

IBM PowerVC Version 2.0 Introduction and Configuration

Sachin P. Deshmukh

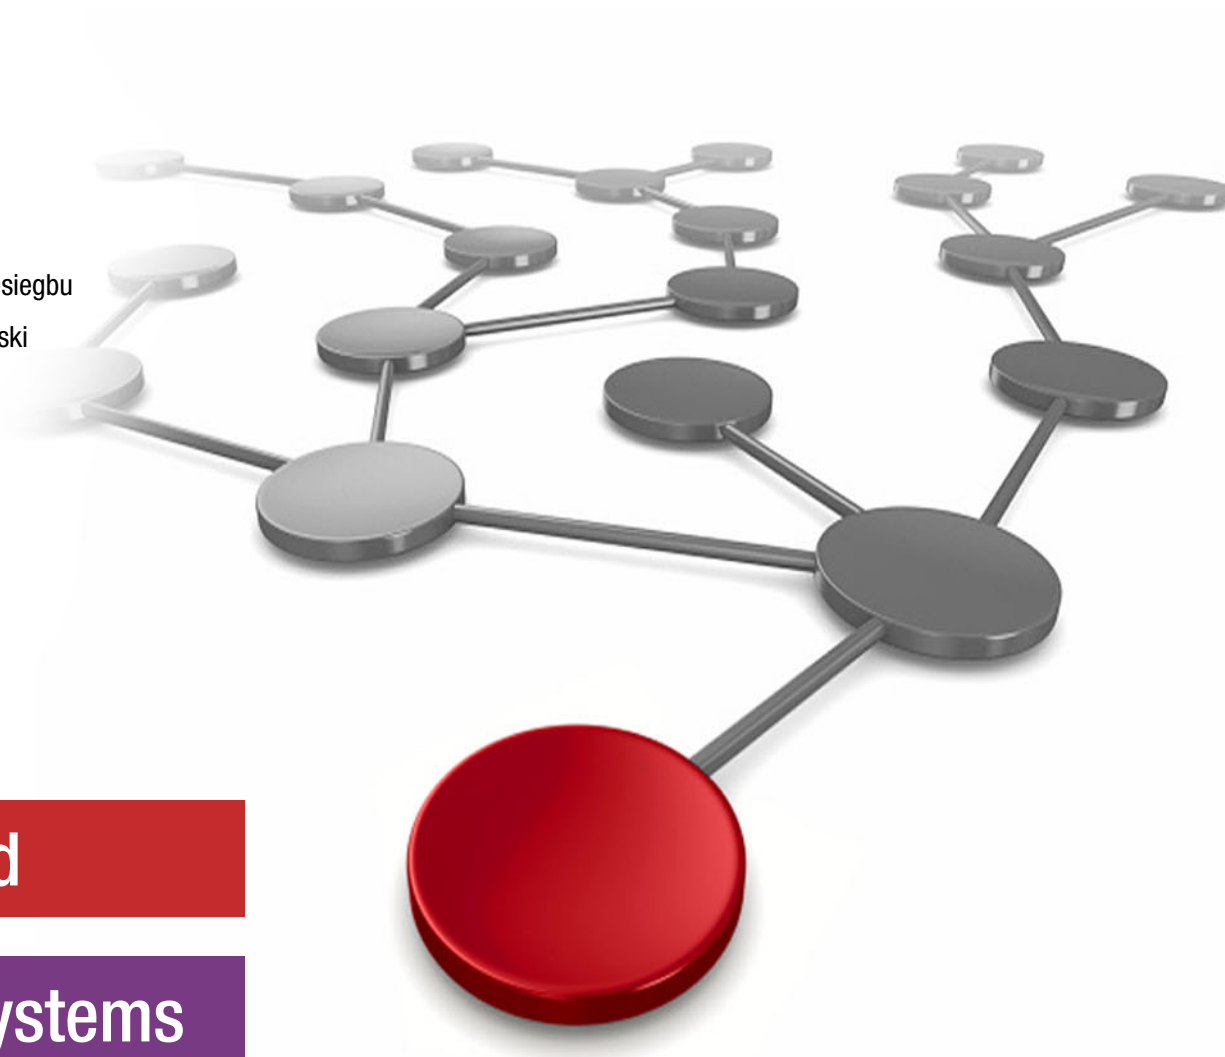
Thierry Huché

Stephen Lutz

Ahmed Mashhour

Christopher Emefiene Osiegbu

Borislav Ivanov Stoymirski



 **Cloud**

Power Systems



IBM Redbooks

**IBM PowerVC Version 2.0 Introduction and
Configuration**

May 2021

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

First Edition (May 2021)

This edition applies to Version 2, Release 0, Modification 0 of IBM PowerVC (product number 5765-VS2) and IBM PowerVC for Private Cloud (product number 5765-VC2).

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

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

Contents

Figures	vii
Tables	xiii
Examples	xv
Notices	xix
Trademarks	xx
Preface	xxi
Authors	xxii
Now you can become a published author, too!	xxiii
Comments welcome	xxiv
Stay connected to IBM Redbooks	xxiv
Chapter 1. IBM PowerVC introduction	1
1.1 IBM PowerVC overview	2
1.1.1 IBM PowerVC functions and advantages	2
1.1.2 IBM PowerVM NovaLink overview	3
1.2 IBM PowerVC editions	7
1.2.1 IBM PowerVC	7
1.2.2 IBM PowerVC for Private Cloud	7
1.3 PowerVC versions and releases	7
1.3.1 IBM PowerVC release to OpenStack edition cross-reference	7
1.3.2 IBM PowerVC 1.4 versions	8
1.3.3 IBM PowerVC Version 2.0.0	9
1.4 IBM PowerVC adoption	11
Chapter 2. IBM PowerVC planning	13
2.1 IBM PowerVC requirements	14
2.1.1 Hardware and software information	14
2.1.2 Hardware and software requirements for IBM PowerVC	14
2.1.3 Other hardware compatibility	19
2.1.4 Web browser supported	19
2.2 IBM PowerVM NovaLink requirements	19
2.2.1 PowerVM NovaLink system requirements	19
2.3 Host and partition management planning	22
2.3.1 Physical server configuration	23
2.3.2 HMC or PowerVM NovaLink planning	23
2.3.3 Virtual I/O Server planning	27
2.4 Placement policies and templates	30
2.4.1 Host groups	30
2.4.2 Placement policies	31
2.4.3 Template types	33
2.4.4 Information that is required for compute template planning	34
2.4.5 Tips for deploy template planning	40
2.5 IBM PowerVC storage access SAN planning	41
2.5.1 vSCSI storage access	41
2.5.2 NPIV storage access	43

2.5.3	Shared storage pools	44
2.6	Storage management planning	45
2.6.1	IBM PowerVC storage terminology	45
2.6.2	Storage templates	49
2.6.3	Storage connectivity group and tags	53
2.6.4	Combining storage connectivity groups, tags and storage port groups	55
2.7	Network management planning	56
2.7.1	Infoblox support	56
2.7.2	Multiple network planning	58
2.7.3	Shared Ethernet Adapter planning	58
2.7.4	Planning Single Root I/O Virtualization networks	60
2.8	Planning users and groups	61
2.8.1	User management planning	61
2.8.2	Projects and role management planning	62
2.9	Security management planning	65
2.9.1	Ports that are used by IBM PowerVC	65
2.9.2	Providing a certificate	66
2.10	Product information	66
Chapter 3. IBM PowerVC installation		69
3.1	Setting up the IBM PowerVC environment	70
3.1.1	Creating the virtual machine to host IBM PowerVC	70
3.1.2	Downloading and installing the Linux operating system	77
3.1.3	Customizing the Linux operating system	78
3.2	Installing IBM PowerVC	81
3.3	Uninstalling IBM PowerVC	86
3.4	Upgrading IBM PowerVC	87
3.4.1	Before you begin	88
3.4.2	Performing an in-place upgrade	89
3.4.3	Upgrading or migrating from earlier versions	95
3.4.4	Migrating PowerVC from ppc64 to ppc64le or x86_64	96
3.4.5	Upgrading to PowerVC for Private Cloud from PowerVC Standard Edition	96
3.5	Updating IBM PowerVC	97
3.5.1	IBM Notifications	98
3.6	IBM PowerVC backup and recovery	98
3.6.1	Backing up IBM PowerVC	98
3.6.2	Recovering IBM PowerVC data	100
3.6.3	Status messages during backup and recovery	101
3.6.4	Considerations about backup and recovery	101
3.7	IBM PowerVC command-line interface	102
3.7.1	Exporting audit data	103
3.8	Virtual machines that are managed by IBM PowerVC	104
3.8.1	IBM AIX virtual machines	104
3.8.2	IBM i virtual machines	104
3.8.3	Linux on Power virtual machines	105
Chapter 4. IBM PowerVC for managing IBM PowerVM		107
4.1	IBM PowerVC graphical user interface	108
4.2	Introduction to the IBM PowerVC setup	118
4.3	Managing resources outside of IBM PowerVC	118
4.3.1	Changes of the LPAR in the HMC	119
4.3.2	Changes in the storage or zoning	119
4.4	Connecting to IBM PowerVC	120

4.5	Host setup	121
4.5.1	Planning for host management - NovaLink	121
4.5.2	Planning for host management - HMC	121
4.5.3	Adding a host	121
4.5.4	Host maintenance mode	124
4.5.5	Host groups	126
4.6	Hardware Management Console management	131
4.6.1	Adding an HMC	131
4.6.2	Changing the HMC credentials	131
4.6.3	Changing the HMC	132
4.6.4	Change management connection, HMC, and PowerVM NovaLink	132
4.7	Storage and SAN fabric setup	133
4.7.1	Adding a storage controller to IBM PowerVC	133
4.7.2	Adding a SAN fabric to IBM PowerVC	134
4.8	VIOS port tags setup	136
4.9	Storage connectivity group setup	138
4.10	Storage template setup	140
4.11	Storage volume setup	143
4.12	Network setup	144
4.12.1	Basic network setup	144
4.12.2	SR-IOV network setup	149
4.13	Compute template setup	154
4.14	Environment verification	156
4.15	Management of virtual machines	157
4.15.1	Virtual machine onboarding	158
4.15.2	Start, stop, and restart of virtual machines	167
4.15.3	Opening a console window	169
4.15.4	Resizing the virtual machine	171
4.15.5	Attaching and detaching of network adapters	174
4.15.6	Attaching, detaching, and cloning of storage volumes	176
4.15.7	Migration of virtual machines	181
4.15.8	VM snapshots	186
4.15.9	Cloning a virtual machine	190
4.15.10	Host maintenance mode	192
4.15.11	Restarting virtual machines remotely from a failed host	196
4.15.12	Resetting the state of a virtual machine	205
4.15.13	Unmanaging a virtual machine	207
4.15.14	Deleting a virtual machine	208
4.16	Images	210
4.16.1	Preparation of a new image	210
4.16.2	Capturing a virtual machine image	224
4.16.3	Import an image from OVA or raw image	226
4.16.4	Define an image with existing volumes	231
4.16.5	Export images into OVA	232
4.16.6	Import and export with backup service	235
4.16.7	Deploying a new virtual machine	239
4.16.8	Deleting images	246
4.17	Users, groups, and roles setup	247
4.17.1	Adding user accounts	247
4.17.2	Updating the user accounts	249
4.17.3	Disabling the root user account from IBM PowerVC	250
4.17.4	Working with roles	250
4.17.5	Setting project policies	251

4.17.6 Lightweight Directory Access Protocol	252
4.17.7 User and Group filtering	253
Chapter 5. IBM PowerVC for Private Cloud	255
5.1 Creating a self-service user	256
5.2 Creating a deploy template	256
5.3 Setting cloud policies	258
5.4 Logging in as a self-service user	260
5.5 Self-service portal graphical user interface	260
5.6 Deploying virtual machines from a template	261
5.6.1 Ownership and expiration dates of VMs	264
Chapter 6. IBM PowerVC integration with Terraform and Ansible	267
6.1 Terraform	268
6.1.1 Install and configure Terraform	268
6.1.2 Create a template for an AIX deployment	270
6.1.3 Specify the Storage Connectivity Group	277
6.1.4 Specify the Storage Template for the boot volume	280
6.2 Ansible automation	283
6.2.1 Install and configure Ansible	283
6.2.2 Create a playbook to retrieve PowerVC informations	284
6.2.3 Create a playbook for an AIX deployment	287
6.2.4 Create an Ansible ready VM with an extra disk and ssh keys	288
Chapter 7. IBM PowerVC in the Hybrid Cloud	293
7.1 Hybrid cloud overview	294
7.2 Red Hat OpenShift Container Platform	294
7.3 IBM PowerVC and IBM Power Systems Virtual Server	308
7.3.1 Migrating VMs to IBM PowerVS	309
7.4 IBM Cloud Pak for Multicloud Management Version 2.2	312
7.4.1 Configure PowerVC as an Openstack Cloud provider	313
7.4.2 Manage VMs with Infrastructure management	316
Appendix A. Useful commands and scripts	317
Useful OpenStack commands	318
Useful powervc-config commands	320
Useful scripts	324
REST APIs example	325
Related publications	331
IBM Redbooks	331
Online resources	331
Help from IBM	337

Figures

1-1	PowerVM NovaLink partition	4
1-2	IBM PowerVC interface	5
1-3	Host Registration window	5
1-4	PowerVM NovaLink diagrams	6
2-1	Changing the logical memory block size	23
2-2	Add redundant HMC	25
2-3	Add HMC details	25
2-4	Secondary HMC available	26
2-5	Enable Active Memory Expansion (AME)	26
2-6	VIOS settings that must be managed by IBM PowerVC	27
2-7	How to change the Maximum Virtual Adapters	28
2-8	HMC Advanced button in HMC Version V9R1M942	28
2-9	Maximum virtual adapters in a VIOS profile	29
2-10	Save configuration changes to profile disable until next activated	29
2-11	Host group sample	30
2-12	Migration of a partition by using a placement policy	33
2-13	Memory region size view on the HMC	36
2-14	IBM PowerVC storage access by using vSCSI	42
2-15	IBM PowerVC storage access by using NPIV	44
2-16	BM PowerVC storage access by using an SSP	45
2-17	List of managed storages	46
2-18	List of managed fabrics	46
2-19	IBM PowerVC Service node	47
2-20	Storage template of an IBM FlashSystem	48
2-21	Storage connectivity groups created in IBM PowerVC	48
2-22	Storage template definition	50
2-23	Selected storage pool with mirroring enabled	52
2-24	Storage port groups page with blue and green port groups created	52
2-25	Fibre Channel port tags	53
2-26	Utilizing storage connectivity groups and storage port groups example	55
3-1	Create Client Partition via HMC Enhanced GUI	71
3-2	Baseboard management controller (BMC) on POWER8	74
3-3	Connect install media in POWER8 BMC	75
3-4	Connect install media in POWER9 BMC	76
3-5	Serial over LAN console on POWER9 BMC system	76
3-6	Checking the IBM PowerVC version via the graphical user interface	88
3-7	Exiting maintenance mode	94
3-8	NovaLink successfully updated	95
4-1	IBM PowerVC Dashboard Overview page	108
4-2	Additional functionalities with IBM PowerVC for Private Cloud 2.0.0	109
4-3	PowerVC 2.0.0 GUI notifications	109
4-4	VM list displays the running task	110
4-5	2-step add storage dialog	110
4-6	Virtual machines tab at the Host edit page	111
4-7	Left navigation bar - expanded	111
4-8	Left navigation bar - collapsed	112
4-9	Download and export the logged messages	112
4-10	Alerts pane	113

4-11	PowerVC Version 2.0.0.0 Dashboard	113
4-12	Inline help pane and tooltip	114
4-13	Pop-up level help	114
4-14	Changing multiple attributes at the same time	115
4-15	VM deployment using the VM list page	115
4-16	Resources column on the VM list page	116
4-17	Context-sensitive logs in VM edit page	116
4-18	Breadcrumb of the trail	117
4-19	The eye icon	117
4-20	IBM PowerVC Login window	120
4-21	First time login	120
4-22	Adding a host in IBM PowerVC	122
4-23	Adding a HMC managed host through existing HMCs or a new connection	122
4-24	Add HMC dialog	123
4-25	Adding a NovaLink managed host to IBM PowerVC	123
4-26	Host list window on IBM PowerVC	124
4-27	Putting a host in maintenance mode	125
4-28	Selecting the action to carry out when setting a host to maintenance mode	126
4-29	Creating a host group	127
4-30	Details page when creating a host group	129
4-31	Selecting hosts when creating a host group	129
4-32	Configuration of DRO	130
4-33	Adding a storage controller	133
4-34	Adding an IBM Storwize storage controller	134
4-35	List of added storage controllers	134
4-36	Adding a SAN fabric to IBM PowerVC	135
4-37	Adding SAN fabric details	135
4-38	Viewing the added SAN fabric	136
4-39	Fibre Channel ports selection	136
4-40	Enter tag	137
4-41	Set connectivity	137
4-42	Fibre Channel port configuration	138
4-43	Storage connectivity group details	139
4-44	Storage connectivity group VIOS redundancy	139
4-45	Disable a storage connectivity group	140
4-46	Hyperswap Storage template details	141
4-47	Hyperswap Storage template Storage pools	142
4-48	Hyperswap Storage template Port groups	142
4-49	Create Hyperswap Storage volume	143
4-50	List available storage volume	144
4-51	Show Hyperswap volume on SVC	144
4-52	Networks in PowerVC	145
4-53	Definition of a basic network	145
4-54	Networks in PowerVC	147
4-55	Used IP addresses of a network	147
4-56	Lock IP addresses	148
4-57	Locked IP addresses in a network	148
4-58	SR-IOV connectivity options	150
4-59	SR-IOV ports of a host	151
4-60	Available SR-IOV ports	152
4-61	Set the name of a physical network	152
4-62	Physical SR-IOV network in PowerVC	153
4-63	SR-IOV ports in the HMC	153

4-64	Physical SR-IOV networks available for the PowerVC network	154
4-65	Create Compute template	155
4-66	Compute template list	155
4-67	Check environment	156
4-68	Check environment result	156
4-69	VM management operations in the VM list	157
4-70	VM management operations in the host list	157
4-71	Onboard a VM	158
4-72	Onboard message	159
4-73	Onboarded VMs	159
4-74	VM details	160
4-75	Enter IP addresses for onboarded VMs	161
4-76	SR-IOV connected network adapter	162
4-77	VM details for an imported VM	162
4-78	Volumes of a virtual machine	163
4-79	Boot volumes	163
4-80	Set storage template	164
4-81	Start a virtual machine	167
4-82	Stop a VM	168
4-83	Restart a VM	169
4-84	NovaLink console	170
4-85	Resizing a VM	171
4-86	Values for VM resizing	172
4-87	Resizing a stopped VM	172
4-88	Resize the memory boundaries for a VM in the HMC	173
4-89	Resize the processor boundaries for a VM in the HMC	173
4-90	Attach network adapter	174
4-91	Select network to be attached	175
4-92	Network successfully attached	175
4-93	Detach a network	176
4-94	Attach a volume to a VM	176
4-95	Attach an existing volume	177
4-96	Attaching new volumes to a VM	177
4-97	Attached volumes in a VM	178
4-98	Detach a volume from a VM	179
4-99	Detach a volume from a VM - confirmation	179
4-100	Clone volumes in the VM volumes tab	180
4-101	Clone volume suffix	180
4-102	Cloned volumes	181
4-103	Migrate selected virtual machines	183
4-104	Migration destination	184
4-105	Virtual machine migration in progress	184
4-106	Virtual machine migration finished	185
4-107	Hard pinned VM	185
4-108	Error when migrating a hard pinned VM	186
4-109	Create snapshot from a VM	187
4-110	Details of the snapshot	187
4-111	Snapshots of a VM	188
4-112	Restore a snapshot	188
4-113	Confirmation to restore the snapshot	189
4-114	Delete a snapshot	189
4-115	Confirmation to delete a snapshot	190
4-116	Clone a VM	190

4-117	VM details for the VM clone	191
4-118	Network details for the clone	191
4-119	Volumes for the clone	192
4-120	VM Clone available	192
4-121	Put host in maintenance mode	193
4-122	VMs to be migrated in maintenance mode	194
4-123	Host in maintenance mode	194
4-124	Exit maintenance mode	195
4-125	Exit maintenance mode confirmation	195
4-126	Recall option	196
4-127	Remote Restart enabled Compute Template	197
4-128	Remote restart attribute in the VM details	198
4-129	Remote Restart state of a VM	199
4-130	Start remote restart	200
4-131	Select VMs to be remote restarted	200
4-132	Host in Remote Restart Rebuilding	201
4-133	VMs during Remote Restart	201
4-134	Remote restart related notifications	202
4-135	Recall VMs	203
4-136	Recall VMs - confirmation	203
4-137	Automated remote restart in a host group	204
4-138	Host included in an automated remote restart task	204
4-139	Exclude a VM from automated remote restart	205
4-140	Reset the virtual machines state	206
4-141	Unmanage a VM	207
4-142	Unmanage a VM - confirmation	207
4-143	Delete a VM	208
4-144	Delete VM - confirmation	209
4-145	Volume for an empty image	211
4-146	Define an empty image - details	211
4-147	Define an empty image - volumes	212
4-148	Deployment of the empty image	212
4-149	VM with empty image is ready	213
4-150	Access the IBM i Open Source Package Management	221
4-151	Connection information for the IBM i system	221
4-152	Installing the IBM i open source environment	222
4-153	IBM i packages to install for cloud-init	222
4-154	Confirmation for IBM i package install	223
4-155	IBM i - cloud-init installed	223
4-156	Capture an image	224
4-157	Details of the image to capture	225
4-158	Captured image is available	225
4-159	Image details	226
4-160	Find the ID of a storage template	228
4-161	Imported image in the GUI of PowerVC	230
4-162	Define an image based on existing volumes	231
4-163	Details for the image to create	231
4-164	Volumes for the image to be defined	232
4-165	Created image ready for use	232
4-166	Select image to be exported	236
4-167	Name for image to be exported	236
4-168	Service nodes	237
4-169	Exported images in the service node	237

4-170	Import an image from backup services	238
4-171	Select image to be imported from backup services	238
4-172	Details of the image import from backup services	239
4-173	Restore of image is running	239
4-174	Deploy the selected image	240
4-175	Image deployment details	240
4-176	SSH keys for injection	242
4-177	VM deployment - Networks tab	243
4-178	Volumes in image to be deployed	244
4-179	Existing volumes to be attached during image deployment	244
4-180	New volumes for a new VM to be deployed	245
4-181	Running deployment	245
4-182	Delete an image	246
4-183	Delete an image - confirmation	247
5-1	Creating a deploy template from the Red Hat Enterprise Linux 8.3 image	256
5-2	The details section when creating a deploy template	257
5-3	image volume tab	257
5-4	Deploy template network tab	258
5-5	List of deploy templates	258
5-6	Setting cloud policies	259
5-7	IBM PowerVC for Private Cloud login window	260
5-8	Self-service portal	260
5-9	Viewing deploy templates from the self-service portal	261
5-10	Deployment of a VM by a self-service user	262
5-11	Viewing deployed requests	262
5-12	Administrator approval page	263
5-13	Self-service user VM list page	263
5-14	Self-service user VM list page	264
5-15	Changing the owner of a VM	264
5-16	Edit the expiration date	265
5-17	Selecting a new expiration date	265
5-18	Administrator approval page	265
5-19	VM with a changed expiration date	266
6-1	Storage connectivity group	277
6-2	Ansible and PowerVC	283
7-1	Create and attach an additional empty volume with 120 GB	295
7-2	Volume successfully attached	295
7-3	Detaching a volume from PowerVC 2.0.0 UI	298
7-4	Creating a new Image for CoreOS	299
7-5	CoreOS image name and operating system	299
7-6	Attaching the volume where the CoreOS was defined	300
7-7	Creating a compute template from PowerVC 2.0.0 UI	301
7-8	Configuring memory resources on a compute template	301
7-9	Configured PowerVC CSI driver on Red Hat OpenShift Container Platform	305
7-10	Listing all Persistent Volumes from OCP UI	307
7-11	Listing all Persistent Volume Claims from OCP UI	307
7-12	Listing the newly created Volume on Storage device	308
7-13	Importing an OVA file from COS to IBM PowerVS	312
7-14	List of boot images in IBM PowerVS	312
7-15	Add a New Cloud Provider	313
7-16	PowerVC cloud provider 1/3	314
7-17	PowerVC cloud provider 2/3	315
7-18	PowerVC cloud provider 3/3	315

7-19 Manage VM.....	316
7-20 Stop VM	316
A-1 REST request to get the tenant IDs	328
A-2 API response example	329
A-3 Get volume ID attribute.....	330
A-4 Volume deletion status message	330

Tables

1-1 IBM PowerVC releases cross-referenced to OpenStack versions	7
2-1 Hardware and OS requirements	15
2-2 Resource requirements for 5 hosts, 2 storage providers, and 2 fabrics	15
2-3 Supported operating systems for VMs on the managed hosts	16
2-4 HMC minimum requirements	16
2-5 Supported VIOS platforms	17
2-6 Supported network hardware and software	17
2-7 Supported Storage and Fabric drivers for PowerVM	18
2-8 Supported storage connectivity options	18
2-9 Supported LDAP server versions	19
2-10 Amount of memory that is needed for a single NovaLink Version 2.0.0 partition	20
2-11 Supported virtualization platforms	21
2-12 Processor compatibility modes	38
2-13 User and Group filters	62
2-14 Available quotas	63
2-15 IBM PowerVC Security Roles	64
3-1 Options for the IBM PowerVC install command	83
3-2 IBM PowerVC users	84
3-3 Available options for the powervc-uninstall command	87
3-4 Options for the custom backup script	90
3-5 Options for powervc-backup command	98
3-6 Options for powervc-restore command	100
3-7 IBM PowerVC available commands	102
3-8 Options for the powervc-audit-export command	103
4-1 VM attributes that can be changed	160
4-2 Information about status icons	164
4-3 Information about the processor specs	165
4-4 Information about the memory specs	165
4-5 Information about the memory specs	166
4-6 Modules and descriptions	214
4-7 Description of the Specifications	226
7-1 OpenShift cluster minimum resource requirements	300
A-1 OpenStack API to remove a volume	329

