure 2-5 Minimal Red Hat OpenShift V4.3 deployment on POWER scale-out architecture

Important: When sizing with an IC922, only four VMs can be supported on a single system for PowerVC GA 1.5.

Figure 2-6 shows an example production level deployment on scale-up (E950) systems.

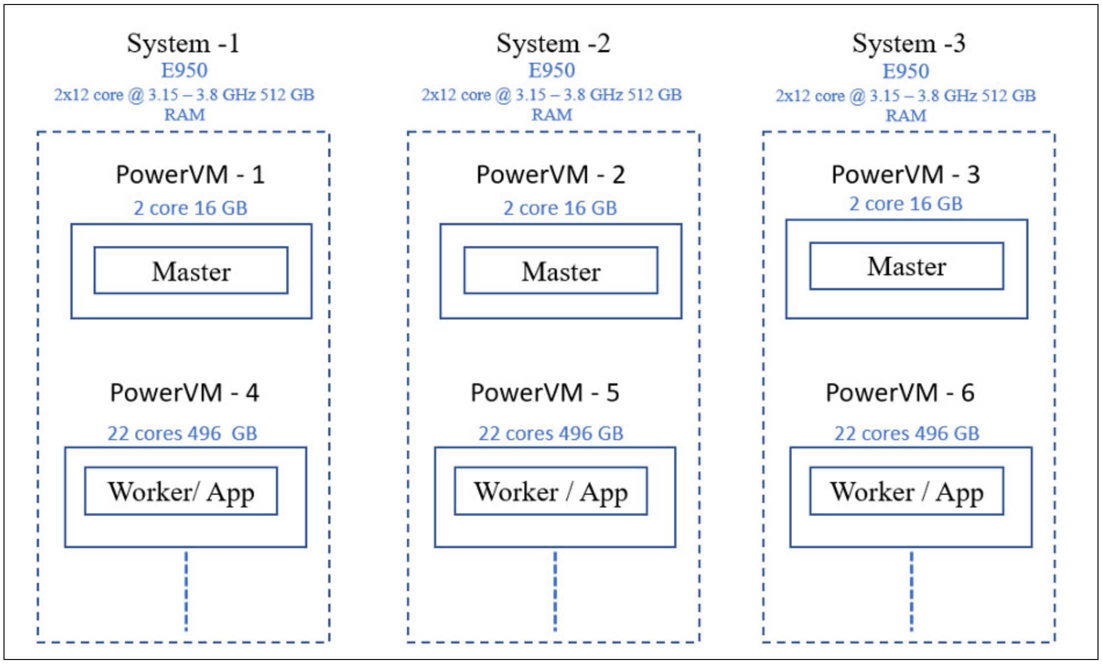

Figure 2-6 Minimal Red Hat OpenShift V4.3 deployment on POWER scale-Up architecture

2.3.3 Installation on restricted networks

With Red Hat OpenShift V4.3, you can perform an installation that does not require an active connection to the internet to obtain software components. You can complete an installation in a restricted network only on an infrastructure that you provision, not with the infrastructure that the installation program provisions; therefore, your platform selection is limited.

To complete a restricted network installation, you must create a registry that mirrors the contents of the Red Hat OpenShift registry and contains the installation media. You can create this mirror on a bastion host, which can access internet and your closed network, or by using other methods that meet your restrictions.

> **Note:** Restricted network installations always use user-provisioned infrastructures. Because of the complexity of the configuration for user-provisioned installations, consider completing a standard user-provisioned infrastructure installation before you attempt a restricted network installation. Completing this test installation can make it easier to isolate and troubleshoot any issues that might arise during your installation in a restricted network.

For more information about supported configurations, see Installing a cluster on IBM Power.

For more information about use cases, see this web page.

2.4 Red Hat OpenShift V4.3 sizing guidelines

This section describes how containers use CPU and the benefits of the use of Power Systems.

Power Systems can have up to eight threads per core, which is a benefit for containers that use CPU that is based on the operating system CPU count, as shown in Example 2-1. You can see that this count is based on the number of threads that are available in the system.

Example 2-1 lscpu output

```
[root@client ~]# lscpu
Architecture:          ppc64le
Byte Order:            Little Endian
CPU(s):                16
On-line CPU(s) list:   0-15
Thread(s) per core:    8
Core(s) per socket:    1
Socket(s):             2
NUMA node(s):          2
Model:                 2.1 (pvr 004b 0201)
Model name:            POWER8 (architected), altivec supported
Hypervisor vendor:     pHyp
Virtualization type:   para
L1d cache:             64K
L1i cache:             32K
L2 cache:              512K
L3 cache:              8192K
NUMA node0 CPU(s):
NUMA node3 CPU(s):     0-3,8-11
```

Four times more containers can be used per core, which maintains the required and limited settings of your YAML files when compared to x86.

For examples of different workloads running on Power Systems two and a half times faster when running eight threads on a PowerVM system when compared to the same number of cores on x86, see the following resources:

- YouTube
- IBM IT Infrastructure web page

Although performance is not a simple 2.5-to-1 comparison and workloads can vary, this section explains how to use this performance advantage.

With Kubernetes, Pods are assigned CPU resources on a CPU thread basis. On deployment and Pod definitions, *CPU* refers to an operating system CPU that maps to a CPU hardware thread.

For an x86 system, an x86 core running with hyperthreading is equivalent to two Kubernetes CPUs. Therefore, when running with x86 hyperthreading, a Kubernetes CPU is equivalent to half of an x86 core.

A PowerVM core can be defined to be 1, 2, 4, or 8 threads with the SMT setting. Therefore, when running on PowerVM with SMT-4, a PowerVM core is equivalent to four Kubernetes CPUs whereas when running with SMT-8, the same PowerVM core is equivalent to eight Kubernetes CPUs. Therefore, when running with SMT-4, a Kubernetes CPU is equivalent to a quarter of a PowerVM core and when running with SMT-8 a Kubernetes CPU is equivalent to one eighth of a PowerVM core.

If your Pod CPU resource was defined to run on x86, you must consider the effects of the POWER's performance advantage and the effects of Kubernetes resources being assigned on a thread basis. For example, for a workload where POWER has a 2X advantage over x86 when running with PowerVM SMT-4, you can assign the same number of Kubernetes CPUs to POWER that you do to x86 to get equivalent performance.

From a hardware perspective, you are assigning the performance of half the number of cores to POWER that you assigned to x86. Whereas for a workload where POWER has a twice the advantage over x86 when running with PowerVM SMT-8, you must assign twice the number of Kubernetes CPUs to POWER that you do to x86 to realize equivalent performance. Although you are assigning twice the number of Kubernetes CPUs to POWER, from a hardware perspective, you are assigning the performance of half the number of cores to POWER that you assigned to x86.

Note: IBM Cloud Pak for Data V3.0.1 defaults to running in SMT-4 mode on POWER. This default can change in the future.

Transforming this abstract concept on core performance that is divided on threads is difficult. A summary is shown in Figure 2-7.

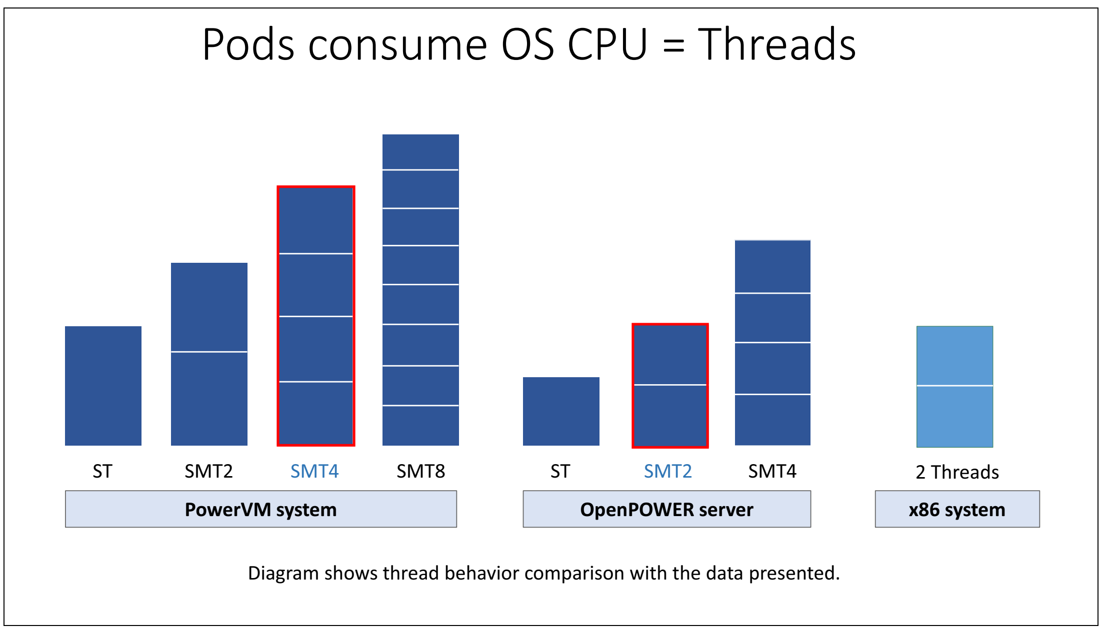

Figure 2-7 Threads on a core

As seen in Figure 2-7, for the workload that is shown, because the PowerVM system can deliver twice the performance of x86 when running with SMT-4, 2 Kubernetes CPUs for PowerVM (which is equivalent to half of a physical core) can deliver the same performance as two Kubernetes CPUs for x86 (which is equivalent to one physical core).

Figure 2-8 shows the container on a Pod limited (throttled) to one CPU (in yellow). You can that see the behavior is different because of the relative performance from the different SMT modes (ppc64le servers).

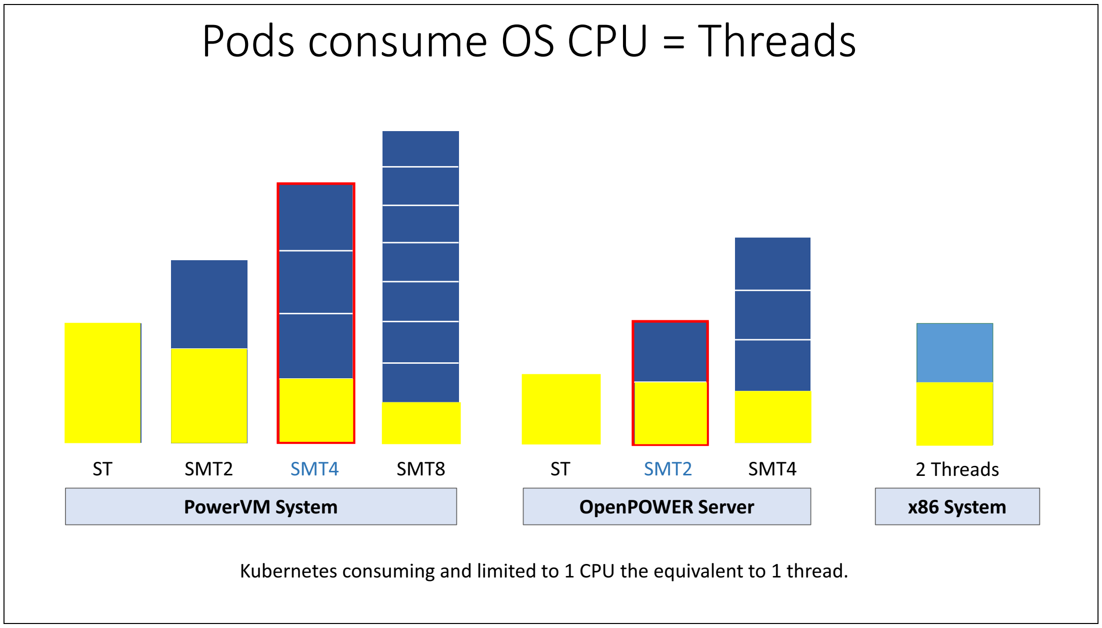

Figure 2-8 Threads on a core, continued

All development performance testing that is done on PowerVM systems use SMT4 and can use half of the cores than x86 that uses the same assumptions. On OpenPower processor systems, all the tests for IBM Cloud Pak for Data are done with SMT2. By using these tests as baseline, we provide directions for sizing guidelines in 2.4.2, “Sizing for IBM Cloud Paks” on page 23.

Appendix A, “Configuring Red Hat CoreOS” on page 73 shows steps for controlling SMT on a node using labels. This label enables different SMT across your cluster for different purposes. You can use the label to select where your application runs depending on needs. By using this method, you can have a *super packing* configuration of nodes, which as nodes that run SMT8 and nodes that run a specific SMT because of application constraints.

2.4.1 Red Hat OpenShift sizing guidelines

Table 2-10 shows the conversion between vCPU, x86 physical cores, and IBM POWER9™ physical cores. Power Systems performance is normally higher than that of x86, but this performance is workload-dependent. For initial sizing, a conservative 1-to-1 mapping of x86 cores to Power Systems Opal only cores and 2-to-1 for Power Systems PowerVM capable were used.

Table 2-10 CPU conversion

	vCPU	x86 Cores (SMT-2)	Physical SMT-2 POWER cores	Physical SMT-4 POWER cores
vCPUs and x86/POWER cores	2	1	1	0.5

First, we highlight the master node resource sizing requirements. The master node sizing requirements depend on the number of worker nodes in the cluster. Table 2-11 shows recommended master node sizing by worker node density.

Table 2-11 Master node initial sizing

Worker nodes	vCPU	Physical SMT-2 POWER cores	Physical SMT-4 POWER cores	Memory (GB)
25	8	4	2	16
100	16	8	4	32
250	32	16	8	64

To determine the number of worker nodes in the cluster, it is important to determine how many Pods are expected to fit per node. This number depends on the application because the application's memory, CPU, and storage requirements must be considered. Red Hat also provided guidelines for the maximum number of Pods per node, which is 250. It is recommended to not exceed this number because it results in lower overall performance.

2.4.2 Sizing for IBM Cloud Paks

This section gives a high-level overview of the sizing guidelines for deploying the various IBM Cloud Paks on Red Hat OpenShift V4.3 on Power Systems. For more information about each IBM Cloud Pak, see Chapter 4, "IBM Cloud Paks" on page 59.

Table 2-12 outlines the number of VMs that are required for deployment and sizing of each VM for IBM Cloud Pak for Multicloud Manager.

Table 2-12 IBM Cloud Pak for Multicloud Manager

Node	Number of VMs	vCPU per node	Memory (GB per node)	Disk (GB per node)
Bastion	1	2	4	100
Master	3	16	32	300
Worker	8	4	16	200
Storage	1	4	16	2x100, 3x200
Total	**13**	**86**	**244**	**3200**

Table 2-13 lists the conversion of the number of vCPUs to physical cores.

Table 2-13 vCPUs to physical cores conversion

vCPU	Physical x86 cores	Physical SMT-2 POWER cores	Physical SMT-4 POWER cores
86	43	43	22

The recommended sizing for installing IBM Cloud Pak for Applications is shown in Table 2-14.

Table 2-14 IBM Cloud Pak for Applications

Node	**Number of VMs**	**vCPU per node**	**Memory (GB per node)**	**Disk (GB per node)**
Master	**3**	**8**	**16**	**200**
Worker	3	8	16	100
Shared Service	1	8	16	100
Total	7	56	128	1000

Table 2-15 lists the conversion of the number of vCPUs to physical cores.

Table 2-15 vCPUs to physical cores conversion

vCPU	**Physical x86 cores**	**Physical SMT-2 POWER cores**	**Physical SMT-4 POWER cores**
56	28	28	14

The recommended sizing for installing IBM Cloud Pak for Integration is listed in Table 2-16.

Table 2-16 IBM Cloud Pak for Integration

Node	**Number of VMs**	**vCPU per node**	**Memory (GB per node)**	**Disk (GB per node)**
Master	3	8	16	200
Worker	8	8	16	200
Storage	1	4	16	2x100, 3x200
Total	**12**	**92**	**192**	**2000**

Table 2-17 lists the conversion of the number of vCPUs to physical cores.

Table 2-17 vCPUs to physical cores conversion

vCPU	**Physical x86 cores**	**Physical SMT-2 POWER cores**	**Physical SMT-4 POWER cores**
100	50	50	25

The recommended sizing for installing the IBM Cloud Pak for Data is shown in Table 2-18. IBM Cloud Pak for Data is the most variable in terms of the sizing because it is highly dependent on the add-ons that are required for a project. The Power Systems release features almost 50 add-on modules for AI (IBM Watson®), Analytics, Dashboarding, Governance, Data Sources, Development Tools, and Storage. Part of this effort must include sizing for add-ons, but for now, the information that is provided Table 2-18 on page 25 is used for a base installation.

Table 2-18 IBM Cloud Pak for Data

Node	Number of VMs	vCPU per node	Memory (GB per node)	Disk (GB per node)
Master	3	8	32	n/a
Worker	3	16	64	200
Storage	1	n/a	n/a	800
Load Balancer	1	8	16	n/a
Total	**8**	**80**	**304**	**1000**

Table 2-19 shows the conversion of the number of vCPUs to physical cores.

Table 2-19 vCPUs to physical cores conversion

vCPU	Physical x86 cores	Physical SMT-2 POWER cores	Physical SMT-4 POWER cores
80	40	40	20

The recommended sizing for installing the IBM Cloud Pak for Automation is shown in Table 2-20.

Table 2-20 IBM Cloud Pak for Automation

Node	Number of VMs	vCPU per node	Memory (GB per node)	Disk (GB per node)
Master	3	8	16	200
Worker	5	8	16	100
Shared Service	1	8	32	400
Total	**7**	**72**	**160**	**1500**

Table 2-21 shows the conversion of the number of vCPUs to physical cores.

Table 2-21 vCPUs to physical cores conversion

vCPU	Physical x86 cores	Physical SMT-2 POWER cores	Physical SMT-4 POWER cores
72	36	36	18

2.5 Storage guidelines

This section discusses the storage options that are available for Red Hat OpenShift V4.3 on Power Systems. Red Hat OpenShift environments with multiple workers often require a distributed shared file system. This requirement stems from the fact that most applications require some sort of persistent store.

In most cases, the developer of the application does not have any control over which worker in the Red Hat OpenShift cluster the application Pod is dispatched. Therefore, regardless of which worker the application can be deployed to, the persistent storage that is required by the application must be available. However, some environments do not require the added complexity that is required to provision a shared, distributed storage environment.

2.5.1 NFS storage

At the time of the writing, the only supported ReadWriteMany (RWX) storage provider for Red Hat OpenShift V4.3 on POWER architecture is an NFS server. You must install the required NFS components on each of the Red Hat OpenShift worker nodes in addition to the node that is designated as the NFS server. For each wanted Red Hat OpenShift persistent volume, a specific NFS export must be created. However, you do not specifically mount the exports on the Red Hat OpenShift workers. The mounts are done by Red Hat OpenShift as needed by the containers when they issue an associated PVC for the predefined persistent volumes. A YAML file that is created to define a static persistent volume based on NFS is shown in Example 2-2.

Example 2-2 Static NFS persistent volume provisioning

```
apiVersion: v1
kind: PersistentVolume
metadata:
  name: nfs-pv01
spec:
  capacity:
    storage: 300Gi
  accessModes:
    - ReadWriteMany
  PersistentVolumeReclainPolicy: Retain
  nfs:
    path:/nfsFileSystem/pv01
    nfsserver.domain.com
    readOnly: False
```

2.5.2 NFS dynamic provisioning

This section uses the provisioner with the deployment files, as seen at this GitHub web page.

You use the `deployment.yaml` file but change it to use the ppc64le image `docker.io/ibmcom/nfs-client-provisioner-ppc64le:latest`, as shown in Example 2-3 on page 27. You must have a preexisting NFS export to use this provisioner.

Change `<NFS_SERVER>` and `<NFS_BASE_PATH>` to match your exported NFS. You can use Spectrum Scale as your NFS server, as described in 2.5.4, “IBM Spectrum Scale and NFS” on page 30.

The use of Spectrum Scale over a kernel NFS includes some advantages; for example, the use of snapshots. This is due to the possibility to use Active File Management for DR and multicloud purposes, and more advanced features that Spectrum Scale provides.

Example 2-3 NFS provisioner deployment file

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: nfs-client-provisioner
  labels:
    app: nfs-client-provisioner
  # replace with namespace where provisioner is deployed
  namespace: default
spec:
  replicas: 1
  strategy:
    type: Recreate
  selector:
    matchLabels:
      app: nfs-client-provisioner
  template:
    metadata:
      labels:
        app: nfs-client-provisioner
    spec:
      serviceAccountName: nfs-client-provisioner
      containers:
        - name: nfs-client-provisioner
          image: docker.io/ibmcom/nfs-client-provisioner-ppc64le:latest
          volumeMounts:
            - name: nfs-client-root
              mountPath: /persistentvolumes
          env:
            - name: PROVISIONER_NAME
              value: fuseim.pri/ifs
            - name: NFS_SERVER
              value: <NFS_SERVER>
            - name: NFS_PATH
              value: <NFS BASE PATH>
      volumes:
        - name: nfs-client-root
          nfs:
            server: <NFS_SERVER>
            path: <NFS BASE PATH>
```

You also need the `rbac.yaml` file, as shown in Example 2-4.

Example 2-4 Role-based access for the provisioner in rbac.yaml file

```
apiVersion: v1
kind: ServiceAccount
metadata:
  name: nfs-client-provisioner
  # replace with namespace where provisioner is deployed
  namespace: default
```

```
---
kind: ClusterRole
apiVersion: rbac.authorization.k8s.io/v1
metadata:
  name: nfs-client-provisioner-runner
rules:
  - apiGroups: [""]
    resources: ["persistentvolumes"]
    verbs: ["get", "list", "watch", "create", "delete"]
  - apiGroups: [""]
    resources: ["persistentvolumeclaims"]
    verbs: ["get", "list", "watch", "update"]
  - apiGroups: ["storage.k8s.io"]
    resources: ["storageclasses"]
    verbs: ["get", "list", "watch"]
  - apiGroups: [""]
    resources: ["events"]
    verbs: ["create", "update", "patch"]
---
kind: ClusterRoleBinding
apiVersion: rbac.authorization.k8s.io/v1
metadata:
  name: run-nfs-client-provisioner
subjects:
  - kind: ServiceAccount
    name: nfs-client-provisioner
    # replace with namespace where provisioner is deployed
    namespace: default
roleRef:
  kind: ClusterRole
  name: nfs-client-provisioner-runner
  apiGroup: rbac.authorization.k8s.io
---
kind: Role
apiVersion: rbac.authorization.k8s.io/v1
metadata:
  name: leader-locking-nfs-client-provisioner
  # replace with namespace where provisioner is deployed
  namespace: default
rules:
  - apiGroups: [""]
    resources: ["endpoints"]
    verbs: ["get", "list", "watch", "create", "update", "patch"]
---
kind: RoleBinding
apiVersion: rbac.authorization.k8s.io/v1
metadata:
  name: leader-locking-nfs-client-provisioner
  # replace with namespace where provisioner is deployed
  namespace: default
subjects:
  - kind: ServiceAccount
    name: nfs-client-provisioner
    # replace with namespace where provisioner is deployed
    namespace: default
```

```
roleRef:
  kind: Role
  name: leader-locking-nfs-client-provisioner
  apiGroup: rbac.authorization.k8s.io
```

For the storage class creation, use the `class.yaml` file, as shown in Example 2-5.

Example 2-5 Storage class definition class.yaml

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: managed-nfs-storage
provisioner: fuseim.pri/ifs # or choose another name, must match deployment's env
PROVISIONER_NAME'
parameters:
  archiveOnDelete: "false"
```

To define the provisioner, apply the `rbac` file, security constrains, deployment, and class, as shown in Example 2-6.