Examples

2-1	Default user and group filter settings	62
3-1	Create the LPAR powervc001 and assign resources	71
3-2	Configure processing units, virtual processors and memory settings	72
3-3	Create and verify the VEA	72
3-4	List and create a logical volume	72
3-5	Map logical volume	72
3-6	List and create the repository	72
3-7	Upload the virtual media disk into the repository	73
3-8	Map the .iso file to the powervc001 LPAR	73
3-9	Power on and boot the LPAR to SMS	73
3-10	yum configuration	79
3-11	Verify yum	79
3-12	List and enable DVD mounted repository	80
3-13	Install prerequisites verification for SUSE Linux Enterprise Server ppc64le	82
3-14	IBM PowerVC 2.0.0 installation	83
3-15	Installation completed	84
3-16	The powervc-services command example	85
3-17	Successful uninstall process	86
3-18	Checking the IBM PowerVC version via the command-line interface	88
3-19	Creating a backup without uninstalling the PowerVC instance	90
3-20	IBM PowerVC upgrade to Version 2.0.0	92
3-21	Upgrading from PowerVC to PowerVC for Private Cloud	97
3-22	IBM PowerVC backup	99
3-23	Active backup of PowerVC without an interruption to services running	100
3-24	IBM PowerVC restore	101
3-25	Maintenance message for logged-in users	101
3-26	Using the powervc-audit command	103
3-27	'IBM Installation Toolkit' installation	105
3-28	Resource Monitoring Control status	105
4-1	Output example of vtmenu	171
4-2	How to change memory and processor resources for a VM in NovaLink	173
4-3	Installation of IBM Service and Productivity tools for Red Hat Enterprise Linux	215
4-4	cloud.cfg additions for Red Hat Enterprise Linux 8.3	216
4-5	Generating a new machine ID in cloud-init for Red Hat Enterprise Linux	217
4-6	Installation of IBM Service and Productivity tools for SUSE Linux Enterprise Server	217
4-7	Generating a new machine ID in cloud-init for SUSE Linux Enterprise Server	218
4-8	Installation of IBM Service and Productivity tools for Ubuntu	218
4-9	Change the PowerVC staging directory for image imports	227
4-10	Find the ID of a storage template	228
4-11	Import an AIX virtual machine image from ESS	228
4-12	Extract the raw image from an OVA	230
4-13	Change the OVA export directory	233
4-14	Find the ID of the image to be exported	233
4-15	Export an image into OVA	234
4-16	Storage path	235
4-17	Add a new group filter	248
4-18	Edit /opt/ibm/powervc/powervcrc environment file	249
4-19	Check the user creation	249

4-20	List the available OpenStack roles	250
4-21	List the OpenStack role arguments.	250
4-22	Check the identity repository filter.	253
4-23	Create a new user group filter.	253
4-24	Change user group filters	254
5-1	Adding a self-service user account called clouduser	256
6-1	Download the Terraform provider	268
6-2	Terraform plugins	268
6-3	Simple Terraform script to test	269
6-4	Terraform init.	269
6-5	Copy powervrc	270
6-6	Retrieve information from PowerVC	271
6-7	Terraform template to create an IBM PowerVC partition	271
6-8	Terraform execution plan	272
6-9	Apply the Terraform template	273
6-10	Inspect current state of a resource	274
6-11	List VMs in PowerVC	275
6-12	Destroy VM with Terraform.	275
6-13	Add Storage Connectivity Group in Flavor	277
6-14	SCG Terraform execution plan	278
6-15	Apply with auto approve option the Terraform plan	279
6-16	Inspect current state of the flavor	280
6-17	Retrieve Storage Template ID	280
6-18	Add Storage Template in Flavor	281
6-19	Storage Template Terraform execution plan	281
6-20	Create a VM using a specific Storage Template.	282
6-21	List resources in terraform state	282
6-22	Installation of Ansible	284
6-23	Configure the OpenStack environment for Ansible.	285
6-24	Playbook to retrieve PowerVC information	285
6-25	Execution output of playbook	286
6-26	VM Deploy playbook	287
6-27	Deploy VM with Ansible playbook.	288
6-28	List VM on PowerVC.	288
6-29	Playbook to create a VM with extra disk and ssh keys	288
6-30	Deploy VM with extra disk and ssh keys.	289
6-31	List VM on PowerVC.	290
6-32	Connect to the deployed VM.	291
7-1	Extract CoreOS image	296
7-2	Convert CoreOS qcow2 to raw image	296
7-3	Rescanning the bus and listing the disk devices	296
7-4	Verifying the raw disk and dd the image.	297
7-5	Rescanning the bus and listing the disk devices	298
7-6	List all nodes.	302
7-7	Creating a new project by using OCP command-line tool.	302
7-8	This command provides access to system service account	302
7-9	This command provides access to system service account default	303
7-10	Downloading and applying scc.yaml file	303
7-11	Downloading, editing and applying ibm-powervc-csi-driver-template.yaml file.	303
7-12	Downloading secret.yaml and replace parameters of secret file.	303
7-13	Applying secret.yaml file	304
7-14	Validating all parameters and default values	304
7-15	Command below will create "a.yaml" that has all the artifacts in it	304

7-16	Installing the OpenShift artifacts	304
7-17	List all CSI nodes (in our case worker-0 and worker-1)	305
7-18	Listing all pod and verifying that there are no errors in the output	305
7-19	Downloading static-pvc.yaml	306
7-20	Applying static-pvc.yaml file	306
7-21	Listing all Persistent Volumes and Persistent Volume Claims	306
7-22	Deleting the Persistent Volume Claim	308
7-23	Exporting the AIX7.2 image to an OVA file	310
7-24	installing and configuring awscli on the IBM PowerVC host server	310
7-25	Accessing the COS and transferring the OVA image to the COS.	311
A-1	Source the OpenStack environment.	318
A-2	Commands to add user to access ibm-default project	318
A-3	Create and list a project	318
A-4	Create users and roles for access to a project.	319
A-5	Change a role.	319
A-6	Set the network domain	320
A-7	List an interface configuration.	320
A-8	List an IP address.	320
A-9	List a host name	320
A-10	List a token expiration.	321
A-11	List a repository	321
A-12	Image subcommand	321
A-13	Storage subcommand	321
A-14	Compute subcommand	322
A-15	Current time to live value command	322
A-16	Data query to obtain the total number of volumes	323
A-17	Data query to obtain the total number of volumes and show up to 1000 entries	323
A-18	The inactivity_timeout command	323
A-19	The token_expiration_warning_time subcommand	323
A-20	List by project, role, and user name	324
A-21	List by user name, project, and role	325
A-22	Token request from the command line	326
A-23	Tenant-id request from command line	326

Notices

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

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

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

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

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

Trademarks

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

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

AIX®	IBM FlashSystem®	PowerHA®
Aspera®	IBM Garage™	PowerVM®
DS8000®	IBM Security™	Redbooks®
FlashCopy®	IBM Spectrum®	Redbooks (logo)  ®
Global Technology Services®	POWER®	Storwize®
IBM®	POWER7®	Tivoli®
IBM Cloud®	POWER8®	XIV®
IBM Cloud Pak®	POWER9™	

The following terms are trademarks of other companies:

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

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

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

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

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

Preface

IBM® Power Virtualization Center (IBM® PowerVC™) is an advanced enterprise virtualization management offering for IBM Power Systems. This IBM Redbooks® publication introduces IBM PowerVC and helps you understand its functions, planning, installation, and setup. It also shows how IBM PowerVC can integrate with systems management tools such as Ansible or Terraform and that it also integrates well into a OpenShift container environment.

IBM PowerVC Version 2.0.0 supports both large and small deployments, either by managing IBM PowerVM® that is controlled by the Hardware Management Console (HMC), or by IBM PowerVM NovaLink. With this capability, IBM PowerVC can manage IBM AIX®, IBM i, and Linux workloads that run on IBM POWER® hardware.

IBM PowerVC is available as a Standard Edition, or as a Private Cloud Edition.

IBM PowerVC includes the following features and benefits:

- ▶ Virtual image capture, import, export, deployment, and management
- ▶ Policy-based virtual machine (VM) placement to improve server usage
- ▶ Snapshots and cloning of VMs or volumes for backup or testing purposes
- ▶ Support of advanced storage capabilities such as IBM SVC vdisk mirroring of IBM Global Mirror
- ▶ Management of real-time optimization and VM resilience to increase productivity
- ▶ VM Mobility with placement policies to reduce the burden on IT staff in a simple-to-install and easy-to-use graphical user interface (GUI)
- ▶ Automated Simplified Remote Restart for improved availability of VMs ifor when a host is down
- ▶ Role-based security policies to ensure a secure environment for common tasks
- ▶ The ability to enable an administrator to enable Dynamic Resource Optimization on a schedule

IBM PowerVC Private Cloud Edition includes all of the IBM PowerVC Standard Edition features and enhancements:

- ▶ A self-service portal that allows the provisioning of new VMs without direct system administrator intervention. There is an option for policy approvals for the requests that are received from the self-service portal.
- ▶ Pre-built deploy templates that are set up by the cloud administrator that simplify the deployment of VMs by the cloud user.
- ▶ Cloud management policies that simplify management of cloud deployments.
- ▶ Metering data that can be used for chargeback.

This publication is for experienced users of IBM PowerVM and other virtualization solutions who want to understand and implement the next generation of enterprise virtualization management for Power Systems.

Unless stated otherwise, the content of this publication refers to IBM PowerVC Version 2.0.0.

Authors

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

Sachin P. Deshmukh is a certified IT Specialist with IBM Global Technology Services®(GTS) in the United States since 1999. He has worked on IBM AIX and the Power Systems platform for 28 years and has 31 years of experience on Unix OS-based systems. He is the Global AIX and Power Systems lead for IBM GTS Delivery and Integrated Operations. His areas of expertise include operating system provisioning and support, IBM PowerHA®, virtualization and Cloud platform. He provides guidance, oversight and assistance to global delivery teams supporting IBM GTS accounts. As member of the Critical Response Team, he works on major incidents, high severity issues, performs Technical Health Checks and Service Reviews. He interacts with automation, design, architecture and support teams for setting delivery standards and creating best practices documentation.

Thierry Huché is an IT Specialist working at the IBM Garage™ for Systems in Montpellier, France, as Power Systems demonstration leader. He has worked at IBM France for 32 years, and have more than 31 years of AIX System Administration and Power Systems experience. His areas of expertise include benchmarking, performance tuning, high availability, Cloud Computing, OpenShift, Ansible, Red Hat. He coauthored four Redbooks and he is speaker for Power topics at the IBM Technical University.

Stephen Lutz is a Certified Leading Technical Sales Professional for Power Systems working at IBM Germany. He holds a degree in Commercial Information Technology from the University of Applied Science Karlsruhe, Germany. He is an IBM POWER9™ champion and has 21 years of experience in AIX, Linux, virtualization, and Power Systems and its predecessors. He provides pre-sales technical support to clients, IBM Business Partners, and IBM sales representatives in Germany. Stephen is an expert in IBM PowerVC. He is involved with many IBM PowerVC client projects in DACH (Germany, Austria, and Switzerland), and has given presentation about IBM PowerVC at the IBM Technical Universities in Europe for several years.

Ahmed Mashhour is a Power Systems Lab Services Consultant at IBM Egypt. He is IBM AIX, Linux, and IBM Tivoli® certified with 15 years of professional experience in IBM AIX and Linux systems. He is an IBM AIX back-end SME who supports several customers in the US, Europe, and the Middle East. His core experience is in IBM AIX, Linux systems, clustering management, virtualization tools, and various Tivoli and database products. He has authored several publications inside and outside IBM, including co-authoring other IBM Redbooks publications. He has hosted IBM AIX, PowerVM, PowerHA, and IBM Spectrum® Scale classes more than 70 times around the world.

Christopher Emefiene Osiegbu is a Level-2 certified IT Specialist at IBM and an SME on Power Cloud solutions. He has been with IBM since 2015, supporting Western Africa accounts. He has been an instructor in several IBM and business partner led events including the IBM Technical University presenting IBM Power Cloud solutions. He has a Master of Science degree in Electrical Engineering from North Carolina A&T state University, USA. His areas of expertise are Power Systems, AIX, Red Hat and Red Hat OpenShift.

Borislav Ivanov Stoymirski is a Power Systems subject matter expert (SME) at IBM Bulgaria. His main area of expertise is solving complex hardware and software issues on IBM Power Systems products, IBM AIX, VIOS, HMC, IBM i, IBM PowerVC, PowerVM NovaLink, PowerVM, Linux on Power Systems servers and OpenShift. He has been with IBM since 2015 providing reactive break-fix, proactive, preventative and cognitive support to several clients in Europe, Africa and Middle East. He holds a Master degree in Computer and Software Engineering by Technical University of Sofia, Bulgaria and Master degree in Ecology and environmental protection by University of Chemical Technology and Metallurgy, Bulgaria. He has authored several publications and hosted technical training programs inside IBM. His areas of interest include Artificial Intelligence, Deep Learning, Machine Learning, Cyber security, Cloud and Quantum computing.

The project that created this publication was managed by:
Scott Vetter, PMP

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

Ritu Agarwal, Jordan Antonov, Ankit Arora, Pradipta Banerjee, Gergana Borisova, Karyn Corneli, Ibrahim Alade Rufai, Mariyan Ivanov, Sabine Jordan, Imranuddin Kazi, Divya Konoor, Christoph Krafft, Doyle J McCoy, Joe Mulholland, Sujeet Pai, Samuel Pavlov, Sri Ram Pisupati, Gautam Prasad, Ramakrishnan R (Techmahindra), Dr. Wolfgang Rother, Udo Sachs (SVA), Yussuf Shaikh, Shyama Venugopal, Russ Young, Jochen Ziegler

Thanks to the authors of the previous editions of this book.

The authors of the first and second edition, *IBM PowerVC Version 1.2.0 and 1.2.1 Introduction and Configuration*, which was published in October 2014, were Bruno Blanchard, Guillermo Corti, Sylvain Delabarre, Ho Jin Kim, Ondrej Plachy, Marcos Quezada, and Gustavo Santos.

The authors of the third edition, *IBM PowerVC Version 1.2.3 Introduction and Configuration*, which was published in October 2015, were Marco Barboni, Guillermo Corti, Benoit Creau, and Liang Hou Xu.

The authors of the fourth edition, *IBM PowerVC Version 1.3.1 Introduction and Configuration Including IBM Cloud® PowerVC Manager*, which was published in September 2016, were Guillermo Corti, Mika Heino and Paul Sonnenberg.

The authors of the fifth edition, *IBM PowerVC Version 1.3.2 Introduction and Configuration*, which was published in January 2017, were Javier Bazan Lazcano and Martin Parrella.

Now you can become a published author, too!

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

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

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

ibm.com/redbooks

- ▶ Send your comments in an email to:

redbooks@us.ibm.com

- ▶ Mail your comments to:

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

Stay connected to IBM Redbooks

- ▶ Find us on LinkedIn:

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



IBM PowerVC introduction

IBM Power Virtualization Center (IBM PowerVC) is an advanced virtualization and cloud management offering for Power Systems servers based on OpenStack technology. This comprehensive virtualization and cloud management offering is simple to install and use, and enables virtual machine (VM) setup and management. IBM PowerVC simplifies the management of the virtualization for Power Systems servers that run on IBM AIX, IBM i, and Linux operating systems.

IBM PowerVC Version 2.0.0 provides many new functional enhancements such as cloning volumes and VMs, snapshots, and storage-related features such as consistency groups and IBM Global Mirror support. Also available is IBM PowerVC for Private Cloud, which includes all the functions of IBM PowerVC plus additional cloud-related features that are described in this publication.

This publication provides introductory and configuration information for IBM PowerVC. After it presents an overview of IBM PowerVC in this chapter, it covers the following topics in subsequent chapters:

- ▶ Planning information in Chapter 2, “IBM PowerVC planning” on page 13.
- ▶ Guidelines in Chapter 3, “IBM PowerVC installation” on page 69.
- ▶ General configuration and setup that are common to all variants of IBM PowerVC in Chapter 4, “IBM PowerVC for managing IBM PowerVM” on page 107.
- ▶ Information that is specific to using IBM PowerVC for Private Cloud in Chapter 5, “IBM PowerVC for Private Cloud” on page 255.
- ▶ PowerVC integration with Ansible and Terraform in Chapter 6, “IBM PowerVC integration with Terraform and Ansible” on page 267.
- ▶ PowerVC in a Hybrid Cloud with OpenShift in Chapter 7, “IBM PowerVC in the Hybrid Cloud” on page 293.

This publication also includes useful information about IBM PowerVM NovaLink and its interaction with IBM PowerVC.

This chapter covers the following topics:

- ▶ 1.1, “IBM PowerVC overview” on page 2
- ▶ 1.2, “IBM PowerVC editions” on page 7
- ▶ 1.3, “PowerVC versions and releases” on page 7
- ▶ 1.4, “IBM PowerVC adoption” on page 11

1.1 IBM PowerVC overview

IBM Power Virtualization Center (PowerVC) is an advanced virtualization and cloud management offering, which is built on OpenStack, that provides simplified virtualization management and cloud deployments for IBM AIX, IBM i, and Linux VMs running on IBM Power Systems. PowerVC is designed to improve administrator productivity and simplify the cloud management of VMs on Power Systems servers. With PowerVC, you can do the following actions:

- ▶ Create VMs and resize the VMs CPU and memory.
- ▶ Attach disk volumes or additional networks to those VMs.
- ▶ Import existing VMs and volumes so that they can be managed by IBM PowerVC.
- ▶ Monitor the use of resources in your environment.
- ▶ Take snapshots of a VM or clone it.
- ▶ Migrate VMs while they are running (live migration between physical servers).
- ▶ Remote restart VMs in case of a server failure.
- ▶ Use advanced storage technologies such as vdisk mirroring or IBM Global mirror.
- ▶ Improve resource usage to reduce capital expense and power consumption.
- ▶ Increase agility and execution to respond quickly to changing business requirements.
- ▶ Increase IT productivity and responsiveness.
- ▶ Simplify Power Systems virtualization management.
- ▶ Accelerate repeatable, error-free virtualization deployments.

IBM PowerVC can manage AIX, IBM i, and Linux-based VMs running under PowerVM virtualization connected to an HMC or using NovaLink. This release supports the scale-out and the enterprise Power Systems servers that are built on IBM POWER8®, IBM POWER9, and subsequent technologies.

1.1.1 IBM PowerVC functions and advantages

Why IBM PowerVC? Why do you need another virtualization management offering? When more than 70% of IT budgets are spent on operations and maintenance, IT clients legitimately expect vendors to focus their new development efforts to reduce IT costs and foster innovation within IT departments.

IBM PowerVC gives IBM Power Systems clients the following advantages:

- ▶ It is deeply integrated with Power Systems.
- ▶ It provides virtualization management tools.
- ▶ It eases the integration of servers that are managed by PowerVM in automated IT environments, such as clouds.
- ▶ It is a building block of IBM Infrastructure as a Service (IaaS), based on Power Systems.

- ▶ PowerVC integrated with other cloud management tool like Ansible, Terraform or OpenShift and can be integrated into orchestration tools like the IBM Cloud Automation Manager (CAM), VMware vRealize or the SAP Landscape Management (LaMa).
- ▶ PowerVC provides also an easy exchange of VM images between private and public clouds.

IBM PowerVC is an addition to the existing PowerVM set of enterprise virtualization technologies that provide virtualization management. It is based on open standards and integrates server management with storage and network management.

Because IBM PowerVC is based on the OpenStack initiative, Power Systems can be managed by tools that are compatible with OpenStack standards. When a system is controlled by IBM PowerVC, it can be managed in one of three ways:

- ▶ By a system administrator using the IBM PowerVC graphical user interface (GUI)
- ▶ By a system administrator that uses scripts containing the IBM PowerVC Representational State Transfer (REST) application programming interfaces (APIs)
- ▶ By higher-level tools that call IBM PowerVC by using standard OpenStack API

The PowerVC offerings are positioned within the available solutions for Power Systems cloud as follows:

- ▶ IBM PowerVC: Advanced Virtualization Management
- ▶ IBM PowerVC for Private Cloud: Basic Cloud
- ▶ IBM Cloud Automation Manager: Advanced Cloud
- ▶ VMware vRealize: Advanced Cloud

PowerVC provides a systems management product that enterprise clients require to manage effectively the advanced features that are offered by IBM premium hardware. It reduces resource use and manages workloads for performance and availability.

1.1.2 IBM PowerVM NovaLink overview

PowerVM NovaLink is a virtualization management paradigm for PowerVM systems and allows for dramatic scale improvements for PowerVM based IBM PowerVC environments. For more information about PowerVM NovaLink and its benefits, see the following blog in the IBM community:

<https://community.ibm.com/community/user/power/blogs/archive-user/2015/11/01/power-vc-NovaLink-overview>

Leveraging the PowerVM NovaLink architecture, IBM PowerVC can significantly increase its scaling for PowerVM based systems. In an existing HMC-managed environment, IBM PowerVC can manage up to 30 hosts and up to 3000 VMs. In a PowerVM NovaLink based environment, IBM PowerVC can manage up to 50 hosts and 10000 VMs. You can use IBM PowerVC to manage your new PowerVM NovaLink systems while still managing your HMC managed systems.

PowerVM NovaLink architecture

PowerVM NovaLink enables highly scalable modern cloud management and deployment for critical enterprise workloads by using a proven PowerVM infrastructure and OpenStack technology. The PowerVM Hypervisor and Virtual I/O Server (VIOS) provide an infrastructure that excels in running enterprise workloads. These items, combined with POWER9 hardware, provide an excellent base on which to build a cloud. However, you cannot operate a cloud without specialized management tools, which is where IBM PowerVC comes into play. This OpenStack-based management offering integrates with PowerVM NovaLink to deliver a PowerVM cloud that is truly scalable, efficient, and simple to manage.

The PowerVM NovaLink architecture changes the virtualization management point for IBM PowerVC. With PowerVM NovaLink, a thin “management” VM exists on the system. This VM runs a Linux operating system and the NovaLink. You can see the thin PowerVM NovaLink partition in Figure 1-1, which is denoted as NVL.

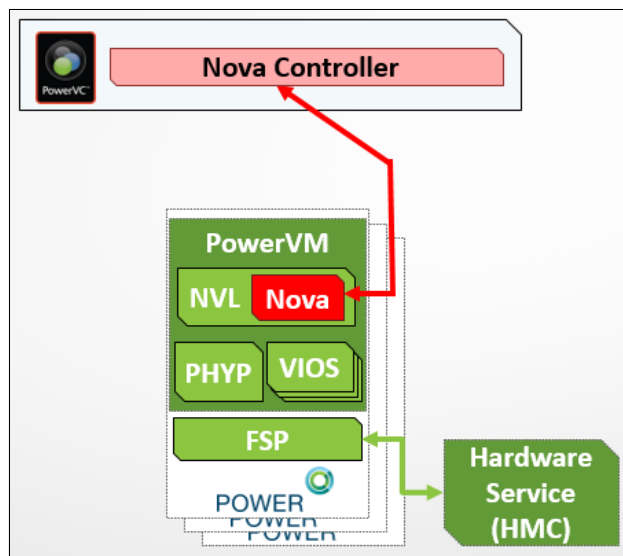


Figure 1-1 PowerVM NovaLink partition

Figure 1-1 shows that the architecture between IBM PowerVC and a PowerVM system is dramatically simplified. The Nova compute processes is run directly on the PowerVM NovaLink thin VM so that the IBM PowerVC can dramatically scale out the number of hosts that it can manage by using this one-to-one link. It also reduces the load on an administrator’s HMC, allowing the hosts to connect more systems to a given HMC than they would otherwise without NovaLink.

In addition, the PowerVM NovaLink code is tuned directly for IBM PowerVC and OpenStack use. This increased efficiency allows IBM PowerVC to scale a single system to 1,000 VMs, double the current 500 VMs per system limitation that exists today. More importantly, it is aligned with the capabilities of the PowerVM platform itself.

PowerVM NovaLink user experience

The integration of PowerVM NovaLink provides a unified PowerVM experience. Whether you choose to have IBM PowerVC managed through PowerVM NovaLink (to take advantage of the scale and speed) or by using the traditional HMC path, IBM PowerVC provides you with a consistent experience.

Figure 1-2 shows the host list in PowerVC, where you can see one NovaLink-connected system and three HMC-connected systems side by side. In general, and administrator will be familiar with most of the current implementation. IBM PowerVC can mix the management types, which means that a single IBM PowerVC can manage some systems through an HMC and others through PowerVM NovaLink.

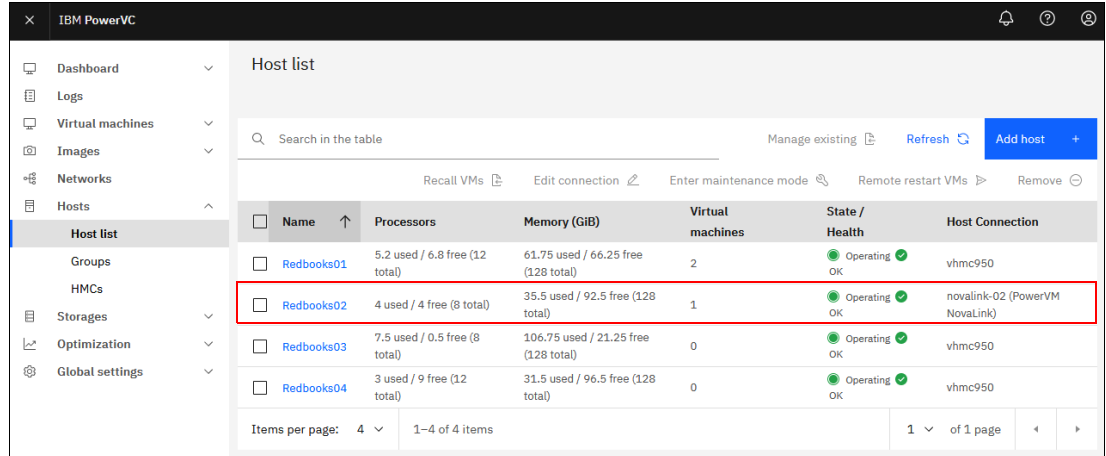


Figure 1-2 IBM PowerVC interface

There are some areas where changes are evident in the user interface. The most obvious one is the Host Registration window. Although host registration for an HMC managed system remains unchanged, there is a path for PowerVM NovaLink host registration. Administrators provide the IP address and credentials of the PowerVM NovaLink VM, which IBM PowerVC uses to register the system. Figure 1-3 shows the Host Registration window.

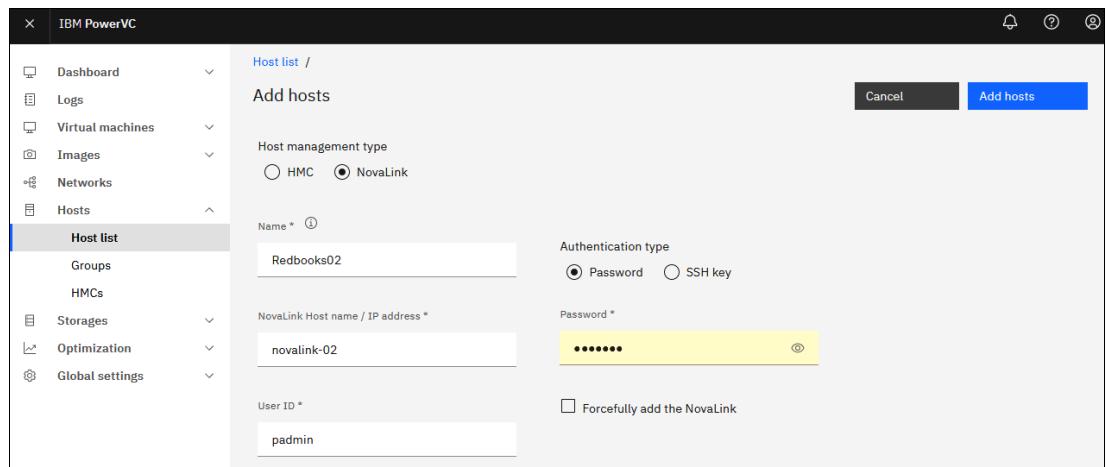


Figure 1-3 Host Registration window

From an architectural point of view, a mixed environment of HMC- and NovaLink-connected systems is shown in Figure 1-4. You can see that the OpenStack nova process, that is in charge of managing a compute node (a host) is either running in PowerVC for the HMC-connected systems (each for one host) or in the NovaLink management partition.

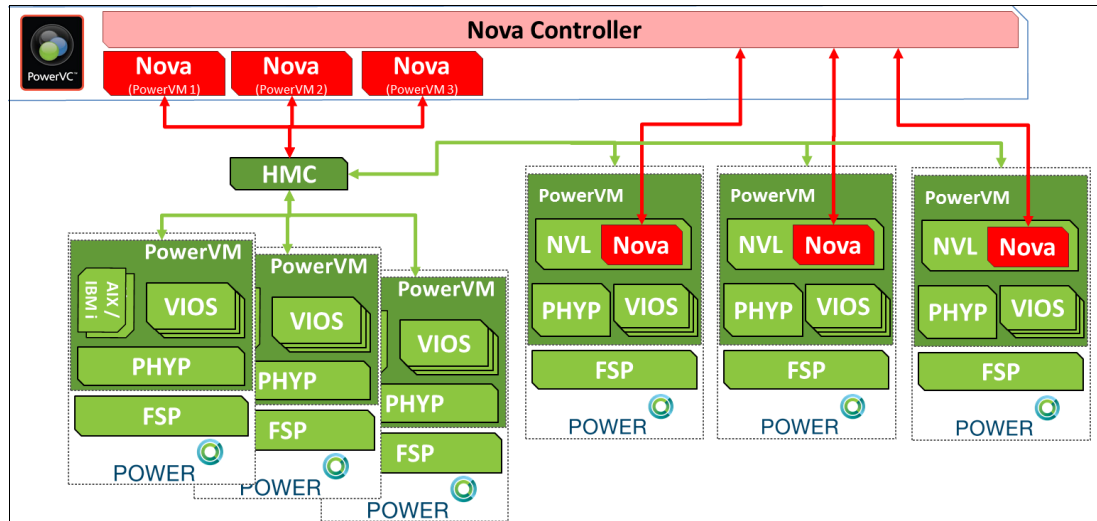


Figure 1-4 PowerVM NovaLink diagrams

PowerVM NovaLink unique features

Besides the higher scalability and performance of NovaLink managed hosts, PowerVC has some features that are only available for NovaLink-connected hosts:

- ▶ Console window out of PowerVC also when the VM has no network connection.
- ▶ Memory utilization graphs in the host details.
- ▶ The Dynamic Resource Optimizer (DRO) can use the memory utilization monitoring to recommend or take actions, such as moving an virtual machine to another host.

Key prerequisites for NovaLink

For IBM PowerVC, the key prerequisites are as follows:

- ▶ IBM PowerVM Edition (5765-VS2) for basic functions, and IBM PowerVM Enterprise Edition (5765-VE3) or IBM PowerVM Linux Edition (5765-VL3). NovaLink is part of these offerings. It is not part of the PowerVC offering.
- ▶ For Remote Restart, POWER8 systems with firmware V8.2 or higher, or POWER9 or later systems are required.
- ▶ PowerVM NovaLink requires systems with a POWER8 processor and Firmware 840 or later, or POWER9 and later systems that are not managed by an HMC.

PowerVM NovaLink provides significant advantages for PowerVM users who want to scale up their environments. NovaLink is highly concurrent and highly scalable, and can reduce infrastructure complexity. The existing IBM PowerVC experience is preserved, enabling administrators to take advantage of these benefits quickly.

1.2 IBM PowerVC editions

IBM offers two editions of PowerVC:

- ▶ IBM PowerVC
- ▶ IBM PowerVC for Private Cloud

1.2.1 IBM PowerVC

IBM PowerVC manages PowerVM servers which are controlled by either the HMC or PowerVM NovaLink. These servers include the scale out as well as the scale up servers. Supported Power systems are IBM POWER7® and later processor classifications.

Either of the two editions mentioned above can be selected during the installation of PowerVC. IBM PowerVC supports the addition of Power systems and other resources that make up the computing environment:

- ▶ Storage devices
- ▶ Fabrics
- ▶ Network segments

IBM PowerVC has all functions of PowerVC with the exception of the self-service portal, which is exclusive to IBM PowerVC for Private Cloud.

1.2.2 IBM PowerVC for Private Cloud

In today's fast paced IT environment, system administrators are responsible for the continuous availability of the system infrastructure. Yet, there is a need to rapidly provision virtual machines for faster go to market or other purposes. Hence, the need for a self-service portal that users can log into and provision virtual machines under policies set by the system administrator.

The self-service portal along with all features available in IBM PowerVC provides a private cloud experience for users of IBM PowerVC for Private Cloud.

1.3 PowerVC versions and releases

This section describes the updates that have been introduced since IBM Power Virtualization Center (IBM PowerVC) Version 1.4.0, with special focus on Version 2.0.0.

1.3.1 IBM PowerVC release to OpenStack edition cross-reference

Table 1-1 cross-references the IBM PowerVC releases to editions of OpenStack.

Table 1-1 IBM PowerVC releases cross-referenced to OpenStack versions

IBM PowerVC release	General availability date	OpenStack edition
V1.4.0	December 2017	Pike
V1.4.1	June 2018	Queens
V1.4.2	December 2018	Queens
V1.4.3	June 2019	Stein

IBM PowerVC release	General availability date	OpenStack edition
V1.4.4	December 2019	Stein
V2.0.0	December 2020	Ussuri

Note: The end of lifecycle support for IBM PowerVC release 1.4.x is April 2022. However new fixes support is provided only for v1.4.4 in addition to v2.0.0, at the time of writing. For the earlier IBM PowerVC 1.4 versions, “how-to-questions” support is provided for existing fixes only.

1.3.2 IBM PowerVC 1.4 versions

The key enhancements and features introduced in each release of the PowerVC 1.4.x versions are discussed in the following subsections.

IBM PowerVC Version 1.4.0

IBM PowerVC Version 1.4.0 provides simplified virtualization management and cloud deployments for IBM AIX, IBM i, and Linux virtual machines (VM) running on IBM Power Systems. It is enhanced with the following features:

- ▶ Image import and export
- ▶ VM live capture
- ▶ PowerVC UI update
- ▶ Flexible SAN zone names
- ▶ Software-defined storage
- ▶ KVM on Power support

IBM PowerVC Version 1.4.1

IBM PowerVC Version 1.4.1 is updated to transform the next generation of IBM POWER9 technology-based enterprise servers into fully-functional private cloud platforms by adding a sophisticated cloud management interface and intuitive self-experience for cloud users. PowerVC 1.4.1 includes the following capabilities:

- ▶ Support for POWER9 servers running the POWERVM hypervisor
- ▶ Streamlined deployment of SAP HANA workloads
- ▶ Integration with IBM Cloud Private
- ▶ Improvements to PowerVC Manager for Software-Defined Infrastructure
- ▶ New zoning capability “Dual-initiator target zoning”

IBM PowerVC Version 1.4.2

IBM PowerVC Version 1.4.2 added the following major enhancements:

- ▶ SR-IOV support over HMC
- ▶ AIX Secure Boot support
- ▶ SDI enhancements
- ▶ Cloud enhancements (copy of deploy or storage templates)
- ▶ Storage: pluggable FC tape driver

IBM PowerVC Version 1.4.3

IBM PowerVC Version 1.4.3 is designed to simplify the management of virtualized and cloud resources in IBM Power Systems environments. PowerVC 1.4.3 provides the following features:

- ▶ Support for all POWER9 enterprise servers
- ▶ Redundant HMC (Auto Failover) support
- ▶ Dell EMC PowerMAX support
- ▶ VMAX REST driver upgrade
- ▶ Redundant Mover Service Partitions (MSP) for Live Partition Mobility (LPM)
- ▶ Image Sharing
- ▶ Cross HMC remote restart
- ▶ Support for EPP 2.0
- ▶ Metering of network and storage traffic

IBM PowerVC Version 1.4.4

IBM PowerVC Version 1.4.4, which is designed to simplify the management of virtual resources in IBM Power Systems environments, has been updated with several enhancements to storage integration, usability, and performance:

- ▶ IBM i license key injection
- ▶ Hitachi GAD support
- ▶ Initiator storage tagging
- ▶ Live Partition Mobility (LPM) VMs to the original system after evacuation
- ▶ Ability to pin VMs to a specific host
- ▶ Ability to dynamically add PowerVC created VM to an HMC user resource role
- ▶ Inactive partition migration
- ▶ Image sharing between projects
- ▶ New restricted administrator role, with no deletion access
- ▶ NovaLink support for multivolume attachment
- ▶ FlexVolume driver support for Red Hat OpenShift

1.3.3 IBM PowerVC Version 2.0.0

IBM PowerVC Version 2.0.0, which is designed to simplify the management of virtual resources in IBM Power Systems environments, has several enhancements to its storage integration, usability, and performance capabilities.

New GUI with improved usability

IBM PowerVC Version 2.0.0 introduces a new user interface that is designed based on the Carbon framework. Carbon is IBM's open source design system for products and digital experiences.

The PowerVC Version 2.0.0 GUI is developed over MERN stack, a web development framework that stands for MongoDB, Express, ReactJS, and NodeJS. At a high level, MongoDB is used as a caching layer, which is used by NodeJS or Express to feed data to the GUI. The data loading on the GUI is much faster than any of the previous releases.

SUSE Linux Enterprise Server support for the PowerVC server

The installation of the PowerVC management server is now supported on SUSE Linux Enterprise Server 15 SP1 and SUSE Linux Enterprise Server 15 SP2. The PowerVC installation continues to be supported on Red Hat Enterprise Linux, with versions 8.2 and 8.3.

Red Hat Enterprise Linux 8 support with Python compatibility

The PowerVC Version 2.0.0 has support for only Python3 on management and NovaLink nodes, in alignment with the Python and OpenStack community.

Consistency group snapshot and restore

An internal consistency group will be created before the snapshot is completed, if more than one volume is captured through a snapshot. Virtual machine (instance) snapshot allows taking of consistent snapshot of the volumes attached to a virtual machine. Multiple snapshot copies can be created for a single virtual machine. One or more volumes can be chosen while creating instance snapshots. These instance snapshots can be further used to restore the virtual machine at any point in time. Virtual machine (instance) snapshot is only supported for IBM Storwize®, PowerMAX and Hitachi storage providers with non-replicated volumes.

Volume cloning for backup purposes

Clone volume feature allows creation of a consistent copy of volumes that are available or attached to a virtual machine.

VM cloning to simplify the redundant deployment of workloads

The clone virtual machine feature allows the creation of new identical virtual machine. A copy of an existing virtual machine with same CPU, memory, availability priority, volumes, etc. can be created.

Support for IBM Global Mirror with IBM Storwize

Starting with PowerVC Version 2.0.0, the IBM Global Mirror functionality of IBM Storwize for volume replication can be used. A secondary storage for global mirror on the storage details page of the user interface or through API can be registered.

Persistent memory virtualization

The persistent memory virtualization feature can be enabled for a virtual machine. When the virtual machine is deployed with a persistent memory volume, during VM restart the persistent memory volume is retained, thus reducing the restart time.

Retype support for existing volumes

The current storage template of a volume to a target storage template can be changed. The volume properties like QoS, provisioning, SLO/workload, etc. can also be changed. Migration of the volume to another pool specified in the target template can be performed.

Dynamic Resource Optimizer (DRO) policy enhancements

Two new DRO policies are available:

- ▶ Combined CPU and memory utilization
- ▶ Shared Processor pool utilization

IBM PowerVC Dynamic Resource Optimizer (DRO) continuously monitors the resource utilization of physical hosts and virtual machines in a cloud environment. Starting with PowerVC Version 2.0.0, DRO can monitor CPU, memory utilization, and shared processor pools. The CPU and Memory Utilization Policy can be selected for the DRO to optimize hosts by monitoring both CPU and memory utilization of the hosts.

Pure Storage

PowerVC supports Pure Storage as an integrated driver. A Pure Storage driver can be added by using the GUI.

PowerMAX updates

The following updates are available:

- ▶ Port group load balancing: A list of port groups on the storage template can be selected. When a virtual machine is deployed using this storage template, the port group with least number of masking views will be used to form the masking view.
- ▶ Host name format: The template for short host name customization for PowerMAX storage provider can be set.
- ▶ Port group template: The template for port group customization for PowerMax storage providers can be set.

Hitachi GAD

Live capture support is now provided for Hitachi GAD.

Multifactor Authentication (MFA) support

Starting with PowerVC Version 2.0.0, an additional authentication mechanism called Time-based One-Time Password (TOTP) is added to provide enhanced security for the users logging in. TOTP along with password must be provided for user authentication, making it a two-factor or multi-factor authentication. The MFA support is available in PowerVC for Private Cloud.

1.4 IBM PowerVC adoption

Two features are useful for a smooth adoption of IBM PowerVC in an existing environment:

- ▶ When IBM PowerVC manages a physical server, it can manage the full set or only a subset of the partitions that are hosted on that server.
- ▶ When IBM PowerVC is adopted in an environment where partitions are already in production, IBM PowerVC can discover the existing partitions and selectively start to manage them.

Therefore, the adoption of IBM PowerVC in an existing environment does not require a major change. It can be a smooth transition that is planned over several days or more.

IBM PowerVC provides a way to check a system to ensure all the prerequisites are met before an installation by running a script. For more information, see section 3.2, “Installing IBM PowerVC” on page 81.

When you are not sure, if PowerVC fits your requirements, test it. You can find a 90 days evaluation copy under the following website:

<https://www-01.ibm.com/marketing/iwm/iwm/web/preLogin.do?source=IPVCv12E>

Note: You need an IBM ID to enter that website. If you don't have an IBM ID, follow the Create an IBM ID link on the login to the website.



IBM PowerVC planning

This chapter describes the key aspects of IBM® Power Virtualization Center (IBM PowerVC) installation planning divided into the following sections:

- ▶ Section 2.1, “IBM PowerVC requirements” on page 14 and 2.2, “IBM PowerVM NovaLink requirements” on page 19 presents the hardware and software requirements for the various components of an IBM PowerVC environment: management station, managed hosts, network, storage area network (SAN), and storage devices.
- ▶ Sections 2.3, “Host and partition management planning” on page 22 through 2.10, “Product information” on page 66 provide detailed planning information for various aspects of the environment’s setup including:
 - 2.4, “Placement policies and templates” on page 30
 - 2.5, “IBM PowerVC storage access SAN planning” on page 41
 - 2.6, “Storage management planning” on page 45
 - 2.7, “Network management planning” on page 56
 - 2.8, “Planning users and groups” on page 61
 - 2.9, “Security management planning” on page 65
 - 2.10, “Product information” on page 66

2.1 IBM PowerVC requirements

This section describes the necessary software and hardware components to implement IBM PowerVC to manage AIX, IBM i, and Linux platforms.

IBM PowerVC for Private Cloud continues to be included along with IBM PowerVC in the IBM PowerVC 2.0.0 installation media. For information about available releases, see this website:

<http://www.ibm.com/software/support/lifecycle/>

In addition to the functionalities offered by IBM PowerVC, IBM PowerVC for Private Cloud provides cloud capabilities such as self-service portal that allows for the provisioning of new virtual machines (VMs) in a PowerVM-based private cloud, without direct system administrator intervention.

2.1.1 Hardware and software information

The following sections describe the hardware, software, and minimum resource requirements for Version 2.0.0 of IBM PowerVC for Private Cloud and IBM PowerVC. For the complete requirements, see the following IBM Knowledge Center:

- ▶ IBM PowerVC Version 2.0.0

https://www.ibm.com/support/knowledgecenter/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/kc_welcome-standard-supermap.html

- ▶ IBM PowerVC for Private Cloud Version 2.0.0

https://www.ibm.com/support/knowledgecenter/SSVSPA_2.0.0/com.ibm.powervc.cloud.help.doc/kc_welcome-cloud-supermap.html

Note: Support for KVM has been deprecated in IBM PowerVC Version 2.0.0.

2.1.2 Hardware and software requirements for IBM PowerVC

The following information provides a consolidated view of the hardware and software requirements for both IBM PowerVC and IBM PowerVC for Private Cloud.

IBM PowerVC management host and managed hosts

The IBM PowerVC architecture supports a single management host for each managed domain. It is not possible to configure redundant IBM PowerVC management hosts that control the same objects.

The VM that hosts the IBM PowerVC management host should be dedicated to this function. No other software or application should be installed on this VM. However, you can install software for the management of this VM, such as monitoring agents and data collection tools for audit or security. Table 2-1 on page 15 lists the IBM PowerVC hardware and software requirements.

Table 2-1 Hardware and OS requirements

Host type	Supported hardware	Supported operating systems
IBM PowerVC management server	<ul style="list-style-type: none"> ▶ ppc64le (Power8 and above) ▶ x86_64 Support for PowerVC installation on ppc64 architecture is being withdrawn.	<ul style="list-style-type: none"> ▶ Power Platform: Red Hat Enterprise Linux 8.2 and Red Hat Enterprise Linux 8.3; SUSE Linux Enterprise Server 15 SP1 and SUSE Linux Enterprise Server 15 SP2 ▶ x86_64 Platform: Red Hat Enterprise Linux 8.2 and Red Hat Enterprise Linux 8.3
Managed Hosts	IBM POWER7 or later processor-based servers. All form factors are included such as chassis, rack, blade and Power Flex. Managed hosts should have a minimum of 4 cores and 8 GB memory.	Refer to Table 2-3 on page 16.

Table 2-2 describes the minimum and recommended resources that are required for IBM PowerVC VMs. In the table, the meaning of the processor capacity row depends on the type of host that is used as the IBM PowerVC management host:

- ▶ If the IBM PowerVC management host is PowerVM, processor capacity refers to either the number of processor units of entitled capacity for shared processors, or the number of dedicated processors.
- ▶ If the IBM PowerVC management host is x86, processor capacity refers to the number of physical cores.

Table 2-2 Resource requirements for 5 hosts, 2 storage providers, and 2 fabrics

Item	Minimum	Recommended				
		Up to 500	501-1000	1001-2000	2001-3000	3001-6000
Number of VMs						
Processor capacity	1	2	4	8	8	12
Virtual CPUs	2	2	4	8	8	12
Memory and swap space (GB)	22	32	35	40	45	55
Disk used (GB)	80	100	120	140	160	180

Note: The resource requirements for PowerVC and PowerVC for Private Cloud will vary when there are more than five hosts, two storage providers and two fabrics. Refer to the 2.1.1, “Hardware and software information” on page 14 section for the links to IBM Knowledge Center.

PowerVC has the following file system space requirements:

/tmp	250 MB
/usr	250 MB
/opt	5 GB
/srv	2 GB
/var	3 GB

It is recommended that 20% of the disk space is assigned to /var.

Guest operating system support

Table 2-3 lists the supported virtual machine (VM) operating systems on the managed hosts.

Table 2-3 Supported operating systems for VMs on the managed hosts

Operating system	Little Endian (LE) or Big Endian (BE)	Version
AIX	BE	7.1 TL0, SP0 7.2 TL0, SP0
IBM i	BE	7.2 TL1 7.3 TR3 7.4 Note: Version 7.2 TR8 or 7.3 TR4 is required for POWER9
Red Hat Enterprise Linux	LE	7.6, Red Hat Enterprise Linux 7.6-ALT, 7.7, 7.8, 7.9 8.0, 8.1, 8.2, and 8.3
SUSE Linux Enterprise Server	LE	12 & 15 SUSE Linux Enterprise Server 15 SP1 and SUSE Linux Enterprise Server 15 SP2
Ubuntu	LE	16.04

Note: The operating system versions shown above are the minimum supported levels by PowerVC. Newer versions may be available from the vendor. It is recommended that systems be maintained at versions that are currently supported. Check the support lifecycle documentation for the respective vendor.

Hardware Management Console

Table 2-4 shows the Hardware Management Console (HMC) version and release requirements to support IBM PowerVC.

Table 2-4 HMC minimum requirements

Item	Requirement
Software level	V9.1.941 V9.2.950
Hardware level	For the HMC hardware that supports the software levels listed above, refer to: https://www.ibm.com/support/pages/recommended-fixes-hmc-code-upgrades

Item	Requirement
HMC memory capacity	Requirements Up to 300 VMs: 4 GB More than 300VMs: 8 GB Recommended Up to 300 VMs: 8 GB More than 300 VMs: 16 GB

As a preferred practice, update to the latest HMC fix pack for the specific HMC release. You can check the fixes for HMC and other IBM products by using the IBM Fix Level Recommendation Tool at the following website:

<https://www14.software.ibm.com/webapp/set2/flrt/home>

You can get the latest fix packages from IBM Fix Central at the following website:

<http://www.ibm.com/support/fixcentral/>

Virtual I/O Server

Table 2-5 includes the Virtual I/O Server (VIOS) version requirements for IBM PowerVC managing PowerVM.

Table 2-5 Supported VIOS platforms

Platform	Requirement
Virtual I/O Server	<ul style="list-style-type: none"> ▶ Version 3.1.0.30 ▶ Version 3.1.1.25 ▶ Version 3.1.2.10

Note: A minimum of 6 GB is required for VIOS 3.1.x.

Network

Table 2-6 lists the network infrastructure that is supported by IBM PowerVC.

PowerVC can communicate with hosts and storage over IPv4 addresses. PowerVC Version 2.0.0 does not support IPv6.

Table 2-6 Supported network hardware and software

Item	Requirement
Network switches	IBM PowerVC does not manage network switches, but it supports network configurations that use virtual LAN (VLAN)-capable switches.
Virtual Networks	PowerVC Supports Shared Ethernet Adapters (SEA) for virtual machine networking and SR-IOV based vNICs on PowerVM.

Storage

Table 2-7 lists the storage systems and drivers, SAN switches and connectivity types that are supported by PowerVC.

Table 2-7 Supported Storage and Fabric drivers for PowerVM

Item	Supported
Storage systems and drivers	IBM Storwize family. IBM FlashSystem® A9000 and A9000R. IBM SAN Volume Controller (SVC). IBM XIV® Storage system. IBM System Storage DS8000®. Hitachi Enterprise Hitachi Block Storage Driver (HBSD) 10.0.0. Dell EMC VNX and Dell EMC PowerMAX (VMAX). Pure Storage. Pluggable (A pluggable storage device is an OpenStack supported storage device.) Dell EMC VNX requires the PowerVC management server running on an x86 host.
SAN switches and Fabric Drivers	Brocade zone OpenStack driver for Ussuri. Cisco Fibre Channel zone driver. VSAN support is included. Pluggable. Any fabric supported by OpenStack driver can be registered with PowerVC. PowerVC supports up to 25 fabrics.
Storage connectivity type	FC attachment through at least one N_PORT Virtualization (NPIV)-capable hosts bus adapter (HBA) on each host. Virtual SCSI (vSCSI). Shared Storage Pools.

Supported storage connectivity options

Each virtual machine is associated with a single connectivity group that manages the virtual machine's connectivity to storage volumes. Each storage connectivity group supports a single connectivity type for volumes in the boot set. Each VM can have only one type of boot connectivity and one type of data connectivity.

Table 2-8 shows the supported storage connectivity options. Both NPIV and vSCSI data volumes are shown supported for a vSCSI boot volume, but you can have either NPIV or vSCSI data volumes, not both. This is because storage connectivity groups only support one type of connectivity. Similarly, you cannot have both shared storage pool volumes and NPIV boot volumes in a VM.

Table 2-8 Supported storage connectivity options

Boot volume	Data volume		
	Shared storage pool	NPIV	vSCSI
Shared storage pool (SSP)	X	X	
NPIV		X	
vSCSI		X	X

Security

PowerVC can be optionally configured to work with an existing Lightweight Directory Access Protocol (LDAP) server.

Table 2-9 shows the supported LDAP software and requirements.

Table 2-9 Supported LDAP server versions

Item	Requirement
Lightweight Directory Access Protocol (LDAP) server (optional)	<ul style="list-style-type: none">▶ OpenLDAP version 2.0▶ Microsoft Active Directory 2016▶ IBM Security™ Directory Suite 8.0.1

2.1.3 Other hardware compatibility

IBM PowerVC is based on OpenStack, so rather than being compatible with specific hardware devices, IBM PowerVC is compatible with drivers that conform to OpenStack standards. They are called *pluggable devices* in IBM PowerVC. Therefore, IBM PowerVC can take advantage of hardware devices that are available from vendors that provide OpenStack-compatible drivers for their products. The level of functionality that the pluggable devices have, will depend on the driver. IBM cannot state the support of other hardware vendors for their specific devices and drivers that are supported by IBM PowerVC, so check with the vendors to learn about their drivers. For more information about pluggable devices, see IBM Knowledge Center:

https://www.ibm.com/support/knowledgecenter/en/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_planning_storage_hmc.html

2.1.4 Web browser supported

IBM PowerVC works on many web browsers:

- ▶ Mozilla Firefox ESR 83

Note: PowerVC does not load if *Ask me every time* is selected as the custom history setting.

- ▶ Google Chrome, Version 86.0.4240.198
- ▶ Microsoft Edge Version 86.0.622.58 (official build) (64-bit) and later
- ▶ Safari, Version 12

2.2 IBM PowerVM NovaLink requirements

PowerVM NovaLink is a software interface that is used for virtualization management. You can install PowerVM NovaLink on a supported PowerVM server. PowerVM NovaLink enables highly scalable modern cloud management and deployment of critical enterprise workloads. PowerVM NovaLink can be used to provision large numbers of Virtual Machines (VMs) on PowerVM servers quickly and at a reduced cost.

2.2.1 PowerVM NovaLink system requirements

For normal operation, PowerVM NovaLink requires specific hardware and software criteria to be met.

Power Systems requirements

PowerVM NovaLink can be installed on POWER8 processor-based servers running firmware level FW840 or later, PowerVM NovaLink can also be installed on POWER9 processor-based servers running firmware level FW910, or later. If the target server does not meet the minimum firmware requirements, it must be updated/upgraded before installing PowerVM NovaLink.

Note: NovaLink levels older than NovaLink 1.0.0.16 Feb 2020 release with partitions running certain SR-IOV capable adapters are NOT supported on POWER9 servers running firmware level FW930 and FW940. Upgrading systems in this configuration is supported only if NovaLink is first updated to NovaLink 1.0.0.16 Feb 2020 release or later. Additional details are available on IBM Support [Fix Central](#) website, in the firmware description file.

IBM PowerVC requirements

For normal operation, IBM PowerVC Version 2.0.0 requires PowerVM NovaLink version 2.0.0.