> **Note:** We are defining the provisioner on the default project. We recommend creating a new project and define the provisioner. Remember to change all namespaces from default to your namespace. Also, remember to change from default to your namespace on the **`add-scc-to-user`** command, as shown in Example 2-6.

Example 2-6 Defining the provisioner

```
[root@client ~]# oc apply -f rbac.yaml
serviceaccount/nfs-client-provisioner created
clusterrole.rbac.authorization.k8s.io/nfs-client-provisioner-runner created
clusterrolebinding.rbac.authorization.k8s.io/run-nfs-client-provisioner created
role.rbac.authorization.k8s.io/leader-locking-nfs-client-provisioner created
rolebinding.rbac.authorization.k8s.io/leader-locking-nfs-client-provisioner
created
[root@client ~]# oc adm policy add-scc-to-user hostmount-anyuid
system:serviceaccount:default:nfs-client-provisioner
securitycontextconstraints.security.openshift.io/hostmount-anyuid added to:
["system:serviceaccount:default:nfs-client-provisioner"]
[root@client ~]# oc apply -f deployment.yaml
deployment.apps/nfs-client-provisioner created
[root@client ~]# oc apply -f class.yaml
storageclass.storage.k8s.io/managed-nfs-storage created
[root@client ~]#
```

You can use the provisioner now.

Guidelines for IBM Cloud Pak for Data storage performance

IBM Cloud Pak for Data relies on dynamic persistent volumes that are provisioned during installation. The method that is described in 2.5.2, “NFS dynamic provisioning” on page 26 is used in this book as the base for IBM Cloud Pak for Data installation.

To ensure correct behavior of your system, check that your NFS exports meet latency and bandwidth specifications. To test latency, check that your result is comparable or better than the result, as shown in Example 2-7.

Example 2-7 Test latency for IBM Cloud Pak for Data

```
[root@client ~]# dd if=/dev/zero of=/mnt/testfile bs=4096 count=1000 oflag=dsync
1000+0 records in
1000+0 records out
4096000 bytes (4.1 MB, 3.9 MiB) copied, 1.5625 s, 2.5 MB/s
```

To test bandwidth, check that your result is comparable or better than the result, as shown in Example 2-8.

Example 2-8 Test bandwidth for IBM Cloud Pak for Data

```
[root@client ~]# dd if=/dev/zero of=/mnt/testfile bs=1G count=1 oflag=dsync
1+0 records in
1+0 records out
1073741824 bytes (1.1 GB) copied, 5.14444 s, 209 MB/s
```

2.5.3 IBM Spectrum Scale

IBM Spectrum Scale is a cluster file system that provides concurrent access to a single file system or set of file systems from multiple nodes. The nodes can be SAN-attached, network-attached, a mixture of SAN-attached and network-attached, or in a shared nothing cluster configuration. This file system enables high-performance access to this common set of data to support a scale-out solution or to provide a high availability platform.

> **Note:** Native Spectrum Scale support for CoreOS is intended for all architectures, s390x, ppc64le and x86_64, and can be a great option for persistent storage.

IBM Spectrum Scale has many features beyond common data access, including data replication, policy-based storage management, and multi-site operations. You can create a cluster of AIX nodes, Linux nodes, Windows server nodes, or a mix of all three. IBM Spectrum Scale can run on virtualized instances, which provide common data access in environments to take advantage of logical partitioning or other hypervisors. Multiple IBM Spectrum Scale clusters can share data within a location or across wide area network (WAN) connections.

For more information about the benefits and features of IBM Spectrum Scale, see IBM Knowledge Center.

2.5.4 IBM Spectrum Scale and NFS

IBM Spectrum Scale provides extra protocol access methods. Providing these extra file and object access methods and integrating them with Spectrum Scale offers several benefits:

- Enables users to consolidate various sources of data efficiently in one global namespace
- Provides a unified data management solution
- Enables efficient space utilization
- Avoids the need to move data unnecessarily because access methods can be different

Protocol access methods that are integrated with Spectrum Scale are NFS, SMB, and Object. Although each of these server functions (NFS, SMB, and Object) uses open source technologies, this integration adds value by providing scaling and high availability by using the clustering technology in Spectrum Scale. The NFS support for IBM Spectrum Scale enables clients to access the Spectrum Scale file system by using NFS clients with their inherent NFS semantics.

For more information about setting up IBM Spectrum Scale as an NFS server, see IBM Knowledge Center.

2.5.5 PowerVC Container Storage Interface driver

The Container Storage Interface (CSI) allows Red Hat OpenShift to use storage from storage backends that implement the CSI interface as persistent storage. The PowerVC CSI driver is a standard for providing storage ReadWriteOnce (RWO) functionality to containers. The PowerVC CSI pluggable driver interacts with PowerVC storage for operations, such as create volumes, delete volumes, and attach or detach volumes.

For more information about the Container Storage Interface in Red Hat OpenShift, see this web page.

For more information about the PowerVC CSI and how to configure it, see IBM Knowledge Center.

2.5.6 Backing up and restoring your Red Hat OpenShift cluster and applications

Many parts of the cluster must be backed up so that it can be restored. This section describes that process at a high level.

Backing up the etcd to restore the cluster state

The state of the cluster can be backed up to the master node. This backup can be used to restore the cluster to this previous state if anything happens to the cluster.

For more information about the procedure to backup the `etcd`, see this web page.

Check the two files that are generated are backed up in a safe place outside of the cluster so they can be retrieved and used if needed.

For more information about the procedure to restore to a previous state, see this web page.

Backing up application consistent persistent volumes

The backup of the persistent volumes is highly dependent on the application being used. Each application likely has a way to correctly backup the data that is on the persistent volume.

Some databases, such as DB2, have their own backup tool (**`db2 backup`**) to create online backups. This tool is also used on the containers. For more information, see IBM Knowledge Center.

If you are using NFS or Spectrum Scale, you can use any other backup tool to backup the directory after the backup completed successfully. Check your backup to a PV that is available externally and has enough space to perform the operation.

Other type of applications that are easily scaled down and scaled up can be scaled down to replicas=0, take a snapshot of the volume (can be a fileset if you are using Spectrum Scale for example), and return the applications back to their former replica number.

Some applications (for example, IBM Cloud Pak for Data) provide tools to quiesce the workload (**`cpdbr quiesce`** and **`cpdbr unquiesce`**). You can also combine this tool to create a backup strategy for your PVC.

Some applications do not need to maintain consistency, and a backup of the PV files is enough.

Backing up crash consistent persistent volumes

If you are using Spectrum Scale, you can always snapshot the file system and back it up.

> **Note:** This type of backup is not advisable if you intend to always maintain an application-consistent method of restore.

Some applications can fail to start this method because they need consistency.

3

Reference installation guide for Red Hat OpenShift V4.3 on Power Systems servers

This chapter describes the installation procedure and dependencies to install Red Hat OpenShift Container Platform V4.3 on Power Systems servers. The procedures that are described here are based on the Red Hat install documentation that uses a user provisioned infrastructure that most enterprise system administrators are used to, and can use to provision to their developers and application environments.

This chapter includes the following topics:

3.1 Introduction

This chapter focuses on the user-provided infrastructure method of installation that targets enterprise customers that have their network infrastructure to provide load balancing, firewall, and DNS.

On all architectures, the control plane runs only Red Hat CoreOS and does not permit the use of Red Hat Linux for this role, unlike the previous Red Hat OpenShift V3.x release. The installation of this operating system depends on a file that is called ignition that acts on the configuration of the CoreOS. Because CoreOS is an immutable operating system, no rpm package management is needed, as you expect on a Red Hat Enterprise Linux operating system.

The minimum number of nodes that are required for a production Red Hat OpenShift cluster are three master nodes and two worker nodes. The only option for worker nodes on ppc64le environment is CoreOS; therefore, you have a flat environment with CoreOS on master and worker nodes. The collection of the master nodes is called *control plane*. A bootstrap node is needed to bring up the control plane, but it is destroyed after the control plane is fully up.

For more information, see web page.

This chapter highlights the differences and uses the BOOTP for the boot process instead of PXE boot as suggested in the documentation for PowerVM LPARs. However, we use the defined PXE boot for bare metal installation.

The minimum resource requirement to start a Red Hat OpenShift V4.3 cluster on Power Systems servers is listed in Table 3-1.

Table 3-1 Minimum resource requirements

Machine	Operating system	vCPU	Virtual RAM	Storage
Bootstrap	Red Hat CoreOS	4	16 GB	120 GB
Control plane	N/A	2	16 GB	120 GB
Compute	N/A	2	16 GB	120 GB

3.2 Red Hat OpenShift V4.3 deployment with internet connection stand-alone installation

This section explains how to use the network boot method to translate the PXE boot that is in the official document to BOOTP installation for PowerVM LPARs. Most customers use BOOTP on NIM.

This section also describes how to create a Linux environment to segregate the NIM server from the Red Hat OpenShift installation server. You can also unify them and use an existing NIM server to be your network installation infrastructure (BOOTP and TFTP), as described in 3.3, “Using NIM as your BOOTP infrastructure” on page 50.

3.2.1 PowerVM configuration for network installation

The client process to install CoreOS by using the network boot is the same as in any AIX NIM installation. For example, this installation covers the installation of Red Hat OpenShift Container Platform V4.3 across three IBM POWER8® or POWER9 L922 with enough hardware for a starter kit for a basic Spectrum Scale environment. Figure 3-1 shows the network connections for a minimal workload with Spectrum Scale as the storage backend for Read-Write-Many (RWX) persistent volumes.

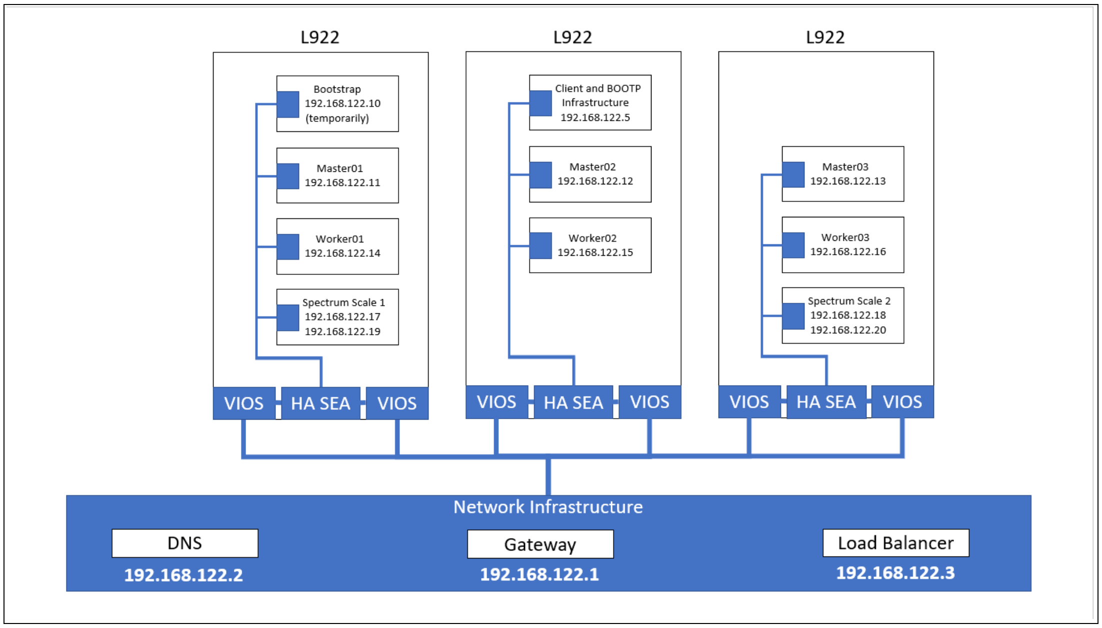

Figure 3-1 PowerVM network architecture example

This test environment uses other servers to act as the DNS and load balancer, considering we do not control the enterprise network resources, only to the Power Systems servers. Because it is a single point of failure, do not use this method on a production environment. Instead, use the highly available network services that are provided by the enterprise networking, infrastructure, and security teams.

The DHCP/BOOTP in this case is another single server. This server works for production environments if the nodes are installed with fixed IP because the server is used during installation only. The installation server can be your NIM server if you have one.

3.2.2 BOOTP infrastructure server for installing CoreOS by using the network

During the initial boot, the machines require a DHCP/TFTP server or a static IP address to be set to establish a network connection to download their ignition configuration files. For this installation, all our nodes feature direct internet access, and include installed DHCP/TFTP on the client server to serve IP addresses. We are also serving HTTP and for this environment, the IP address of this server is 192.168.122.5. We used a ppc64le Red Hat Enterprise Linux V7.7 for simplicity of configuration of the grub network environment.

This document does not describe in great deal setting up a boot environment. For more information about how this process is done, see Red Hat's *Configuring Network Boot on IBM Power Systems Using GRUB2*.

Complete the following steps to set up a boot environment:

1. Install a DHCP server, a TFTP server, and an HTTP server by using YUM.
2. Follow the procedure that is shown in Example 3-1 to create a GRUB2 network boot directory inside the tftp root.

Example 3-1 Creating netboot directory

```
[root@jinete ~]# grub2-mknetdir --net-directory=/var/lib/tftpboot
Netboot directory for powerpc-ieee1275 created. Configure your DHCP server to
point to /var/lib/tftpboot/boot/grub2/powerpc-ieee1275/core.elf
```

3. Configure the grub load for each server on `/var/lib/tftpboot/boot/grub2/grub.cfg`. Example 3-2 shows entries for `/var/lib/tftpboot/boot/grub2/grub.cfg`.

Example 3-2 Example for each entry on /var/lib/tftpboot/boot/grub2/grub.cfg

```
if [ ${net_default_mac} == 52:54:00:af:db:b6 ]; then
default=0
fallback=1
timeout=1
menuentry "Bootstrap CoreOS (BIOS)" {
echo "Loading kernel Bootstrap"
linux "/rhcos-4.3.18-ppc64le-installer-kernel-ppc64le" rd.neednet=1
ip=192.168.122.10::192.168.122.1:255.255.255.0:bootstrap.ocp4.ibm.lab:env2:none
nameserver=192.168.122.2 console=tty0 console=ttyS0 coreos.inst=yes
coreos.inst.install_dev=vda
coreos.inst.image_url=http://192.168.122.5:8080/rhcos-4.3.18-ppc64le-metal.ppc64le
.raw.gz coreos.inst.ignition_url=http://192.168.122.5:8080/bootstrap.ign
echo "Loading initrd"
initrd  "/rhcos-4.3.18-ppc64le-installer-initramfs.ppc64le.img"
}
fi

if [ ${net_default_mac} == 52:54:00:02:23:c7 ]; then
default=0
fallback=1
timeout=1
menuentry "Master1 CoreOS (BIOS)" {
echo "Loading kernel Master1"
linux "/rhcos-4.3.18-ppc64le-installer-kernel-ppc64le" rd.neednet=1
ip=192.168.122.11::192.168.122.1:255.255.255.0:master1.ocp4.ibm.lab:env2:none
nameserver=192.168.122.2 console=tty0 console=ttyS0 coreos.inst=yes
coreos.inst.install_dev=vda
coreos.inst.image_url=http://192.168.122.5:8080/rhcos-4.3.18-ppc64le-metal.ppc64le
.raw.gz coreos.inst.ignition_url=http://192.168.122.5:8080/master.ign
echo "Loading initrd"
initrd  "/rhcos-4.3.18-ppc64le-installer-initramfs.ppc64le.img"
}
fi

if [ ${net_default_mac} == 52:54:00:68:5c:7c ]; then
default=0
fallback=1
timeout=1
menuentry "Worker1 CoreOS (BIOS)" {
echo "Loading kernel Worker1"
```

```
linux "/rhcos-4.3.18-ppc64le-installer-kernel-ppc64le" rd.neednet=1
ip=192.168.122.14::192.168.122.1:255.255.255.0:worker1.ocp4.ibm.lab:env2:none
nameserver=192.168.122.2 console=tty0 console=ttyS0 coreos.inst=yes
coreos.inst.install_dev=vda
coreos.inst.image_url=http://192.168.122.5:8080/rhcos-4.3.18-ppc64le-metal.ppc64le
.raw.gz coreos.inst.ignition_url=http://192.168.122.5:8080/worker.ign
echo "Loading initrd"
initrd  "/rhcos-4.3.18-ppc64le-installer-initramfs.ppc64le.img"
}
fi
```

4. Example 3-2 on page 36 shows the installation image and the ignition file point to an HTTP server. Use `httpd` and change the port to run on 8080. Leave in the directory the default `/var/www/html`.

> **Note:** Remember to repeat the entries for all nodes because these entries are only examples for each node type to bootstrap the cluster. You must at least three masters and two worker nodes. In our case, we have three masters and two worker nodes. The `env2` entry on the IP parameter is the one in our environment and changes with the virtual ID you assign to it on the LPAR.

   Example 3-2 on page 36 shows some referenced files. These installation files are available to download from this web page.

   You find the client, the installer, and the CoreOS assets. Remember to retrieve all of the referenced files in `grub.cfg: rhcos*metal.ppc64le.raw.gz`, `rhcos*initramfs.ppc64le.img`, and `rhcos*kernel-ppc64le`.

   The `rhcos*metal.ppc64le.raw.gz` file must be placed on the root directory of the HTTP server, and `rhcos*initramfs.ppc64le.img` and `rhcos*kernel-ppc64le` must be on the TFTP root directory.

5. Install `dhcpd` and configure `dhcp.conf` to match your configuration, as shown in Example 3-3.

Example 3-3 Contents of /etc/dhcp/dhcp.conf

```
default-lease-time 900;
max-lease-time 7200;
subnet 192.168.122.0 netmask 255.255.255.0 {
        option routers 192.168.122.1;
        option subnet-mask 255.255.255.0;
        option domain-search "ocp4.ibm.lab";
        option domain-name-servers 192.168.122.2;
        next-server 192.168.122.5;
        filename "boot/grub2/powerpc-ieee1275/core.elf";
}

allow bootp;
option conf-file code 209 = text;
host bootstrap {
        hardware ethernet 52:54:00:af:db:b6;
        fixed-address 192.168.122.10;
        option host-name "bootstrap.ocp4.ibm.lab";
        allow booting;
}
host master1 {
```

```
        hardware ethernet 52:54:00:02:23:c7;
        fixed-address 192.168.122.11;
        option host-name "master1.ocp4.ibm.lab";
        allow booting;
}
host master2 {
        hardware ethernet 52:54:00:06:c2:ee;
        fixed-address 192.168.122.12;
        option host-name "master2.ocp4.ibm.lab";
        allow booting;
}
host master3 {
        hardware ethernet 52:54:00:df:15:3e;
        fixed-address 192.168.122.13;
        option host-name "master3.ocp4.ibm.lab";
        allow booting;
}
host worker1 {
        hardware ethernet 52:54:00:07:b4:ec;
        fixed-address 192.168.122.14;
        option host-name "worker1.ocp4.ibm.lab";
        allow booting;
}
host worker2 {
        hardware ethernet 52:54:00:68:5c:7c;
        fixed-address 192.168.122.15;
        option host-name "worker2.ocp4.ibm.lab";
        allow booting;
}

host worker3 {
        hardware ethernet 52:54:00:68:ac:7d;
        fixed-address 192.168.122.16;
        option host-name "worker3.ocp4.ibm.lab";
        allow booting;
}
```

The MAC address must match `grub.cfg` and `dhcpd.conf` files for a correct installation.

> **Note:** The MAC addresses that are used differ because they are normally randomly generated upon LPAR definition. You can use your preferred method to define the LPARs and collect the MAC address from the configuration to use on these files.

3.2.3 Network infrastructure prerequisites for Red Hat OpenShift Container Platform installation in a production environment

The network prerequisites can be found in the Red Hat OpenShift Container Platform installation documentation on Power Systems at this web page.

Note: Use highly available enterprise services that available in your infrastructure to provide the services that are described in this section (load balancer, DNS, and network connectivity). This process is normally done before the installation by the enterprise networking, security, and active directory teams.

This section shows copies of the documentation tables for easier reference. Check for updates in the source documentation.

3.2.4 Creating a network infrastructure for a test environment

If you want to perform a test or get acquainted with Red Hat OpenShift without creating a full network infrastructure for it, follow the procedures to create the prerequisites if permitted on your network.

Important: Do not follow this procedure for production environments unless you created a full network infrastructure.

Provisioning DNS with dnsmasq on test environments

Table 3-2 is from the Red Hat OpenShift documentation and is included here for reference. It lists that needed DNS entries that are used by the cluster nodes and Pods.

Table 3-2 DNS entries

Component	Record	Description
Kubernetes external API	api.<cluster_name>.<base_domain>.	This DNS A/AAAA or CNAME record must point to the load balancer for the control plane machines. This record must be resolvable by both clients external to the cluster and from all the nodes within the cluster.
Kubernetes internal API	api-int.<cluster_name>.<base_domain>.	This DNS A/AAAA or CNAME record must point to the load balancer for the control plane machines. This record must be resolvable from all the nodes within the cluster. The API server must be able to resolve the worker nodes by the host names that are recorded in Kubernetes. If it cannot resolve the node names, proxied API calls can fail, and you cannot retrieve logs from Pods.

Component	Record	Description
Routes	*.apps.<cluster_name>.<base_domain>.	A wildcard DNS A/AAAA or CNAME record that points to the load balancer that targets the machines that run the Ingress router Pods, which are the worker nodes by default. This record must be resolvable by both clients external to the cluster and from all the nodes within the cluster.
etcd Name Record	etcd-<index>.<cluster_name>.<base_domain>.	Red Hat OpenShift Container Platform requires DNS A/AAAA records for each etcd instance to point to the control plane machines that host the instances. The etcd instances are differentiated by <index> values, which start with 0 and end with n-1, where n is the number of control plane machines in the cluster. The DNS record must resolve to a unicast IPv4 address for the control plane machine, and the records must be resolvable from all the nodes in the cluster.
etcd Service Record	_etcd-server-ssl._tcp.<cluster_name>.<base_domain>.	For each control plane machine, Red Hat OpenShift Container Platform also requires an SRV DNS record for etcd server on that machine with priority 0, weight 10, and port 2380. A cluster that uses three control plane machines requires the following records: # _service._proto.name.TTL class SRV priority weight port target. _etcd-server-ssl._tcp.<cluster_name>.<base_domain>. 86400 IN SRV 0 10 2380 etcd-0.<cluster_name>.<base_domain> _etcd-server-ssl._tcp.<cluster_name>.<base_domain>. 86400 IN SRV 0 10 2380 etcd-1.<cluster_name>.<base_domain> _etcd-server-ssl._tcp.<cluster_name>.<base_domain>. 86400 IN SRV 0 10 2380 etcd-2.<cluster_name>.<base_domain>

To provision the DNS, enter all the details as shown in Table 3-2 on page 39 in your DNS. Ensure it forwards the requests to the network DNS servers (in this case, represented by 10.124.0.1 and 10.124.0.2). In this way, all other addresses can also be resolved when pulling containers from the registries. Example 3-4 shows our `dnsmasq` configuration file.