PowerVM NovaLink partition requirements

PowerVM NovaLink requires its own partition on the managed system. The PowerVM NovaLink Version 2.0.0 partition requires the following system resources:

- ▶ 0.5 shared processors that are uncapped with a non-zero weight and two virtual processors.
- ▶ 8 GB of memory is required for Red Hat Enterprise Linux 8.2. The minimum required memory capacity for Ubuntu 18.04 is 6.5 GB, the memory capacity can be lowered to 2.5 GB after installation. See Table 2-10 for the memory requirements when scaling VMs.
- ▶ At least 30 GB of vSCSI storage (LV, PV, or NPIV).
- ▶ A virtualized network that is bridged through SEA.
- ▶ Minimum virtual slots that are set to 200 or higher.

Table 2-10 Amount of memory that is needed for a single NovaLink Version 2.0.0 partition

Number of VMs	Up to 250	251 - 500	More than 500
Red Hat Enterprise Linux 8.2 - Memory needed in a standard environment (in GB)	8	14	18
Ubuntu 18.04 - Memory needed in a standard environment (in GB)	2.5	5	10

PowerVM NovaLink considerations:

- ▶ When the PowerVM NovaLink environment is installed on a new managed system, the PowerVM NovaLink installer creates the PowerVM NovaLink partition automatically. In such cases, the PowerVM NovaLink installer creates the PowerVM NovaLink partition on a new managed system, the PowerVM NovaLink installer always uses virtualized storage that is provisioned from the Virtual I/O Server. The installer allocates logical volumes from the VIOS rootvg to the PowerVM NovaLink partition.
- ▶ If the PowerVM NovaLink installer is set to use I/O redundancy, the storage for the PowerVM NovaLink partition is automatically mirrored to accomplish redundancy by using RAID 1.

- If the PowerVM NovaLink software is installed on a HMC-managed system, the HMC is used to create a Linux logical partition (LPAR) and define its associated amounts of resources. When the HMC is used to create the Linux logical partition (LPAR), the `powervm_mgmt_capable` flag must be set to true.

The default OS installed by the PowerVM NovaLink media is Ubuntu 16.04.1 LTS. Red Hat Enterprise Linux version 7.3, or later, is also supported.

NovaLink 2.0.0 is also supported on Linux partitions running Ubuntu 18.04 with latest fixes applied, as well as Red Hat Enterprise Linux 8.2 and Red Hat Enterprise Linux 8.3 (LE) (valid as of the date when this document was created). The installer provides an option to install Red Hat Enterprise Linux once the required setup and configuration steps are complete.

Note: Software-defined infrastructure (SDI) is no longer supported with NovaLink 2.0.0 valid as of the date when this document was created.

Supported operating systems for hosted logical partitions

PowerVM NovaLink supports all operating systems that are supported on the machine type and model of the managed system.

Virtual I/O Server partition requirements

Table 2-11 shows the PowerVC, PowerVM NovaLink and Virtual I/O Server minimum required versions.

Table 2-11 Supported virtualization platforms

PowerVC	NovaLink	Platform	VIOS
1.4.2.0	1.0.0.12	POWER8 and POWER9	Version 3.1.0.10 Version 2.2.6.23 Version 2.2.5.10 ^a
1.4.2.1	1.0.0.13		
1.4.3.0	1.0.0.14	POWER8 and POWER9	Version 2.2.6.41 with mandatory iFix IJ16853 applied Version 3.1.0.21 with mandatory iFix IJ16854 applied
1.4.3.1	1.0.0.15		
1.4.4.0	1.0.0.16	POWER8 and POWER9	Version 2.2.6.51 Version 3.1.0.21 with mandatory iFix IJ16854 applied Version 3.1.0.30 Version 3.1.1.10 Version 3.1.1.25 Version 3.1.2.0
1.4.4.1	1.0.0.16		
1.4.4.2	1.0.0.16		
2.0.0.0	2.0.0.0	POWER8 and POWER9	Version 3.1.0.30 Version 3.1.1.25 Version 3.1.2.10

a. PowerVM Virtual I/O Server compatible only with POWER8 systems

Before installing PowerVM NovaLink, review the following considerations to determine if any of any are applicable for your environment:

- ▶ If you install the PowerVM NovaLink environment on a new managed system, configure one disk with at least 60 GB of storage for each Virtual I/O Server instance that you plan to create on the system. You can configure the system's internal serial-attached Small Computer System Interface (SAS) disk units, alternatively disk units residing within SAN (Storage Area Network) can be used as well.
- ▶ If you create two instances of Virtual I/O Server, create each disk on a separate SAS controller or Fibre Channel (FC) card to accomplish redundancy. Otherwise, the resource requirements for Virtual I/O Servers that are installed by the PowerVM NovaLink installer are the same as the resource requirements for Virtual I/O Servers that are not installed by the PowerVM NovaLink.

Reliable Scalable Cluster Technology (RSCT) for Resource Monitoring Control (RMC) connections

To enable IPv6 link-local address support for Resource Monitoring Control (RMC) connections, update the Reliable Scalable Cluster Technology (RSCT) packages on AIX and Linux logical partitions to be at version 3.2.1.0 or later.

IBM PowerVC requirements

IBM PowerVC Version 1.3 (or later) is required to manage a PowerVM NovaLink host. IBM PowerVC Version 2.0.0 is required to manage a host that has PowerVM NovaLink version 2.0.0. Support for Software Defined Networking (SDN) is discontinued with NovaLink 2.0.0. Also, make sure that your NovaLink version 2.0.0 native Linux host meets the following levels:

- ▶ Ubuntu 18.04 with latest fixes applied (recommended)
- ▶ Red Hat Enterprise Linux 8.2 and Red Hat Enterprise Linux 8.3 (LE)

Hardware Management Console requirements

HMC Version 8.4.0 service pack 1, or later, is required to co-manage a system with PowerVM NovaLink.

2.3 Host and partition management planning

When planning your IBM Power Virtualization Center (PowerVC) environment, consider your host and I/O management strategy, number of hosts' limitations and VMs that can be managed by IBM PowerVC, and the benefits from using multiple VIOs.

2.3.1 Physical server configuration

If you plan to use Live Partition Mobility (LPM), you must ensure that all servers are configured with the same logical memory block size. The logical memory block size can be changed from the Advanced System Management Interface (ASMI) as shown in Figure 2-1.

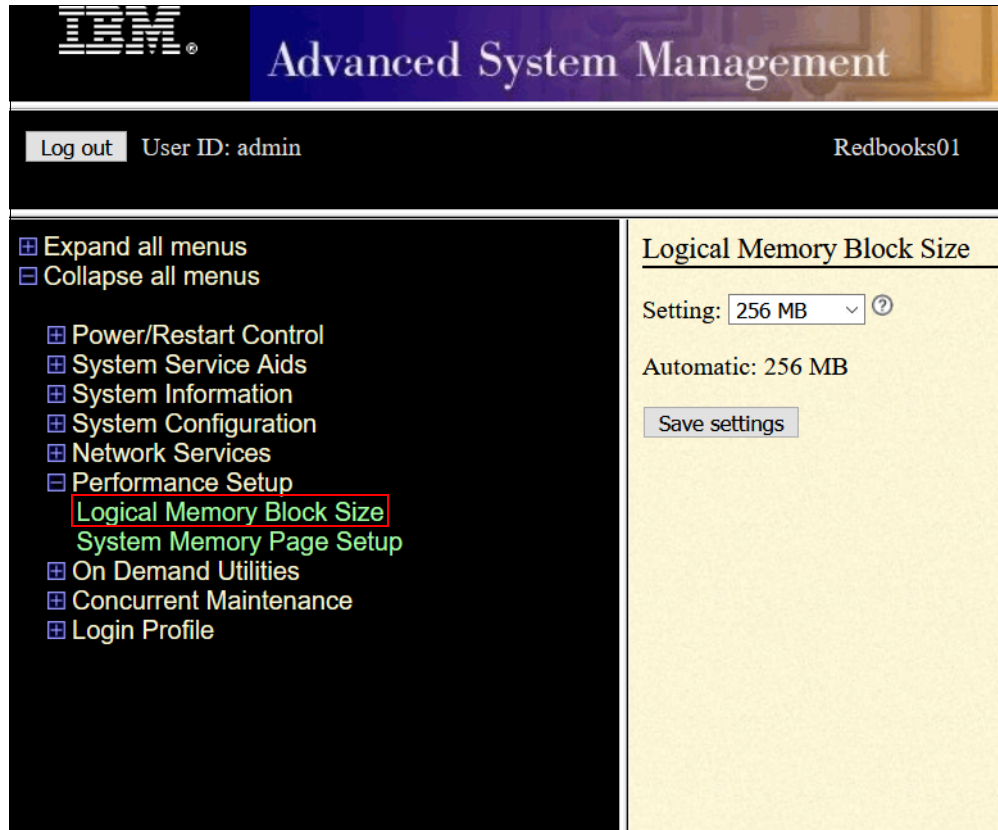


Figure 2-1 Changing the logical memory block size

2.3.2 HMC or PowerVM NovaLink planning

Data centers can contain hundreds of hosts and thousands of VMs. For IBM PowerVC Version 2.0.0, the following maximum numbers are recommended:

- ▶ IBM PowerVC version 2.0.0 managing PowerVM hosts by an HMC:
 - A maximum of 30 HMC-managed hosts are supported.
 - Each host can have up to 1000 virtual machines on it.
 - A maximum of 3000 virtual machines can be on all of the hosts combined.
 - Each HMC can manage up to 2000 virtual machines.
- ▶ IBM PowerVC Version 2.0.0 managing PowerVM hosts by using PowerVM NovaLink:
 - Up to 50 NovaLink-managed hosts are supported.
 - Up to 10000 virtual machines and 20000 volumes can be on all of the NovaLink-managed hosts combined.

- Up to 1000 virtual machines (NovaLink, Virtual I/O Servers, or client workloads) per PowerVM host are supported (This number is valid as of the date when this document was released, the maximum number is also dependent on the PowerVM system firmware versions). For more information on scaling your system, see *IBM PowerVM Best Practices*, SG24-8062.
- ▶ When your IBM PowerVC version 2.0.0 environment consists of both HMC-managed and NovaLink-managed hosts:
 - Up to 50 hosts are supported, 30 of which can be HMC-managed.
 - A maximum of 3000 virtual machines can be on all of the hosts combined.

Notes:

- ▶ No hard limitations exist in IBM PowerVC. These are the suggested values from performance standpoint only.
- ▶ KVM is no longer supported on PowerVC 2.0.0 and NovaLink 2.0.0 (valid as of the date when this document was created).

You must consider how to partition your HMC, and PowerVM NovaLink in subsets, where each is managed by a IBM PowerVC management host.

Advanced system planning and installation typically use redundant HMCs to manage the hosts. Support for redundant HMCs was added with the release of PowerVC Version 1.2.3. Support for redundant HMCs was enhanced with version 1.4.2.0, which allowed the user to manage a host via more than one HMC, where the first registered HMC would be the primary HMC for the host.

If there is another HMC managing the host and it is added to PowerVC, it is set as the secondary HMC. In case the primary HMC fails, PowerVC automatically fails over the host connection to the secondary HMC. Alternatively, in case of a planned HMC outage, you can manually switch the connection to the Secondary HMC, so that the secondary takes over the role of primary HMC. See Figure 2-2 on page 25 for how to add redundant HMC in PowerVC 2.0.0.

Note: It is recommended to create a separate HMC user, dedicated for PowerVC to authenticate to the HMC.

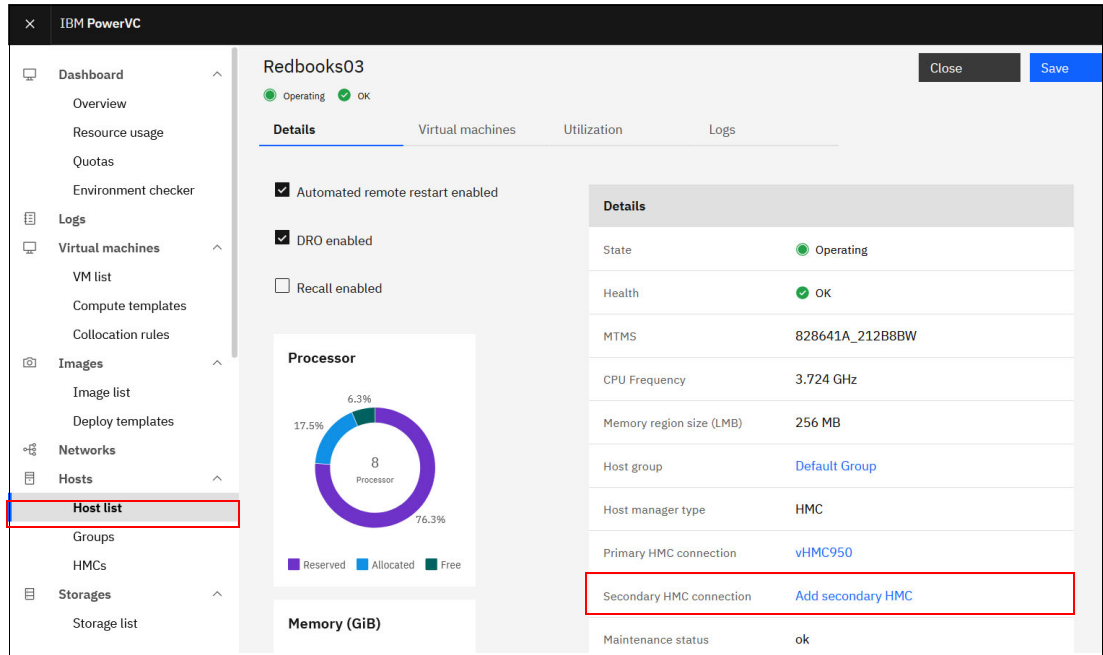


Figure 2-2 Add redundant HMC

1. To add a redundant HMC, select your host from the host list menu and click on **Add secondary HMC** located under Details on Secondary HMC connection field (Figure 2-3).

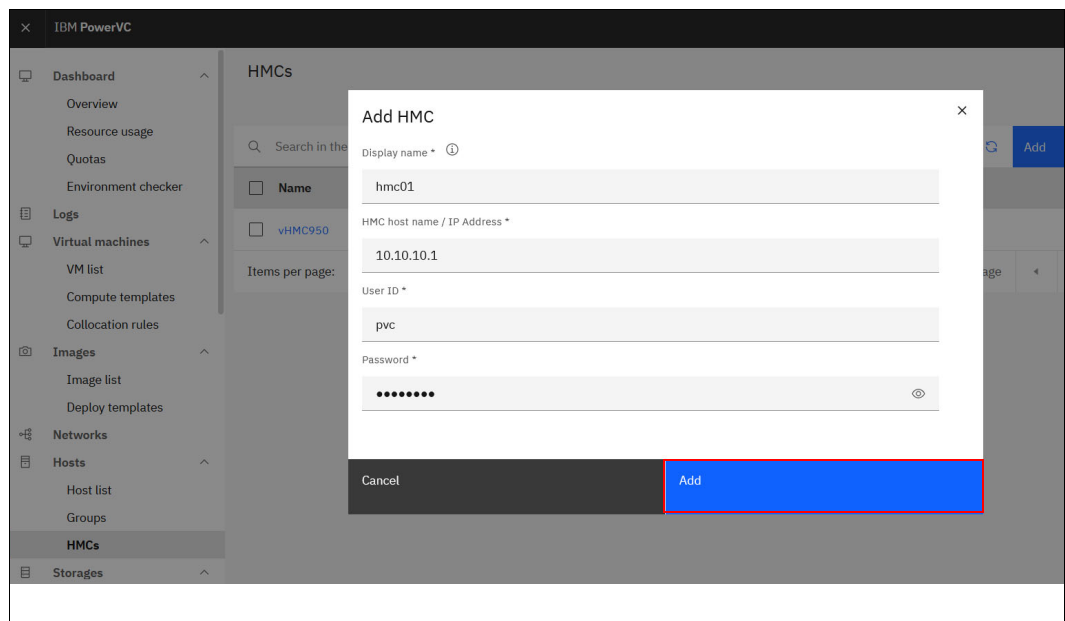


Figure 2-3 Add HMC details

2. Enter a unique name that will be displayed in PowerVC, HMC IP address, user ID, and password for authentication.

Note: PowerVC does not support management of OpenPOWER systems through HMC.

Now the secondary HMC was successfully added as shown in Figure 2-4.

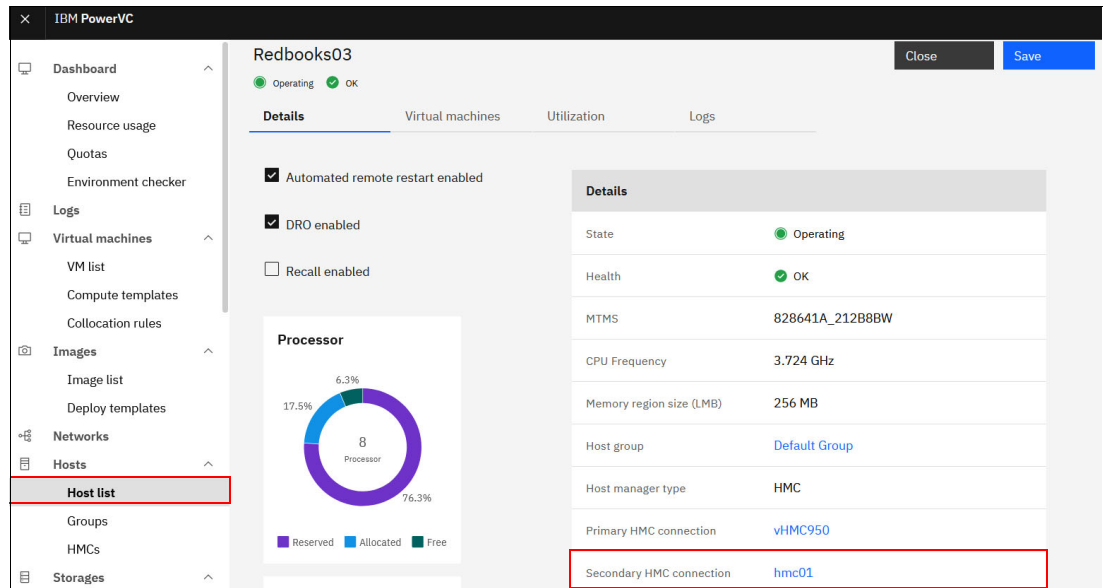


Figure 2-4 Secondary HMC available

Planning for AME

Active Memory Expansion (AME) improves memory usage for virtual machines. It expands the physical memory that is assigned to a virtual machine by allowing the operating system to effectively use more memory than it was assigned. The hypervisor compresses the least used memory pages to achieve the expansion. The increase in available memory is called the expansion factor. To enable AME, use an advanced compute template. See Figure 2-5 on for how to enable Active Memory Expansion (AME) to expand the effective memory capacity of an AIX virtual machine beyond its physical allocation.

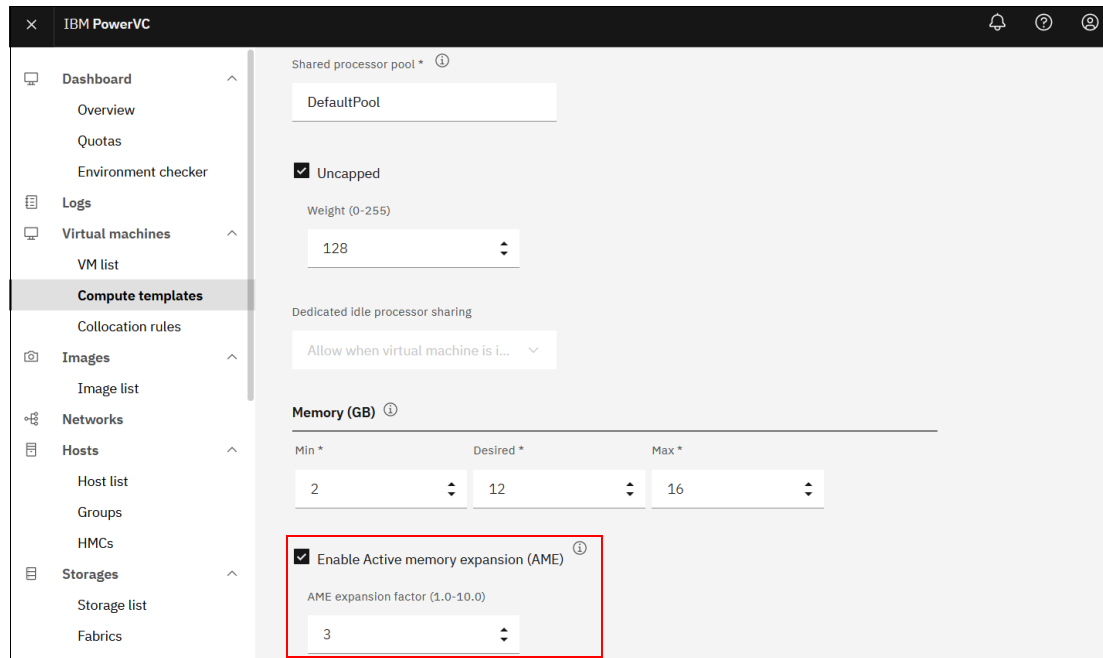


Figure 2-5 Enable Active Memory Expansion (AME)

AME Requirements in PowerVC 2.0.0

Your system must meet the following requirements to use AME with PowerVC:

- ▶ PowerVC Version 2.0.0
- ▶ AIX Version 6.1, or later
- ▶ POWER7, or later

2.3.3 Virtual I/O Server planning

IBM PowerVC supports more than one VIOS server.

Consider a second VIOS to provide redundancy and I/O connectivity resilience to the hosts. Use two VIOSs to avoid outages to the hosts when you must perform maintenance, updates, or changes in the VIOS configuration.

If you plan to make partitions mobile, define the VIOS that provides the mover service on all hosts, and ensure that the **Mover service partition** option is enabled in the profile of these VIOSs. Save configuration changes to profile must be set to Enable on all VIOSs. On the HMC, verify the settings of all VIOSs, as shown in Figure 2-6.

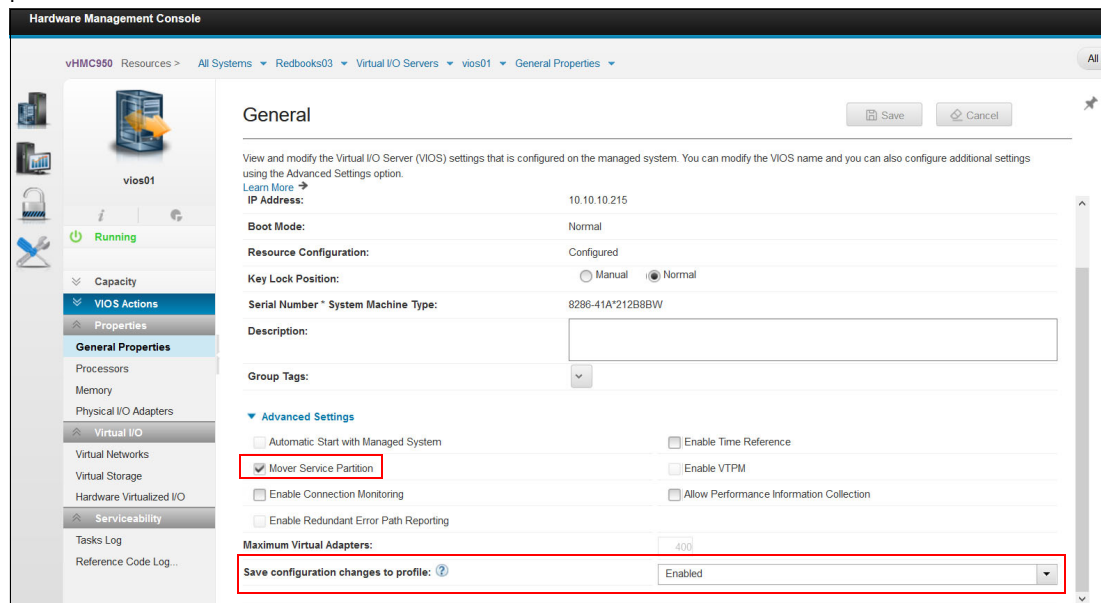


Figure 2-6 VIOS settings that must be managed by IBM PowerVC

Important: Configure the maximum number of virtual resources (virtual adapters) for the VIOS to at least 200. This setting provides sufficient resources on your hosts while you create and migrate VMs throughout your environment. Otherwise, IBM PowerVC indicates a warning during the verification process.

Changing the maximum virtual adapters in a Virtual I/O Server

The preferred way to change the Maximum Virtual Adapters in the HMC running at version 9 is using the Enhanced GUI:

1. Select **Resources** → **All Systems** then click the name of the desired managed server.
2. The partitions view is displayed. Under Power VM click **Virtual I/O Servers**, select the VIOS.

3. Click **Actions** → **View Virtual I/O Server Properties**.
4. Now click **Advanced** (see Figure 2-7), that is displayed on the upper right-hand corner.
5. Check the Maximum Virtual Adapters field. In order to edit that field, the target VIOS must be powered off in advance (Figure 2-7).

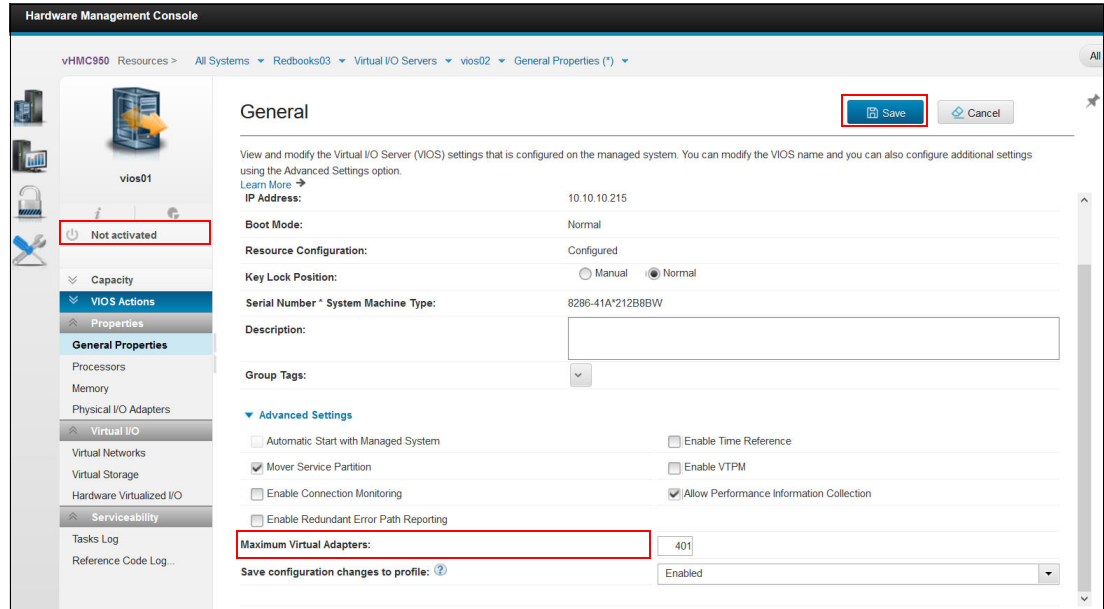


Figure 2-7 How to change the Maximum Virtual Adapters

Note: On certain HMC versions, the Advanced settings tab is located on the upper right-hand side corner and is represented by the Advanced button. See Figure 2-8, for an example of an HMC running version V9R1M942 (MH01876).

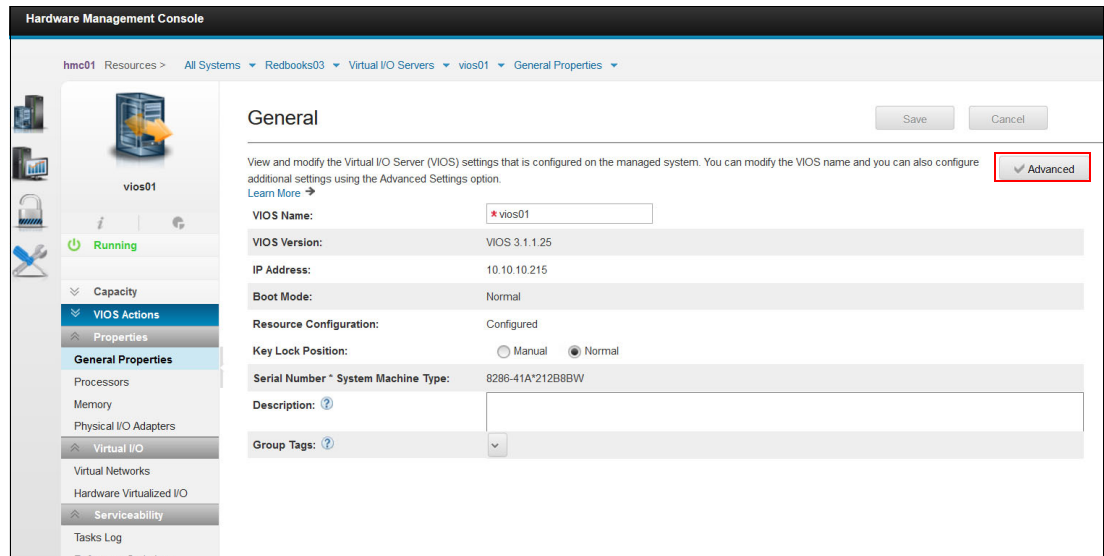


Figure 2-8 HMC Advanced button in HMC Version V9R1M942

Another way to change the Maximum Virtual Adapters in HMC Version 9:

1. Select **Resources** → **All Systems**, then click the name of the desired managed server to open the Partitions view.
2. Under PowerVM click **Virtual I/O Servers**.
3. Now click the name of the desired VIOS.
4. Under VIOS Actions click **Profiles** → **Manage Profiles**, select the desired profile.
5. Then under Actions click **Edit** and switch to the Virtual Adapters tab.
6. Now replace the value in the Maximum virtual adapters field with a new value and click **OK** to save it. See Figure 2-9 for an example.

Logical Partition Profile Properties: default_profile @ vios01

General Processors Memory I/O **Virtual Adapters** Power Controlling Settings

Actions ▼

Virtual resources allow for the sharing of physical hardware between logical partitions. The current virtual adapter settings are listed below.

Maximum virtual adapters : * 400

Number of virtual adapters : 9

Select	Type	Adapter ID	Server/Client Partition	Partner Adapter	Required
<input type="checkbox"/>	Ethernet	2	N/A	N/A	No
<input type="checkbox"/>	Server Fibre Channel	6	powervc00(6)	4	No
<input type="checkbox"/>	Server SCSI	3	vHMC950(3)	3	No
<input type="checkbox"/>	Server SCSI	4	powervc01(4)	3	No
<input type="checkbox"/>	Server SCSI	5	powervc00(6)	2	No
<input type="checkbox"/>	Server SCSI	7	ibmi74-for-ca-1e1441eb-00000015(2)	2	No
<input type="checkbox"/>	Server SCSI	8	ubuntu-16.04--fe8af493-0000001c(8)	2	No
<input type="checkbox"/>	Server Serial	0	Any Partition	Any Partition Slot	Yes
<input type="checkbox"/>	Server Serial	1	Any Partition	Any Partition Slot	Yes

Total: 9 Filtered: 9 Selected: 0

OK Cancel Help

Figure 2-9 Maximum virtual adapters in a VIOS profile

Updating the profile settings will set Save configuration changes to profile to Disable until next activated as shown in Figure 2-10.

Update Profile with Sync profile suspended till next activation/apply - vios01

sync_current_profile is turned on. Attempting to update the profile will set the sync_current_profile capability to 'sync suspended till next activation/apply'
Do you want to continue?

Yes No

Figure 2-10 Save configuration changes to profile disable until next activated

Note: Save the current configuration of the Virtual I/O Server to a new partition profile. To save the current configuration to a new partition profile, you must be a super administrator, service representative, operator, or product engineer.

2.4 Placement policies and templates

One goal of IBM PowerVC is to simplify the management of VMs and storage by providing the automated creation of partitions, virtual storage disks, and the automated placement of partitions on physical hosts. This automation replaces the manual steps that are needed when you use PowerVM directly. In the manual steps, you must create disks, select all parameters that define each partition to deploy, and configure the mapping between the storage units and the partitions in the VIOs.

This automation is performed by using placement policies and various templates.

2.4.1 Host groups

Use host groups to group hosts logically regardless of the features that they might share. For example, the hosts do not need the same architecture, network configuration, or storage. Host groups have these important features:

- ▶ Every host must be in a host group.

Any hosts that do not belong to a user-defined host group are members of the default host group. The default host group cannot be deleted.

- ▶ VMs are kept within the host group.

A VM can be deployed to a specific host or to a host group. After deployment, if that VM is migrated, it must always be migrated within the host group. Also advanced features of PowerVC like the Dynamic Resource Optimizer (DRO), Maintenance mode, or Automated Simplified Remote Restart only operate inside a host group.

- ▶ Placement policies are associated with host groups.

Every host within a host group is subject to the host group's placement policy. The default placement policy is *striping*.

An enterprise client can group its hosts to meet different business needs, for example, for test, development, and production, as shown in Figure 2-11. With different placement policies, even with different hardware, the client can archive at different service levels.

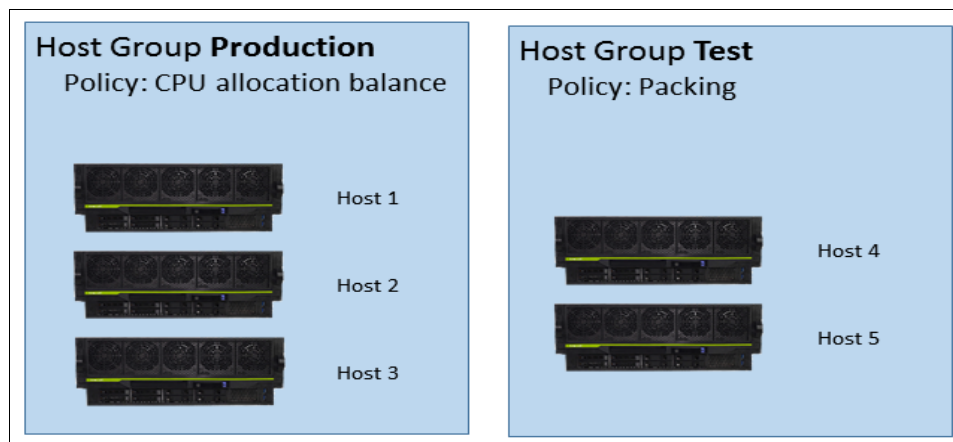


Figure 2-11 Host group sample

2.4.2 Placement policies

If you want to deploy a new partition, you can indicate to IBM PowerVC the host on which you want to create this partition. You can also ask IBM PowerVC to identify the hosts on which the partitions will best fit in a host group, based on a policy that matches your business needs. If you ask IBM PowerVC to identify the hosts on which the partitions will best fit in a host group, it compares the requirements of the partitions with the availability of resources on the possible set of target hosts. IBM PowerVC considers the selected placement policy to make a choice.

IBM PowerVC offers six policies to deploy VMs:

- ▶ **Striping placement policy**

The striping placement policy distributes your VMs evenly across all of your hosts. For each deployment, IBM PowerVC determines the hosts with sufficient processing units and memory to meet the requirements of the VM. Other factors for determining eligible hosts include the storage and network connectivity that are required by the VM. From the group of eligible hosts, IBM PowerVC chooses the host that contains the fewest number of VMs and places the VM on that host.

- ▶ **Packing placement policy**

The packing placement policy places VMs on a single host until its resources are fully used, and then it moves on to the next host. For each deployment, IBM PowerVC determines the hosts with sufficient processing units and memory to meet the requirements of the VM. Other factors for determining eligible hosts include the storage and network connectivity that are required by the VM. From the group of eligible hosts, IBM PowerVC chooses the host that contains the most VMs and places the VM on that host. After the resources on this host are fully used, IBM PowerVC moves on to the next eligible host that contains the most VMs.

This policy can be useful when you deploy large partitions on small servers. For example, you must deploy four partitions that require eight, nine, and seven cores on two servers, each with 16 cores. If you use the striping policy, the first two partitions are deployed on the two servers, which leaves only eight free cores on each. IBM PowerVC cannot deploy the 9-core partition because an LPM operation must be performed before the 9-core partition can be deployed.

By using the packing policy, the first two 8-core partitions are deployed on the first hosts, and IBM PowerVC can then deploy the 9-core and 7-core partitions on the second host. This example is simplistic, but it illustrates the difference between the two policies: The striping policy optimizes performance, and the packing policy optimizes human operations.

- ▶ **CPU utilization balance placement policy**

This placement policy places VMs on the host with the lowest CPU utilization in the host group. The CPU utilization is computed as a running average over the last 15 minutes.

- ▶ **CPU allocation balance placement policy**

This placement policy places VMs on the host with the lowest percentage of its CPU that is allocated post-deployment or after relocation.

For example, consider an environment with two hosts:

- Host 1 has 16 total processors, four of which are assigned to VMs.
- Host 2 has four total processors, two of which are assigned to VMs.

Assume that the user deploys a VM that requires one processor. Host 1 has $(4+1)/16$, or $5/16$ of its processors that are allocated. Host 2 has $(2+1)/4$, or $3/4$ of its processors that are allocated. Therefore, the VM is scheduled to Host 1.

► Memory utilization balanced

This placement policy places virtual machines on the host that has the lowest memory utilization in the host group. The memory utilization is computed as a running average over the last 15 minutes.

- HMC managed hosts do not accurately report their memory utilization (only NovaLink managed systems). Therefore, host groups that use this policy should not contain HMC managed hosts. If there are any HMC managed hosts in the host group, PowerVC always targets the HMC hosts for placement because their utilization is recorded as 0.
- All virtual machines on a PowerVC host should have RMC running for the most accurate memory utilization estimates.

► Memory allocation balanced placement policy

This placement policy places VMs on the host with the lowest percentage of its memory that is allocated post-deployment or after relocation.

For example, consider an environment with two hosts:

- Host 1 has 24 GB total memory, 11 GB of which are assigned to VMs.
- Host 2 has 8 GB total memory, 2 GB of which are assigned to VMs.

Assume that the user deploys a VM that requires 1 GB of total memory. Host 1 has $(11+1)/24$, or $1/2$ of its memory that is allocated. Host 2 has $(2+1)/8$, or $3/8$ of its memory that is allocated. Therefore, the VM is scheduled to Host 2.

Note: A default placement policy change does not affect existing VMs. It affects only new VMs that are deployed after the policy setting is changed. Therefore, changing the placement policy for an existing environment does not result in moving existing partitions.

Tip: The following settings might increase the throughput and decrease the duration of deployments:

- Use the striping policy rather than the packing policy.
- Limit the number of concurrent deployments to match the number of hosts.

When a new host is added to the host group that is managed by IBM PowerVC, if the placement policy is set to the striping mode, new VMs are deployed on the new host until it catches up with the existing hosts. IBM PowerVC allocates partitions only on this new host until the resources use of this host is about the same as on the previously installed hosts.

When a new partition is deployed, the placement algorithm uses several criteria to select the target server for the deployment, such as availability of resources and access to the storage that is needed by the new partitions. By design, the IBM PowerVC placement policy is deterministic. Therefore, the considered resources are the amounts of processing power and memory that are needed by the partition, as defined in the partition profile (virtual processors, entitlement, and memory). Dynamic resources, such as I/O bandwidth, are not considered, because they result in a non-deterministic placement algorithm.

Note: The placement policies are predefined. You cannot create your own policies.

The placement policy can also be used when you migrate a VM. Figure 2-12 shows the IBM PowerVC user interface for migrating a partition. Use this interface to select between specifying a specific target or letting IBM PowerVC select a target according to the current placement policy.

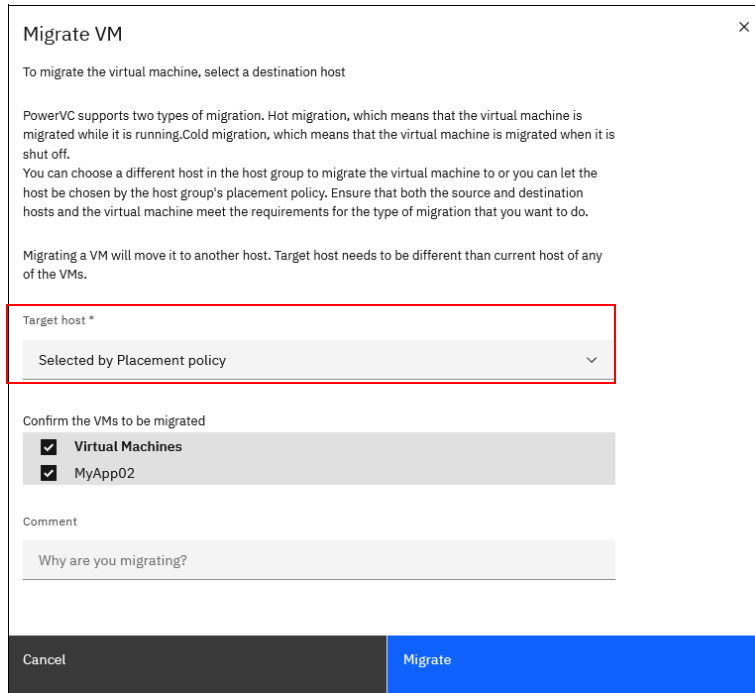


Figure 2-12 Migration of a partition by using a placement policy

2.4.3 Template types

Rather than define all characteristics for each partition or each storage unit that must be created, the usual way to create them in IBM PowerVC is to instantiate these objects from a template that was previously defined. The amount of effort that is needed to define a template is similar to the effort that is needed to define a partition or storage unit. Therefore, reusing templates saves significant effort for the system administrator, who must deploy many objects.

IBM PowerVC provides a GUI to help you create or customize templates. Templates can be easily defined to accommodate your business needs and your IT environment.

Three types of templates are available:

Compute templates These templates are used to define processing units, memory, and disk space that are needed by a partition. Compute templates also define additional details of a VM, such as remote restartable, Compatibility mode or Secure boot. They are described in 2.4.4, “Information that is required for compute template planning” on page 34.

Deploy templates These templates are used in PowerVC for Private Cloud to allow authorized self-service users to quickly, easily, and reliably deploy an image. They are described in 2.4.5, “Tips for deploy template planning” on page 40.

Storage templates These templates are used to define storage settings, such as a specific volume type, storage pool, and storage provider. They are described in 2.6.2, “Storage templates” on page 49.

Use the templates to deploy new VMs. This approach propagates the values for all of the resources into the VMs. The templates accelerate the deployment process and create a baseline for standardization.

2.4.4 Information that is required for compute template planning

The IBM PowerVC management host provides six predefined compute templates (sizes from tiny to xxlarge). Your redefined templates can be edited and removed. You can also create your own templates.

Before you create templates, plan for the amount of resources that you need for the different classes of partitions. For example, different templates can be used for partitions that are used for development, test, and production, or you can have different templates for database servers, application servers, and web servers.

The following information about the attributes of a compute template helps your planning efforts regarding compute templates:

- ▶ **Template name**

The name to use for the template.

- ▶ **Virtual processors**

The number of virtual processors. A VM usually performs best if the number of virtual processors is close to the number of processing units that is available to the VM. You can specify the following values:

Minimum The smallest number of virtual processors that you accept for deploying a VM.

Desired The number of virtual processors that you want for deploying a VM.

Maximum The largest number of virtual processors that you allow when you resize a VM. This value is the upper limit to resize a VM dynamically. When it is reached, you need to power off the VM, edit the profile, change the maximum to a new value, and restart the VM.

- ▶ **Use shared processors**

If checked, a VM with shared processors will be deployed. If not checked, the VM gets dedicated processor that belong to just that LPAR.

The following attributes are only visible if Use shared processor is checked, meaning the template creates a shared processor VM:

- ▶ **Processing units**

Number of entitled processing units. A processing unit is the minimum amount of processing resource that the VM can use. For example, a value of 1 (one) processing unit corresponds to 100% use of a single physical processor.

Processing units are split between virtual processors, so a VM with two virtual processors and one processing unit appears to the VM user as a system with two processors, each running at 50% speed.

You can specify the following values:

- Minimum** The smallest number of processing units that you accept for deploying a VM. If this value is not available, the deployment does not occur.
- Desired** The number of processing units that you want for deploying a VM. The deployment occurs with a number of processing units that is less than or equal to the wanted value and greater than or equal to the minimum value.
- Maximum** The largest number of processing units that you allow when you resize a VM. This value is the upper limit to which you can resize dynamically. When it is reached, you must power off the VM, edit the profile, change the maximum value to a new value, and restart the VM.

► Shared processor pool

PowerVC supports multiple shared processor pools. This allows you to share a group of processors between multiple virtual machines. You can group your applications together and set a shared processor pool size limit on the total number of processing units for each pool, which limits the software license exposure for that pool.

The default is DefaultPool.

► Uncapped

If checked, this template creates uncapped VMs that can use processing units that are not being used by other VMs, up to the number of virtual processors that is assigned to the uncapped VM.

If not checked, this template creates capped VMs can use only the number of processing units that are assigned to them.

► Weight (0 - 255)

The Weight attribute is only available for uncapped VMs.

If multiple uncapped VMs require unused processing units, the uncapped weights of the uncapped VMs determine the ratio of unused processing units that are assigned to each VM. For example, an uncapped VM with an uncapped weight of 200 receives two processing units for every processing unit that is received by an uncapped VM with an uncapped weight of 100.

Important: Processing units and virtual processor are values that work closely and must be calculated carefully. For more information about virtual processor and processing units, see *IBM PowerVM Virtualization Managing and Monitoring*, SG24-7590.

The following attribute is only visible if Use shared processor is unchecked, meaning the template creates a dedicated processor VM:

► Idle sharing

This setting enables this VM to share its dedicated processors with other VMs running in shared processor (also known as a *dedicated donating* partition).

The possible selections are:

- Allow when virtual machine is inactive
- Allow when virtual machine is active
- Allow always
- Never

The rest of the attributes are not related to shared or dedicated processor usage.

► **Memory (GB)**

Amount of memory, expressed in GB. The value for memory must be a multiple of the memory region size that is configured on your host. The minimum value is 16 MB. To see the region size for your host, open the Properties window for the selected host on the HMC, and then open the Memory tab to view the memory region size. Figure 2-13 shows an example.

You can specify the following values:

- Minimum** The smallest amount of memory that you want for deploying a VM. If the value is not available, the deployment does not occur.
- Desired** The total memory that you want in the VM. The deployment occurs with an amount of memory less than or equal to the wanted amount and greater than or equal to the minimum amount that is specified.
- Maximum** The largest amount of memory that you allow when you resize a VM. This value is the upper limit to resize a VM dynamically. When it is reached, you must power off the VM, edit the profile, change the maximum to a new value, and restart the VM.

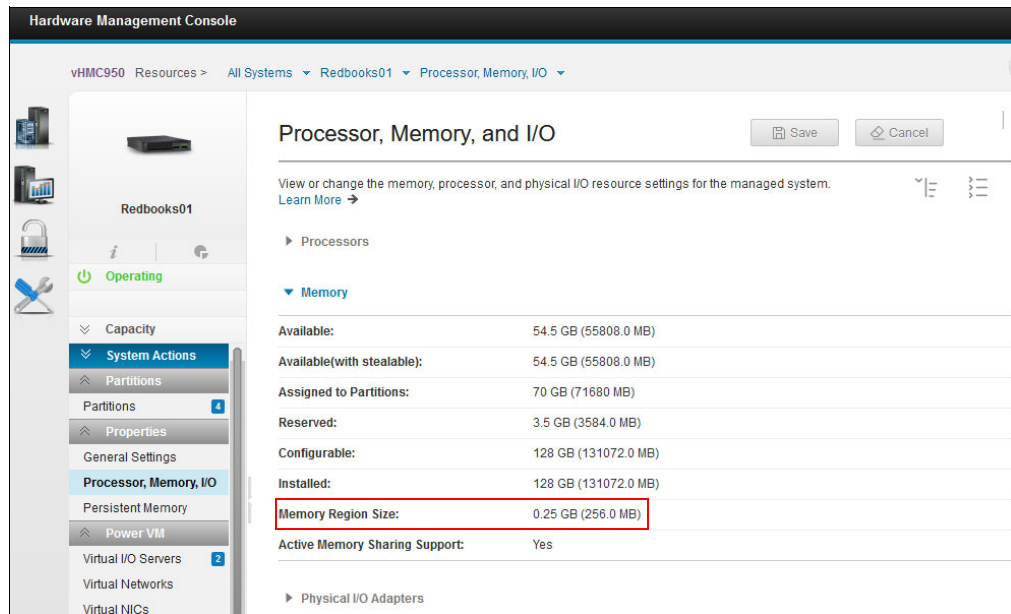


Figure 2-13 Memory region size view on the HMC

► **Enable Active memory expansion (AME)**

AME is an AIX only feature.

Active Memory Expansion (AME) improves memory usage for virtual machines. It expands the physical memory that is assigned to a virtual machine by allowing the operating system to effectively use more memory than it was assigned. The hypervisor compresses the least used memory pages to achieve the expansion. The increase in available memory is called the expansion factor.

- ▶ AME expansion factor

The expansion factor can only be set, when AME is enabled / checked in the template. A factor of 1.5 means that 50% memory expansion will be provided. You can set the factor from 1 to 10. A factor of 1.0 mean AME is enabled, but effectively not running. If AME is enabled in a VM, the factor can be changed while the VM is running.

Note: Compressing memory needs processor resources. When setting the value to high and the VM need much more memory than physically available, this can cause a lot of processor overhead.

For more information about AME, consult IBM Knowledge Center:

https://www.ibm.com/support/knowledgecenter/en/ssw_aix_72/performance/intro_ame_process.html

The attributes discussed until now, are on the Details tab of a compute template. No let's discuss the attributes from the Miscellaneous tab.

- ▶ Enable virtual machine remote restart

Users can remote restart a VM on another host easily if the current host fails. This feature enhanced the availability of applications in addition to the solutions that are based on IBM PowerHA and LPM.

Note: This function is based on the PowerVM simplified remote restart function and is supported only by POWER8 an POWER9 servers. For the requirements of remote restart, see IBM Knowledge Center:

https://www.ibm.com/support/knowledgecenter/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_remote_restart_hmc.html

- ▶ Enable performance information collection

Enable the operating system on a partition to collect performance information.

- ▶ Availability priority

To avoid shutting down mission-critical workloads when your server firmware unconfigures a failing processor, set availability priorities for the VMs (0 - 255). A VM with a failing processor can acquire a replacement processor from a VM with a lower availability priority. The acquisition of a replacement processor allows the VM with the higher availability priority to continue running after a processor failure.

► Compatibility mode

Select the processor compatibility that you need for your VM. Table 2-12 describes each compatibility mode and the servers on which the VMs that use each mode can operate.

Table 2-12 Processor compatibility modes

Processor compatibility mode	Description	Supported servers
POWER6	Use the POWER6 processor compatibility mode to run operating system versions that use all of the standard features of the POWER6 processor. POWER6 can still be selected, but POWER6 server are no longer supported by PowerVC.	VMs that use the POWER6 processor compatibility mode can run servers that are based on POWER6, IBM POWER6+, POWER7, or POWER8 processors.
POWER6+	Use the POWER6+ processor compatibility mode to run operating system versions that use all of the standard features of the POWER6+ processor. POWER6+ can still be selected, but POWER6+ server are no longer supported by PowerVC.	VMs that use the POWER6+ processor compatibility mode can run on servers that are based on POWER6+, POWER7, or POWER8 processors.
POWER7, including POWER7+	Use the POWER7 processor compatibility mode to run operating system versions that use all of the standard features of the POWER7 processor.	VMs that use the POWER7 processor compatibility mode can run servers that are based on POWER7, POWER8 or POWER9 processors.
POWER8	Use the POWER8 processor compatibility mode to run operating system versions that use all of the standard features of the POWER8 processor.	VMs that use the POWER8 processor compatibility mode can run servers that are based on POWER8 or POWER9 processors.
POWER9_Base	The operating system version on the logical partition is set to use all the standard features of the POWER9 processor enabled by the firmware at level FW910.	VMs that use POWER9 servers with firmware level 910 or later

Processor compatibility mode	Description	Supported servers
POWER9	<p>In POWER9 mode, all the features of POWER9 are available, as introduced by Firmware version 940.</p> <p>New features with FW 940 are for example:</p> <ul style="list-style-type: none"> ▶ Support for the External Interrupt Virtualization Engine (XIVE) ▶ User Mode NX Acceleration Enablement for PowerVM ▶ Extended support for PowerVM Firmware Secure Boot. <p>For more information visit the following website: https://www.ibm.com/support/pages/processor-compatibility-modes-ibm-power9-based-systems</p>	VMs that use POWER9 servers with firmware level 940 or later
Default	<p>The default processor compatibility mode is a preferred processor compatibility mode that enables the hypervisor to determine the current mode for the VM. When the preferred mode is set to Default, the hypervisor sets the current mode to the most fully featured mode that is supported by the operating environment. In most cases, this mode is the processor type of the server on which the VM is activated. For example, assume that the preferred mode is set to Default and the VM is running on a POWER9 processor-based server. The operating environment supports the POWER9 processor with all capabilities, so the hypervisor sets the current processor compatibility mode to POWER9.</p>	The servers on which VMs with the preferred processor compatibility mode of Default can run depend on the current processor compatibility mode of the VM. For example, if the hypervisor determines that the current mode is POWER9, the VM can run on servers that are based on POWER9 processors.

▶ Secure boot

Secure boot is a PowerVM feature introduced in POWER9 systems. On supported operating systems, the kernel and applications will have their digital signatures verified before being allowed to run. This attribute can be enabled or disabled.

▶ Physical page table ratio

The Physical Page Table (PPT) ratio. The PPT ratio is the ratio of the maximum memory to the size of the Physical Page Table. It controls the size of the page table that is used by the hypervisor when performing live partition migration. The larger the PPT, the more entries are available for use by the partition for mapping virtual addresses to physical real addresses. If the PPT is set too small and the partition is running workloads, performance in the partition can decline during live partition migration. If this ratio is too large, more memory is reserved for the PPT with no performance gain. A higher ratio reserves less memory than a lower ratio. For example, a ratio of 1:4096 reserves less memory for the PPT than a ratio of 1:64 does.