Example 3-4 dnsmasq configuration file

```
address=/client.ocp4.ibm.lab/192.168.122.5

address=/bootstrap.ocp4.ibm.lab/192.168.122.10
ptr-record=10.122.168.192.in-addr.arpa,bootstrap.ocp4.ibm.lab

address=/master1.ocp4.ibm.lab/192.168.122.11
address=/etcd-0.ocp4.ibm.lab/192.168.122.11
srv-host=_etcd-server-ssl._tcp.ocp4.ibm.lab,etcd-0.ocp4.ibm.lab,2380
ptr-record=11.122.168.192.in-addr.arpa,master1.ocp4.ibm.lab

address=/master2.ocp4.ibm.lab/192.168.122.12
address=/etcd-1.ocp4.ibm.lab/192.168.122.12
srv-host=_etcd-server-ssl._tcp.ocp4.ibm.lab,etcd-1.ocp4.ibm.lab,2380
ptr-record=12.122.168.192.in-addr.arpa,master2.ocp4.ibm.lab

address=/master3.ocp4.ibm.lab/192.168.122.13
address=/etcd-2.ocp4.ibm.lab/192.168.122.13
srv-host=_etcd-server-ssl._tcp.ocp4.ibm.lab,etcd-2.ocp4.ibm.lab,2380
ptr-record=13.122.168.192.in-addr.arpa,master3.ocp4.ibm.lab

address=/worker1.ocp4.ibm.lab/192.168.122.14
ptr-record=14.122.168.192.in-addr.arpa,worker1.ocp4.ibm.lab

address=/worker2.ocp4.ibm.lab/192.168.122.15
ptr-record=15.122.168.192.in-addr.arpa,worker2.ocp4.ibm.lab

address=/worker3.ocp4.ibm.lab/192.168.122.16
ptr-record=16.122.168.192.in-addr.arpa,worker3.ocp4.ibm.lab

address=/api.ocp4.ibm.lab/192.168.122.3
address=/api-int.ocp4.ibm.lab/192.168.122.3
address=/.apps.ocp4.ibm.lab/192.168.122.3

# Listen on lo and env2 only
bind-interfaces
interface=lo,env2
server=10.124.0.1
server=10.124.0.2
```

Provision load balancing with haproxy on test environments

The next prerequisite is the load balancing. We use haproxy to meet this prerequisite for the test environments, as shown in Table 3-3 on page 42.

Table 3-3 Load balancing entries requirement

Port	Machines	Internal	External	Description
6443	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane.	X	X	Kubernetes API server.
22623	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane.	X		Machine config server.
443	The machines that run the Ingress router Pods, compute, or worker, by default.	X	X	HTTPS traffic.
80	The machines that run the Ingress router Pods, compute, or worker, by default.	X	X	HTTP traffic.

Example 3-5 shows the configuration file of what the installation documentation describes when haproxy is used to implement it.

Example 3-5 haproxy configuration file

```
global
    log         127.0.0.1 local2

    chroot      /var/lib/haproxy
    pidfile     /var/run/haproxy.pid
    maxconn     4000
    user        haproxy
    group       haproxy
    daemon

    # turn on stats unix socket
```

```
    stats socket /var/lib/haproxy/stats
defaults
    mode                    http
    log                     global
    option                  httplog
    option                  dontlognull
    option http-server-close
    option forwardfor       except 127.0.0.0/8
    option                  redispatch
    retries                 3
    timeout http-request    10s
    timeout queue           1m
    timeout connect         10s
    timeout client          1m
    timeout server          1m
    timeout http-keep-alive 10s
    timeout check           10s
    maxconn                 3000

frontend openshift-api
    bind *:6443
    default_backend openshift-api
    mode tcp
    option tcplog
backend openshift-api
    balance source
    mode tcp
    server ocp43-bootstrap 192.168.0.10:6443 check
    server ocp43-master01 192.168.0.11:6443 check
    server ocp43-master02 192.168.0.12:6443 check
    server ocp43-master03 192.168.0.13:6443 check

frontend openshift-configserver
    bind *:22623
    default_backend openshift-configserver
    mode tcp
    option tcplog
backend openshift-configserver
    balance source
    mode tcp
    server ocp43-bootstrap 192.168.0.10:22623 check
    server ocp43-master01 192.168.0.11:22623 check
    server ocp43-master02 192.168.0.12:22623 check
    server ocp43-master03 192.168.0.13:22623 check

frontend openshift-http
    bind *:80
    default_backend openshift-http
    mode tcp
    option tcplog
backend openshift-http
    balance source
    mode tcp
    server ocp43-worker01 192.168.0.14:80 check
    server ocp43-worker02 192.168.0.15:80 check
```

```
    server ocp43-worker02 192.168.0.16:80 check
frontend openshift-https
    bind *:443
    default_backend openshift-https
    mode tcp
    option tcplog
backend openshift-https
    balance source
    mode tcp
    server ocp43-worker01 192.168.0.14:443 check
    server ocp43-worker02 192.168.0.15:443 check
    server ocp43-worker03 192.168.0.16:443 check
```

After your configuration is complete, confirm that all services are started. Be aware that you might need to change SELinux Boolean configurations to get haproxy to serve on any port.

3.2.5 Installing Red Hat OpenShift on PowerVM LPARs by using the BOOTP network installation

This section uses the BOOTP process to install CoreOS. Red Hat OpenShift Container Platform is installed with all the configurations to be passed to CoreOS installation with the bootstrap ignition file.

Check you downloaded the Red Hat OpenShift installation package, the Red Hat OpenShift client, and the binaries are on your path (you must decompress the packages you download). Also, confirm that you downloaded the CoreOS assets to build your BOOTP infrastructure and the files are correctly placed as directed. On this same page, you find your pull secret that you need to configure your `install-config.yaml` file.

The pull secret is tied to your account and you can use your licenses. At the time of this writing, we were tied to a 60-day trial. If you intend to maintain the cluster for more than 60 days, check that you have a valid subscription.

Complete the following steps to install a Power Systems cluster until the point you have your YAML file ready. Do not create the ignition files at this moment. Create the sample file as shown at this web page.

Our YAML file is shown in Example 3-6.

Example 3-6   install-config.yaml file

```
apiVersion: v1
baseDomain: ibm.lab
compute:
- hyperthreading: Enabled
  name: worker
  replicas: 0
controlPlane:
  hyperthreading: Enabled
  name: master
  replicas: 3
metadata:
  name: ocp4
networking:
  clusterNetwork:
```

```
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  networkType: OpenShiftSDN
  serviceNetwork:
  - 172.30.0.0/16
platform:
  none: {}
fips: false
pullSecret: '<PULL SECRET HERE>'
sshKey: '<SSH PUBLIC KEY SECRET HERE>'
```

If you need proxy to access the internet, complete the process that is described at this web page.

The creation of the manifests that are a set of YAML files that are used to create the ignition files that configure the CoreOS installation. The creation of the manifests is done by using the YAML file that you prepared (see Example 3-6 on page 44), and the **openshift-install** command that was downloaded from this web page.

Place the `openshift-installer` and `install-config.yaml` files into a single directory and create a backup of the `install-config.yaml` because it is deleted after use. For any installation, these three files are the only files in the directory. Our installation directory is shown in Example 3-7.

Example 3-7 Install directory

```
[root@client install]# ls
install-config.yaml install-config.yaml.bak openshift-install
[root@client install]#
```

After you prepare the install directory, run the manifest creation, as shown in Example 3-8.

Example 3-8 Manifest creation

```
[root@client install]# ./openshift-install create manifests
INFO Consuming Install Config from target directory
WARNING Making control-plane schedulable by setting MastersSchedulable to true for
Scheduler cluster settings
[root@client install]# ls -la
total 318028
drwxr-xr-x.  4 root root       163 Jun  1 08:26 .
drwxrwxrwt. 11 root root      4096 Jun  1 08:18 ..
-rw-r--r--.  1 root root      3549 Jun  1 08:18 install-config.yaml.bak
drwxr-x---.  2 root root      4096 Jun  1 08:26 manifests
drwxr-x---.  2 root root      4096 Jun  1 08:26 openshift
-rwxr-xr-x.  1 root root 325462976 Jun  1 08:21 openshift-install
-rw-r--r--.  1 root root     23999 Jun  1 08:26 .openshift_install.log
-rw-r-----.  1 root root    153751 Jun  1 08:26 .openshift_install_state.json
```

As shown in Example 3-8, a structure is created. Per the documentation, change the manifests/`cluster-scheduler-02-config.yml` files to mark the masters not schedulable, as shown in Example 3-9 on page 46.

Example 3-9 Change cluster-scheduler-02-config.yml

```
apiVersion: config.openshift.io/v1
kind: Scheduler
metadata:
  creationTimestamp: null
  name: cluster
spec:
  mastersSchedulable: false
  policy:
    name: ""
status: {}
```

Now, create the ignition files, as shown in Example 3-10.

Example 3-10 Creating ignition files

```
[root@client install]# ./openshift-install create ignition-configs
INFO Consuming OpenShift Install (Manifests) from target directory
INFO Consuming Worker Machines from target directory
INFO Consuming Openshift Manifests from target directory
INFO Consuming Common Manifests from target directory
INFO Consuming Master Machines from target directory
[root@client install]# ls -la
total 319416
drwxr-xr-x.  3 root root       219 Jun  1 12:20 .
drwxrwxrwt. 11 root root      4096 Jun  1 08:18 ..
drwxr-x---.  2 root root        50 Jun  1 12:20 auth
-rw-r-----.  1 root root    293828 Jun  1 12:20 bootstrap.ign
-rw-r--r--.  1 root root      3549 Jun  1 08:18 install-config.yaml.bak
-rw-r-----.  1 root root      1829 Jun  1 12:20 master.ign
-rw-r-----.  1 root root        96 Jun  1 12:20 metadata.json
-rwxr-xr-x.  1 root root 325462976 Jun  1 08:21 openshift-install
-rw-r--r--.  1 root root     73018 Jun  1 12:20 .openshift_install.log
-rw-r-----.  1 root root   1226854 Jun  1 12:20 .openshift_install_state.json
-rw-r-----.  1 root root      1829 Jun  1 12:20 worker.ign
```

Copy the ignition files (`bootstrap.ign`, `master.ign`, and `worker.ign`) to your httpd file server.

> **Note:** You need the files inside the auth directory to access your cluster. Copy the file `auth/kubeconfig to /root/.kube/config`. Find the password for the console inside the `kubeadmin-password` file.

Verifying all components to install your cluster

Complete the following steps to verify that all components are ready to install your cluster:

1. Confirm that the web server is up on the correct port (for example, 8080).
2. Check that the CoreOS raw image and the ignition files include the correct permissions.
3. Check that TFTP is started.
4. Confirm that the initramfs and kernel files are available.
5. Check that you created the grub netboot directory inside your tftp server root. Also, that you created the `grub.cfg` file and it is correctly configured.
6. Confirm that DHCP points to the tftp server and the correct boot file.

7. Check that DNS has all needed entries, including the etcd entries.
8. Confirm that the LoadBalancer correctly points to the control-plane (22623 and 6443), and to the worker nodes (80 and 443).

Installing your CoreOS operating systems on the LPARs

Follow the minimum requirements to install the LPARs. Start the boot from the bootstrap. The first time that you boot, blank disks are assigned; therefore, it is not necessary to change the boot list.

For more information about the Power Systems boot, see Appendix B, "Booting from System Management Services" on page 85.

Use a network boot of your choice to install the servers. For more information, see 3.3, "Using NIM as your BOOTP infrastructure" on page 50.

After you start all servers, wait for the bootstrap to complete. This section does not discuss troubleshoot during the boot process. To get information about the bootstrap completion, use the command that is shown in Example 3-11.

Example 3-11 Wait for bootstrap completion

```
[root@client install]# ./openshift-install wait-for bootstrap-complete
INFO Waiting up to 30m0s for the Kubernetes API at
https://api.ocp4.ibm.lab:6443...
INFO API v1.16.2 up
INFO Waiting up to 30m0s for bootstrapping to complete...
INFO It is now safe to remove the bootstrap resources
```

After completion, it is safe to shut down the bootstrap server because it is no longer needed in the cluster lifecycle, and everything is done by using the control plane. Even adding nodes is done without a bootstrap.

To get configurations ready for use, we show how to change CoreOS parameters.

The supported way of making configuration changes to CoreOS is through the machineconfig objects (see the login message of CoreOS), as shown in Example 3-12.

Example 3-12 Message when logging in to CoreOS

```
Red Hat Enterprise Linux CoreOS 43.81.202005200338.0
  Part of OpenShift 4.3, RHCOS is a Kubernetes native operating system
  managed by the Machine Config Operator (`clusteroperator/machine-config`).

WARNING: Direct SSH access to machines is not recommended; instead,
make configuration changes via `machineconfig` objects:

https://docs.openshift.com/container-platform/4.3/architecture/architecture-rhcos.
html

---
Last login: Sun May 31 11:58:27 2020 from 20.0.1.37
[core@worker1 ~]$
```

For more information about how to create machineconfig objects and how to add them, see Appendix A, "Configuring Red Hat CoreOS" on page 73.

Important: The configuration that is shown in Appendix A, “Configuring Red Hat CoreOS” on page 73 is important when VIOS Shared Ethernet Adapters (SEA) is used, and also for IBM Cloud Pak for Data. It also simplifies SMT management across the cluster, which makes it possible to run different levels of SMT to take advantage of the parallel threading offered by the POWER processor. However, it also makes it possible for applications that are sensitive to context switching to run at their best. If you do not apply this configuration, you can end up with authentication operator in unknown state if using SEA, as in our case because of our network configuration. The **oc apply -f <FILE>** command is used to apply the files.

With time, the range of operations start and become available, as shown in Example 3-13.

Example 3-13 Waiting for all cluster operators to become available

```
[root@client install]# oc get co
NAME                                  VERSION AVAILABLE PROGRESSING DEGRADED SINCE
authentication                        4.3.19  True      False       False    1h
cloud-credential                      4.3.19  True      False       False    1h
cluster-autoscale                     4.3.19  True      False       False    1h
console                               4.3.19  True      False       False    1h
dns                                   4.3.19  True      False       False    1h
image-registry                        4.3.19  True      False       False    1h
ingress                               4.3.19  True      False       False    1h
insights                              4.3.19  True      False       False    1h
kube-apiserver                        4.3.19  True      False       False    1h
kube-controller-manager               4.3.19  True      False       False    1h
kube-scheduler                        4.3.19  True      False       False    1h
machine-api                           4.3.19  True      False       False    1h
machine-config                        4.3.19  True      False       False    1h
marketplace                           4.3.19  True      False       False    1h
monitoring                            4.3.19  True      False       False    1h
network                               4.3.19  True      False       False    1h
node-tuning                           4.3.19  True      False       False    1h
openshift-apiserver                   4.3.19  True      False       False    1h
openshift-controller-manager          4.3.19  True      False       False    1h
openshift-samples                     4.3.19  True      False       False    1h
operator-lifecycle-manager            4.3.19  True      False       False    1h
operator-lifecycle-manager-catalog    4.3.19  True      False       False    1h
service-ca 4.3.19                     4.3.19  True      False       False    1h
service-catalog-apiserver             4.3.19  True      False       False    1h
service-catalog-controller-manager    4.3.19  True      False       False    1h
storage                               4.3.19  True      False       False    1h
```

The process to check when the cluster is ready can also be monitored by using the **openshift-install** command, as shown in Example 3-14.

Example 3-14 Waiting for installation completion

```
[root@client install]# ./openshift-install wait-for install-complete
INFO Waiting up to 30m0s for the cluster at https://api.ocp4.ibm.lab:6443 to
initialize...
INFO Waiting up to 10m0s for the openshift-console route to be created...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/root/install/auth/kubeconfig'
```

```
INFO Access the OpenShift web-console here:
https://console-openshift-console.apps.ocp4.ibm.lab
INFO Login to the console with user: kubeadmin, password: 5hIjT-FkY3u-SbMBI-Fzan7
```

3.2.6 Configuring the internal Red Hat OpenShift registry

Red Hat OpenShift has an internal registry. IBM Cloud Pak for Data documentation mentions the use of this registry to simplify operations. Therefore, this section shows how to configure it.

Export the NFS server as shown in the documentation that is available at this web page.

Example 3-15 shows a sample YAML file that describes how to define a static persistent volume in the NFS storage backend.

Example 3-15 YAML file to define persistent value with NFS

```
apiVersion: v1
kind: PersistentVolume
metadata:
  name: imageregistry
spec:
  capacity:
    storage: 100Gi
  accessModes:
  - ReadWriteMany
  nfs:
    path: /ocpgpfs/imageregistry
    server: ces.ocp4.ibm.lab
  persistentVolumeReclaimPolicy: Retain
```

Apply the file by issuing the **`oc apply -f <FILE>`** command, as shown in Example 3-16.

Example 3-16 Apply the file

```
[root@client ~]# oc apply -f imageregistrypv.yaml
persistentvolume/imageregistry created
[root@jinete nfs]# oc get pv
NAME          CAPACITY  ACCESS MODES   RECLAIM POLICY   STATUS      CLAIM
imageregistry 100Gi     RWX            Retain           Available
```

Now, enable the internal image registry by changing the `managementState` parameter from `Deleted` to `Managed`, as shown in Example 3-17. Issue the **`oc edit configs.imageregistry.operator.openshift.io`** command to open a text editor and make the change.

Example 3-17 Edit Image registry operator

```
apiVersion: imageregistry.operator.openshift.io/v1
kind: Config
metadata:
  creationTimestamp: "2020-05-27T17:04:58Z"
  finalizers:
  - imageregistry.operator.openshift.io/finalizer
  generation: 4
  name: cluster
```

```
  resourceVersion: "7473646"
  selfLink: /apis/imageregistry.operator.openshift.io/v1/configs/cluster
  uid: 913f5985-a50f-4878-b058-72eae57fc8a5
spec:
  defaultRoute: true
  disableRedirect: false
  httpSecret: 7ab51007dbad0462dfeb89f5f1d97edbfc42782c534bf0911ab77497bd285357055f
92129469643cdf507d3f15ed1ab38e5cd4e0c8b4bf71aea9a4f3b531c39c
  logging: 2
  managementState: Managed
  proxy:
    http: ""
    https: ""
    noProxy: ""
  readOnly: false
  replicas: 1
  requests:
    read:
      maxInQueue: 0
      maxRunning: 0
      maxWaitInQueue: 0s
    write:
      maxInQueue: 0
      maxRunning: 0
      maxWaitInQueue: 0s
  storage:
    pvc:
      claim:
```

Exposing the internal registry

To push content to the internal registry, you can optionally enable the default route, as shown in Example 3-18.

Example 3-18 Enable the default route

```
[root@client ~]# oc patch configs.imageregistry.operator.openshift.io/cluster
--patch '{"spec":{"defaultRoute":true}}' --type=merge
config.imageregistry.operator.openshift.io/cluster patched
```

3.3 Using NIM as your BOOTP infrastructure

Network Installation Management (NIM) is used to boot, install, and maintain AIX operating systems up to date. NIM uses TFTP and BOOTP protocols to boot and install AIX servers from the network. This section describes how to configure a NIM server running on an AIX operating system to boot and install a CoreOS LPAR from the network.

For more information about NIM, see IBM Knowledge Center.

Tip: This step-by-step guide can also be used with Red Hat Linux V7 and later versions.

3.3.1 Prerequisites and components

This section describes the prerequisites that must be met.

Required steps to set up the NIM environment and boot

Complete the following steps to set up the NIM environment:

1. Install NIM master files.
2. Install and configure Apache web server for AIX.
3. Copy the ignition files and the CoreoOS bare metal file to the http server root directory.
4. Configure tftp and bootp.
5. Copy the CoreOS image files to tftp directory.
6. Copy the ignition files to the Apache web server root directory.
7. Netboot the LPAR in the Power Systems Hardware Management Console (HMC).

Installing NIM master files

For NIM server installation in AIX, the `nim.master` and `nim.spot` files from AIX installation image must be installed. For more information about how to install and configure NIM, see IBM Knowledge Center.

Installing Apache web server in AIX

The Apache RPM package `httpd` for AIX can be downloaded (see Example 3-19) and manually installed with all of the package dependencies from the IBM Toolbox for Linux Applications web page.

Example 3-19 Download and install apache for AIX from aixtoolbox IBM site

```
# wget --no-check-certificate
https://public.dhe.ibm.com/aix/freeSoftware/aixtoolbox/RPMS/ppc/httpd/httpd-2.4.41-1.aix6.1
.ppc.rpm
--2020-05-25 20:59:09--
https://public.dhe.ibm.com/aix/freeSoftware/aixtoolbox/RPMS/ppc/httpd/httpd-2.4.41-1.aix6.1
.ppc.rpm
Resolving public.dhe.ibm.com... 170.225.15.112
Connecting to public.dhe.ibm.com|170.225.15.112|:443... connected.
WARNING: cannot verify public.dhe.ibm.com's certificate, issued by 'CN=GeoTrust RSA CA
2018,OU=www.digicert.com,O=DigiCert Inc,C=US':
Self-signed certificate encountered.
HTTP request sent, awaiting response... 200 OK
Length: 4279776 (4.1M) [text/plain]
Saving to: 'httpd-2.4.41-1.aix6.1.ppc.rpm'
httpd-2.4.41-1.aix6.1.ppc.rpm
100%[=====================================================>]   4.08M   310KB/s    in 27s
2020-05-25 20:59:37 (154 KB/s) - 'httpd-2.4.41-1.aix6.1.ppc.rpm' saved [4279776/4279776]

[ocp43nimserver@root:/bigfs:] rpm -ivh httpd-2.4.41-1.aix6.1.ppc.rpm
error: Failed dependencies:
apr >= 1.5.2-1 is needed by httpd-2.4.41-1.ppc
apr-util >= 1.5.4-1 is needed by httpd-2.4.41-1.ppc
expat >= 2.2.6 is needed by httpd-2.4.41-1.ppc
libapr-1.so is needed by httpd-2.4.41-1.ppc
libaprutil-1.so is needed by httpd-2.4.41-1.ppc
libgcc >= 6.3.0-1 is needed by httpd-2.4.41-1.ppc
liblber.a(liblber-2.4.so.2) is needed by httpd-2.4.41-1.ppc
libldap.a(libldap-2.4.so.2) is needed by httpd-2.4.41-1.ppc
libpcre.a(libpcre.so.1) is needed by httpd-2.4.41-1.ppc
openldap >= 2.4.40 is needed by httpd-2.4.41-1.ppc
```

```
pcre >= 8.42 is needed by httpd-2.4.41-1.ppc
```

Example 3-19 on page 51 shows that package dependencies must be installed first before the httpd package can be installed. Another option is to install the YUM package manager for AIX and let YUM resolve and install the dependencies.

Installing with YUM

Other installation methods can be by installing YUM for AIX and then use YUM to install the Apache httpd package. This process can be the preferred method because YUM can automatically download and install all package dependencies for Apache.

> **YUM installation AIX:** For more information, see this web page.

After YUM is installed, you can use it to install the Apache web server and all package dependencies, as shown in Example 3-20.

Example 3-20 Install Apache web server in AIX using YUM

```
#yum search httpd
============== N/S Matched: http ==========================================
httpd.ppc : Apache HTTP Server
httpd-devel.ppc : Development tools for the Apache HTTP server.
httpd-manual.ppc : Documentation for the Apache HTTP server.
libnghttp2.ppc : A library implementing the HTTP/2 protocol
libnghttp2-devel.ppc : Files needed for building applications with libnghttp2
mod_http2.ppc : Support for the HTTP/2 transport layer
.......... output removed ..........
....................................
Name and summary matches only, use "search all" for everything.

#yum install httpd.ppc
Setting up Install Process
Resolving Dependencies
--> Running transaction check
---> Package httpd.ppc 0:2.4.41-1 will be installed
Installed:
httpd.ppc 0:2.4.41-1
Dependency Installed:
apr.ppc 0:1.5.2-1     apr-util.ppc 0:1.5.4-1    cyrus-sasl.ppc 0:2.1.26-3
libiconv.ppc 0:1.14-2   libstdc++.ppc 0:8.3.0-2   libunistring.ppc 0:0.9.9-2
libxml2.ppc 0:2.9.9-1
ncurses.ppc 0:6.2-1   openldap.ppc 0:2.4.48-1   pcre.ppc 0:8.43-1
xz-libs.ppc 0:5.2.4-1
Dependency Updated:
bzip2.ppc 0:1.0.8-2   expat.ppc 0:2.2.9-1   gettext.ppc 0:0.19.8.1-5   glib2.ppc
0:2.56.1-2   info.ppc 0:6.6-2   libgcc.ppc 0:8.3.0-2   readline.ppc 0:8.0-2
sqlite.ppc 0:3.28.0-1
Complete!
```

Basic Apache configuration and service start

Edit the Apache configuration file as shown in Example 3-21 to change the service port that listens and verifies the root directory.