If a value other than default is selected, hosts are filtered during deploys, resizes, and migrations to exclude hosts that do not support specifying the PPT ratio. However, this filtering is not applied during remote restarts.

This setting is particularly useful for SAP HANA workloads. The recommended setting for SAP HANA workloads is 1:1024.

- ▶ Check affinity score check during migrations

This attribute applies to POWER9 hosts only. When this value is true, the affinity score for the partition on the destination host is calculated before migration. If migrating to the proposed target host would result in a lower affinity score, the migration fails.

The recommended value for SAP HANA workloads is true. The default value is false.

- ▶ Use SAP HANA recommended Values

If this attribute is checked, the physical page table ratio and check affinity score can not be selected. In this case the physical page table ratio will be set to 1:1024 and the Check affinity score to true.

- ▶ Persistent memory volumes

Persistent memory is a new concept available with POWER9 servers. Persistent memory will be available to the LPAR as volumes that are persistent when the VM is shutdown. Persistent memory volumes can be defined as part of a compute template, or individually in the deployment of a new VM.

To Add a persistent memory volume, click **Add a persistent memory** to add a new line. Then enter the size and check Affinity if you want the physical used memory near the physical used processors. To delete a persistent memory volume click the trash symbol.

Special considerations:

- Users cannot perform migrate or remote restart operations on the virtual machine with persistent memory volume.
- Virtual machines with persistent volumes are hard pinned to the host.
- Users cannot unpin a virtual machine with persistent memory volume.
- PowerVC allows only DRAM based persistent memory volumes.

To use persistent memory for SAP HANA implementations, have a look into the paper SAP HANA and PowerVM Virtual Persistent Memory under following link:

<https://www.ibm.com/support/pages/system/files/inline-files/vPMEM-SAPHANA-White paper-V1.2.pdf>

2.4.5 Tips for deploy template planning

Use the following information when planning deploy templates

- ▶ Deploy templates are only available in PowerVC for Private Cloud.
- ▶ Administrators can configure image deployment properties and save them as a deploy template. A deploy template includes an images that includes the operating system and may also include data volumes. Additionally a deploy template has the necessary information to create a VM quickly, including the deployment target, storage connectivity group, compute template, and so on. This reduces a lot of resource details from self-server users.

For more information on PowerVC Private Cloud see Chapter 5, “IBM PowerVC for Private Cloud” on page 255.

2.5 IBM PowerVC storage access SAN planning

IBM PowerVC and IBM PowerVC for Private Cloud can manage different storage volumes types, which can be attached to virtual servers. The virtual servers can access their storage by using either of three protocols:

- ▶ Normal vSCSI
- ▶ NPIV
- ▶ vSCSI to SSP

These storage volumes can be integrated devices or pluggable devices. A minimum configuration of the SAN and storage is necessary before IBM PowerVC can use them. Since IBM PowerVC creates virtual disks on storage devices, but these devices must be set up first, you must perform the following actions before you use IBM PowerVC:

- ▶ Configuration of the FC fabric for the IBM PowerVC environment must be planned first including cable attachments, SAN fabrics, and redundancy. It is always recommended to create at least two independent fabrics to provide SAN redundancy.
- ▶ IBM PowerVC provides storage for VMs through the VIOS.

The storage is accessed by using NPIV, vSCSI, or an SSP that uses vSCSI. The VIOS and SSP must be configured before IBM PowerVC can manage them.

Be aware of the following information when working with Virtual Fabrics:

- ▶ PowerVC supports a total of 25 total fabrics. Each Virtual Fabric counts toward this total.
- ▶ The number of Virtual Fabrics that you can create depends on the switch model.
- ▶ The switch must be at the appropriate fabric operating system level.
- ▶ The SAN switch administrator user ID and password must be set up. These IDs are used by IBM PowerVC.
- ▶ The storage controller administrator user ID and password must be set up.

2.5.1 vSCSI storage access

Before using virtual vSCSI attached storage in your environment, ensure that your environment is configured correctly and be aware of the below considerations:

- ▶ The supported multipathing software driver solutions are AIX path control module (PCM) and EMC PowerPath. The number of volumes that can be attached to a virtual machine is not limited by PowerVC. However, each virtual machines operating system limits how many volumes its vSCSI driver supports.
- ▶ PowerVM supports migration only between Virtual I/O Servers that use the same multipathing software driver solution.

Before using vSCSI attached storage, perform the following steps:

1. Turn off SCSI reserves for volumes being discovered on all the Virtual I/O Servers being used for vSCSI connections. This is required for live partition mobility (LPM) operations and for dual VIOSs. For the IBM Storwize family, XIV, IBM System Storage DS8000, and EMC using the AIX PCM model, you must run the following command on every Virtual I/O Server where vSCSI operations will be run:

```
$ chdef -a reserve_policy=no_reserve -c disk -s fcp -t mpioosdisk
```

2. If using DS8000, you must run the following command on every Virtual I/O Server where vSCSI operations will be run:


```
$ chdef -a reserve_policy=no_reserve -c disk -s fcp -t aixmpiods8k
```
3. If using Hitachi ODM attributes in the Virtual I/O Server, you must run the following command on every Virtual I/O Server where vSCSI operations will be run:


```
$ chdef -a reserve_policy=no_reserve -c disk -s fcp -t htcvspmpio
```
4. Zoning between the Virtual I/O Server and the storage device ports must be configured to allow vSCSI environments to be imported easily and also to allow you to use many fabrics with vSCSI.

Figure 2-14 shows how VMs in IBM PowerVC access storage by using vSCSI.

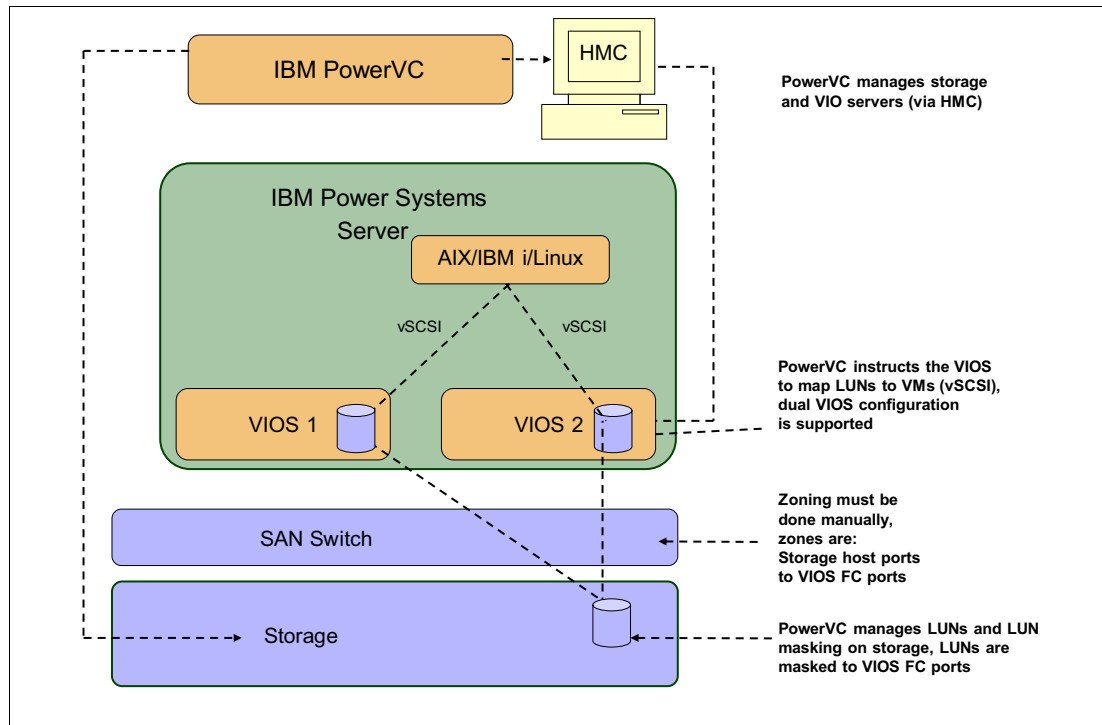


Figure 2-14 IBM PowerVC storage access by using vSCSI

vSCSI best practices

It is important to understand the following information when working with virtual SCSI (vSCSI) attached storage for the best practices purposes:

- ▶ After a volume is attached to a virtual machine, log onto that virtual machine and discover the disk by running the appropriate configuration manager command per the operating system type:
 - For IBM AIX operating system, run the following command as root:


```
# cfgmgr
```
 - For On Linux, run the following command as root:


```
# ls /sys/class/scsi_host/ | while read host ; do echo "- - -" > /sys/class/scsi_host/$host/scan ; done
```
- ▶ Before you detach a volume from a virtual machine, that virtual machine should stop using the disk device for the volume. If a file system is mounted on the disk device on the virtual machine, that file system should first be unmounted and the disk device closed.

- ▶ After you detach a volume from the virtual machine, a disk device is left on the virtual machine. This device can be removed by running the appropriate command:
 - For IBM AIX operating system, run the following command as root:


```
# rmdev -Rdl hdiskX
```
 - For On Linux, run the following command as root:


```
# echo 1 > /sys/block/<device-name>/device/delete
```
- ▶ It is always recommended to turn off disconnected ports. Using disconnected Fibre Channel ports for vSCSI operations can cause volume attach or detach operations to take longer because the Virtual I/O Servers will try to discover the attached volume on the disconnected ports. To turn off ports, follow these steps:
 - a. From the PowerVC user interface, select **Storages** → **FC ports**.
 - b. For each port you want to turn off, select **None** for Connectivity.
 - c. After making all necessary changes, click **Save**.
- ▶ The default setting for Virtual I/O Server pathing is failover, which the Virtual I/O Server uses a single path until it fails. So it is recommended to tell the Virtual I/O Server to use all available paths to send I/O requests by running the command:


```
$ chdef -a algorithm=round_robin -c PCM -s friend -t fcpoother
```
- ▶ For fast failure detection in the Fibre Channel fabric, run the following command on each Virtual I/O Server for each Fibre Channel adapter:


```
$ chdev -l fscsiX -a fc_err_recov=fast_fail
```
- ▶ To detect when Fibre Channel cables are moved, the following command should be run on each Virtual I/O Server for each Fibre Channel adapter:


```
$ chdev -l fscsiX -a dyntrk=yes
```

2.5.2 NPIV storage access

PowerVC performs automatic actions to manage the flow of storage from physical storage LUNs to VMs. The following are list of the actions:

- ▶ Access to the SAN from VMs is configured on VIOSs by using an FC adapter pair and NPIV by running the **vfcmmap** command to maps the virtual fiber adapters to physical adapters.
- ▶ LUNs are provisioned on a supported storage controller. The following are storage providers that are supported to work with PowerVC:
 - EMC VNX and EMC PowerMax (VMAX)
 - Enterprise Hitachi Block Storage Driver (HBSD)
 - Hitachi Global-Active Device (GAD)
 - IBM System Storage DS8000
 - IBM Storwize family
 - IBM XIV
 - Pure Storage

See the following link for more details:

https://www.ibm.com/support/knowledgecenter/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_storage_providers.html

- ▶ LUNs are masked to VM virtual FC ports.

- ▶ SAN zoning is adjusted so VMs have access from their virtual FC ports to the storage controller host ports. Changes in zoning are performed automatically by IBM PowerVC.
- ▶ LUNs are viewed as logical devices in VMs.

Figure 2-15 shows how VMs in IBM PowerVC access storage by using NPIV.

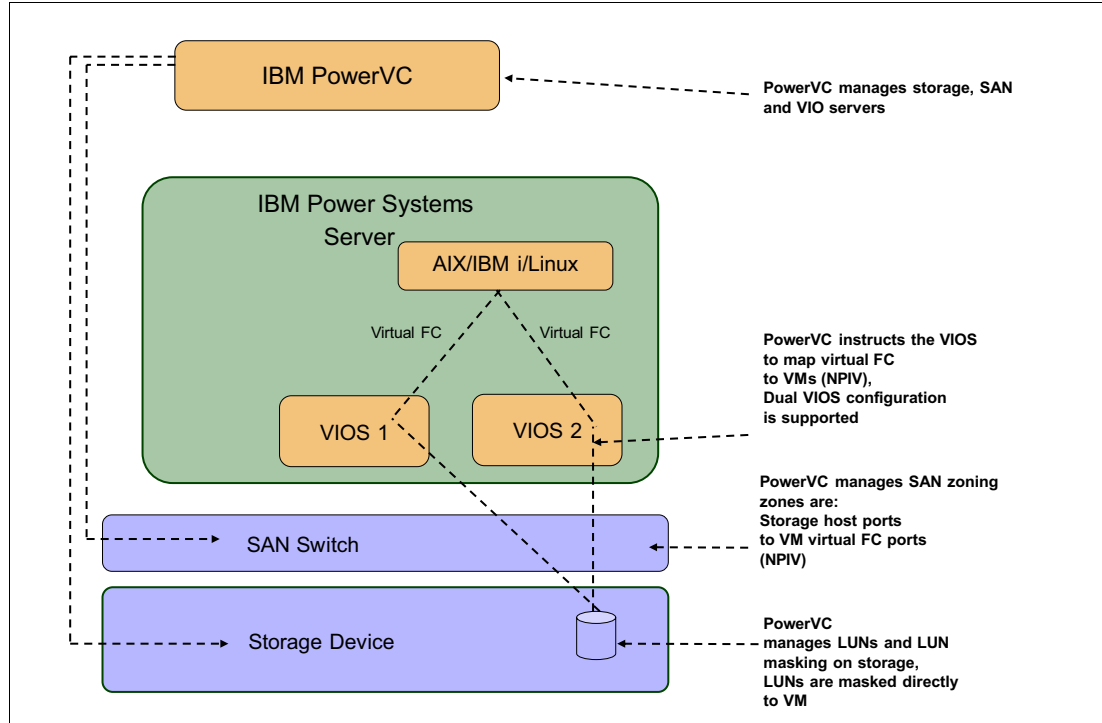


Figure 2-15 IBM PowerVC storage access by using NPIV

2.5.3 Shared storage pools

Shared storage pools allow a set of SAN volumes on one or more supported storage controllers to be managed as a clustered storage device from any Virtual I/O Server that is in the cluster. Those Virtual I/O Servers share access to aggregated logical volumes and present this aggregated storage space as a single pool of storage, optionally divided into separate tiers for Quality of Service (QoS) needs. PowerVC can manage and allocate storage volumes from this pool as it does from Fibre Channel SAN providers.

A cluster consists of up to 24 Virtual I/O Servers with a shared storage pool that provides distributed storage access to the Virtual I/O Servers in the cluster.

Below is the flow of storage management from physical storage LUNs to VMs in IBM PowerVC:

- ▶ The SSP is configured manually by creating the SSP cluster, inclusion of VIOSs in the cluster, and additions of disk to the pool.

See the following two links for more details:

<https://www.ibm.com/support/pages/creating-simple-ssp-among-two-vio-servers>

https://www.ibm.com/support/knowledgecenter/9119-MME/p8eew/p8eew_shared_storage.htm

- ▶ IBM PowerVC discovers the SSP when it discovers the VIOSs.
- ▶ IBM PowerVC can create additional logical units (LUs) in the SSP when it creates a VM.

- ▶ IBM PowerVC instructs the VIOS to map the SSP LUs for the VIO clients partitions that access them through vSCSI devices. They can be shown by several SSP commands and from VIOS `lsmmap` command as well.
- ▶ SSPs are supported on hosts that are managed either by HMC or NovaLink.

Figure 2-16 shows how VMs in IBM PowerVC access SSP storage using vSCSI.

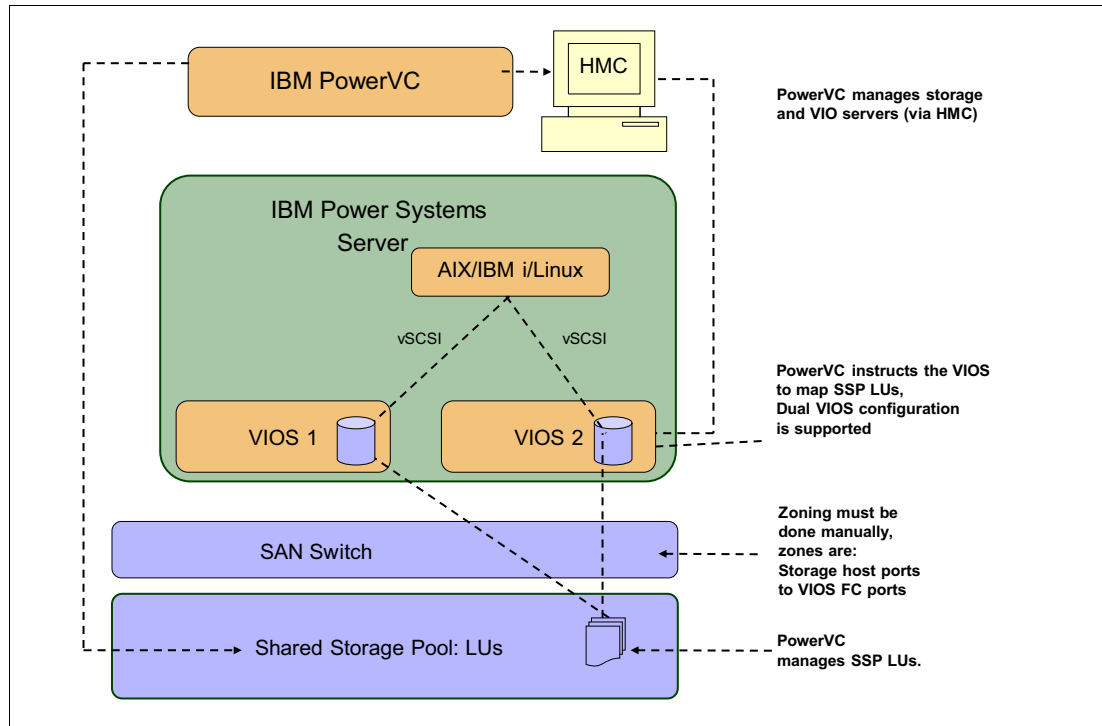


Figure 2-16 IBM PowerVC storage access by using an SSP

2.6 Storage management planning

IBM PowerVC offers a platform for aspects of the enterprise infrastructure to be managed. Supported storage sub systems can be added and managed by IBM PowerVC. Functions such as creating, attaching, detaching, sharing, deleting and cloning of volumes can be performed through IBM PowerVC. IBM PowerVC requires IP connectivity to the storage provider or a REST API server that is connected to the storage system in order to manage the storage. IBM PowerVC uses the OpenStack cinder block storage service to interact with storage devices.

2.6.1 IBM PowerVC storage terminology

The following sections will explain the storage terminology in PowerVC.

Storage Provider

A system that provides volumes in IBM PowerVC is termed a storage provider. Currently, the supported storage providers in IBM PowerVC 2.0 include the following:

- ▶ EMC VNX and EMC PowerMax (VMAX)
- ▶ Enterprise Hitachi Block Storage Driver (HBSD)

- ▶ Hitachi Global-Active Device (GAD)
- ▶ IBM System Storage DS8000
- ▶ IBM FlashSystem family
- ▶ IBM XIV
- ▶ Pure Storage

Figure 2-17 shows the managed storage provider and how its access by clicking on **Storage list** under Storages in the drop-down list.

<input type="checkbox"/>	Name	Hostname or IP Address	Available space (GiB)	Total space (GiB)	State / Health	Volumes	Type
<input type="checkbox"/>	V7000	9.114.181.189	14045.00	21216.00	Running OK	8	IBM Storwize

Items per page: 1 | 1-1 of 1 items | 1 of 1 page

Figure 2-17 List of managed storages

Fabric

The Storage area network (SAN) switches are the link between the servers and the SAN storage providers named *fabrics*. Fabrics must be initially configured before being managed by IBM PowerVC. Figure 2-18 shows the fabrics list page and the managed fabrics.

<input type="checkbox"/>	Name	Hostname or IP Address	Designation	Fabric type
<input type="checkbox"/>	Fabric-A	9.114.181.33	A	Brocade
<input type="checkbox"/>	Fabric-B	9.114.181.34	B	Brocade

Items per page: 2 | 1-2 of 2 items | 1 of 1 page

Figure 2-18 List of managed fabrics

Service node

This is the IBM PowerVC management server showing an instance of the Cinder backup service running. Service nodes are used when exporting and importing images to and from other service nodes. The service node of IBM PowerVC is shown in Figure 2-19.

Name	IP Address	State	Health
Local management node	9.47.76.117	Running	OK

Items per page: 1 | 1-1 of 1 items | 1 of 1 page

Figure 2-19 IBM PowerVC Service node

Volumes

These are disks or LUNs that can either be boot or data volumes. These volumes are created from storage pools of storage providers.

Storage templates

Storage templates define the properties of a volume created within IBM PowerVC. These properties include the definition of the type of volume, whether it is a thin, generic or compressed volume. The template also defines properties such as the storage pool the volume is created from, and the port groups to use when zoning the volume to a virtual machine managed by IBM PowerVC. Figure 2-20 shows the properties of a storage template.

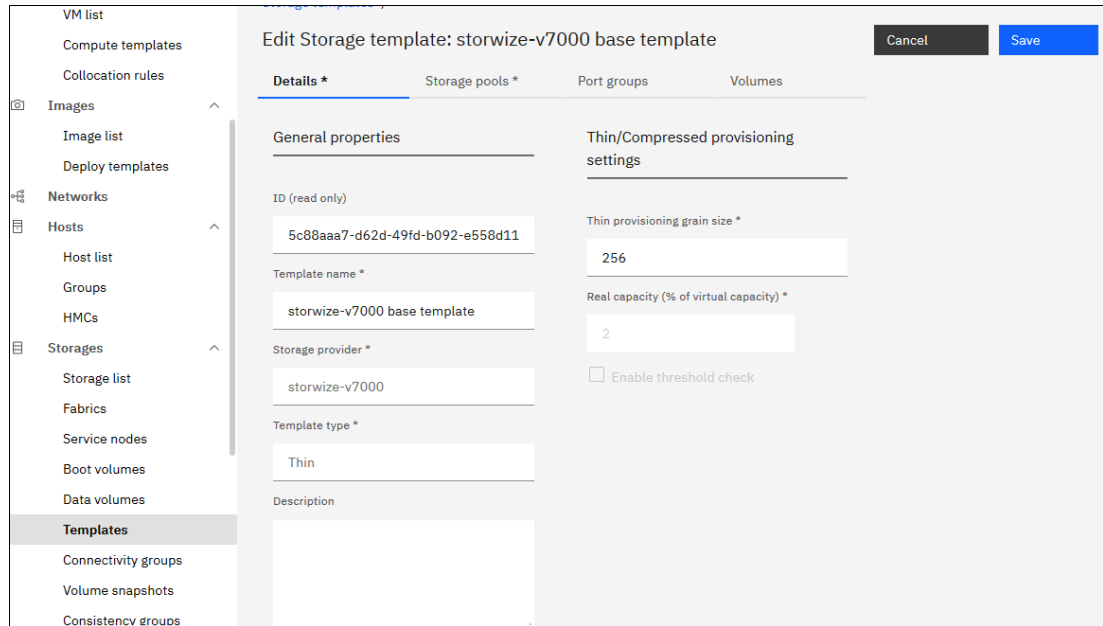


Figure 2-20 Storage template of an IBM FlashSystem

Storage connectivity groups

Storage connectivity groups is a concept that is local to IBM PowerVC. They are a logical grouping of resources that connect to the storage. Using storage connectivity groups, redundancy rules at the Virtual I/O Server layer and optionally at the fabric layer are defined for connections to the storage provider. Figure 2-21 shows created storage connectivity groups.

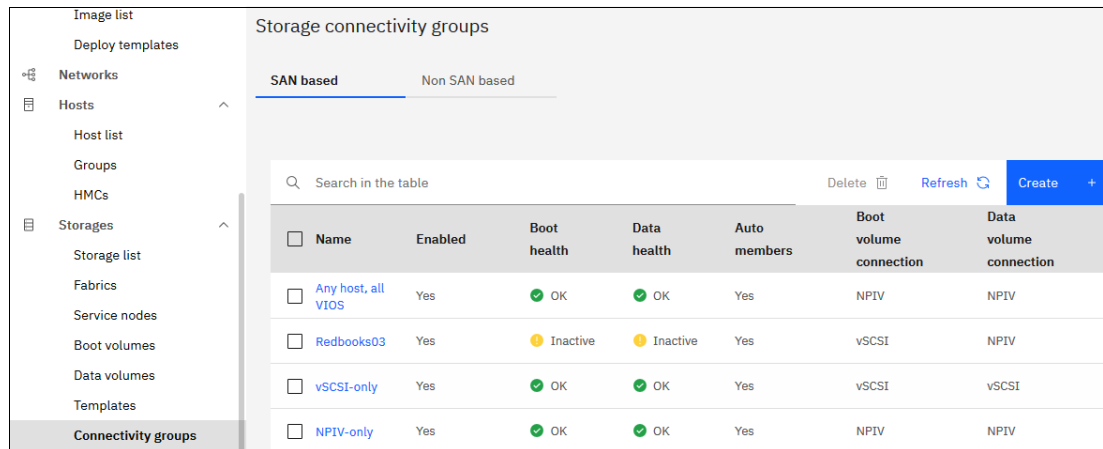


Figure 2-21 Storage connectivity groups created in IBM PowerVC

Fibre Channel port tags

Fibre Channel port tags are strings placed on Fibre Channel ports from a host system. Storage connectivity groups can be configured to only connect through tagged Fibre Channel ports. Fibre Channel ports are mainly tagged for two reasons, separate workload and create redundancy on the Fibre Channel adapter, Virtual I/O servers and the fabric layer.

Volume snapshots

Volume snapshots provide the capability of a non disruptive backup of disks that will be taken at a point in time. These volumes can be used to boot up virtual machines at a point in time.

Consistency groups

Consistency groups are groups of volumes with consistent data which are used to create a point in time snapshot or used to create a consistent copy of volumes. IBM PowerVC offers three types of consistency groups.

1. Consistent Group Snapshot Enabled
2. Consistent Group Replication Enabled
3. Replication Enabled VMRM

2.6.2 Storage templates

When a storage provider is managed by IBM PowerVC, a default storage template is created. This template can be modified or more templates can be created to suit a purpose. However, every storage provider must have one default template which can either be the one created by IBM PowerVC after managing the storage provider or a user defined template. A storage templates speeds up the process of creating volumes by utilizing predefined properties. If a storage template is not used by any volume, you can delete and modify it.

However, if the storage template is in use by existing volumes you cannot delete the template. The name, whether it is the default template for that storage provider and the port groups are the properties that can be altered. If you want to change a storage template that is already being used, you can make a copy and update the new version. Preexisting storage volumes managed by IBM PowerVC do not have an associated storage template. Although after the volumes are managed, they can be assigned a storage template using the set storage template function. In the succeeding section is an example of a storage template for the IBM FlashSystem.

Storage template definition

Figure 2-22 shows the definition of a new storage template.