The Apache default configuration file is `/opt/freeware/etc/httpd/conf/httpd.conf`.

Example 3-21 Apache config file modifications

```
#vi /opt/freeware/etc/httpd/conf/httpd.conf
Listen 8080
#
# DocumentRoot: The directory out of which you will serve your
# documents. By default, all requests are taken from this directory, but
# symbolic links and aliases may be used to point to other locations.
#
DocumentRoot "/var/www/htdocs"
<Directory "/var/www/htdocs">
#
# Possible values for the Options directive are "None", "All",
# or any combination of:
#   Indexes Includes FollowSymLinks SymLinksifOwnerMatch ExecCGI MultiViews
#
# Note that "MultiViews" must be named *explicitly* --- "Options All"
# doesn't give it to you.
#
# The Options directive is both complicated and important.  Please see
# http://httpd.apache.org/docs/2.4/mod/core.html#options
# for more information.
#
Options Indexes FollowSymLinks
#
# AllowOverride controls what directives may be placed in .htaccess files.
# It can be "All", "None", or any combination of the keywords:
#   AllowOverride FileInfo AuthConfig Limit
#
AllowOverride None
#
# Controls who can get stuff from this server.
#
Require all granted
</Directory>
```

Copying ignition files and CoreOS boot file to the Apache root directory

After generating the ignition files as described in Example 3-10 on page 46, copy these files to the Apache root directory in the NIM server.

> **Ignition files:** Because no Red Hat OpenShift installer is available for AIX, the ignition files must be created by using the `openshift-install` script in a Linux Red Hat or MacOS operating system.

Copy the `bootstrap.ign`, `master.ign`, and `worker.ign` files to the NIM server, as shown in Example 3-22 on page 54.

Copy the `rhcos-4.3.18-ppc64le-metal.ppc64le.raw.gz` file to the Apache root directory, as shown in Example 3-22.

Example 3-22 Copy the ignition files and CoreOS boot file to the NIM server

```
deploy01# scp /root/openshift-install/*.ign ocp43nimserver:/var/www/htdocs/
deploy01# scp rhcos-4.3.18-ppc64le-metal.ppc64le.raw.gz ocp43nimserver:/var/www/htdocs/
```

Example 3-23 shows the contents of the Apache RootFS directory.

Example 3-23 Apache RootFS directory content

```
ls -la /var/www/htdocs
total 1517040
drwxr-xr-x    2 root     system          256 May 28 23:20 .
drwxr-xr-x    6 root     system          256 May 26 11:09 ..
-rwxrwxr-x    1 root     system       295410 May 28 21:44 bootstrap.ign
-rwxrwxr-x    1 root     system           45 Jun 11 2007  index.html
-rwxrwxr-x    1 root     system         1829 May 28 21:44 master.ign
-rwxrwxr-x    1 root     system    776412494 May 28 21:43
rhcos-4.3.18-ppc64le-metal.ppc64le.raw.gz
-rwxrwxr-x    1 root     system         1829 May 28 21:44 worker.ign
```

Start and verify the Apache service

Start the httpd service to publish the content in the rootfs, as shown in Example 3-24.

Example 3-24 Apache start service command

```
#/opt/freeware/sbin/apachectl -k start
AH00558: httpd: Could not reliably determine the server's fully qualified domain
name, using 172.16.140.120. Set the 'ServerName' directive globally to suppress
this message
```

NOTE: The AH00558: ServerName Warning message can be ignored.

Verify that Apache is running and listening on port 8080, as shown in Example 3-25.

Example 3-25 Apache server status

```
#ps -ef | grep -i http
apache  5636450 14221670   0 18:42:49      -  0:00 /opt/freeware/sbin/httpd -k start
apache  5701922 14221670   0 18:42:49      -  0:00 /opt/freeware/sbin/httpd -k start
apache  8323540 14221670   0 18:42:49      -  0:00 /opt/freeware/sbin/httpd -k start
apache  9765306 14221670   0 18:42:49      -  0:00 /opt/freeware/sbin/httpd -k start
root 14221670        1   0 18:42:49      -  0:00 /opt/freeware/sbin/httpd -k start
apache 15073660 14221670   0 18:42:49      -  0:00 /opt/freeware/sbin/httpd -k start

#netstat -an | grep 8080
tcp        0      0  *.8080                *.*                    LISTEN
```

Bootp and tftp configuration

The `/etc/bootptab` file must be manually updated to permit network communication to the LPAR you want to boot from the network.

Add the following information according to your configuration to the end of the `bootptab` file, as shown in Example 3-26:

- LPAR IP address: 172.16.140.95
- LPAR HW address: fa698d138b20
- NIM server IP address: 172.16.140.120
- Network mask 255.255.255.0

Example 3-26 bootptab file content

```
bootstrap:bf=/tftpboot/coreos/boot/grub2/powerpc-ieee1275/core.elf:ip=172.16.140.9
5:ht=ethernet:ha=fa698d138b20:sa=172.16.140.120:sm=255.255.255.0:
```

Note: After the installation is completed, the bootptab lines must be removed manually.

Copy all files from `powerpc-ieee1275` directory (including the `core.elf` from the Red Hat deploy01 LPAR) to the NIM server, as shown in Example 3-27.

Example 3-27 ieee1275 files copy to the NIM server

```
deploy01#scp /var/lib/tftpboot/boot/grub2/powerpc-ieee1275/*
ocp43nimserver:/tftpboot/coreos/boot/grub2/powerpc-ieee1275/
```

Grub file configuration

CoreOS uses grub as the boot file configuration. Example 3-28 shows the boot file only to netboot the bootstrap LPAR.

Example 3-28 grub2 file options

```
cat /tftpboot/coreos/boot/grub2/powerpc-ieee1275/grub.cfg
set default=0
set fallback=1
set timeout=10
echo "Loading kernel bootstrap test"
menuentry "bootstrap CoreOS (BIOS)" {
echo "Loading kernel bootstrap"
insmod linux
linux   "/tftpboot/coreos/rhcos-4.3.18-installer-kernel-ppc64le"
systemd.journald.forward_to_console=1  rd.neednet=1
ip=172.16.140.95::172.16.140.1:255.255.255.0:bootstrap.demolab.uy.ibm.com::none
nameserver=172.16.140.70      console=tty0   console=ttyS0  coreos.inst=yes
coreos.inst.install_dev=sda
coreos.inst.image_url=http://172.16.140.120:8080/rhcos-4.3.18-ppc64le-metal.ppc64l
e.raw.gz       coreos.inst.ignition_url=http://172.16.140.120:8080/bootstrap.ign
echo "Loading initrd"
initrd "/tftpboot/coreos/rhcos-4.3.18-installer-initramfs.ppc64le.img"
}
```

Configuring tftp access

The last config file to change is the `/etc/tftpaccess.ctl` file that permits tftp process to some directories on the server, as shown in Example 3-29.

Example 3-29 tftpaccess configuration file to permit boot process

```
cat /etc/tftpaccess.ctl
# NIM access for network boot
allow:/tftpboot/coreos
```

Netboot from the HMC

Finally, the LPAR can be started and booted from the network. This process can be done from the HMC command line (as shown in Example 3-30) or by using the HMC graphical user interface as described in Appendix B, "Booting from System Management Services" on page 85.

Consider the following points:

- `-S` is the NIM server IP address
- `-G` is the default gateway
- `-C` is the LPAR IP address
- `-K` is the network mask

Example 3-30 LPAR netboot command in the HMC

```
lpar_netboot  -t ent -s auto -d auto -D -S 172.16.140.120 -G 172.16.140.1 -C 
172.16.140.95 -K 255.255.255.0 ocp43bootstrap default_profile dc4p824
```

> **Tip:** If more verbose debug is needed for the **`lpar_netboot`** command, the **`–x`** flag can be specified.

3.4 Installing on scale-out servers bare metal

The ideal container implementation creates an isolation layer for workloads better than server virtualization, which saves memory and CPU that are needed to maintain different workloads. PowerVM can provide you with a layer of separation better than x86 virtualization because of its tight integration with the hardware. However, you can also deploy bare metal servers when you want large worker nodes.

This section describes how to add a bare metal server (also known by Opal or PowerNV servers) to your PowerVM cluster. This example can be the most complex case because you are mixing different ways to provision nodes into your cluster. The entire cluster can also be PowerNV only.

Petitboot is the operating system bootloader for scale-out PowerNV systems and is based on Linux kexec. Petitboot can use PXE boot to ease the process booting with a simple configuration file. It can load any operating system image that supports the Linux kexec reboot mechanism, such as Linux and FreeBSD. Petitboot can load images from any device that can be mounted by Linux, and can also load images from the network by using the HTTP, HTTPS, NFS, SFTP, and TFTP protocols.

You can still use the same DHCP, HTTP, and TFTP servers that are used for the KVM installation and with a few changes provide the necessary environment to also install PowerNV servers, as shown in Example 3-31.

Example 3-31 dhcp.conf file for mixed PowerVM and PowerNV environment

```
default-lease-time 900;
max-lease-time 7200;
subnet 192.168.122.0 netmask 255.255.255.0 {
        option routers 192.168.122.1;
        option subnet-mask 255.255.255.0;
        option domain-search "ocp4.ibm.lab";
        option domain-name-servers 192.168.122.1;
        next-server 192.168.122.5;
        filename "boot/grub2/powerpc-ieee1275/core.elf";
}
allow bootp;
option conf-file code 209 = text;
.
.
.
host powervmworker {
        hardware ethernet 52:54:00:1a:fb:b6;
        fixed-address 192.168.122.30;
        option host-name "powervmworker.ocp4.ibm.lab";
        next-server 192.168.122.5;
        filename "boot/grub2/powerpc-ieee1275/core.elf";
        allow booting;
}
host powernvworker {
        hardware ethernet 98:be:94:73:cd:78;
        fixed-address 192.168.122.31;
        option host-name "powernvworker.ocp4.ibm.lab";
        next-server 192.168.122.5;
        option conf-file "pxelinux/pxelinux.cfg/98-be-94-73-cd-78.cfg";
        allow booting;
}
```

After changing the `dhcp.conf` file, restart the dhcpd server. The PXE boot does not point to the `grub.cfg` file created for the PowerVM hosts. Instead, it points to the configuration file that is stated on the dhcpd entry.

> **Note:** We identified the config file to be `pxelinux/pxelinux.cfg/98-be-94-73-cd-78.cfg`. Therefore, the file must be on your tftp server; in our case, `/var/lib/tftpboot/pxelinux/pxelinux.cfg/98-be-94-73-cd-78.cfg`.

Create the configuration file as shown in Example 3-32.

Example 3-32 Configuration file example for PXE boot of a PowerNV worker node

```
DEFAULT pxeboot
TIMEOUT 20
PROMPT 0
LABEL pxeboot
KERNEL http://192.168.122.5:8080/rhcos-4.3.18-ppc64le-installer-kernel-ppc64le
```

```
APPEND
initrd=http://192.168.122.5:8080/rhcos-4.3.18-ppc64le-installer-initramfs.ppc64le.
img rd.neednet=1 ip=dhcp nameserver=192.168.122.1 console
```

3.5 NVMe 4096 block size considerations

At the time of this writing, no 4 K Block size image is available for CoreOS. If you want to install an NVMe with block size of 4096, you first must format it to 512, as shown in Example 3-33.

Example 3-33 Formatting NVMe to 512 block size

```
[root@localhost ~]# nvme format /dev/nvme0n1 -b 512
Success formatting namespace:1
[root@localhost ~]# nvme list ns
Node             SN                   Model
Namespace Usage                      Format           FW Rev
---------------- -------------------- ----------------------------------------
--------- -------------------------- ---------------- --------
/dev/nvme0n1     S46JNE0M700050       PCIe3 1.6TB NVMe Flash Adapter III x8    1
1.60  TB /   1.60  TB    512   B +  0 B   MA24MA24
```

3.6 Offline Red Hat OpenShift V4.3 deployment

If you do not have an internet connection to your Enterprise (even by way of Proxy), you can still install Red Hat OpenShift Container Platform on Power Systems.

Find a bastion host that can accesses the internet that your cluster can reach. This server is not used as a router or any other network service other than the registry for the Red Hat OpenShift Platform. Therefore, it is not in the production path and is used for maintenance only.

For more information about this concept, see the online documentation on the Red Hat website.

Note: The most important part of this documentation is where it shows you how to configure the mirror registry on the bastion node (see this web page).

These instructions assume that you perform the instructions on a x86 architecture. If you are performing it on a ppc64le, the only change is to change the registry container image to one supported in ppc64le when performing Step 5 of the section at this web page.

4

IBM Cloud Paks

This chapter describes the IBM Cloud Paks offerings that are available to use with Red Hat OpenShift. This chapter also describes our experiences installing IBM Cloud Pak for Data, and points the user to the documentation with more information about supported services and prerequisites.

This chapter includes the following topics:

4.1 Introduction

The IT world is primarily divided into two major segments. The first segment is new development and the second is *run* operations where new features and enhancements are deployed to production environments. Both phases have their own set of requirements, constrains, and challenges. This is the situation during migration to the cloud. Cloud migration process for a production environment is significantly different from a development environment. This leads to new terminologies, such as cloud migration and cloud modernization.

Key performance indexes include the following examples:

- Easy to manage orchestration service
- Improved security
- Efficiency in governance
- Reduced cost to maximize the return on the investment (ROI)

IBM Cloud Paks are enterprise-ready, containerized middleware and common software solutions that are hosted on Kubernetes-based platforms that give clients an open, faster, and more secure way to move core business applications to any Cloud. This full stack, converged infrastructure includes a virtualized cloud hosting environment that helps to extend applications to the Cloud.

IBM Cloud Paks features the following benefits:

- Market ready

  IBM Cloud Paks with Red Hat OpenShift is a flexible combination to ensure faster deployment with high scalability. API-based micro services ensure faster adoption of changes. Cloud Native applications are quickly developed by using container and microservices that can take advantage of capabilities of middleware database through DevOps practices.

- Run anywhere

  IBM Cloud Paks are portable. They can run on-premises, on public clouds, or in an integrated system.

IBM Cloud Paks are certified by IBM with up-to-date software to provide full-stack support, from hardware to applications.

4.2 IBM Cloud Paks

IBM Cloud Paks are built on the Kubernetes-based portable platform that use a common container platform from Red Hat OpenShift. This enterprise-ready containerized software solution provides several key benefits to different segments of users (see Figure 4-1 on page 61) during the following phases:

- Build: Packaged with open platform components to take advantage of several API services available from different sources. This phase is easy to build and distribute. Developers take advantage of a single application environment that is configured with all required tools for planning the modernization process to build Cloud native API and runtime platforms to deploy the solution.
- Move: Run-anywhere model allows same software to be ported to on-premises or private or public cloud that is based on the client's requirement. This phase also is a transition from a large monolithic application to an API-based micro service model.

- Run: Available built-in services, such as logging, monitoring, metering, security, identity management, and image registry. Each business is unique with a set of key business values that can lead the application to operate from on-premises, private, and public clouds. The unified Kubernetes platform provides such flexibility to deploy and run the same application in any wanted platform.

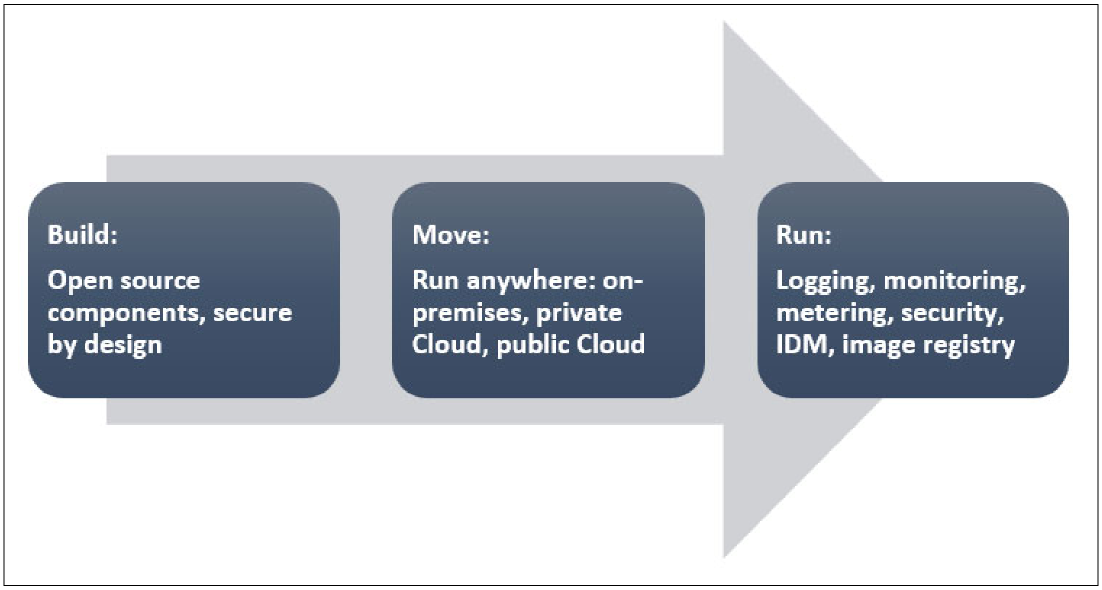

Figure 4-1 IBM Cloud Paks helps users looking to enable workloads

4.3 IBM Cloud Paks offerings

This section describes the different IBM Cloud Paks and capabilities.

4.3.1 IBM Cloud Pak for Application

This offering accelerates the modernization and building of applications by using built-in developer tools and processes. This offering includes support for analyzing applications and guiding the application owner through the modernization journey. In addition, it supports cloud-native development microservices functions and serverless computing. Customers can quickly build cloud apps, although IBM middleware clients gain the most straightforward path to modernization.

Transformation of a traditional application is one the key features in this scope. Development, testing, and redeployment are some of the phases where most of the effort and challenges are experienced in the traditional development model. The Agile DevOps-based development model is the potential solution to break this challenge.

To complement the Agile development process, IBM Cloud Pak for Application extends Kubernetes features for consistent and faster development cycle, which helps IBM clients to build cost-optimized, smarter applications.

4.3.2 IBM Cloud Pak for Data

This offering unifies and simplifies the collection, organization, and analysis of data. Enterprises can turn data into insights through an integrated cloud-native architecture. IBM Cloud Pak for Data is extensible and can be customized to a client's unique data and AI landscapes through an integrated catalog of IBM, open source, and third-party microservices add-ons.

> **Note:** Although all IBM Cloud Paks are intended to be supported on the ppc64le architecture, at the time of this writing, IBM Cloud Pak for Data was readily available for us to test. Therefore, we worked more with this IBM Cloud Pak specifically.

For more information about IBM Cloud Pak for Data, see 4.4, “IBM Cloud Pak for Data” on page 63.

4.3.3 IBM Cloud Pak for Integration

This offering Integrates applications, data, cloud services, and APIs. It also supports the speed, flexibility, security, and scale that is required for modern integration and digital transformation initiatives. It includes a pre-integrated set of capabilities, which includes API lifecycle, application and data integration, messaging and events, high-speed transfer, and integration security.

Personalize customer experience is the primary business focus that needs an integrated view of all scattered data. IBM Cloud Pak for Integration facilitates rapid integration of data along with security, compliance, and version capability. If featured the following key capabilities:

- API lifecycle management
- Application and data integration
- Enterprise messaging
- Event streaming
- High-speed data transfer
- Secure gateway

4.3.4 IBM Cloud Pak for Automation

This offering transforms business processes, decisions, and content. A pre-integrated set of essential software enables clients to easily design, build, and run intelligent automation applications at scale. The following major KPIs are featured:

- Improved efficiently and productivity
- Enhanced customer experience
- Operational insight

IBM Cloud Pak for Automation helps to automate business operations with an integrated platform. Kubernetes makes it easier to configure, deploy, and manage containerized applications. It is compatible with all types of projects small and large to deliver better end-to-end customer journeys with improved governing of content and processes.

Automation capabilities empower to work more effectively in the following cases:

- With limited work force to manage higher workload from new applications or services, rising customer demand, or seasonal fluctuations.
- Create enhanced and personalized customer experiences that increase loyalty by drawing insights instantly from multiple sources of information.

- When you need to scale operations to help maximize revenue and customer service.

4.3.5 IBM Cloud Pak for Multicloud Management

This offering provides consistent visibility, governance, and automation across a range of multicloud management capabilities, such as infrastructure management, application management, multicluster management, edge management, and integration with existing tools and processes.

4.4 IBM Cloud Pak for Data

IBM Cloud Pak for Data is a native cloud solution that provides a single unified interface for your team to connect to your data no matter where it lives and manage it, find it, and use it for analysis.

Data management features the following capabilities:

- Connect to data
- Data governance
- Identify wanted data
- Data transformation for analytics

User access and collaboration capability includes the following features:

- Single unified interface
- Services to data management and collaboration
- Data readiness for ready use in analytics
- Access and connection are built-in features

Functionality with fully integrated data and AI platform features the following capabilities:

- Data collection
- Data organization
- Data analysis
- Infuse AI into the business
- Support multicloud architecture

The following benefits are realized:

- Cost saving
- Built-in services
- Integrated tools
- Scope for customize solution

4.4.1 IBM Cloud Pak for Data features and enhancements

IBM Cloud Pak for Data V3.0.1 is the exclusive offering for Red Hat OpenShift. Clients can selectively install and activate required services only. IBM Cloud Pak for Data features several enhancements for usability and more integrated service orientation, as shown in Figure 4-2.