Storage templates /

Create storage template

Details * Storage pools * Port groups

General properties

Template name *

Storage provider *

Template type *

Description

Use this storage template as default

Use all available WWPNs

Auto expand

Thin/Compressed provisioning settings

Thin provisioning grain size *

Real capacity (% of virtual capacity) *

Enable threshold check

Warning threshold (%)

Figure 2-22 Storage template definition

The following are defined in a storage template:

- Template name** The template name is the name used to identify the storage template in IBM PowerVC. This name is local to IBM PowerVC.
- Storage provider** The storage provider refers to the storage system that volumes utilizing this template will be provisioned from. A storage template cannot be used by more than one storage provider.
- Template type** Three different volume types can be provisioned in IBM PowerVC: Thin, thick (generic) or compressed volumes. Thin and compressed volume types have additional settings that can be set through the storage template, such as:
- Real capacity** Real capacity is the percentage of size of the provisioned volume that is actually provisioned at the time of volume creation. This percentage determines the real storage capacity of the volume at the time of creation. For example, if a 200 GB disk is thin provisioned using the values in Figure 2-22, after provisioning the actual size of the disk is 4 GB.

Warning threshold Warning threshold is a percentage of the virtual capacity. Once the percentile value is reached, an alert is issued. For example, if a 200 GB disk is thin provisioned using the values in Figure 2-22 on page 50 (i.e. warning threshold 80%), when the real capacity of the volume reaches 160 GB, a notification will appear in the IBM PowerVC logs.

Description Description is an optional text description of the storage template. In this description, the purpose of the template or information about the port grouping can be documented here.

Checkboxes on the details tab:

- The Use this storage template as default ensures that at the time of provisioning volumes from the selected storage provider, this storage template will be used as the default.
- The Use all available WWPN ensures IBM PowerVC uses all available WWPNs from all of the I/O groups in the storage controller to attach the volume to the VM.
- The Auto expand ensures thin provisioned volumes do not go offline when they reach their virtual capacity. As a thin-provisioned volume uses more of its capacity, this feature maintains a fixed amount of unused real capacity, which is called the contingency capacity.
- The Select I/O groups select the I/O group to which to add the volume. For the SAN Volume Controller, the maximum I/O groups that are supported is four.
- The Throttle I/O is used to achieve a better distribution of storage controller resources by limiting the volume of I/O processed by the storage controller at various levels. IOPS and bandwidth throttle limits can be set. Throttling can be set at a volume level, host, host cluster and storage pool level. The throttling check box on the storage template sets the throttling at the storage volume level.
- Flash copy rate controls the rate at which updates are propagated from a source volume to a target volume. IBM FlashCopy® mapping copy rate values can range from 128 KBps (10) to 2 GBps (150) and can be changed when the FlashCopy mapping is in any state. The default value on IBM FlashSystem systems is 50 (2 MBps). The user can determine the priority that is given to the background copy process and adjust.
- Disable fast formatting of disk.

Storage pools Each storage template can have only one storage pool from which volumes are created.

Enable mirroring When checked, you must select another pool for volume mirroring. The volume that is created has one more copy in the mirroring pool. IBM FlashSystem systems clients can use two pools from two different back-end storage devices to provide high availability.

Figure 2-23 shows the powervc_data pool selected with mirroring enabled.

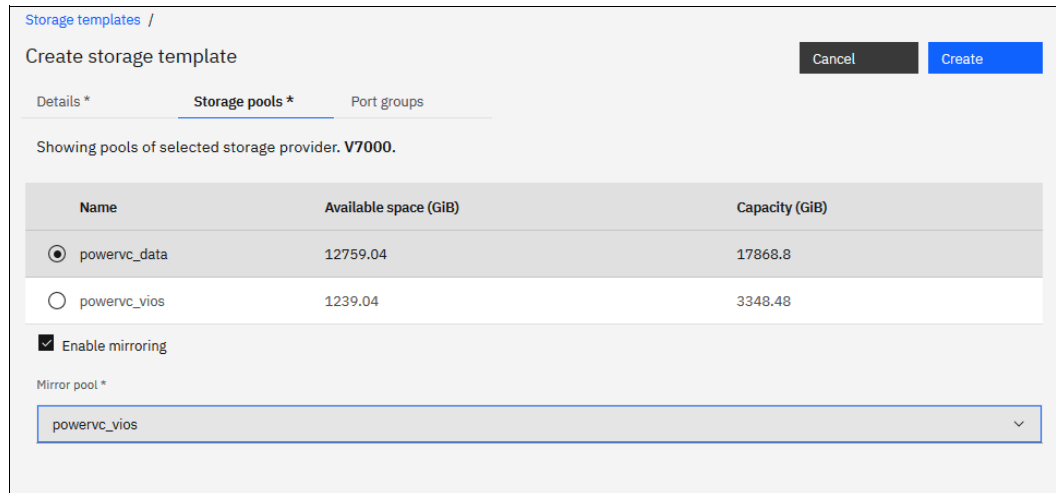


Figure 2-23 Selected storage pool with mirroring enabled

Port groups

Storage port groups provide the capability to balance workload across the storage array and improve redundancy when zoning NPIV-attached volumes to virtual machines. When multiple port groups are defined on a storage template, deployed VMs utilize port groups in an iterative fashion, ensuring balance and not based on I/O metrics from the array. Figure 2-24 shows two port groups created, the blue and green port groups. If two VMs are deployed using this storage template, VM1 will use one port group and VM2 will use the other port group, creating a balance.

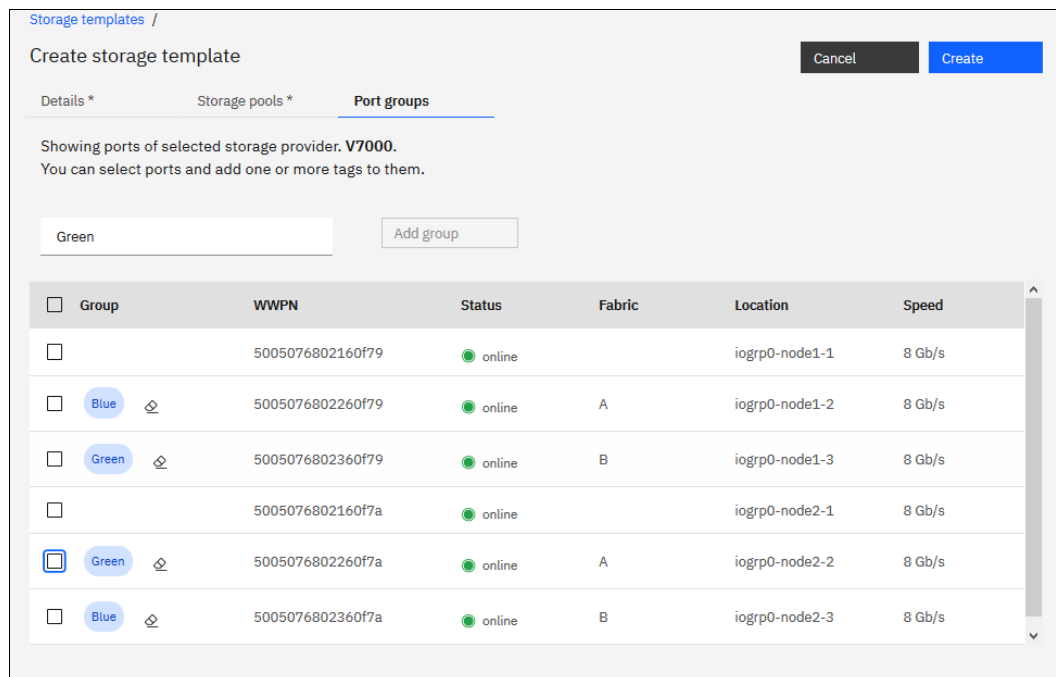


Figure 2-24 Storage port groups page with blue and green port groups created

2.6.3 Storage connectivity group and tags

The basic use of storage connectivity groups is to isolate storage traffic. Storage connectivity groups can be required for various reasons. Popular reasons are as follows:

- ▶ Ensuring workloads are separated based on function, for example production, test and development workloads. In this scenario, Fibre Channel ports are tagged based on their use.
- ▶ Ensuring specific Virtual I/O Server pairs are used when deploying a VM.
- ▶ Ensuring node or I/O drawer redundancy.
- ▶ Ensuring virtual machines are deployed on only certain hosts.
- ▶ Ensuring specific virtual machines can be migrated to only certain hosts.

The default storage connectivity groups are created for NPIV, vSCSI or for shared storage pools (SSP) depending on the situation.

The default storage connectivity group for NPIV connectivity is created when IBM PowerVC first initializes, and as resources are managed in the environment, they are added to this group when applicable.

Default groups that allow vSCSI connectivity are created only when an existing VM with vSCSI connectivity is managed in the environment. A default storage connectivity group that is specific for an SSP is created when the SSP is first automatically managed into the environment. The default storage connectivity groups can be disabled but not deleted.

When a new virtual machine is deployed from an image, you must specify a storage connectivity group. The virtual machine will be deployed to a hosts that satisfies the storage connectivity group settings.

Fibre Channel port tags are vital in ensuring adapter and port level redundancy, assuming a configuration such a Power E980 system with Fibre Channel adapters split between two I/O drawers. It is essential to have virtual machine's virtual Fibre Channel adapters, mapped to physical Fibre Channel ports on both I/O drawers in order to provide adapter and port level redundancy. This can be ensured using Fibre Channel port tags. Figure 2-25 shows two ports tagged with SCG1.

Fibre channel ports										Save		
Search in the table										Set port tag	Set connectivity	Refresh
<input type="checkbox"/>	Name	Host	↑	VIOS	WWPN	Status	Available connections	Tag	Connectivity	Fabric		
<input type="checkbox"/>	fcs0	Redbooks01		vios08	10000090faa094e5	OK	59	SCG1	Any	B		
<input type="checkbox"/>	fcs1	Redbooks01		vios07	10000090faa92980	Offline			Any	None		
<input type="checkbox"/>	fcs1	Redbooks01		vios08	10000090faa094e6	Offline			Any	None		
<input type="checkbox"/>	fcs0	Redbooks01		vios07	10000090faa9297f	OK	59	SCG1	Any	A		

Figure 2-25 Fibre Channel port tags

Assuming fcs0 and fcs1 are on the same Fibre Channel adapter assigned to vios7, and it is located on separate I/O drawer relative to the adapter with ports fcs0 and fcs1 assigned to vios8. Creating a storage connectivity group CG1, and restricting all virtual machine deployments to Fibre Channel ports with the tag SCG1 will ensure Fibre Channel port and adapter layer high availability for the deployed virtual machines. This is the foundation for workload separation in IBM PowerVC.

If we tag other ports on another VIOS pair on a different host with the same SCG1 and ensure the VIOS pair is part of CG1, during live migrations, virtual machines deployed with the storage connectivity group CG1, will migrate to the VIOS pair that is part of CG1 and virtual machines will be mapped to the ports tagged SCG1. This ensures priority workloads are migrated to priority ports.

Working with initiator port groups

Initiator port groups (IPGs) define the set of VIOS ports to be used for volume attachment when using NPIV storage. This feature provides an IBM PowerVC administrator with further flexibility in specifying VIOS ports for volume attachment. It also affords a virtual machine the capability to scale on the amount of volume attachment. IPGs boot and data volumes can be attached to separate ports. IPGs are defined in storage connectivity groups.

Multiple IPGs can be defined for a single storage connectivity group, but it must contain ports from all VIOS members of the storage connectivity group. During live migration, the ports within the same IPG are selected on the target host. A VIOS port can only be a member of one IPG and once a virtual machine is associated with a shared storage connectivity group, you cannot edit the IPG. IPGs can be used in the storage template to define an exact path for attached volumes by matching IPGs and storage port groups. This is possible in the storage template for IBM storages (DS8000 XIV/A9000), PowerMax and Hitachi.

2.6.4 Combining storage connectivity groups, tags and storage port groups

By using the storage connectivity groups, Fibre Channel port tags and storage port group functions, you can tailor them to the specific needs of virtual machines and ensure high availability in different layers of the infrastructure: port, adapter, VIOS, fabric and the storage port layer. Figure 2-26 shows an example of two virtual machines, one for production and the other one for testing. Their paths to storage has been separated using storage connectivity groups and storage port groups.

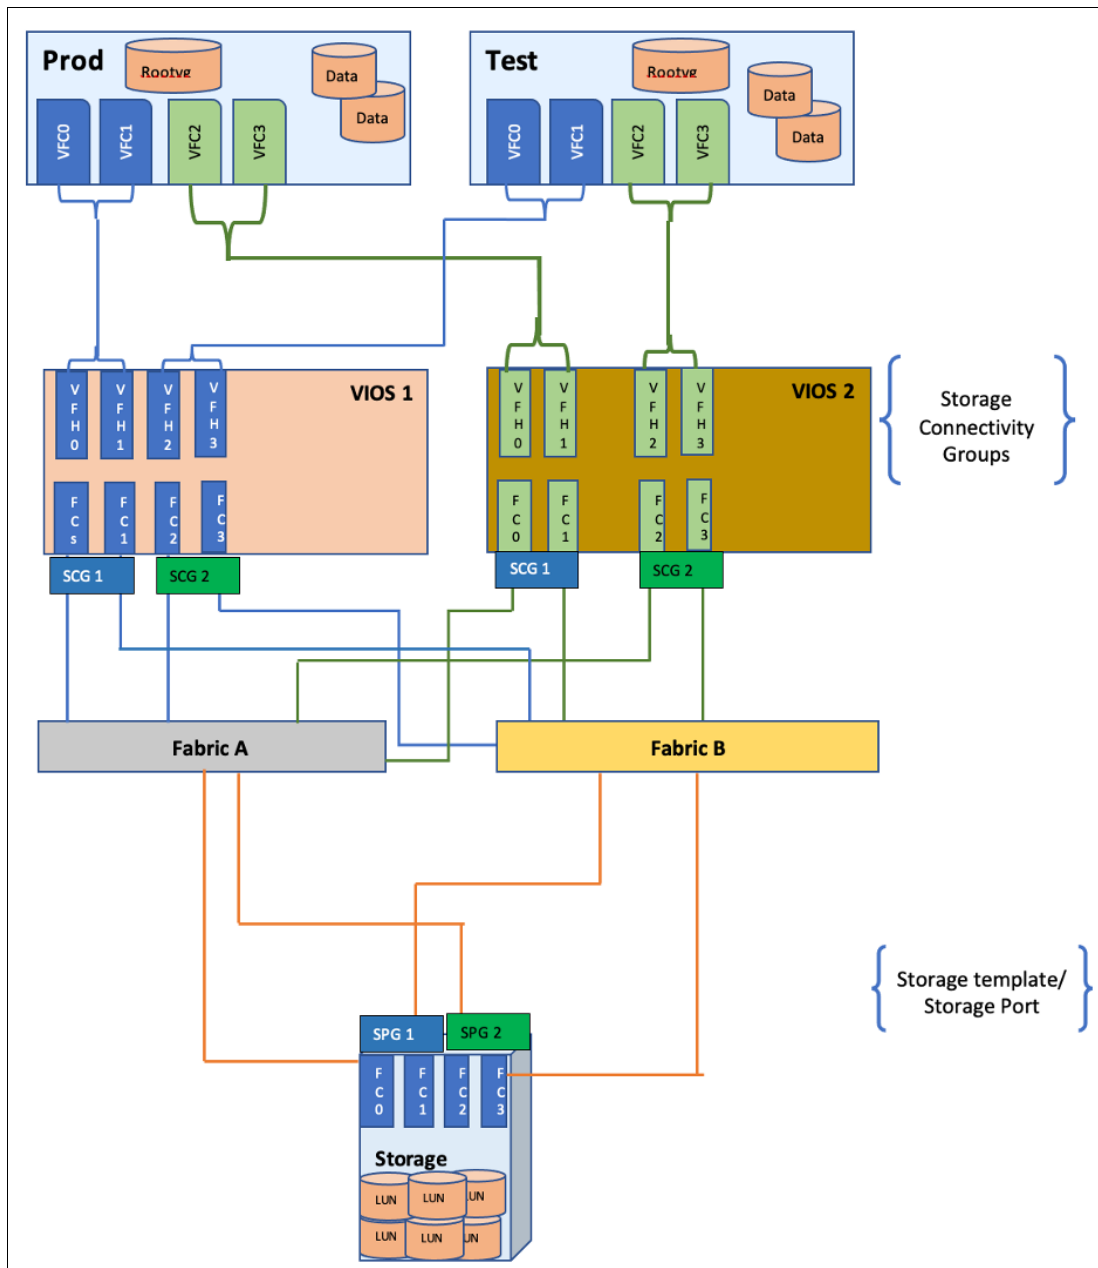


Figure 2-26 Utilizing storage connectivity groups and storage port groups example

Using the IBM PowerVC GUI the following groups is been made to achieve the workload separation:

- ▶ VIOS 1 and VIOS 2 ports fcs0 and fcs1 have been labelled as SCG1.
- ▶ VIOS 1 and VIOS 2 ports fcs2 and fcs3 have been labelled as SCG2.
- ▶ Two storage connectivity groups have been created with the following rules:
 - VIOS redundancy: At least two VIO servers must be mapped to a VM for storage connectivity.
 - Fabric redundancy: Every fabric per VIOS.
- ▶ Two storage port groups were defined on the same storage template with the WWPN of fc0 and fc1 belonging to SPG1 and fc2 and fc3 belonging to SPG2. When multiple port groups are defined on a storage template, deployed VMs utilize port groups in an iterative fashion ensuring balance.

2.7 Network management planning

A network represents a set of Layer 2 and 3 network specifications, such as how your network is subdivided with VLANs. It also provides information about the subnet mask, gateway, and other characteristics. When you are deploying an image, you choose one or more existing networks to apply to the new virtual machine. Setting up networks beforehand reduces the amount of information that you need to input during each deployment and helps to ensure a successful deployment.

During deploy time, the network with the Primary Network flag provides the system default gateway address. You can add more networks to segregate and manage the network traffic. If using vNIC, note that a maximum of 32 vNICs can be attached to a virtual machine.

IBM PowerVC supports IP addresses by using hardcoded `/etc/hosts` or Domain Name Server (DNS)-based host name resolution. IBM PowerVC also supports Dynamic Host Configuration Protocol (DHCP) or static IP address assignment. For DHCP, an external DHCP server is required to provide the address on the VLANs of the objects that are managed by IBM PowerVC.

2.7.1 Infoblox support

Typically, when a virtual machine is deployed or deleted, you must manually create or delete the DNS record. However, you can optionally register an Infoblox vNIOS server into PowerVC. Infoblox then automatically updates the DNS records when virtual machines are deployed or deleted; including details such as the IP address, virtual machine name.

Infoblox Version 6.8 and later is supported. Running `create_ea_defs.py` requires Version 7.2 or later. This command is used to create environment variables in Infoblox. Previous versions of Infoblox require the user interface to create environment variables.

To register Infoblox, go to the Network page, click **Configure DNS Authority** and add network details. The details are saved as a zone name in the Infoblox appliance. Make sure that you have valid Infoblox server credentials for configuring the DNS authority.

Note: The value that is provided for hostname or IP address must match the name of the grid member in the Infoblox DNS configuration authority.

You need to configure several attributes that act as environment variables in Infoblox to generate zones and DNS records. Set the following values in the Members section of the Infoblox appliance:

- ▶ Default Domain Name Pattern: {instance_name} - Neutron network name that is used to create the authoritative zone if not present in Infoblox.
- ▶ Default Host Name Pattern: {network_name} - Virtual machine name that is used to create a DNS record.
- ▶ DNS Support: True.
- ▶ Admin Network Deletion: True - Deletes the zone when a network is deleted.

Note: To automatically configure the above settings, run `$ python create_ea_defs.py` from the PowerVC management system. This file is in the following directory:
`/usr/lib/python2.7/site-packages/networking_infoblox/tools/create_ea_defs.py`

When you deploy a virtual machine, PowerVC adds the virtual machine name as the hostname and when you delete a virtual machine PowerVC deletes the DNS record in the Infoblox appliance. If two virtual machines have identical names, Infoblox creates a duplicate record with a separate IP address allotted to each of the virtual machines. In this situation, DNS resolution might not work as expected.

Considerations

You need to understand these considerations before adding Infoblox into PowerVC:

1. When you migrate a virtual machine, DNS reconfiguration is not needed. This is because the hostname is set by using cloud-init during the initial deployment on source host, the same hostname is used on the destination host.
2. During upgrade or update operations, DNS configuration is not required for existing virtual machines. DNS configuration is only performed for virtual machines that are deployed post DNS configuration.
3. When a virtual machine is unmanaged, no DNS record is updated.
4. When a virtual machine is managed or remanaged, you must manually update the IP address to create a DNS record.

Limitations

Consider the following limitations while working with vNICs, performing virtual machine activities, or adding a network.

1. When a vNIC is attached or detached from a virtual machine, no DNS records are created or removed.
2. After registering Infoblox support, you should run `sync_neutron_to_infoblox.py` to add existing networks to Infoblox. Currently, when you run the sync tool, zones are not created because Infoblox does not support Keystone v3. In this case, DNS records will not be written in Infoblox when a virtual machine is deployed.
3. You cannot create a duplicate network with existing subnet details.
4. When you delete virtual machine or network details on PowerVC, Infoblox does not delete the DNS records if they were manually created.

2.7.2 Multiple network planning

Each virtual machine that you deploy must have one or more networks. Using multiple networks provide allow you to separate traffic. If you have multiple projects, you can create shared networks from the Networks page.

PowerVM Enterprise Edition generally uses three types of networks when you are deploying virtual machines.

Data network

Data network provides the route over which workload traffic is sent. At least one data network is required for each virtual machine, and more than one data network is allowed.

Management network

Management network type of network is optional but highly recommended to provide a higher level of function with the virtual machine. A management network provides the Remote Monitoring and Control (RMC) connection between the management console and the client LPAR. Several PowerVC features, such as live migration and dynamic LPAR, add or remove NIC require an active RMC connection between the management console (HMC or the NovaLink partition) and the virtual machine.

In a NovaLink environment, the system will try to use an internal virtual switch (named MGMTSWITCH) to provide the RMC connections. This internal management network requires that images be created with rsct 3.2.1.0-15216 or later. Virtual machines are not required to have a dedicated management network, but having one provides beneficial advanced features.

PowerVC provides the ability for you to connect to a management network, but you must first set up the networking on the switches and the shared Ethernet adapter to support it.

Live Partition Migration network

This optional network provides the route over which migration data is sent from one host to another. Creating a separate network for migration data helps you to control and prioritize your network traffic. For example, you can specify a higher or lower priority for the migration data as compared with standard data or management traffic. If you do not want to use a separate network for Live Partition Migration (LPM), you can reuse an existing data network connection or a management network connection.

2.7.3 Shared Ethernet Adapter planning

If you plan to use Shared Ethernet Adapters for your virtual machine networking, the Shared Ethernet Adapters must be created outside of PowerVC. The configuration of each host's Shared Ethernet Adapters determines how networks treat each host.

When you create a network in PowerVC, a Shared Ethernet Adapter is automatically chosen from each registered host. The Shared Ethernet Adapter is chosen based on the VLAN that you specified when you defined the network. You can always change the Shared Ethernet Adapter to which a network maps or remove the mapping altogether. However, consider the automatic assignment when you set up your networks if you do not want to change many settings later.

For each host, you can change the network mapping of the Shared Ethernet Adapter. You can also opt to not map the network to any of the host's Shared Ethernet Adapters. If you want virtual machines that use a network to not reside on a particular host, do not assign to the network any Shared Ethernet Adapters for that host.

The Shared Ethernet Adapter is chosen based on the VLAN that you specified when you defined the network. The Shared Ethernet Adapter that is chosen as the default is the one with the same network VLAN as the new network. If no such Shared Ethernet Adapter exists, the adapter with the lowest primary VLAN ID (PVID) that is in an available state is chosen.

Note the following regarding Shared Ethernet Adapters:

- ▶ If there are no usable Shared Ethernet Adapters on any host for a specific VLAN ID, you are directed to choose a different VLAN ID.
- ▶ If a Shared Ethernet Adapter is set to Do Not Use, you can select it. However, you cannot use it in a deploy until it is not set to Do Not Use.
- ▶ If the status for the Shared Ethernet Adapter is Unavailable, the RMC connection may be down. The connection must be fixed before you can select this adapter.

Certain configurations might assure the assignment of a particular Shared Ethernet Adapter to a network. For example, assume that you create a new network in PowerVC and choose the PVID of the Shared Ethernet Adapter or one of the additional VLANs of the primary virtual Ethernet adapter as the VLAN. In this case, the chosen Shared Ethernet Adapter must back the network; no other options are made available.

Note these considerations if you change the Shared Ethernet Adapters after the initial configuration:

- ▶ If you create a network, deploy virtual machines to use it, and then change the Shared Ethernet Adapter to which that network is mapped, your workloads are impacted. At a minimum, the network experiences a short outage while the reconfiguration takes place.
- ▶ If you modify a network to use a different Shared Ethernet Adapter and that existing VLAN is already deployed by other networks, those other networks also move to the new adapter. To split a single VLAN across multiple Shared Ethernet Adapters, you must have separate virtual switches assigned to each of those Shared Ethernet Adapters.

Multiple servers and switches

In PowerVC, a host can have multiple dual VIOS pairs. A dual VIOS setup promotes redundancy, accessibility, and serviceability. It can enhance virtual I/O client partition performance and allows you easily expand hardware or add new functions. It also offers load balancing capabilities for Multipath I/O (MPIO) and multiple shared Ethernet adapter configurations.

Note the following regarding network creation and modification:

- ▶ If you create a network, deploy virtual machines to use it, and then change the Shared Ethernet Adapter to which that network is mapped, your workloads are impacted. At a minimum, the network experiences a short outage while the reconfiguration takes place.
- ▶ If you modify a network to use a different Shared Ethernet Adapter and that existing VLAN is already deployed by other networks, those other networks also move to the new adapter.

2.7.4 Planning Single Root I/O Virtualization networks

You can deploy virtual machines that leverage Single Root Input/Output Virtualization (SR-IOV). SR-IOV supports pass-through of Ethernet data from guest virtual machines directly to hardware. This improves performance by allowing data to pass directly from guest virtual machines to physical adapters with minimal processing between, allowing a guest virtual machine to achieve near wire-speed Ethernet performance. SR-IOV also supports some additional configuration options, such as Quality of Service (QoS) for enforcing bandwidth allocations to guest virtual machines.

A given SR-IOV adapter can have multiple physical ports, connected to external switches. Each physical port is divided into logical ports. These logical ports are connected to a virtual machine for network connectivity. These logical ports allow a single physical hardware device to appear as multiple devices to guest virtual machines.

SR-IOV versus Shared Ethernet adapters

In PowerVC, without SR-IOV, you have an SEA and a virtual Ethernet adapter on the Virtual I/O Server (VIOS). These adapters connect to a physical network adapter and to a client network adapter on the virtual machine. This setup allows you to segment your network using VLAN IDs and allows you virtualize your network hardware; providing migration and failover support. An SEA environment also supports higher virtual machine density. However, all network traffic is routed through the SEA on the VIOS, which adds processing cycles.