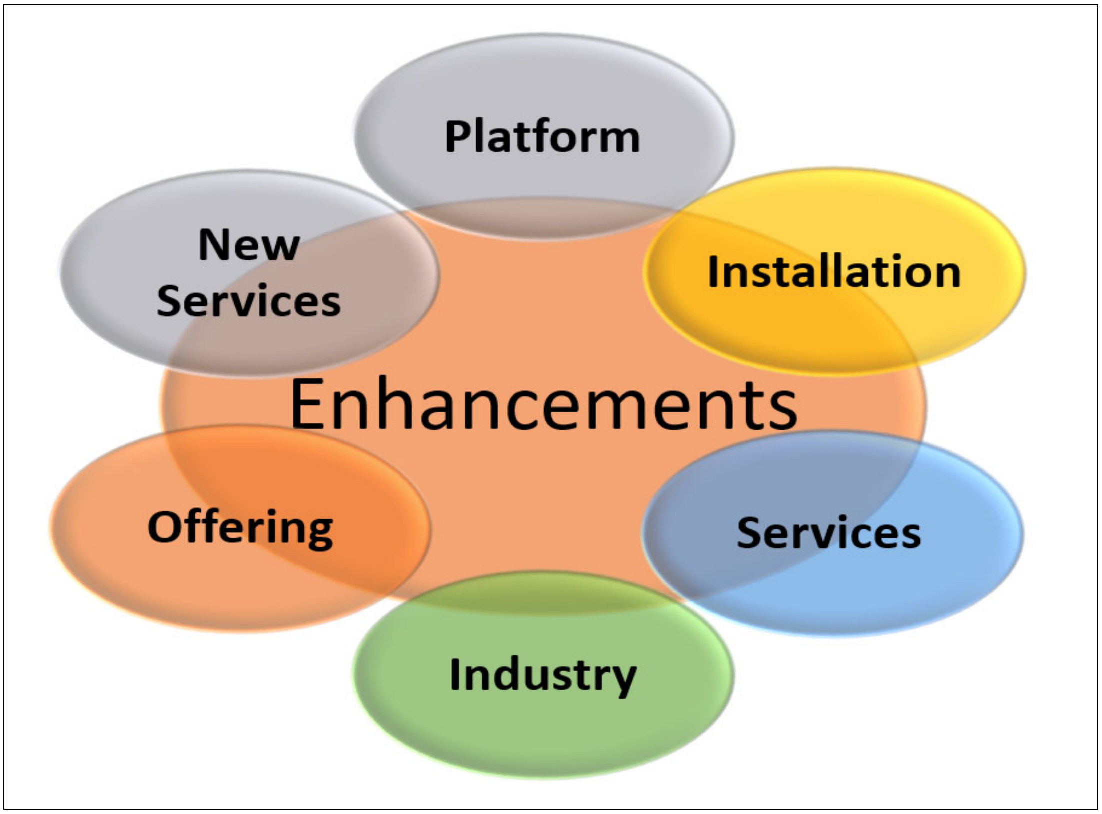

Figure 4-2 IBM Cloud Pak for Data enhancements

These enhancements (see Figure 4-2) are classified into the following sections:

- Platform:
  - Modular services-based platform setup for efficient and optimized use of resources.
  - Built-in dashboard with meaningful KPI for better monitoring, metering, and serviceability.
  - Open extendable platform with new age Platform and Service APIs.
- Installation:
  - Simplified installation
  - Red Hat OpenShift V4.3
- Service:
  - Data processing and analytics are some of the key enhancements.
  - Advanced integration with IBM DataStage® and IBM Db2®.
  - Advanced predictive and analytical models that use IBM SPSS® and Streams Watson suits.

- Industry

  Available industry examples to some solutions for various industries. For more information, see this web page.

- Offering (IBM Cloud Pak for Data):
  - DataStage Edition
  - Watson Studio Premium
  - Db2
- New services categories:
  - AI
  - Integration with Spark, Hadoop
  - Developer tools like Jupyter Notebook with Python V3.6, RStudio V3.6

IBM Cloud Pak for Data emerged as an extensible and highly scalable data and AI platform. It is built on a comprehensive foundation of unified code, data, and AI services that can use multiple Cloud native services. It also is flexible enough to adopt customizations to address specific business needs through extended service catalog.

The catalog features the following services:

- AI: Consists of several Watson libraries, tools, and studios.
- Analytics: Services include Trusted Predictive and analytical platforms, such as Cognos, SPSS, Dashboards, and Python.
- Data governance: Consists of IBM InfoSphere® and Watson libraries.
- Data store service: Industry solution and storage.

4.4.2 IBM Cloud Pak for Data requirements

The installation requirements can change without notice. It is recommended that you refer to the latest documentation when planning and performing your installation.

The minimum resource recommendations that are described in this publication are for guidance only. Consult with your IBM Sales representative for recommendations that are based on your specific needs.

To check that your persistent volume provider complies with latency and throughput requirements for IBM Cloud Pak for Data, follow the storage test that is described at this web page.

IBM Cloud Pak for Data includes many services to ensure that you have a complete solution available to you. For more information about the services that supported on the ppc64le platform and their requirements, see this web page.

> **Note:** The time on all of the nodes must be synchronized within 500 ms. Check that you have the NTP correctly set. Use the machine configuration method that is described in Appendix A, "Configuring Red Hat CoreOS" on page 73 to configure `/etc/chrony.conf` to point to the correct NTP server.

4.4.3 Installing IBM Cloud Pak for Data on Power Systems servers

For more information about the prerequisites that must be met to install IBM Cloud Pak for Data on Red Hat OpenShift Container Platform V4.3, see IBM Knowledge Center.

> **Note:** Appendix A, "Configuring Red Hat CoreOS" on page 73 provides steps to configure these settings by using machine config objects and tuned operator. Ensure that you apply these settings before installing IBM Cloud Pak for Data.

Before installing, you need a registry to hold IBM Cloud Pak for Data images. For simplicity, the internal registry of Red Hat OpenShift can be configured as described in 3.2.6, "Configuring the internal Red Hat OpenShift registry" on page 49.

Downloading the installation and repository file

Complete the process that is described at IBM Knowledge Center to check that you have the necessary files to install IBM Cloud Pak for Data.

This case uses the `repo.yaml` file and the `cpd-ppc64le` client. Example 4-1 shows the `repo.yaml` file.

Example 4-1 repo.yaml file

```
registry:
  - url: <URL>
    username: <USERNAME>
    apikey: [Get you entitlement key here-
https://myibm.ibm.com/products-services/containerlibrary]
    name: base-registry
fileservers:
  - url: <URL>
```

You need the entitlement key to install IBM Cloud Pak for Data.

Offline installation method

The assembly package contains a service and can be installed on Red Hat OpenShift. The basic offline installation method for all services consists of the following steps:

1. Download the assembly on a server that includes access to the internet and the internal registry.
2. Push downloaded images to the internal registry.
3. Apply the admin setup of the namespace for the assembly.
4. Install the assembly into the namespace.

The base service is the lite assembly that contains the infrastructure and the core part of IBM Cloud Pak for Data. To download the assembly, configure your `repo.yaml` file and issue the command as shown in Example 4-2.

Example 4-2 Downloading the assembly

```
[root@client ~]# ./cpd-ppc64le preloadImages -a lite --action download -r
./repo.yaml --download-path=/repository/lite --arch ppc64le --verbose
[v] [2020-06-02 06:37:35-0186] Path is verified: /repository/lite
[v] [2020-06-02 06:37:35-0311] Running action mode: download
[v] [2020-06-02 06:37:35-0541] {CustomYaml:./repo.yaml ClusterRegistry:{Mode:1
PullPrefix: PullUsername: PullKey: PushPrefix: PushUsername: PushKey: PushCert:}
Assembly:{Assembly:lite Versio
n: Arch:ppc64le DownloadPath:/repository/lite Namespace: StorageClass: LoadPath:
HiddenForceUpgrade:false} AdmApplyRun:false AdmForce:false
AskPushRegistryPassword:false AskPullRegistryPassw
```

```
ord:false TransferImage:false InsecureTLS:false IsUpgrade:false IsPatch:false
IsScale:false IsUninstall:false IsAdmSetup:false IsCRStatus:false
CpdinstallImage:cpdoperator:v3.0.1-13 CpdinitI
mage:cpd-operator-init:v1.0.1 InstallDryRun:false TillerImage:cp4d-tiller:v2.16.6
BaseRegistry:base-registry OperatorSCC:cpd-zensys-scc GlobalOverridePath:
FileServerCert: SilentInstall:fals
e IgnoreLicenses:false UninstallDryRun:false UpgDryRun:false DisplayPatches:false
DisplayScale:false AvailableUpdates:false ShowAll:false ScaleSize: Platform:oc
IngressHost: K8SConfig:<nil>
K8SClient:<nil> OptionalModulesList: AllOptionalModules:false SkipSICheck:false
MultiInstanceName: TetheredNamespace: LegacyDownload:true NoAppVerCheck:false
MaxImageDownloadRetry:5 NoManife
stRefresh:false}

[INFO] [2020-06-02 06:37:35-0651] Parsing custom YAML file ./repo.yaml
[INFO] [2020-06-02 06:37:35-0751] Overwritten default download settings using
./repo.yaml
.
.
.
Copying blob
sha256:87c6ac3990f2a6debb4e67ba3651030fe0b722975019a9f8079a1687882c1cde
Copying blob
sha256:0817eba40402f0bee8a2e8d3eb9dc27051d7f1758ba85d7409b01ae8e7be0e38
Copying blob
sha256:5f3688d5514c88858a861f52d2b9631e8c621bc01cee126002f3ead1255265ec
Copying blob
sha256:f5c8abbe4f91deb321e8c3e21eda810bc4337f621330c63a9b4c035de71affc5
Copying blob
sha256:29b01004f92ddaa2eacf7a2a11f018e6e6440020afe1a1446744aa9d64e32e92
Copying config
sha256:ca3fa91dc9dbc9b84a7dd2c151957f2a7732cf80d36eda757d4a327cf5723fac
Writing manifest to image destination
Storing signatures
[INFO] [2020-06-02 06:43:26-0721] The image
/repository/lite/images/zen-data-sorcerer-----v3.0.1.0-ppc64le-42.tar in the local
directory validated
[INFO] [2020-06-02 06:43:26-0768] All images are now downloaded successfully, the
directory is ready to perform offline push.
```

After you download it, push the images to your registry, as shown in Example 4-3.

Example 4-3 Pushing images to the internal registry

```
[root@client ~]# oc login https://api.ocp4.ibm.lab:6443
Authentication required for https://api.ocp4.ibm.lab:6443 (openshift)
Username: kubeadmin
Password:
Login successful.

You have access to 82 projects, the list has been suppressed. You can list all
projects with 'oc projects'

Using project "default".
[root@client ~]# ./cpd-ppc64le preloadImages --action push
--load-from=/repository/lite
```

```
--transfer-image-to=default-route-openshift-image-registry.apps.ocp4.ibm.lab/zen
--target-registry-username=kubeadmin --target-registry-password=$(oc whoami -t) -a
lite --arch ppc64le -v 3.0.1 --insecure-skip-tls-verify
[INFO] [2020-06-03 05:59:12-0351] --download-path not specified, creating
workspace cpd-ppc64le-workspace
[INFO] [2020-06-03 05:59:12-0356] Enter credentials for target registry
default-route-openshift-image-registry.apps.ocp-ppc64le-test-fae390.redhat.com/zen

[INFO] [2020-06-03 05:59:12-0361] 0 file servers and 0 registries detected from
current configuration.
[INFO] [2020-06-03 05:59:12-0365] Server configure files validated

[INFO] [2020-06-03 05:59:12-0370] Assembly version is validated

*** Parsing assembly data and generating a list of charts and images for download
***

[INFO] [2020-06-03 05:59:12-0380] Assembly data validated
[INFO] [2020-06-03 05:59:13-0062] The category field of module 0010-infra is not
specified in its manifest file, assuming default type 'helm-chart'
[INFO] [2020-06-03 05:59:13-0063] The category field of module 0015-setup is not
specified in its manifest file, assuming default type 'helm-chart'
[INFO] [2020-06-03 05:59:13-0064] The category field of module 0020-core is not
specified in its manifest file, assuming default type 'helm-chart'
---------------------------------------------------------------------------------
List of charts required by assembly:

MODULE         VERSION     ARCHITECTURE     CHART                        STATUS
0010-infra     3.0.1       ppc64le          0010-infra-3.0.1.tgz         Downloaded
0015-setup     3.0.1       ppc64le          0015-setup-3.0.1.tgz         Downloaded
0020-core      3.0.1       ppc64le          0020-zen-base-3.0.1.tgz      Downloaded
.
.
.
[INFO] [2020-06-03 05:59:20-0316] Pushing image
zen-data-sorcerer:v3.0.1.0-ppc64le-42 from file
/repository/lite/images/zen-data-sorcerer-----v3.0.1.0-ppc64le-42.tar (17/17)
Getting image source signatures
Copying blob
sha256:b042988a8be982f5747f4ebc010f18f0f06dcac5a5ee4e4302904944a930665c
Copying blob
sha256:9ca4978d2b98bdfd2b55077e7aa15c53e76f1e3d3fb083f59ead330f9c8b201b
Copying blob
sha256:3e17c9d6bc285418e1bcefd2fb1278c7d9fcb5c6380bfeea3bd64ed0bafccae2
Copying blob
sha256:b47ddfaa51172f9d7d0c7cf9c59b06d86bed0dfe4eb120afd450c799043f3bb1
Copying blob
sha256:a39cb5f9504dec02f899b0eea1d778c290a7c0b725b6dedb2fb34a0debda415f
Copying blob
sha256:f209e77e229cfbefa77bb0c09a6c22ae0449d6dc46fc72e01036e8ca26ba87af
Copying blob
sha256:7336912f24443d080ffb5b320ca56529e8286534c4e6108fbf0bc65ec6f5a9ff
Copying blob
sha256:935e174f20de07a5ea843b3352961834d0181d023fe4d90c9082ecac37bfa659
```

```
Copying blob
sha256:35b9a4278afae0bb7ea1cfc563615482e67477c80b94fb0e903cddc6d8542ffa
Copying blob
sha256:80f8b4ea1b3dce9aaa6a280befc5234a5a389a10985098fca822d67269f7cc8f
Copying config
sha256:ca3fa91dc9dbc9b84a7dd2c151957f2a7732cf80d36eda757d4a327cf5723fac
Writing manifest to image destination
Storing signatures
[INFO] [2020-06-03 05:59:20-0564] All images are now loaded to your registry, the
registry is ready to perform offline install
```

The next step is to apply the admin setup, as shown in Example 4-4.

Example 4-4 Applying admin setup on the namespace

```
[root@client ~]# ./cpd-ppc64le adm --assembly lite --arch ppc64le --version 3.0.1
--namespace zen --load-from /repository/lite --apply
[INFO] [2020-06-03 05:59:37-0525] 0 file servers and 0 registries detected from
current configuration.
[INFO] [2020-06-03 05:59:37-0530] Server configure files validated

[INFO] [2020-06-03 05:59:37-0540] Assembly version is validated

* Parsing assembly data and generating a list of charts and images for download *
[INFO] [2020-06-03 05:59:37-0548] Assembly data validated
[INFO] [2020-06-03 05:59:38-0046] The category field of module 0010-infra is not
specified in its manifest file, assuming default type 'helm-chart'
[INFO] [2020-06-03 05:59:38-0048] The category field of module 0015-setup is not
specified in its manifest file, assuming default type 'helm-chart'
[INFO] [2020-06-03 05:59:38-0050] The category field of module 0020-core is not
specified in its manifest file, assuming default type 'helm-chart'
.
.
.
[INFO] [2020-06-03 05:59:46-0201]
securitycontextconstraints.security.openshift.io/cpd-zensys-scc added to:
["system:serviceaccount:zen:cpd-admin-sa"]

[INFO] [2020-06-03 05:59:46-0203] Finished executing requests

*** Executing add Cluster Role to SA requests ***

[INFO] [2020-06-03 05:59:46-0210] Finished executing requests

[INFO] [2020-06-03 05:59:46-0213] Admin setup executed successfully
```

Finally, you are ready to install the assembly for the lite service, as shown in Example 4-5 on page 69.

Example 4-5 Installing the lite assembly

```
[root@client ~]# ./cpd-ppc64le -a lite --arch ppc64le -c managed-nfs-storage -n
zen --cluster-pull-prefix=image-registry.openshift-image-registry.svc:5000/zen
--cluster-pull-username=kubeadmin --cluster-pull-password=$(oc whoami -t)
--insecure-skip-tls-verify --load-from=/repository/lite
```

```
[INFO] [2020-06-03 06:01:31-0974] --download-path not specified, creating
workspace cpd-ppc64le-workspace
[INFO] [2020-06-03 06:01:33-0154] Detected root certificate in kube config.
Ignoring --insecure-skip-tls-verify flag
[INFO] [2020-06-03 06:01:33-0179] Detected root certificate in kube config.
Ignoring --insecure-skip-tls-verify flag
[INFO] [2020-06-03 06:01:33-0599] 0 file servers and 0 registries detected from
current configuration.
[INFO] [2020-06-03 06:01:33-0601] Server configure files validated

[INFO] [2020-06-03 06:01:33-0603] Version for assembly is not specified, using the
latest version '3.0.1' for assembly lite

[INFO] [2020-06-03 06:01:33-0604] Verifying the CR apiVersion is expected. This
process could take up to 30 minutes
*** Parsing assembly data and generating a list of charts and images for download
***

[INFO] [2020-06-03 06:01:33-0778] Assembly data validated
[INFO] [2020-06-03 06:01:33-0949] The category field of module 0010-infra is not
specified in its manifest file, assuming default type 'helm-chart'
[INFO] [2020-06-03 06:01:33-0949] The category field of module 0015-setup is not
specified in its manifest file, assuming default type 'helm-chart'
[INFO] [2020-06-03 06:01:33-0950] The category field of module 0020-core is not
specified in its manifest file, assuming default type 'helm-chart'
.
.
.
2020-06-03 06:10:01.15394324 -0400 EDT m=+509.315738297
Module                                  Arch       Version             Status
0010-infra                              ppc64le    3.0.1               Ready
0015-setup                              ppc64le    3.0.1               Ready
0020-core                               ppc64le    3.0.1               Ready
[INFO] [2020-06-03 06:10:01-0993] Access the web console at
https://zen-cpd-zen.apps.ocp4.ibm.lab

*** Initializing version configmap for assembly lite ***

[INFO] [2020-06-03 06:10:02-0249] Assembly configmap update complete

[INFO] [2020-06-03 06:10:02-0251] *** Installation for assembly lite completed
successfully ***
```

Your IBM Cloud Pak for Data is ready and accessible. Access the web console URL by using a web browse as the installer installation instruction shows in Example 4-5. You can now log in to the console and see the window, as shown in Figure 4-3 on page 71.

Note: The default user name is `admin` and the password is `password`.

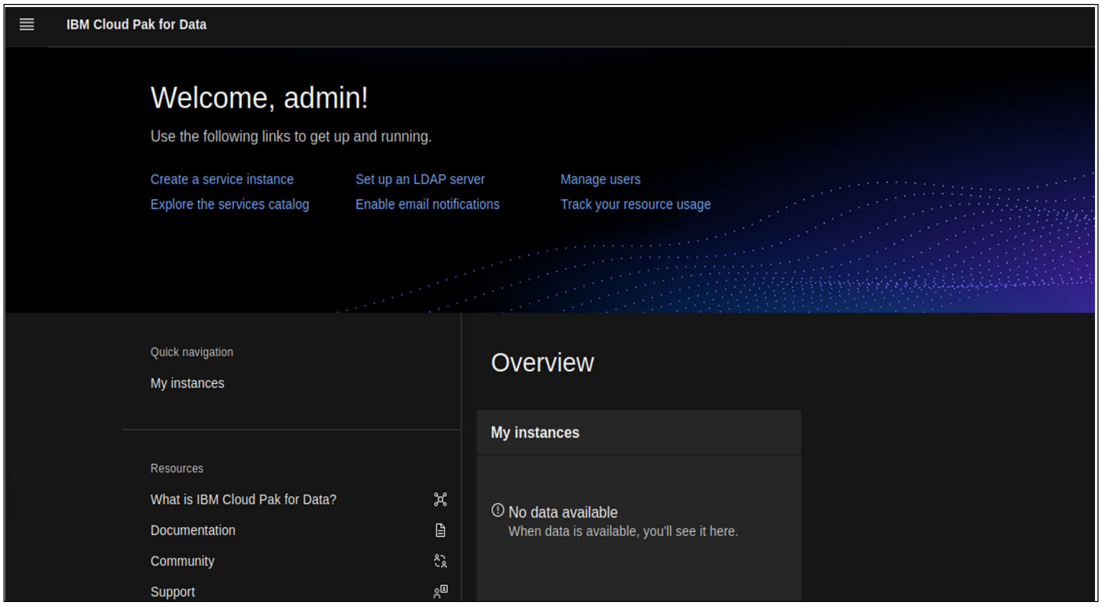

Figure 4-3 IBM Cloud Pak for Data: Welcome window

You can repeat the steps from Example 4-2 on page 66 - Example 4-4 on page 69 to install other available assemblies, including the following examples:

- aiopenscale
- cde
- db2oltp
- db2wh
- dods
- hadoop
- hadoop-addon
- lite
- rstudio
- spark
- spss
- spss-modeler
- wml
- wsl

Remember to change `lite` for the service that you want to install.

4.4.4 IBM Cloud Pak for Data backup

Backups on most IBM Cloud Pak for Data services are done quiescing the applications by scaling the applications down to 0, as described at IBM Knowledge Center.

If you use Spectrum Scale as storage for IBM Cloud Pak for Data persistent volumes that use Cluster Export Services, you can use the snapshot feature to enable the backup.

Note: You must be aware of any exception because you might want some online backup features for specific applications. For example, IBM Db2 and Db2 Warehouse provide a method for online backup inside the container. For more information, see the following IBM Knowledge Center web pages:

- Backing up and restoring Db2
- Backing up and restoring Db2 Warehouse

A

Configuring Red Hat CoreOS

This appendix shows how to create a machineconfig configuration that is beneficial to use with IBM Power Systems and helps with the requirements for IBM Cloud Pak for Data.

The changes that are provided in this appendix automatically set:

- Correct kernel value for `slub_max_order` parameter for IBM Cloud Pak for Data
- Correct number of open files and pids on `crio.conf` for IBM Cloud Pak for Data
- Automatic setting for SMT with labels
- Correct sysctl values for IBM Cloud Pak for Data

This appendix includes the following topics:

CoreOS machine configuration management and machineconfig objects

When you log in to the CoreOS by way of SSH, the message says to use the machineconfig objects, as shown in Example A-1.

Example A-1 Message of the day on a regular CoreOS node

```
[root@client ~]# ssh core@worker1
Red Hat Enterprise Linux CoreOS 43.81.202004201335.0
  Part of OpenShift 4.3, RHCOS is a Kubernetes native operating system
  managed by the Machine Config Operator (`clusteroperator/machine-config`).

WARNING: Direct SSH access to machines is not recommended; instead,
make configuration changes via `machineconfig` objects:
  https://docs.openshift.com/container-platform/4.3/architecture/architecture-rhco
s.html

---
Last login: Tue Jun  2 12:00:39 2020 from 10.17.201.99
[core@worker1 ~]$
```

Machine objects can be created by using YAML files and enable systemd services or change files on disk. The basic YAML structure is shown in Example A-2.

Example A-2 Basic structure of a machineconfig object

```
apiVersion: machineconfiguration.openshift.io/v1
kind: MachineConfig
metadata:
  labels:
    machineconfiguration.openshift.io/role: master
  name: <OBJECT NAME>
spec:
  kernelArguments:
    - '<PARAMETER>=<VALUE>'
  config:
    ignition:
      version: 2.2.0
    storage:
      files:
        - path: /path/to/file
          overwrite: true
          mode: 0700
          filesystem: root
          contents:
            source: data:text/plain;charset=utf-8;base64,<BASE64 CONTENT>
systemd:
      units:
        - name: <SERVICENAME>.service
          enabled: true
```

We use the base64 content for text files because it eases the occurrence of any special characters and non-printable characters resulting in a single stream.

For our example, we configure the crio daemon, create a service that watches for an SMT label on the node, and apply the wanted SMT configuration by using that label. We also set the `slub_max_order` kernel parameter to 0.

The crio changes are necessary because some workloads include busy containers, and need more pids and open files than the allocated by default. We do not paste the full `crio.conf` because we changed only one parameter from the default file `pids_limit = 1024` and added a ulimit parameter, as shown in Example A-3.

Example A-3 Changes on the crio.conf default file

```
pids_limit = 12288

# Adding nofiles for CP4D
default_ulimits = [
        "nofile=66560:66560"
]
```

We create two files to control the SMT across the cluster, as shown in Example A-4. For more information about how Kubernetes uses CPU, see 2.4.1, “Red Hat OpenShift sizing guidelines” on page 22.

Example A-4 Systemd file to run powersmt service

```
[root@client ~]# ssh core@ultra cat /etc/systemd/system/powersmt.service
[Unit]
Description=POWERSMT
After=network-online.target

[Service]
ExecStart="/usr/local/bin/powersmt"

[Install]
WantedBy=multi-user.target
```

Example A-4 shows the reference to the file `/usr/local/bin/powersmt`. We need to create that file to set the parameters correctly, as shown in Example A-6 on page 76.

You must transform the file to change any special characters and spaces (that are meaningful in YAML files) to a simple string in a single line. The process is shown in Example A-5.

> **Note:** During the project, we attempted different ways to create the YAML file and use base64 was the easiest way to overcome all situations on the YAML file definition. It is for this reason that we used this format and we encourage users to use this format as well.

Example A-5 Transforming a file in base64 string

```
[root@client ~]# cat /etc/systemd/system/powersmt.service |base64 |xargs echo |sed
's/ //g'
W1VuaXRdCkRlc2NyaXBOaW9uPVBPVOVSUO1UCkFmdGVyPW5ldHdvcmstb25saW5lLnRhcmdldAoKW1Nlcn
ZpY2VdCkV4ZWNTdGFydDOiL3Vzci9sb2NhbC9iaW4vcG93ZXJzbXQiCgpbSW5zdGFsbFOKV2FudGVkQnk9
bXVsdGktdXNlci5OYXJnZXQK
```

> **Note:** We do not show this process for other files, but if you want to use base64 for the source of the file contents, repeat the process for any file you created.

Example A-6 shows a bash loop that uses the **lscpu** command output and the SMT label that is defined for the node where it is running to decide whether to get CPU offline or online. This script checks the value that is always smaller than 2 for x86 architectures, 4 for OpenPower, and 8 for Power Systems servers with PowerVM capabilities.

Example A-6 Service executable for powersmt service

```
#!/bin/bash
export PATH=/root/.local/bin:/root/bin:/sbin:/bin:/usr/local/sbin:/usr/local/bin:/
usr/sbin:/usr/bin
export KUBECONFIG=/var/lib/kubelet/kubeconfig
COREPS=$(/bin/lscpu | /bin/awk -F: ' $1 ~ /^Core\(s\) per socket$/ {print $2}'|/bi
n/xargs)
SOCKETS=$(/bin/lscpu | /bin/awk -F: ' $1 ~ /^Socket\(s\)$/ {print $2}'|/bin/xargs)
let TOTALCORES=$COREPS*$SOCKETS
MAXTHREADS=$(/bin/lscpu | /bin/awk -F: ' $1 ~ /^CPU\(s\)$/ {print $2}'|/bin/xargs)
let MAXSMT=$MAXTHREADS/$TOTALCORES
CURRENTSMT=$(/bin/lscpu | /bin/awk -F: ' $1 ~ /^Thread\(s\) per core$/ {print $2}'
|/bin/xargs)

while :
do
  ISNODEDEGRADED=$(/bin/oc get node $HOSTNAME -o yaml |/bin/grep machineconfigurat
ion.openshift.io/reason |/bin/grep "unexpected on-disk state validating")
  SMTLABEL=$(/bin/oc get node $HOSTNAME -L SMT --no-headers |/bin/awk '{print $6}'
)
  if [[ -n $SMTLABEL ]]
    then
      case $SMTLABEL in
        1) TARGETSMT=1
      ;;
        2) TARGETSMT=2
      ;;
        4) TARGETSMT=4
      ;;
        8) TARGETSMT=8
      ;;
        *) TARGETSMT=$CURRENTSMT ; echo "SMT value must be 1, 2, 4, or 8 and small
er than Maximum SMT."
      ;;
      esac
    else
      TARGETSMT=$MAXSMT
  fi

  if [[ -n $ISNODEDEGRADED ]]
    then
      touch /run/machine-config-daemon-force
  fi

  CURRENTSMT=$(/bin/lscpu | /bin/awk -F: ' $1 ~ /^Thread\(s\) per core$/ {print $2
}'|/bin/xargs)

  if [[ $CURRENTSMT -ne $TARGETSMT ]]
    then
      INITONTHREAD=0
```

```
      INITOFFTHREAD=$TARGETSMT
      if [[ $MAXSMT -ge $TARGETSMT ]]
        then
          while [[ $INITONTHREAD -lt $MAXTHREADS ]]
          do
            ONTHREAD=$INITONTHREAD
            OFFTHREAD=$INITOFFTHREAD

            while [[ $ONTHREAD -lt $OFFTHREAD ]]
            do
              /bin/echo 1 > /sys/devices/system/cpu/cpu$ONTHREAD/online
              let ONTHREAD=$ONTHREAD+1
            done
            let INITONTHREAD=$INITONTHREAD+$MAXSMT
            while [[ $OFFTHREAD -lt $INITONTHREAD ]]
            do
              /bin/echo 0 > /sys/devices/system/cpu/cpu$OFFTHREAD/online
              let OFFTHREAD=$OFFTHREAD+1
            done
            let INITOFFTHREAD=$INITOFFTHREAD+$MAXSMT
          done
        else
          echo "Target SMT must be smaller or equal than Maximum SMT supported"
      fi
  fi
  /bin/sleep 30
done
```

The machineconfig file that is used after including all of these alterations is shown in Example A-7.

> **Important:** The file that is shown in Example A-7 applies the configuration to all worker nodes. This configuration does not need to be applied to the master nodes. If you must apply this configuration, change `worker` to `master` on both metadata stanza occurrences, and create two different YAML configuration files. Both configuration files must be applied to on all masters and workers nodes.
>
> Do *not* use this example for your installation; instead, build your own file. You must at minimum update the base64 strings from the ones in your `crio.config`. Remember that configuration entries can change over time; for example, the `pause_image`.
>
> By default, the nodes automatically reboot in rolling fashion, one-by-one, after you apply the configuration.

Example A-7 Machine configuration operator 99_openshift-machineconfig_99-worker-ibm.yaml file

```
apiVersion: machineconfiguration.openshift.io/v1
kind: MachineConfig
metadata:
  labels:
    machineconfiguration.openshift.io/role: worker
  name: 99-worker-ibm
spec:
  kernelArguments:
    - 'slub_max_order=0'
```

```
  config:
    ignition:
      version: 2.2.0
    storage:
      files:
        - path: /usr/local/bin/powersmt
          overwrite: true
          mode: 0700
          filesystem: root
          contents:
            source: data:text/plain;charset=utf-8;base64,IyEvYmluL2Jhc2gKZXhwb3J0I
FBBVEg9L3Jvb3QvLmxvY2FsL2Jpbjovcm9vdC9iaW46L3NiaW46L2JpbjovdXNyL2xvY2FsL3NiaW46L3V
zci9sb2NhbC9iaW46L3Vzci9zYmluOi91c3IvYmluCmV4cG9ydCBLVUJFQ09ORklHPS92YXIvbGliL2t1Y
mVsZXQva3ViZWNvbmZpZwpDT1JFUFM9JCgvYmluL2xzY3B1IHwgL2Jpbi9hd2sgLUY6ICcgJDEgfiAvXkN
vcmVcKHNcKSBwZXIgc29ja2V0JC8ge3ByaW50ICQyfSd8L2Jpbi94YXJncykKU09DS0VUUzOkKC9iaW4vb
HNjcHUgfCAvYmluL2F3ayAtRjogJyAkMSB+IC9eU29ja2V0XChzXCkkLyB7cHJpbnQgJDJ9J3wvYmluL3h
hcmdzKQpsZXQgVE9UQUxDT1JFUzOkQO9SRVBTKiRTT0NLRVRTCk1BWFRIUkVBRFM9JCgvYmluL2xzY3B1I
HwgL2Jpbi9hd2sgLUY6ICcgJDEgfiAvXkNQVVwoc1wpJC8ge3ByaW50ICQyfSd8L2Jpbi94YXJncykKbGV
OIE1BWFNNVDOkTUFYVEhSRUFEUy8kVE9UQUxDT1JFUwpDVVJSRU5UU01UPSQoL2Jpbi9sc2NwdSB8IC9ia
W4vYXdrIC1GOiAnICQxIH4gL15UaHJlYWRcKHNcKSBwZXIgY29yZSQvIHtwcmludCAkMnOnfC9iaW4veGF
yZ3MpCgp3aGlsZSA6CmRvCiAgSVNOOTORFREVHUkFERUQ9JCgvYmluL29jIGdldCBub2RlICRIT1NUTkFNR
SAtbyB5YW1sIHwvYmluL2dyZXAgbWFjaGluZWNvbmZpZ3VyYXRpb24ub3BlbnNoaWZ0LmlvL3JlYXNvbiiB
8L2Jpbi9ncmVwICJ1bmV4cGVjdGVkIG9uLWRpc2sgc3RhdGUgdmFsaWRhdGluZyIpCiAgUO1UTEFCRUw9J
CgvYmluL29jIGdldCBub2RlICRIT1NUTkFNRSAtTCBTTVQgLS1uby1oZWFkZXJzIHwvYmluL2F3ayAne3B
yaW5OICQ2fScpCiAgaWYgW1sgLW4gJFNNVExBQkVMIF1dCiAgICBOaGVuCiAgICAgIGNhc2UgJFNNVExBQ
kVMIGluCiAgICAgICAgMSkgVEFSROVUU01UPTEKICAgICAgOzsKICAgICAgICAyKSBUQVJHRVRTTVQ9Mgo
gICAgICA7OwogICAgICAgIDQpIFRBUkdFVFNNVDOOCiAgICAgIDs7CiAgICAgICAgOCkgVEFSROVUU01UP
TgKICAgICAgOzsKICAgICAgICAqKSBUQVJHRVRTTVQ9JENVUlJFTlRTTVQgOyBlY2hvICJTTVQgdmFsdWU
gbXVzdCBiZSAxLCAyLCAOLCBvciA4IGFuZCBzbWFsbGVyIHRoYW4gTWF4aW11bSBTTVQuIgogICAgICA7O
wogICAgICBlc2FjCiAgICBlbHNlCiAgICAgIFRBUkdFVFNNVDOkTUFYU01UCiAgZmkKCiAgaWYgW1sgLW4
gJElTTk9ERURFR1JBREVEIF1dCiAgICBOaGVuCiAgICAgIHRvdWNoIC9ydW4vbWFjaGluZS1jb25maWctZ
GFlbW9uLWZvcmNlCiAgZmkKCiAgQ1VSUkVOVFNNVDOkKC9iaW4vbHNjcHUgfCAvYmluL2F3ayAtRjogJyA
kMSB+IC9eVGhyZWFkXChzXCkgcGVyIGNvcmUkLyB7cHJpbnQgJDJ9J3wvYmluL3hhcmdzKQoKICBpZiBbW
yAkQ1VSUkVOVFNNVCAtbmUgJFRBUkdFVFNNVCBdXQogICAgdGhlbgogICAgICBJTklUT05USFJFQUQ9MAo
gICAgICBJTklUOZGVEhSRUFEPSRUQVJHRVRTTVQKICAgICAgaWYgW1sgJE1BWFNNVCAtZ2UgJFRBUkdFV
FNNVCBdXQogICAgICAgIHRoZW4KICAgICAgICAgIHdoaWxlIFtbICRJTklUO5USFJFQUQgLWxOICRNQVh
USFJFQURTIF1dCiAgICAgICAgICBkbwogICAgICAgICAgICBPTlRIUkVBRDOkSU5JVE9OVEhSRUFECiAgI
CAgICAgICAgIE9GRlRIUkVBRDOkSU5JVE9GRlRIUkVBRAoKICAgICAgICAgICAgd2hpbGUgW1sgJE9OVEh
SRUFEIC1sdCAkTOZGVEhSRUFEIF1dCiAgICAgICAgICAgIGRvCiAgICAgICAgICAgICAgL2Jpbi9lY2hvI
DEgPiAvc3lzL2RldmljZXMvc3lzdGVtL2NwdS9jcHUkTO5USFJFQUQvb25saW5lCiAgICAgICAgICAgICA
gbGVOIE9OVEhSRUFEPSRPTlRIUkVBRCsxCiAgICAgICAgICAgIGRvbmUKICAgICAgICAgICAgbGVOIElOS
VRPTlRIUkVBRDOkSU5JVE9OVEhSRUFEKyRNQVhTTVQKICAgICAgICAgICAgd2hpbGUgW1sgJE9GRlRIUkV
BRCAtbHQgJElOSVRPTlRIUkVBRCBdXQogICAgICAgICAgICBkbwogICAgICAgICAgICAgIC9iaW4vZWNob
yAwID4gL3N5cy9kZXZpY2VzL3N5c3RlbS9jcHUvY3B1JE9GRlRIUkVBRC9vbmxpbmUKICAgICAgICAgICA
gICBsZXQgTOZGVEhSRUFEPSRPRkZUSFJFQUQrMQogICAgICAgICAgICBkb25lCiAgICAgICAgICAgIGxld
CBJTklUOZGVEhSRUFEPSRJTklUOZGVEhSRUFEKyRNQVhTTVQKICAgICAgICAgIGRvbmUKICAgICAgICB
lbHNlCiAgICAgICAgICBlY2hvICJUYXJnZXQgU01UIG11c3QgYmUgc21hbGxlciBvciBlcXVhbCBOaGFuI
E1heGltdW0gU01UIHN1cHBvcnRlZCIKICAgICAgZmkKICBmaQogIC9iaW4vc2xlZXAgMzAKZG9uZQo=
        - path: /etc/systemd/system/powersmt.service
          overwrite: true
          mode: 0644
          filesystem: root
          contents:
```

```
            source: data:text/plain;charset=utf-8;base64,W1VuaXRdCkRlc2NyaXB0aW9uP
VBPV0VSU01UCkFmdGVyPW5ldHdvcmstb25saW5lLnRhcmdldAoKW1NlcnZpY2VdCkV4ZWNTdGFydD0iL3V
zci9sb2NhbC9iaW4vcG93ZXJzbXQiCgpbSW5zdGFsbF0KV2FudGVkQnk9bXVsdGktdXNlci50YXJnZXQK
        - path: /etc/crio/crio.conf
          overwrite: true
          mode: 0644
          filesystem: root
          contents:
            source: data:text/plain;charset=utf-8;base64,W2NyaW9dCgojIFRoZSBkZWZhd
WxOIGxvZyBkaXJlY3Rvcnkgd2hlcmUgYWxsIGxvZ3Mgd2lsbCBnbyB1bmxlc3MgZGlyZWNObHkgc3BlY2l
maWVkIGJ5CiMgdGhlIGt1YmVsZXQuIFRoZSBsb2cgZGlyZWNOb3J5IHNwZWNpZmllZCBtdXNOIGJlIGFuI
GFic29sdXRlIGRpcmVjdG9yeS4KbG9nX2RpciA9ICIvdmFyL2xvZy9jcmlvL3BvZHMiCgojIExvY2FOaW9
uIGZvciBDUkktTyB0byBsYXkgZG93biBOaGUgdmVyc2lvbiBmaWxlCnZlcnNpb25fZmlsZSA9ICIvdmFyL
2xpYi9jcmlvL3ZlcnNpb24iCgojIFRoZSBjcmlvLmFwaSBOYWJsZSBjb25OYWlucyBzZXROaW5ncyBmb3I
gdGhlIGt1YmVsZXQvZ1JQQyBpbnRlcmZhY2UuCltjcmlvLmFwaV0KCiMgUGF0aCBObyBBRl9MTONBTCBzb
2NrZXQgb24gd2hpY2ggQ1JJLU8gd2lsbCBsaXNOZW4uCmxpc3RlbiA9ICIvdmFyL3J1bi9jcmlvL2NyaW8
uc29jayIKCiMgSG9zdCBJUCBjb25zaWRlcmVkIGFzIHRoZSBwcmltYXJ5IElQIHRvIHVzZSBieSBDUkktT
yBmb3IgdGhpbmdzIHN1Y2ggYXMgaG9zdCBuZXR3b3JrIElQLgpob3NOX2lwID0gIiIKCiMgSVAgYWRkcmV
zcyBvbiB3aGljaCBOaGUgc3RyZWFtIHNlcnZlciB3aWxsIGxpc3Rlbi4Kc3RyZWFtX2FkZHJlc3MgPSAiI
goKIyBUaGUgcG9ydCBvbiB3aGljaCBOaGUgc3RyZWFtIHNlcnZlciB3aWxsIGxpc3Rlbi4Kc3RyZWFtX3B
vcnQgPSAiMTAwMTAiCgojIEVuYWJsZSBlbmNyeXBOZWQgVExTIHRyYW5zcG9ydCBvZiBOaGUgc3RyZWFtI
HNlcnZlci4Kc3RyZWFtX2VuYWJsZV9ObHMgPSBmYWxzZQoKIyBQYXRoIHRvIHRoZSB4NTA5IGNlcnRpZml
jYXRlIGZpbGUgdXNlZCBObyBzZXJ2ZSBOaGUgZW5jcnlwdGVkIHNOcmVhbS4gVGhpcwojIGZpbGUgY2FuI
GNoYW5nZSwgYW5kIENSSS1PIHdpbGwgYXVOb21hdGljYWxseSBwaWNrIHVwIHRoZSBjaGFuZ2VzIHdpdGh
pbiA1CiMgbWludXRlcy4Kc3RyZWFtX3Rsc19jZXJOIDOgIiIKCiMgUGF0aCBObyBOaGUga2V5IGZpbGUgd
XNlZCBObyBzZXJ2ZSBOaGUgZW5jcnlwdGVkIHNOcmVhbS4gVGhpcyBmaWxlIGNhbgojIGNoYW5nZSBhbmQ
gQ1JJLU8gd2lsbCBhdXRvbWFOaWNhbGx5IHBpY2sgdXAgdGhlIGNoYW5nZXMgd2lOaGluIDUgbWludXRlc
y4Kc3RyZWFtX3Rsc19rZXkgPSAiIgoKIyBQYXRoIHRvIHRoZSB4NTA5IENBKHMpIGZpbGUgdXNlZCBObyB
2ZXJpZnkgYW5kIGF1dGhlbnRpY2FOZSBjbGllbnQKIyBjb21tdW5pY2FOaW9uIHdpdGggdGhlIGVuY3J5c
HRlZCBzdHJlYW0uIFRoaXMgZmlsZSBjYW4gY2hhbmdlIGFuZCBDUkktTyB3aWxsCiMgYXVOb21hdGljYWx
seSBwaWNrIHVwIHRoZSBjaGFuZ2VzIHdpdGhpbiA1IG1pbnVOZXMuCnNOcmVhbV9ObHNfY2EgPSAiIgoKI
yBNYXhpbXVtIGdycGMgc2VuZCBtZXNzYWdlIHNpemUgaW4gYnlOZXMuIElmIG5vdCBzZXQgb3IgPDOwLCB
OaGVuIENSSS1PIHdpbGwgZGVmYXVsdCBObyAxNiAqIDEwMjQgKiAxMDI0LgpncnBjX21heF9zZW5kX21zZ
19zaXplIDOgMTY3NzcyMTYKCiMgTWF4aW11bSBncnBjIHJlY2VpdmUgbWVzc2FnZSBzaXplLiBJZiBub3Q
gc2V0IG9yIDw9IDAsIHRoZW4gQ1JJLU8gd2lsbCBkZWZhdWxOIHRvIDE2ICogMTAyNCAqIDEwMjQuCmdyc
GNfbWF4X3JlY3ZfbXNnX3NpemUgPSAxNjc3NzIxNgoKIyBUaGUgY3Jpby5ydW50aW1lIHRhYmxlIGNvbnR
haW5zIHNldHRpbmdzIHBlcnRhaW5pbmcgdG8gdGhlIE9DSSBydW50aW1lIHVzZWQKIyBhbmQgb3B0aW9uc
yBmb3IgaG93IHRvIHNldCB1cCBhbmQgbWFuYWdlIHRoZSBPQ0kgcnVudGltZS4KW2NyaW8ucnVudGltZV0
KCiMgZGVmYXVsdF9ydW50aW1lIGlzIHRoZSBfbmFtZV8gb2YgdGhlIE9DSSBydW50aW1lIHRvIGJlIHVzZ
WQgYXMgdGhlIGRlZmF1bHQuCiMgVGhlIG5hbWUgaXMgbWF0Y2hlZCBhZ2FpbnNOIHRoZSBydW50aW1lcyB
tYXAgYmVsb3cuCmRlZmF1bHRfcnVudGltZSA9ICJydW5jIgoKIyBJZiBOcnVlLCBOaGUgcnVudGltZSB3a
WxsIG5vdCB1c2UgcGl2b3Rfcm9vdCwgYnVOIGluc3RlYWQgdXNlIE1TX01PVkUuCm5vX3Bpdm9OIDOgZmF
sc2UKCiMgUGF0aCBObyBOaGUgY29ubW9uIGJpbmFyeSwgdXNlZCBmb3IgbW9uaXRvcmluZyBOaGUgT0NJI
HJ1bnRpbWUuCiMgV2lsbCBiZSBzZWFyY2hlZCBmb3IgdXNpbmcgJFBBVEggaWYgZW1wdHkuCmNvbm1vbiA
9ICIvdXNyL2xpYmV4ZWMvY3Jpby9jb25tb24iCgojIENncm91cCBzZXROaW5nIGZvciBjb25tb24KY29ub
W9uX2Nncm91cCA9ICJwb2QiCgojIEVudmlyb25tZW5OIHZhcmlhYmxlIGxpc3QgZm9yIHRoZSBjb25tb24
gcHJvY2VzcywgdXNlZCBmb3IgcGFzc2luZyBuZWNlc3NhcnkKIyBlbnZpcm9ubWVudCB2YXJpYWJsZXMgd
G8gY29ubW9uIG9yIHRoZSBydW50aW1lLgpjb25tb25fZW52IDOgWwogICAgIlBBVEg9L3Vzci9sb2NhbC9
zYmluOi91c3IvbG9jYWwvYmluOi91c3Ivc2JpbjovdXNyL2Jpbjovc2JpbjovYmluIiwKXQoKIyBJZiBOc
nVlLCBTRUxpbnV4IHdpbGwgYmUgdXNlZCBmb3IgcG9kIHNlcGFyYXRpb24gb24gdGhlIGhvc3QuCnNlbGl
udXggPSBOcnVlCgojIFBhdGggdG8gdGhlIHNlY2NvbXAuanNvbiBwcm9maWxlIHdoaWNoIGlzIHVzZWQgY
XMgdGhlIGRlZmF1bHQgc2VjY29tcCBwcm9maWxlCiMgZm9yIHRoZSBydW50aW1lLiBJZiBub3Qgc3BlY2l
maWVkLCBOaGVuIHRoZSBpbnRlcm5hbCBkZWZhdWxOIHNlY2NvbXAgcHJvZmlsZQojIHdpbGwgYmUgdXNlZ
C4Kc2VjY29tcF9wcm9maWxlIDOgIiIKCiMgVXNlZCBObyBjaGFuZ2UgdGhlIG5hbWUgb2YgdGhlIGRlZmF
1bHQgQXBwQXJtb3IgcHJvZmlsZSBvZiBDUkktTy4gVGhlIGRlZmF1bHQKIyBwcm9maWxlIG5hbWUgaXMgI
```

mNyaW8tZGVmYXVsdC0iIGZvbGxvd2VkIGJ5IHRoZSB2ZXJzaW9uIHN0cmluZyBvZiBDUkktTy4KYXBwYXJ
tb3JfcHJvZmlsZSA9ICJjcmlvLWRlZmF1bHQiCgojIENncm91cCBtYW5hZ2VtZW50IGltcGxlbWVudGF0a
W9uIHVzZWQgZm9yIHRoZSBydW50aW1lLgpjZ3JvdXBfbWFuYWdlciA9ICJzeXN0ZW1kIgoKIyBMaXN0IG9
mIGRlZmF1bHQgY2FwYWJpbGl0aWVzIGZvciBjb250YWluZXJzLiBJZiBpdCBpcyBlbXB0eSBvciBjb21tZ
W50ZWQgb3V0LAojIG9ubHkgdGhlIGNhcGFiaWxpdGllcyBkZWZpbmVkIGluIHRoZSBjb250YWluZXJzIGp
zb24gZmlsZSBieSB0aGUgdXNlci9rdWJlCiMgd2lsbCBiZSBhZGRlZC4KZGVmYXVsdF9jYXBhYmlsaXRpZ
XMgPSBbCiAgICAiQ0hPV04iLAogICAgIkRBQ19PVkVSUklERSIsCiAgICAiRlNFVElEIiwKICAgICJGT1d
ORVIiLAogICAgIk5FVF9SQVciLAogICAgIlNFVEdJRCIsCiAgICAiU0VUVUlEIiwKICAgICJTRVRQQ0FQI
iwKICAgICJORVRfQklORF9TRVJWSUNFIiwKICAgICJTWVNfQ0hST09UIiwKICAgICJLSUxMIiwKXQoKIyB
MaXN0IG9mIGRlZmF1bHQgc3lzY3Rscy4gSWYgaXQgaXMgZW1wdHkgb3IgY29tbWVudGVkIG91dCwgb25se
SB0aGUgc3lzY3RscwojIGRlZmluZWQgaW4gdGhlIGNvbnRhaW5lciBqc29uIGZpbGUgYnkgdGhlIHVzZXI
va3ViZSB3aWxsIGJlIGFkZGVkLgpkZWZhdWx0X3N5c2N0bHMgPSBbCl0KCiMgTGlzdCBvZiBhZGRpdGlvb
mFsIGRldmljZXMuIHNwZWNpZmllZCBhcwojICI8ZGV2aWNlLW9uLWhvc3Q+OjxkZXZpY2Utb24tY29udGF
pbmVyPjo8cGVybWlzc2lvbnM+IiwgZm9yIGV4YW1wbGU6ICItLWRldmljZT0vZGV2L3NkYzovZGV2L3h2Z
GM6cndtIi4KI0lmIGl0IGlzIGVtcHR5IG9yIGNvbW1lbnRlZCBvdXQsIG9ubHkgdGhlIGRldmljZXMKIyB
kZWZpbmVkIGluIHRoZSBjb250YWluZXIganNvbiBmaWxlIGJ5IHRoZSB1c2VyL2t1YmUgd2lsbCBiZSBhZ
GRlZC4KYWRkaXRpb25hbF9kZXZpY2VzID0gWwpdCgojIFBhdGggdG8gT0NJIGhvb2tzIGRpcmVjdG9yaWV
zIGZvciBhdXRvbWF0aWNhbGx5IGV4ZWN1dGVkIGhvb2tzLgpob29rc19kaXIgPSBbCiAgICAiL2V0Yy9jb
250YWluZXJzL29jaS9ob29rcy5kIiwKXQoKIyBMaXN0IG9mIGRlZmF1bHQgbW91bnRzIGZvciBlYWNoIGN
vbnRhaW5lci4gKipEZXByZWNhdGVkOioqIHRoaXMgb3B0aW9uIHdpbGwgKIyBiZSByZW1vdmVkIGluIGZ1d
HVyZSB2ZXJzaW9ucyBpbiBmYXZvciBvZiBkZWZhdWx0X21vdW50c19maWxlLgpkZWZhdWx0X21vdW50cyA
9IFsKXQoKIyBNYXhpbXVtIG51bWJlciBvZiBwcm9jZXNzZXMgYWxsb3dlZCBpbiBhIGNvbnRhaW5lci4Kc
Glkc19saW1pdCA9IDEyMjg4CgojIEFkZGluZyBub2ZpbGVzIGZvciBDUkRECmRlZmF1bHRfdWxpbWl0cyA
9IFsKICAgICAgICAibm9maWxlPTY2NTYwOjY2NTYwIgpdCgojIE1heGltdW0gc2l6ZWQgYWxsb3dlZCBmb
3IgdGhlIGNvbnRhaW5lciBsb2cgZmlsZS4gTmVnYXRpdmUgbnVtYmVycyBpbmRpY2F0ZQojIHRoYXQgbm8
gc2l6ZSBsaW1pdCBpcyBpbXBvc2VkLiBJZiBpdCBpcyBwb3NpdGl2ZSwgaXQgbXVzdCBiZSA+PSA4MTkyI
HRvCiMgbWF0Y2gvZXhjZWVkIGNvbm1vbidzIHJlYWQgYnVmZmVyLiBUaGUgZmlsZSBpcyB0cnVuY2F0ZWQ
gYW5kIHJlLW9wZW5lZCBzbyB0aGUKIyBsaW1pdCBpcyBuZXZlciBleGNlZWRlZC4KbG9nX3NpemVfbWF4I
D0gLTEKCiMgV2hldGhlciBjb250YWluZXIgb3V0cHV0IHNob3VsZCBiZSBsb2dnZWQgdG8gam91cm5hbGQ
gaW4gYWRkaXRpb24gdG8gdGhlIGt1YmVyZW50ZXMgbG9nIGZpbGUKbG9nX3RvX2pvdXJuYWxkID0gZmFsc
2UKCiMgUGF0aCB0byBkaXJlY3RvcnkgaW4gd2hpY2ggY29udGFpbmVyIGV4aXQgZmlsZXMgYXJlIHdyaXR
0ZW4gdG8gYnkgY29ubW9uLgpjb250YWluZXJfZXhpdHNfZGlyID0gIi92YXIvcnVuL2NyaW8vZXhpdHMiC
gojIFBhdGggdG8gZGlyZWN0b3J5IGZvciBjb250YWluZXIgYXR0YWNoIHNvY2tldHMuCmNvbnRhaW5lcl9
hdHRhY2hfc29ja2V0X2RpciA9ICIvdmFyL3J1bi9jcmlvIgoKIyBUaGUgcHJlZml4IHRvIHVzZSBmb3Igd
GhlIHNvdXJjZSBvZiB0aGUgYmluZCBtb3VudHMuCmJpbmRfbW91bnRfcHJlZml4ID0gIiIKCiMgSWYgc2V
0IHRvIHRydWUsIGFsbCBjb250YWluZXJzIHdpbGwgcnVuIGluIHJlYWQtb25seSBtb2RlLgpyZWFkX29ub
HkgPSBmYWxzZQoKIyBDaGFuZ2VzIHRoZSB2ZXJib3NpdHkgb2YgdGhlIGxvZ3MgYmFzZWQgb24gdGhlIGx
ldmVsIGl0IGlzIHNldCB0by4gT3B0aW9ucwojIGFyZSBmYXRhbCwgcGFuaWMsIGVycm9yLCB3YXJuLCBpb
mZvLCBhbmQgZGVidWcuIFRoaXMgb3B0aW9uIHN1cHBvcnRzIGxpdmUKIyBjb25maWd1cmF0aW9uIHJlbG9
hZC4KbG9nX2xldmVsID0gImVycm9yIgoKIyBUaGUgVUlEIG1hcHBpbmdzIGZvciB0aGUgdXNlciBuYW1lc
3BhY2Ugb2YgZWFjaCBjb250YWluZXIuIEEgcmFuZ2UgaXMKIyBzcGVjaWZpZWQgaW4gdGhlIGZvcm0gY29
udGFpbmVyVUlEOkhvc3RVSUQ6U2l6ZS4gTXVsdGlwbGUgcmFuZ2VzIG11c3QgYmUKIyBzZXBhcmF0ZWQgY
nkgY29tbWEuCnVpZF9tYXBwaW5ncyA9ICIiCgojIFRoZSBHSUQgbWFwcGluZ3MgZm9yIHRoZSB1c2VyIG5
hbWVzcGFjZSBvZiBlYWNoIGNvbnRhaW5lci4gQSByYW5nZSBpcwojIHNwZWNpZmllZCBpbiB0aGUgZm9yb
SBjb250YWluZXJHSUQ6SG9zdEdJRDpTaXplLiBNdWx0aXBsZSByYW5nZXMgbXVzdCBiZQojIHNlcGFyYXR
lZCBieSBjb21tYS4KZ2lkX21hcHBpbmdzID0gIiIKCiMgVGhlIG1pbmltYWwgYW1vdW50IG9mIHRpbWUga
W4gc2Vjb25kcyB0byB3YWl0IGJlZm9yZSBpc3N1aW5nIGEgdGltZW91dAojIHJlZ2FyZGluZyB0aGUgcHJ
vcGVyIHRlcm1pbmF0aW9uIG9mIHRoZSBjb250YWluZXIuCmN0cl9zdG9wX3RpbWVvdXQgPSAwCgojIE1hb
mFnZU5ldHdvcmtOU0xpZmVjeWNsZSBkZXRlcm1pbmVzIHdoZXRoZXIgd2UgcGluIGFuZCByZW1vdmUgbmV
0d29yayBuYW1lc3BhY2UKIyBhbmQgbWFuYWdlIGl0cyBsaWZlY3ljbGUuCm1hbmFnZV9uZXR3b3JrX25zX
2xpZmVjeWNsZSA9IGZhbHNlCgojIFRoZSAiY3Jpby5ydW50aW1lLnJ1bnRpbWVzIiB0YWJsZSBkZWZpbmV
zIGEgbGlzdCBvZiBPQ0kgY29tcGF0aWJsZSBydW50aW1lcy4KIyBUaGUgcnVudGltZSB0byB1c2UgaXMgc
Glja2VkIGJhc2VkIG9uIHRoZSBydW50aW1lX2hhbmRsZXIgcHJvdmlkZWQgYnkgdGhlIENSSS4KIyBJZiB
ubyBydW50aW1lX2hhbmRsZXIgaXMgcHJvdmlkZWQsIHRoZSBydW50aW1lIHdpbGwgYmUgcGlja2VkIGJhc
2VkIG9uIHRoZSBsZXZlbAojIG9mIHRydXN0IG9mIHRoZSB3b3JrbG9hZC4gRWFjaCBlbnRyeSBpbiB0aGU

gdGFibGUgc2hvdWxkIGZvbGxvdyB0aGUgZm9ybWF0OgojCiNbY3Jpby5ydW50aW1lLnJ1bnRpbWVzLnJ1b
nRpbWUtaGFuZGxlcl0KIyAgcnVudGltZV9wYXRoID0gIi9wYXRoL3RvL3RoZS9leGVjdXRhYmxlIgojICB
ydW50aW1lX3R5cGUgPSAib2NpIgojICBydW50aW1lX3Jvb3QgPSAiL3BhdGggvdG8vdGhlL3Jvb3QiCiMKI
yBXaGVyZToKIyAtIHJ1bnRpbWUtaGFuZGxlcjogbmFtZSB1c2VkIHRvIGlkZW50aWZ5IHRoZSBydW50aW1
lCiMgLSBydW50aW1lX3BhdGggKG9wdGlvbmFsLCBzdHJpbmcpOiBhYnNvbHV0ZSBwYXRoIHRvIHRoZSByd
W50aW1lIGV4ZWN1dGFibGUgaW4KIyAgIHRoZSBob3N0IGZpbGVzeXN0ZW0uIElmIG9taXR0ZWQsIHRoZSB
ydW50aW1lLWhhbmRsZXIgaWRlbnRpZmllciBzaG91bGQgbWF0Y2gKIyAgIHRoZSBydW50aW1lIGV4ZWN1d
GFibGUgbmFtZSwgYW5kIHRoZSBydW50aW1lIGV4ZWN1dGFibGUgc2hvdWxkIGJlIHBsYWNlZAojICAgaW4
gJFBBVEguCiMgLSBydW50aW1lX3R5cGUgKG9wdGlvbmFsLCBzdHJpbmcpOiB0eXBlIG9mIHJ1bnRpbWUsI
G9uZSBvZjogIm9jaSIsICJ2bSIuIElmCiMgICBvbWl0dGVkLCBhbiAib2NpIiBydW50aW1lIGlzIGFzc3V
tZWQuCiMgLSBydW50aW1lX3Jvb3QgKG9wdGlvbmFsLCBzdHJpbmcpOiByb290IGRpcmVjdG9yeSBmb3Igc
3RvcmFnZSBvZiBjb250YWluZXJzCiMgICBzdGF0ZS4KW2NyaW8ucnVudGltZS5ydW50aW1lcy5ydW5jXQp
ydW50aW1lX3BhdGggPSAiIgpydW50aW1lX3R5cGUgPSAib2NpIgpydW50aW1lX3Jvb3QgPSAiL3J1bi9yd
W5jIgoKIyBDUkktTyByZWFkcyBpdHMgY29uZmlndXJlZCByZWdpc3RyaWVzIGRlZmF1bHRzIGZyb20gdGh
lIHN5c3RlbSB3aWRlCiMgY29udGFpbmVycy1yZWdpc3RyaWVzLmNvbmYoNSkgbG9jYXRlZCBpbiAvZXRjL
2NvbnRhaW5lcnMvcmVnaXN0cmllcy5jb25mLiBJZgojIHlvdSB3YW50IHRvIG1vZGlmeSBqdXN0IENSSS1
PLCB5b3UgY2FuIGNoYW5nZSB0aGUgcmVnaXN0cmllcyBjb25maWd1cmF0aW9uIGluCiMgdGhpcyBmaWxlL
iBPdGhlcndpc2UsIGxlYXZlIGluc2VjdXJlX3JlZ2lzdHJpZXMgYW5kIHJlZ2lzdHJpZXMgY29tbWVudGV
kIG91dCB0bwojIHVzZSB0aGUgc3lzdGVtJ3MgZGVmYXVsdHMgZnJvbSAvZXRjL2NvbnRhaW5lcnMvcmVna
XN0cmllcy5jb25mLgpbY3Jpby5pbWFnZV0KCiMgRGVmYXVsdCB0cmFuc3BvcnQgZm9yIHB1bGxpbmcgaW1
hZ2VzIGZyb20gYSByZW1vdGUgY29udGFpbmVyIHN0b3JhZ2UuCmRlZmF1bHRfdHJhbnNwb3J0ID0gImRvY
2tlcjovLyIKCiMgVGhlIHBhdGggdG8gYSBmaWxlIGNvbnRhaW5pbmcgY3JlZGVudGlhbHMgbmVjZXNzYXJ
5IGZvciBwdWxsaW5nIGltYWdlcyBmcm9tCiMgc2VjdXJlIHJlZ2lzdHJpZXMuIFRoZSBmaWxlIGlzIHNpb
WlsYXIgdG8gdGhhdCBvZiAvdmFyL2xpYi9rdWJlbGV0L2NvbmZpZy5qc29uCmdsb2JhbF9hdXRoX2ZpbGU
gPSAiL3Zhci9saWIva3ViZWxldC9jb25maWcuanNvbiIKCiMgVGhlIGltYWdlIHVzZWQgdG8gaW5zdGFud
GlhdGUgaW5mcmEgY29udGFpbmVycy4KIyBUaGlzIG9wdGlvbiBzdXBwb3J0cyBsaXZlIGNvbmZpZ3VyYXR
pb24gcmVsb2FkLgpwYXVzZV9pbWFnZSA9ICJxdWF5LmlvL29wZW5zaGlmdC1yZWxlYXNlLWRldi9vY3Atd
jQuMC1hcnQtZGV2QHNoYTI1NjpjMmU5YmRhNmY0YmRmODI4Y2I4ZmRiNTA3Mjk0ZDJiMDEyMTk1NjhmZGR
jMDg1YTMwODM3ZGM2Nzg2MzUwM2NjIgoKIyBUaGUgcGF0aCB0byBhIGZpbGUgY29udGFpbmluZyBjcmVkZ
W50aWFscyBzcGVjaWZpYyBmb3IgcHVsbGluZyB0aGUgcGF1c2VfaW1hZ2UgZnJvbQojIGFib3ZlLiBUaGU
gZmlsZSBpcyBzaW1pbGFyIHRvIHRoYXQgb2YgL3Zhci9saWIva3ViZWxldC9jb25maWcuanNvbgojIFRoa
XMgb3B0aW9uIHN1cHBvcnRzIGxpdmUgY29uZmlndXJhdGlvbiByZWxvYWQuCnBhdXNlX2ltYWdlX2F1dGh
fZmlsZSA9ICIvdmFyL2xpYi9rdWJlbGV0L2NvbmZpZy5qc29uIgoKIyBUaGUgY29tbWFuZCB0byBydW4gd
G8gaGF2ZSBhIGNvbnRhaW5lciBzdGF5IGluIHRoZSBwYXVzZWQgc3RhdGUuCiMgV2hlbiBleHBsaWNpdGx
5IHNldCB0byAiIiwgaXQgd2lsbCBmYWxsYmFjayB0byB0aGUgZW50cnlwb2ludCBhbmQgY29tbWFuZAojI
HNwZWNpZmllZCBpbiB0aGUgcGF1c2UgaW1hZ2UuIFdoZW4gY29tbWVudGVkIG91dCwgaXQgd2lsbCBmYWx
sYmFjayB0byB0aGUKIyBkZWZhdWx0OiAiL3BhdXNlIi4gVGhpcyBvcHRpb24gc3VwcG9ydHMgbGl2ZSBjb
25maWd1cmF0aW9uIHJlbG9hZC4KcGF1c2VfY29tbWFuZCA9ICIvdXNyL2Jpbi9wb2QiCgojIFBhdGggdG8
gdGhlIGZpbGUgd2hpY2ggZGVjaWRlcyB3aGF0IHNvcnQgb2YgcG9saWN5IHdlIHVzZSB3aGVuIGRlY2lka
W5nCiMgd2hldGhlciBvciBub3QgdG8gdHJ1c3QgYW4gaW1hZ2UgdGhhdCB3ZSd2ZSBwdWxsZWQuIEl0IGl
zIG5vdCByZWNvbW1lbmRlZCB0aGF0CiMgdGhpcyBvcHRpb24gYmUgdXNlZCwgYXMgdGhlIGRlZmF1bHQgY
mVoYXZpb3Igb2YgdXNpbmcgdGhlIHN5c3RlbS13aWRlIGRlZmF1bHQKIyBwb2xpY3kgKGkuZS4sIC9ldGM
vY29udGFpbmVycy9wb2xpY3kuanNvbikgaXMgbW9zdCBvZnRlbiBwcmVmZXJyZWQuIFBsZWFzZQojIHJlZ
mVyIHRvIGNvbnRhaW5lcnMtcG9saWN5Lmpzb24oNSkgZm9yIG1vcmUgZGV0YWlscy4Kc2lnbmF0dXJlX3B
vbGljeSA9ICIiCgojIENvbnRyb2xzIGhvdyBpbWFnZSB2b2x1bWVzIGFyZSBoYW5kbGVkLiBUaGUgdmFsa
WQgdmFsdWVzIGFyZSBta2RpciwgYmluZCBhbmQKIyBpZ25vcmU7IHRoZSBsYXR0ZXIgd2lsbCBpZ25vcmU
gdm9sdW1lcyBlbnRpcmVseS4KaW1hZ2Vfdm9sdW1lcyA9ICJta2RpciIKCiMgVGhlIGNyaW8ubmV0d29ya
yB0YWJsZSBjb250YWluZXJzIHNldHRpbmdzIHBlcnRhaW5pbmcgdG8gdGhlIG1hbmFnZW1lbnQgb2YKIyB
DTkkgcGx1Z2lucy4KW2NyaW8ubmV0d29ya10KIyBQYXRoIHRvIHRoZSBkaXJlY3Rvcnkgd2hlcmUgQ05JI
GNvbmZpZ3VyYXRpb24gZmlsZXMgYXJlIGxvY2F0ZWQuCm5ldHdvcmtfZGlyID0gIi9ldGMva3ViZXJuZXR
lcy9jbmkvbmV0LmQvIgoKIyBQYXRocyB0byBkaXJlY3RvcmllcyB3aGVyZSBDTkkgcGx1Z2luIGJpbmFya
WVzIGFyZSBsb2NhdGVkLgpwbHVnaW5fZGlycyA9IFsKICAgICIvdmFyL2xpYi9jbmkvYmluIiwKXQoKIyB
BIG5lY2Vzc2FyeSBjb25maWd1cmF0aW9uIGZvciBQcm9tZXRoZXVzIGJhc2VkIG1ldHJpY3MgcmV0cmlld
mFsCltjcmlvLm1ldHJpY3NdCgojIEdsb2JhbGx5IGVuYWJsZSBvciBkaXNhYmxlIG1ldHJpY3Mgc3VwcG9

```
ydC4KZW5hYmxlX21ldHJpY3MgPSBOcnVlCgojIFRoZSBwb3J0IG9uIHdoaWNoIHRoZSBtZXRyaWNzIHNlc
nZlciB3aWxsIGxpc3Rlbi4KbWVOcmljc19wb3J0ID0gOTUzNwo=
    systemd:
      units:
        - name: powersmt.service
```


enabled: true

Check the contents of the base64 hash by using the **base64 -d** command to check that the configuration you are applying is the one to which you intended. You can always create a hash with a different configuration when you have a different requirement.

After you created the YAML file, apply the configuration as shown in Example A-8.

Example A-8 Applying `99_openshift-machineconfig_99-worker-ibm.yaml` *configuration*

```
[root@client ~]# oc apply -f 99_openshift-machineconfig_99-worker-ibm.yaml
machineconfig.machineconfiguration.openshift.io/99-worker-ibm created
```

Note: By default, the nodes automatically reboot in rolling fashion, one-by-one, after you apply the configuration.

Using CoreOS tuned operator to apply the sysctl parameters

To ensure that certain microservices run correctly, you must verify the following kernel parameters:

- Virtual memory limit (`vm.max_map_count`)
- Message limits (`kernel.msgmax, kernel.msgmnb, and kernel.msgmni`)
- Shared memory limits (`kernel.shmmax, kernel.shmall, and kernel.shmmni`)
- Semaphore limits (`kernel.sem`)

These settings are required for all deployments. Example A-9 assumes that you have worker nodes with 64 GB of RAM. If the worker nodes have 128 GB of RAM each, double the values for the `kernel.shm*` values.

Example A-9 64 GB worker node of a tuned.yaml file

```
apiVersion: tuned.openshift.io/v1
kind: Tuned
metadata:
  name: cp4d-wkc-ipc
  namespace: openshift-cluster-node-tuning-operator
spec:
  profile:
  - name: cp4d-wkc-ipc
    data: |
      [main]
      summary=Tune IPC Kernel parameters on OpenShift Worker Nodes running WKC Pod
s
      [sysctl]
      kernel.shmall = 33554432
      kernel.shmmax = 68719476736
      kernel.shmmni = 16384
      kernel.sem = 250 1024000 100 16384
```

```
      kernel.msgmax = 65536
      kernel.msgmnb = 65536
      kernel.msgmni = 32768
      vm.max_map_count = 262144
  recommend:
  - match:
    - label: node-role.kubernetes.io/worker
    priority: 10
    profile: cp4d-wkc-ipc
```

To apply the parameters, follow Example A-10.

Example A-10 Applying tuned configuration

```
[root@client ~]# oc apply -f tuned.yaml
tuned.tuned.openshift.io/cp4d-wkc-ipc created
```

B

Booting from System Management Services

This appendix describes how to boot a partition from the System Management Services (SMS) menu by selecting the boot device for booting from network, and then defining the boot device order to boot the newly installed operating system.

The SMS services helps view information about your system or partition, and perform tasks, such as changing the boot list and setting the network parameters. These SMS menus can be used for AIX or Linux logical partitions.

This appendix includes the following topics:

Entering SMS mode

When booting an LPAR, it is possible to enter the SMS menu from the startup console (see Figure B-1), or directly when booting from the Hardware Management Console (HMC).

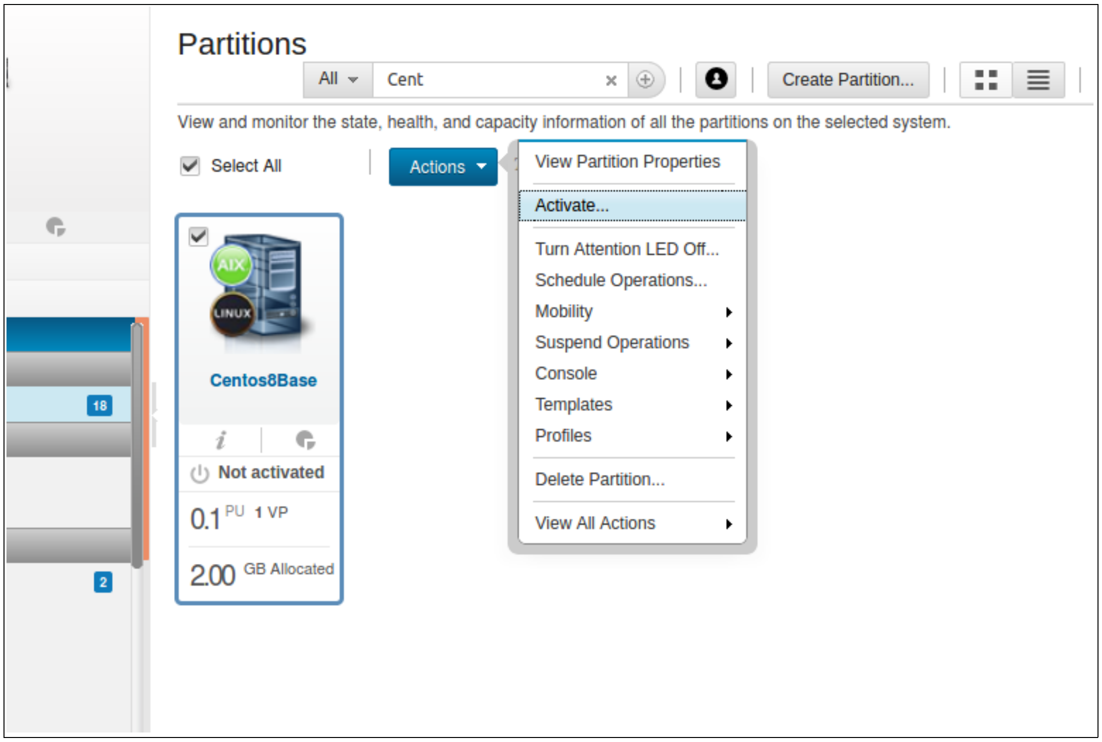

Figure B-1 HMC GUI partitions view

Option 1: Boot directly to SMS from the HMC

To enter SMS from the HMC web GUI, choose the SMS option in the Advanced Settings Boot Mode window, as shown in Figure B-2.

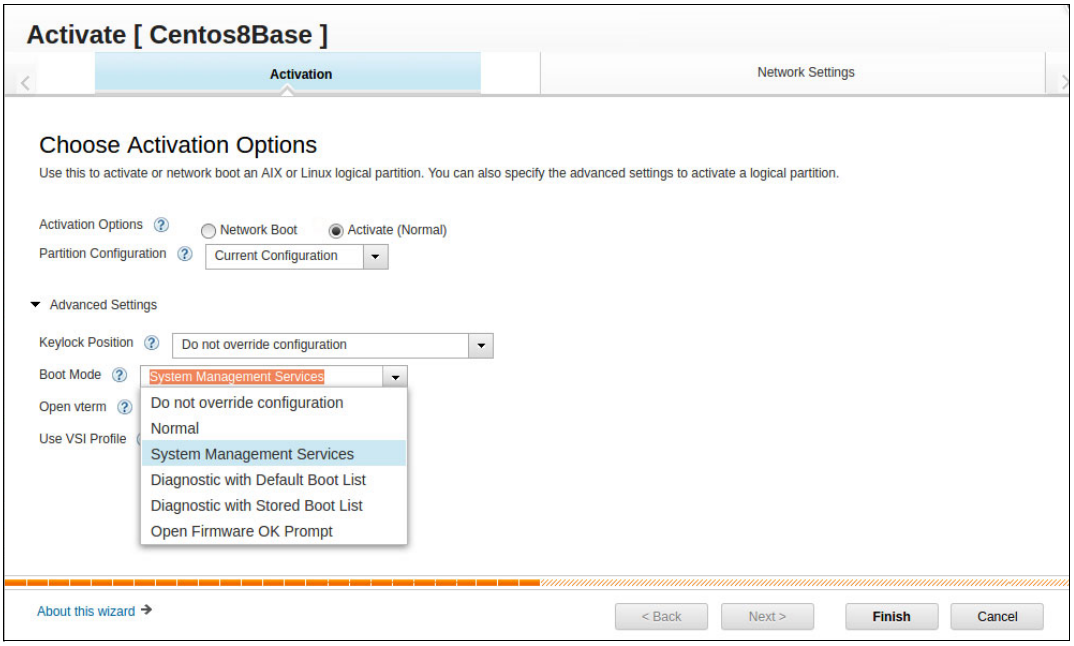

Figure B-2 Choose Activation Options: Advanced Settings window

Alternatively, you can run the following command from the HMC CLI:

```
user@hmc:~> chsysstate -r lpar -m <managed-system> -o on -f <profile> -b sms -n
<lpar_name>
```

Option 2: Enter SMS from the startup console

From the startup console, enter SMS by pressing **1** the first time this window appears, as shown in Figure B-3.

```
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM

          1 = SMS Menu                          5 = Default Boot List
          8 = Open Firmware Prompt              6 = Stored Boot List


     Memory      Keyboard     Network     Speaker
```

Figure B-3 Startup console

Configuring booting from network

After you are in the SMS menu, choose **option number 2: Setup Remote IPL (Initial Program Load)**, as shown in Figure B-4.

```
PowerPC Firmware
Version FW840.40 (SV840_147)
SMS (c) Copyright IBM Corp. 2000,2016 All rights reserved.
-------------------------------------------------------------------------------
Main Menu
1.   Select Language
2.   Setup Remote IPL (Initial Program Load)
3.   I/O Device Information
4.   Select Console
5.   Select Boot Options









-------------------------------------------------------------------------------
Navigation Keys:

                                           X = eXit System Management Services
-------------------------------------------------------------------------------
Type menu item number and press Enter or select Navigation key:
```

Figure B-4 SMS: Main Menu selection pane

Figure B-5 shows the LPAR attached network cards. Select the network card to be used for network booting. In this case, only one card is available.

```
PowerPC Firmware
Version FW840.40 (SV840_147)
SMS (c) Copyright IBM Corp. 2000,2016 All rights reserved.
-------------------------------------------------------------------------------
NIC Adapters
     Device                          Location Code                Hardware
                                                                  Address
1.  Interpartition Logical LAN      U8408.E8E.219283W-V5-C5-T1  4aed2bc15b05










 -------------------------------------------------------------------------------
Navigation keys:
M = return to Main Menu
ESC key = return to previous screen         X = eXit System Management Services
 -------------------------------------------------------------------------------
Type menu item number and press Enter or select Navigation key:
```

Figure B-5 SMS: NIC Adapters pane

Select the IP version to use for network booting, as shown in Figure B-6. In our example, we choose **IPv4**.

```
PowerPC Firmware
Version FW840.40 (SV840_147)
SMS (c) Copyright IBM Corp. 2000,2016 All rights reserved.
-------------------------------------------------------------------------------
Select Internet Protocol Version.

1.   IPv4 - Address Format 123.231.111.222
2.   IPv6 - Address Format 1234:5678:90ab:cdef:1234:5678:90ab:cdef










-------------------------------------------------------------------------------
Navigation keys:
M = return to Main Menu
ESC key = return to previous screen         X = eXit System Management Services
-------------------------------------------------------------------------------
Type menu item number and press Enter or select Navigation key:
```

Figure B-6 SMS: Select Internet Protocol Version pane

Select the **BOOTP** option, as shown in Figure B-7.

```
PowerPC Firmware
Version FW840.40 (SV840_147)
SMS (c) Copyright IBM Corp. 2000,2016 All rights reserved.
-------------------------------------------------------------------------------
Select Network Service.
1.   BOOTP
2.   ISCSI









-------------------------------------------------------------------------------
Navigation keys:
M = return to Main Menu
ESC key = return to previous screen         X = eXit System Management Services
-------------------------------------------------------------------------------
Type menu item number and press Enter or select Navigation key:
```

Figure B-7 SMS: Select Network Service selection pane

Select IP Parameters (see Figure B-8), and set the LPAR address, the bootp server address, and the network gateway and mask.

```
 PowerPC Firmware
 Version FW840.40 (SV840_147)
 SMS (c) Copyright IBM Corp. 2000,2016 All rights reserved.
-------------------------------------------------------------------------------
 Network Parameters
Interpartition Logical LAN: U8408.E8E.219283W-V5-C2-T1
 1.   IP Parameters
 2.   Adapter Configuration
 3.   Ping Test
 4.   Advanced Setup: BOOTP









 -------------------------------------------------------------------------------
 Navigation keys:
 M = return to Main Menu
 ESC key = return to previous screen         X = eXit System Management Services
 -------------------------------------------------------------------------------
 Type menu item number and press Enter or select Navigation key:
```

Figure B-8   SMS: Network Parameters pane

After you are done, press **M** to return to the main menu, as shown in Figure B-9.

```
 PowerPC Firmware
 Version FW840.40 (SV840_147)
 SMS (c) Copyright IBM Corp. 2000,2016 All rights reserved.
-------------------------------------------------------------------------------
 IP Parameters
Interpartition Logical LAN: U8408.E8E.219283W-V5-C2-T1
 1.   Client IP Address                    [000.000.000.000]
 2.   Server IP Address                    [000.000.000.000]
 3.   Gateway IP Address                   [000.000.000.000]
 4.   Subnet Mask                          [000.000.000.000]









 -------------------------------------------------------------------------------
 Navigation keys:
 M = return to Main Menu
 ESC key = return to previous screen         X = eXit System Management Services
 -------------------------------------------------------------------------------
 Type menu item number and press Enter or select Navigation key:
```

Figure B-9 SMS: IP Parameters pane

In the main menu, select **option number 5, Select Boot Options**, as shown in Figure B-10.

```
PowerPC Firmware
Version FW840.40 (SV840_147)
SMS (c) Copyright IBM Corp. 2000,2016 All rights reserved.
-------------------------------------------------------------------------------
Main Menu
1.   Select Language
2.   Setup Remote IPL (Initial Program Load)
3.   I/O Device Information
4.   Select Console
5.   Select Boot Options









-------------------------------------------------------------------------------
Navigation Keys:

                                             X = eXit System Management Services
-------------------------------------------------------------------------------
Type menu item number and press Enter or select Navigation key:
```

Figure B-10   SMS: Main Menu pane

In the boot options menu (see Figure B-11), select **option number 1, Select Install/Boot Device.**

```
PowerPC Firmware
Version FW840.40 (SV840_147)
SMS (c) Copyright IBM Corp. 2000,2016 All rights reserved.
-------------------------------------------------------------------------------
Multiboot
1.   Select Install/Boot Device
2.   Configure Boot Device Order
3.   Multiboot Startup <OFF>
4.   SAN Zoning Support









-------------------------------------------------------------------------------
Navigation keys:
M = return to Main Menu
ESC key = return to previous screen         X = eXit System Management Services
-------------------------------------------------------------------------------
Type menu item number and press Enter or select Navigation key:
```

Figure B-11 SMS: Multiboot pane

Then, select **option number 4, Network**, as shown in Figure B-12.

```
PowerPC Firmware
Version FW840.40 (SV840_147)
SMS (c) Copyright IBM Corp. 2000,2016 All rights reserved.
-------------------------------------------------------------------------------
Select Device Type
1.   Tape
2.   CD/DVD
3.   Hard Drive
4.   Network
5.   List all Devices









-------------------------------------------------------------------------------
Navigation keys:
M = return to Main Menu
ESC key = return to previous screen         X = eXit System Management Services
-------------------------------------------------------------------------------
Type menu item number and press Enter or select Navigation key:
```

Figure B-12 SMS: Select Device Type pane

Select **option number 1, BOOTP**, as shown in Figure B-13.

```
PowerPC Firmware
Version FW840.40 (SV840_147)
SMS (c) Copyright IBM Corp. 2000,2016 All rights reserved.
-------------------------------------------------------------------------------
Select Network Service.
1.   BOOTP
2.   ISCSI











-------------------------------------------------------------------------------
Navigation keys:
M = return to Main Menu
ESC key = return to previous screen         X = eXit System Management Services
-------------------------------------------------------------------------------
Type menu item number and press Enter or select Navigation key:
```

Figure B-13 SMS: Select Network Service pane

Next, select the network card you that configured for network booting, as shown in Figure B-14.

```
PowerPC Firmware
Version FW840.40 (SV840_147)
SMS (c) Copyright IBM Corp. 2000,2016 All rights reserved.
-------------------------------------------------------------------------------
Select Device
Device  Current  Device
Number  Position  Name
1.        -      Interpartition Logical LAN
       ( loc=U8408.E8E.219283W-V5-C5-T1 )










-------------------------------------------------------------------------------
Navigation keys:
M = return to Main Menu
ESC key = return to previous screen         X = eXit System Management Services
-------------------------------------------------------------------------------
Type menu item number and press Enter or select Navigation key:
```

Figure B-14 SMS: Select Device pane

Select option **number 2, Normal Mode Boot**, as shown in Figure B-15.

```
 PowerPC Firmware
 Version FW840.40 (SV840_147)
 SMS (c) Copyright IBM Corp. 2000,2016 All rights reserved.
-------------------------------------------------------------------------------
 Select Task

Interpartition Logical LAN
    ( loc=U8408.E8E.219283W-V5-C5-T1 )

 1.   Information
 2.   Normal Mode Boot
 3.   Service Mode Boot








 -------------------------------------------------------------------------------
 Navigation keys:
 M = return to Main Menu
 ESC key = return to previous screen         X = eXit System Management Services
 -------------------------------------------------------------------------------
 Type menu item number and press Enter or select Navigation key:
```

Figure B-15 SMS: Select Task pane

Exit SMS. Confirm exiting SMS, as shown in Figure B-16. This action boots the operating system from the network, installs the operating system, and reboots.

```
PowerPC Firmware
Version FW840.40 (SV840_147)
SMS (c) Copyright IBM Corp. 2000,2016 All rights reserved.
-------------------------------------------------------------------------------
Are you sure you want to exit System Management Services?
1.   Yes
2.   No















-------------------------------------------------------------------------------
Navigation Keys:

                                             X = eXit System Management Services
-------------------------------------------------------------------------------
Type menu item number and press Enter or select Navigation key:
```

Figure B-16   SMS: Exit System Management Services pane

Configuring boot device order

After the reboot starts, you must enter again the SMS menu to change the primary boot device, or the LPAR boots from the network again (see Figure B-17).

```
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM

          1 = SMS Menu                          5 = Default Boot List
          8 = Open Firmware Prompt              6 = Stored Boot List


     Memory      Keyboard     Network     Speaker
```

Figure B-17 SMS: Change the default boot option

In SMS mode again, select **option number 5, Select Boot Options**, as shown in Figure B-18.

```
PowerPC Firmware
Version FW840.40 (SV840_147)
SMS (c) Copyright IBM Corp. 2000,2016 All rights reserved.
-------------------------------------------------------------------------------
Main Menu
1.   Select Language
2.   Setup Remote IPL (Initial Program Load)
3.   I/O Device Information
4.   Select Console
5.   Select Boot Options









-------------------------------------------------------------------------------
Navigation Keys:

                                      X = eXit System Management Services
-------------------------------------------------------------------------------
Type menu item number and press Enter or select Navigation key:
```

Figure B-18 SMS: Main Menu pane

In the boot options menu, select **option number 2, Configure Boot Device Order**, as shown in Figure B-19.

```
PowerPC Firmware
Version FW840.40 (SV840_147)
SMS (c) Copyright IBM Corp. 2000,2016 All rights reserved.
-------------------------------------------------------------------------------
Multiboot
1.   Select Install/Boot Device
2.   Configure Boot Device Order
3.   Multiboot Startup <OFF>
4.   SAN Zoning Support










-------------------------------------------------------------------------------
Navigation keys:
M = return to Main Menu
ESC key = return to previous screen         X = eXit System Management Services
-------------------------------------------------------------------------------
Type menu item number and press Enter or select Navigation key:
```

Figure B-19 SMS: Multiboot pane

Select **option number 1, Select 1st Boot Device**, to set the first boot device, as shown in Figure B-20.

```
PowerPC Firmware
Version FW840.40 (SV840_147)
SMS (c) Copyright IBM Corp. 2000,2016 All rights reserved.
-------------------------------------------------------------------------------
Configure Boot Device Order
1.   Select 1st Boot Device
2.   Select 2nd Boot Device
3.   Select 3rd Boot Device
4.   Select 4th Boot Device
5.   Select 5th Boot Device
6.   Display Current Setting
7.   Restore Default Setting







-------------------------------------------------------------------------------
Navigation keys:
M = return to Main Menu
ESC key = return to previous screen          X = eXit System Management Services
-------------------------------------------------------------------------------
Type menu item number and press Enter or select Navigation key:
```

Figure B-20   SMS: Configure Boot Device Order pane

Select **option number 6, List All Devices**, as shown in Figure B-21.

```
PowerPC Firmware
Version FW840.40 (SV840_147)
SMS (c) Copyright IBM Corp. 2000,2016 All rights reserved.
-------------------------------------------------------------------------------
Select Device Type
1.   Tape
2.   CD/DVD
3.   Hard Drive
4.   Network
5.   None
6.   List All Devices









 -------------------------------------------------------------------------------
Navigation keys:
M = return to Main Menu
ESC key = return to previous screen         X = eXit System Management Services
 -------------------------------------------------------------------------------
Type menu item number and press Enter or select Navigation key:
```

Figure B-21 SMS: Select Device Type pane

Select the device where you installed the operating system. Figure B-22 shows four hard disk drives that are the same because of the multipath configuration. Select one of the hard drives.

```
PowerPC Firmware
Version FW840.40 (SV840_147)
SMS (c) Copyright IBM Corp. 2000,2016 All rights reserved.
-------------------------------------------------------------------------------
Select Device
Device  Current  Device
Number  Position  Name
1.         -       Interpartition Logical LAN
       ( loc=U8408.E8E.219283W-V5-C5-T1 )
2.         1       103 GB   FC Harddisk
       ( loc=U8408.E8E.219283W-V5-C3-T1-W5001738033270140-L0001000000000000 )
3.         -       103 GB   FC Harddisk
       ( loc=U8408.E8E.219283W-V5-C3-T1-W5001738033270141-L0001000000000000 )
4.         -       103 GB   FC Harddisk
       ( loc=U8408.E8E.219283W-V5-C3-T1-W5001738033270150-L0001000000000000 )
5.         -       103 GB   FC Harddisk
       ( loc=U8408.E8E.219283W-V5-C3-T1-W5001738033270161-L0001000000000000 )


 -------------------------------------------------------------------------------
Navigation keys:
M = return to Main Menu      N = Next page of list
ESC key = return to previous screen         X = eXit System Management Services
-------------------------------------------------------------------------------
Type menu item number and press Enter or select Navigation key:
```

Figure B-22 SMS: Select Device pane

Select **option number 2, Set Boot Sequence: Configure as 1st Boot Device**, as shown in Figure B-23. This action sets the device that is selected as the first boot device option.

```
 PowerPC Firmware
 Version FW840.40 (SV840_147)
 SMS (c) Copyright IBM Corp. 2000,2016 All rights reserved.
-------------------------------------------------------------------------------
 Select Task

103 GB   FC Harddisk
    ( loc=U8408.E8E.219283W-V5-C3-T1-W5001738033270140-L0001000000000000 )

 1.   Information
 2.   Set Boot Sequence: Configure as 1st Boot Device









 -------------------------------------------------------------------------------
 Navigation keys:
 M = return to Main Menu
 ESC key = return to previous screen         X = eXit System Management Services
 -------------------------------------------------------------------------------
 Type menu item number and press Enter or select Navigation key:
```

Figure B-23 SMS: Select Task pane

Figure B-24 shows the current boot sequence. Enter x to exit SMS.

```
PowerPC Firmware
Version FW840.40 (SV840_147)
SMS (c) Copyright IBM Corp. 2000,2016 All rights reserved.
-------------------------------------------------------------------------------
Current Boot Sequence
1.  103 GB   FC Harddisk
   ( loc=U8408.E8E.219283W-V5-C3-T1-W5001738033270140-L0001000000000000 )
2.     None
3.     None
4.     None









-------------------------------------------------------------------------------
Navigation keys:
M = return to Main Menu
ESC key = return to previous screen         X = eXit System Management Services
-------------------------------------------------------------------------------
Type menu item number and press Enter or select Navigation key:
```

Figure B-24 SMS: Current Boot Sequence pane

Confirm that you want to exit SMS and boot normally.

C

Additional material

This paper refers to additional material that can be downloaded from the internet as described in the following sections.

Locating the web material

The web material that us associated with this paper is available in softcopy on the internet from the IBM Redbooks web server:

`ftp://www.redbooks.ibm.com/redbooks/REDP5599`

Alternatively, you can go to the IBM Redbooks website:

ibm.com/redbooks

Search for REDP5599, select the title, and then, click **Additional materials** to open the directory that corresponds with the IBM Redpaper form number, REDP5599.

Using the web material

The additional web material that accompanies this paper includes the following files:

File name	*Description*
REDP5599AdditionalMaterials.zip	Zipped YAML Files

System requirements for downloading the web material

The web material requires the following system configuration:

Hard disk space:	100 MB minimum
Operating System:	Windows, Linux or macOS
Processor:	i3 or higher
Memory:	1024 MB or higher

Downloading and extracting the web material

Create a subdirectory (folder) on your workstation, and extract the contents of the web material `.zip` file into this folder.

Related publications

The publications that are listed in this section are considered particularly suitable for a more detailed discussion of the topics that are covered in this paper.

IBM Redbooks

The following IBM Redbooks 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:

- *Red Hat OpenShift and IBM Cloud Paks on IBM Power Systems: Volume 1*, SG24-8459
- *NIM from A to Z in AIX 5L*, SG24-7296
- *IBM Power System E950: Technical Overview and Introduction*, REDP-5509
- *IBM Power Systems LC921 and LC922: Technical Overview and Introduction*, REDP-5495
- *IBM Power System IC922: Technical Overview and Introduction*, REDP-5584

You can search for, view, download, or order these documents and other Redbooks, Redpapers, Web Docs, draft, and additional materials, at the following website:

ibm.com/redbooks

Online resources

The following websites are also relevant as further information sources:

- IBM PowerVC:

  https://www.ibm.com/us-en/marketplace/powervc
- IBM PowerVC Standard Edition V1.4.4:

  https://www.ibm.com/support/knowledgecenter/SSXK2N_1.4.4/com.ibm.powervc.standard.help.doc/powervc_whats_new_hmc.html
- Network Installation Management:

  https://www.ibm.com/support/knowledgecenter/ssw_aix_72/install/nim_intro.html
- IBM Redbooks highlighting POWER9 processor-based technology:

  http://www.redbooks.ibm.com/Redbooks.nsf/pages/power9?Open

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

REDP-5599-00

ISBN 0738459070

Printed in U.S.A.

Get connected

Redbooks®

ibm.com/redbooks