SR-IOV with vNIC provides a separation of the control plane and data plane for Ethernet within a host. Therefore, an SR-IOV environment performs better because the VIOS is only used to set up and manage the communication channel and provide failover and migration support. SR-IOV does not scale to hundreds or thousands of virtual machines per host. Instead, it is used to set up a few very fast virtual machines.

You can use VLAN IDs to segment your network whether you are using SR-IOV or SEA, but virtual machines cannot be migrated in an SR-IOV environment.

Redundancy support

When deploying a virtual machine that uses SR-IOV networks, PowerVC creates a vNIC adapter for that virtual machine. If you select Redundant adapter when deploying a virtual machine, the vNIC adapter created for the virtual machine includes multiple logical ports. These logical ports are isolated in order to provide redundancy by using multiple physical ports, multiple SR-IOV adapters, and multiple VIOSs if available.

Quality of service

When you create an SR-IOV network, you can use the Virtual NIC capacity field to specify the minimum bandwidth of the network. If that capacity is not available when deploying a virtual machine, the deploy is not allowed.

Requirements

The requirements for implementing this is:

- ▶ POWER8 or later system.
- ▶ The adapter must be in SR-IOV mode.
 - On NovaLink hosts, use `pvmct1 sriov` to view and update the adapter mode.
 - On HMC hosts, in the HMC user interface, modify the SR-IOV adapter, then choose Shared mode.
- ▶ The operating system on the virtual machine must be supported.

Note: Ensure the slot on your server is capable for an SR-IOV adapter.

See more details about SR-IOV and vNIC adapters in the following links:

<https://community.ibm.com/community/user/power/viewdocument/vNIC-and-vNIC-failover-faqs?CommunityKey=71e6bb8a-5b34-44da-be8b-277834a183b0&tab=librarydocuments>
<https://www.ibm.com/support/pages/node/6357843>

Restrictions

- ▶ You cannot directly connect the SR-IOV adapter, physical ports, or logical ports to a virtual machine.
- ▶ Activation engine is not supported.

2.8 Planning users and groups

The following sections describe the planning that is required for users and groups.

2.8.1 User management planning

When you install IBM PowerVC, it is configured to use the security features of the operating system on the management host by default. This configuration sets the root operating system user account as the only initially available account with access to the IBM PowerVC server.

Upon installation of IBM PowerVC, a new operating system group named *powervc-filter* is created. The root user account gets added to this group by default. IBM PowerVC has visibility only to the user and group accounts that are part of the *powervc-filter* group. The other operating system users and groups are not exposed to PowerVC unless they are added to the *powervc-filter* group.

As a preferred practice, create at least one system administrator user account to replace the root user account as the IBM PowerVC management administrator. For more information, see section 4.17.1, “Adding user accounts” on page 247. After a new administrator ID is defined, remove the IBM PowerVC administrator rights from the root user ID, as explained in 4.17.3, “Disabling the root user account from IBM PowerVC” on page 250.

Important: IBM PowerVC also requires user IDs that are defined in `/etc/passwd` that must not be modified, such as `nova`, `neutron`, `keystone`, and `cinder`. All of these users use OpenStack and they must not be changed or deleted.

For security purposes, you cannot connect remotely to these user IDs. These users are configured with the login shell `/sbin/nologin`.

User account planning is important to define standard accounts and the process and requirements for managing these accounts. An IBM PowerVC management host can take advantage of user accounts that are managed by the Linux operating system security tools or can be configured to use the services that are provided by LDAP.

IBM PowerVC does not create users or groups in the underlying operating system. PowerVC backups include information about the configured user and group filters. If operating system users and groups are configured differently when the backup is restored, it may lead to administration issues.

Table 2-13 describes the available attributes to use when working with user and group filters.

Table 2-13 User and Group filters

Attribute name	Description
User filter	Limits which users are visible to PowerVC. The default is “(memberOf=powervc-filter)”.
Group filter	Limits which groups are visible to PowerVC. The default is “(name=powervc-filter)”.

A freshly installed PowerVC system displays the default values, as shown in Example 2-1.

Example 2-1 Default user and group filter settings

```
#powervc-config identity repository
Type: os
User filter: (memberOf=powervc-filter)
Group filter: (name=powervc-filter)
```

2.8.2 Projects and role management planning

This section describes the settings that are required for each user and group to operate and perform actions and work with projects.

Managing projects

A *project*, sometimes referred to as a *tenant*, is a unit of ownership. VMs, volumes, images, and networks belong to a specific project. Only users with a role assignment for a given project can work with the resources belonging to that project. At the time of installation, the `ibm-default` project is created, but IBM PowerVC also supports the creation of more projects for resource segregation.

To work with projects, an admin can login to the `ibm-default` project and click Projects from the configuration page.

You can also use the `openstack project` command to manage projects as needed. As a OpenStack administrator, you can create, delete, list, set, and show projects:

- ▶ Create a project by running the following command:
`openstack project create project-name`
- ▶ Delete an existing project by running the following command:
`openstack project delete project-name`
- ▶ List projects by running the following command:
`openstack project list`
- ▶ Set project properties (name, or description) by running the following commands:
`openstack project set --name <name> project-name`
`openstack project set --description <description> project-name`
- ▶ Display project details by running the following command:
`openstack project show project-name`

After you create a project, you must grant at least one user a role on that project.

Project quotas

Project quotas sets limits on the various types of resources within each project. Administrators can edit, enable, and disable the quotas. Project quotas are set from the Project quotas tab of the user interface in the Dashboard menu.

Notes:

- ▶ When a quota is disabled, that resource is unlimited.
- ▶ You can set a quota to be smaller than its current value. The quota is considered exceeded in this case. PowerVC does not change the effective resource usage, but subsequent requests for resources will fail.

Table 2-14 provides the quotas that can be set per project.

Table 2-14 Available quotas

Quota	Description	Default
Collocation Rules	The total number of collocation rules allowed	25
External IP addresses	The maximum number of external (floating) IP addresses that can be assigned in the project	100
Injected files	The total number of injected files allowed for a project. The data is injected at the time of VM provisioning	5
Injected File Content (Bytes)	The maximum size of each injected file that is allowed in the project	10,240
Injected File Path (Bytes)	The maximum length of each injected file path	255
Memory (GB)	The total memory that can be used across all virtual machines in the project	40000 GB
Per Volume (GB)	The maximum amount of storage that can be allocated to each volume in the project in GB	Unlimited (disabled)
Processing Units	The total number of entitled processing units of all virtual machines within the project	5500
Snapshots	The total number of volume snapshots that are allowed in the project	100,000
Virtual Machines	The total number of virtual machines that are allowed in the project	5500
Virtual Processors	The total number of virtual processors (cores) allowed across all virtual machines in the project	55000
Volume Backup (GB)	The total amount of storage for volume backups allowed per project	15,000
Volume Backups	The number of volume backups allowed per project	30
Volume Groups	The number of volume groups allowed per project	200

Quota	Description	Default
Volume Storage (GB)	The total amount of disk space that can be used across all volumes within the project	10,000,000
Volumes	The total number of volumes that can be part of the project	100,000

Managing roles

Roles are assigned to a user or group. They are inherited by all users in that group. A user or group can have more than one role, allowing them to perform any action that at least one of their roles allows.

Roles are used to specify what actions users can perform. Table 2-15 shows the available roles and actions each role is allowed to perform.

Table 2-15 IBM PowerVC Security Roles

Role	Action
Administrator (admin)	Users with this role can perform all tasks and have access to all resources.
Administrator assistant (admin_assist)	Users with this role can perform create and edit tasks but do not have privileges to perform remove or delete operations. The admin_assist user can perform all virtual machine, image, and volume lifecycle operations except Delete.
Deployer (deployer)	Users with this role can perform the following tasks: <ul style="list-style-type: none"> ▶ Deploying a virtual machine from an image ▶ Viewing all resources except users and groups
Image manager (Image_manager)	Users with this role can perform the following tasks: <ul style="list-style-type: none"> ▶ Creating, capturing, importing, or deleting an image ▶ Editing description of an image ▶ Viewing all resources except users and groups
Storage manager (storage_manager)	Users with this role can perform the following tasks: <ul style="list-style-type: none"> ▶ Creating, deleting, or resizing a volume ▶ Viewing all resources except users and groups
Viewer (viewer)	Users with this role can view resources and the properties of resources, but can perform no tasks. They cannot view users and groups.
Virtual Machine Manager (vm_manager)	Users with this role can perform the following tasks: <ul style="list-style-type: none"> ▶ Deploying a virtual machine from an image ▶ Deleting, resizing, starting, stopping, or restarting a virtual machine ▶ Attaching or detaching volume ▶ Snapshot and restore a volume ▶ Attaching or detaching network interface ▶ Editing details of a deployed virtual machine ▶ Viewing all resources except users and groups ▶ Creating, attaching, detaching, and deleting floating IP addresses
Virtual machine user (vm_user)	Users with this role can perform the following tasks: <ul style="list-style-type: none"> ▶ Starting, stopping, or restarting a virtual machine ▶ Viewing all resources except users and groups

Role assignments are specific to a project. Users can log in to only one project at a time in the IBM PowerVC user interface. If they have a role on multiple projects, they can switch to one of those other projects without having to log out and log back in. When users log in to a project, they see only resources, messages, and other information, that belong to that project. They cannot see or manage resources that belong to a project where they have no role assignment. There is one exception to this rule. The admin role can operate across projects in many cases. Be mindful of this when handing out admin role assignments.

Important: OpenStack does not support moving resources from one project to another project. You can move volumes by unmanaging them and then remanaging them in the new project, but it is not possible to perform the same action for VMs because the network on which that VM depends is tied to the original project.

2.9 Security management planning

IBM PowerVC provides security services that support a secure environment and, in particular, the following security features:

- ▶ Starting with BM PowerVC Version 2.0.0, an additional authentication mechanism called Time-based One-Time Password (TOTP) has been added to provide enhanced security for the users logging into IBM PowerVC. For a user to be authenticated, TOTP along with a password must be provided by the user.
- ▶ PowerVC uses the HSTS, X-XSS-Protection, and X-Content-Type-Options type HTTP security response headers.
- ▶ Signing packages adds an extra level of trustworthiness towards a product. IBM PowerVC ships both RPM packages and Debian packages with its installer.

2.9.1 Ports that are used by IBM PowerVC

Information about the ports that are used by IBM PowerVC management hosts for inbound and outbound traffic is on the following IBM Knowledge Center websites:

- ▶ Ports used on the management server

https://www.ibm.com/support/knowledgecenter/bs/SSVSPA_2.0.0/com.ibm.powervc.cloud.help.doc/powervc_planning_security_firewall_cloud.html?view=kc#powervc_planning_security_firewall_hmc__d264e87

- ▶ Ports used by IBM PowerVC on the management server

https://www.ibm.com/support/knowledgecenter/bs/SSVSPA_2.0.0/com.ibm.powervc.cloud.help.doc/powervc_planning_security_firewall_cloud.html?view=kc#powervc_planning_security_firewall_hmc__d264e567

- ▶ Ports used by PowerVM NovaLink managed host

https://www.ibm.com/support/knowledgecenter/bs/SSVSPA_2.0.0/com.ibm.powervc.cloud.help.doc/powervc_planning_security_firewall_cloud.html?view=kc#powervc_planning_security_firewall_hmc__d264e734

2.9.2 Providing a certificate

The IBM PowerVC management server uses a self-signed X.509 certificate, by default, to secure its web interface and REST APIs. Because that self-signed certificates can be created by anyone, they are not trusted by client's web browsers automatically. In order to improve security, a certificate signed by a certificate authority should be used to replace the default self-signed certificate. Expiring or revoked certificates will also need to be replaced.

The web interface and REST APIs use the private key and certificate at the following locations:

- ▶ /etc/pki/tls/private/powervc.key
- ▶ /etc/pki/tls/certs/powervc.crt

The process to replace the existing certificates can be found here:

https://www.ibm.com/support/knowledgecenter/SSVSPA_2.0.0/com.ibm.powervc.cloud.help.doc/powervc_certificate_cloud.html

2.10 Product information

For additional planning information, see the following resources.

IBM support

IBM Support is your gateway to technical support tools and resources that are designed to help you save time and simplify support. IBM Support can help you find answers to questions, download fixes, troubleshoot, submit and track problem cases, and build skills:

<https://www.ibm.com/support/home/>

Learn and stay informed about the transformation of IBM Support, including new tools, new processes, and new capabilities, by going to the IBM Support Insider:

<https://www.ibm.com/support/insider>

IBM Support Guide

IBM Support gives you an advantage by helping you drive success with your IBM products and services across cloud, on-premises, and hybrid cloud platforms:

<https://www.ibm.com/support/pages/node/733923>

Offering Information

Product information is available on the IBM Offering Information website:

https://www-01.ibm.com/common/ssi/SearchResult.wss?request_locale=en&MPPEFFTR=CONTENTS&MPPEFSCH=5765-VC2&MPPEFSRT=2

Packaging

This offering is delivered through IBM My Entitled System Support Site (ESS) as an electronic download. There is no physical media.

It is possible to obtain the electronic version at IBM My Entitled System Support Site:

<https://www.ibm.com/servers/eserver/ess/ProtectedServlet.wss>

Click **My entitled software** in the left pane, and then click **My entitled software**.

Software maintenance

The IBM Agreement for Acquisition of Software Maintenance (Z125-6011) applies for Subscription and Support and does not require client signatures.

Licenses under the IBM International Program License Agreement (IPLA) provide for support with ongoing access to releases and versions of the program. IBM includes one year of Software Subscription and Support (also referred to as Software Maintenance) with the initial license acquisition of each program acquired. The initial period of Software Subscription and Support can be extended by the purchase of a renewal option, if available. Two charges apply: a one-time license charge for use of the program and an annual renewable charge for the enhanced support that includes telephone assistance (voice support for defects during normal business hours), as well as access to updates, releases, and versions of the program as long as support is in effect.

IBM Enterprise Support and Preferred Care

IBM System Storage or Power Systems Hardware and Software Support Services provide around-the-clock integrated hardware and software services backed by our global support infrastructure, product expertise and proprietary analytics tools:

<https://www.ibm.com/support/pages/node/738889>

Licensing

IBM International Program License Agreement including the License Information document and Proof of Entitlement (PoE) govern your use of the program. PoEs are required for all authorized use.

This software license includes Software Subscription and Support (also referred to as Software Maintenance).

Hardware requirements

Any IBM system that includes an IBM POWER7+, and later, processor.

Software requirements

- ▶ PowerVM Standard Edition (5765-VS3) for basic functions, and PowerVM Enterprise Edition (5765-VE3) or PowerVM PowerLinux Edition (5765-VL3) for full function.
- ▶ Firmware v8.2, or higher, is required for the new Remote Restart function for PowerVM that is managed by PowerVC.

The program's specifications and specified operating environment information may be found in documentation accompanying the program, if available, such as a readme file, or other information published by IBM, such as an announcement letter. Documentation and other program content may be supplied only in the English language.



IBM PowerVC installation

This chapter describes the installation of IBM Power Virtualization Center (IBM PowerVC) and IBM PowerVC for Private Cloud.

It covers the following topics:

- ▶ 3.1, “Setting up the IBM PowerVC environment” on page 70
- ▶ 3.2, “Installing IBM PowerVC” on page 81
- ▶ 3.3, “Uninstalling IBM PowerVC” on page 86
- ▶ 3.4, “Upgrading IBM PowerVC” on page 87
- ▶ 3.5, “Updating IBM PowerVC” on page 97
- ▶ 3.6, “IBM PowerVC backup and recovery” on page 98
- ▶ 3.7, “IBM PowerVC command-line interface” on page 102
- ▶ 3.8, “Virtual machines that are managed by IBM PowerVC” on page 104

3.1 Setting up the IBM PowerVC environment

IBM PowerVC Version 2.0.0 can be installed on Red Hat Enterprise Linux Version 8.2 and 8.3 (either the ppc64le or x86_64 platform), and SUSE Linux Enterprise Server 15 SP1 and SP2 (ppc64le platform only).

Note: Unlike the Hardware Management Console (HMC), IBM PowerVC is not a stand-alone appliance. It must be installed on top of the operating system. The general trend is moving toward Little Endian (LE), and support for PowerVC installation on ppc64 architecture is being withdrawn with Version 1.4.4. On the IBM POWER Systems, there are two supported options to choose from: Red Hat Enterprise Linux (LE) or SUSE Linux Enterprise Server (LE). You must have a valid Linux license to use the operating system and a valid license to use IBM PowerVC.

Before you install IBM PowerVC, first install the operating system (Red Hat Enterprise Linux or SUSE Linux Enterprise Server) on the management virtual machine (VM) or management host. IBM PowerVC requires several additional packages to be installed. These packages are automatically installed if you have configured and enabled a valid Linux repository. For additional information, see the relevant sections in the IBM Knowledge Center documentation for IBM PowerVC:

- ▶ Installing Red Hat Enterprise Linux on the management server or host:

https://www.ibm.com/support/knowledgecenter/SSXK2N_2.0.0/com.ibm.powervc.standa.rd.help.doc/powervc_installing_rhel_mgmt_vm_hmc.html

- ▶ Installing SUSE Linux on the management server or host:

https://www.ibm.com/support/knowledgecenter/SSXK2N_2.0.0/com.ibm.powervc.standa.rd.help.doc/powervc_installing_sles_mgmt_vm.html

Important: The management VM must be dedicated to IBM PowerVC and the operating system on which it runs. Do not install other software on it.

To set up the management hosts, complete the following tasks located in the sections that are provided:

1. Section 3.1.1, “Creating the virtual machine to host IBM PowerVC” on page 70 (only if you plan to install IBM PowerVC in a virtualized environment)
2. Section 3.1.2, “Downloading and installing the Linux operating system” on page 77
3. Section 3.1.3, “Customizing the Linux operating system” on page 78

3.1.1 Creating the virtual machine to host IBM PowerVC

IBM PowerVC can be installed on a VM running on IBM Power System ppc64le or the x86_64 platform.

To install the management VM on a platform that matches your needs, select one of the following methods outlined in the following topics:

- ▶ “Create the VM by using the HMC Enhanced GUI” on page 71
- ▶ “Create the VM by using the NovaLink CLI” on page 71
- ▶ “Install a VM by using BMC on OPAL systems” on page 73
- ▶ “Create the VM by using Red Hat KVM or Ubuntu KVM” on page 77
- ▶ “Create the management VM on an x86_64 machine” on page 77

Create the VM by using the HMC Enhanced GUI

Complete the following steps to create the VM using the HMC Enhanced GUI:

1. In the left navigation pane, open Resources and click **All Systems**.
2. Click the name of the desired managed system (the Partitions view is loaded), then click the **Create Partition** button on the upper right-hand side corner. This is shown in Figure 3-1.

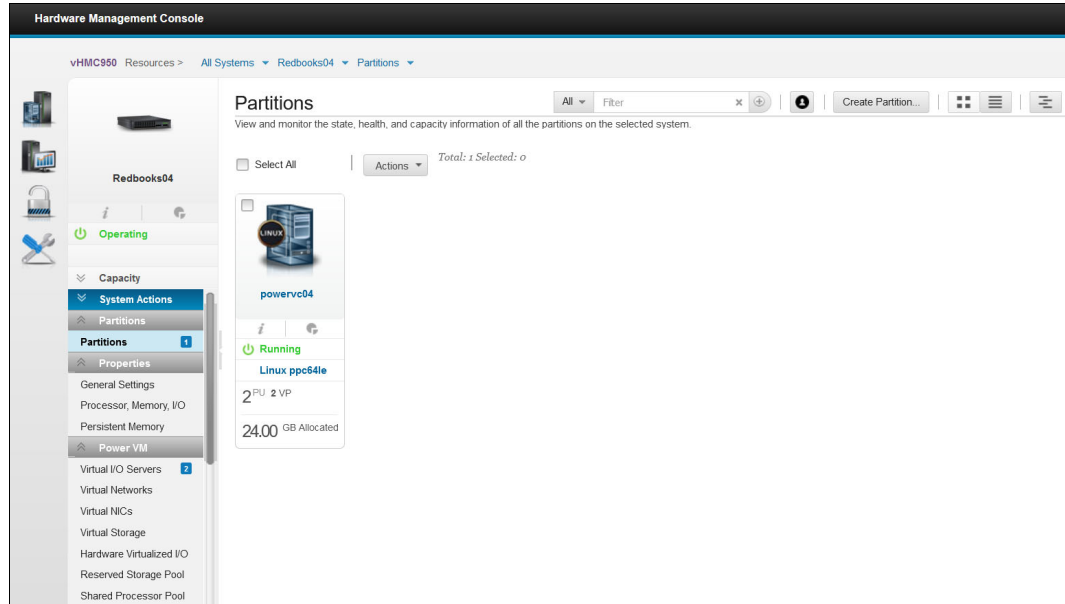


Figure 3-1 Create Client Partition via HMC Enhanced GUI

3. Follow the steps in the Create Partition wizard to create a single logical partition (LPAR) and then configure the partition.

Create the VM by using the NovaLink CLI

Complete the following steps to create the VM using the NovaLink CLI:

1. List all LPARs on the server by using the `pvmctl` command:

```
# pvmctl lpar list
```
2. Create the LPAR and assign processor and memory size, and re-list to verify as shown in Example 3-1.

Example 3-1 Create the LPAR `powervc001` and assign resources

```
# pvmctl lpar create --name powervc001 --proc-type shared --proc 4 --sharing-mode uncapped --type AIX/Linux --mem 8192 --proc-unit 1
```

```
#pvmctl lpar list  
Logical Partitions
```

Name	ID	State	RMC	Env	Ref Code	Mem	CPU	Ent
novalink	1	running	----	AIX/Linux	Linux ppc64le	8192	2	1.0
powervc001	4	not activated	inactive	AIX/Linux	00000000	8192	4	1.0

3. Configure the processing units, virtual processors and memory settings (minimum, maximum, desired) as shown in Example 3-2.

Example 3-2 Configure processing units, virtual processors and memory settings

```
# pvmctl lpar update -i name=powervc001 --set-fields
SharedProcessorConfiguration.min_units=0.1
SharedProcessorConfiguration.max_units=2
SharedProcessorConfiguration.desired_units=1
SharedProcessorConfiguration.min_virtual=2
SharedProcessorConfiguration.max_virtual=6
SharedProcessorConfiguration.desired_virtual=4
PartitionMemoryConfiguration.max=8192 PartitionMemoryConfiguration.min=1024
PartitionMemoryConfiguration.desired=8192
```

4. To list the Virtual Ethernet Adapters (VEA) configuration on the system, use the following command:

```
# pvmctl vea list
```

5. To create and verify the VEA, use the commands shown in Example 3-3.

Example 3-3 Create and verify the VEA

```
# pvmctl vea create --pvid=1 --vswitch=ETHERNET0 --parent-id name=powervc001
# pvmctl vea list
```

6. Use the commands shown in Example 3-4 to list and create a logical volume.

Example 3-4 List and create a logical volume

```
# pvmctl lv list
# pvmctl lv create -p name=vios03 -c name=rootvg --size 90 --name pvc2lv
```

7. Map the created logical volume via vscsi with LPAR powervc001 and re-list to verify the mapping, by issuing commands shown in Example 3-5.

Example 3-5 Map logical volume

```
# pvmctl scsi list
# pvmctl scsi create --type lv --lpar name=powervc001 --stor-id name=pvc2lv
--parent-id name=vios03
# pvmctl scsi list
```

8. List and create repository if needed. See Example 3-6. In this case we used rootvg of VIOS server vios03.

Example 3-6 List and create the repository

```
# pvmctl repo list
# pvmctl repo create -sp rootvg -size 20G --parent-id name=vios03
```

9. Upload the virtual media disk into the repository as shown in Example 3-7.

Example 3-7 Upload the virtual media disk into the repository

```
# pvmctl vom upload --vios vios03 --data <path to .iso image of Red Hat 8.3>
--name rhel83
```

```
# pvmctl repo list
```

Virtual Media Repositories

VIOS	VG	VMR Size (GB)	Free (GB)	OM Name	OM Size (GB)	rw
vios03	rootvg	20.0	12.4404	rhel83	7.5596	rw
vios04	rootvg	20.0	20.0			

10. Map the .iso file to the powervc001 LPAR as shown in Example 3-8.

Example 3-8 Map the .iso file to the powervc001 LPAR

```
# pvmctl scsi create --type vopt --lpar name=powervc001 --stor-id rhel83 -p
name=vios03
```

11. Boot the newly created LPAR powervc001 to SMS (System Management Services) to start the installation process from the media (Example 3-9).

Example 3-9 Power on and boot the LPAR to SMS

```
# chsysstate -r lpar -o on -n powervc001 -b sms
# mkvterm -p powervc001
```

For additional details on usage and syntax of **pvmctl** command, use the link below:

https://www.ibm.com/support/knowledgecenter/9009-42A/p9eig/p9eig_cli.htm

Install a VM by using BMC on OPAL systems

In an OPAL-based Power system the command-line utilities are the **ipmitool** and the **openbmctool**. HTML-based management tools that are called BMC GUI and OpenBMC GUI.

For the installation, use the following steps:

1. On the Baseboard management controller (BMC) on POWER8, click on **Launch** next to Remote control to start JViewer as shown in Figure 3-2.

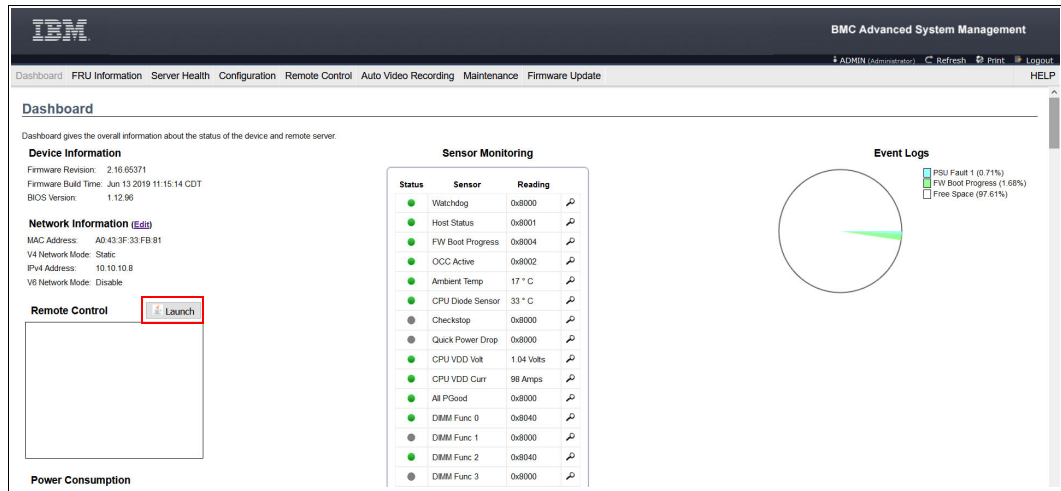


Figure 3-2 Baseboard management controller (BMC) on POWER8

- To load the Linux installation media within the POWER8 BMC JViewer click **Media -> Virtual Media Wizard**. There click the Browse button next to CD/DVD Media: I. Now click the **Connect CD/DVD** button as shown in Figure 3-3.

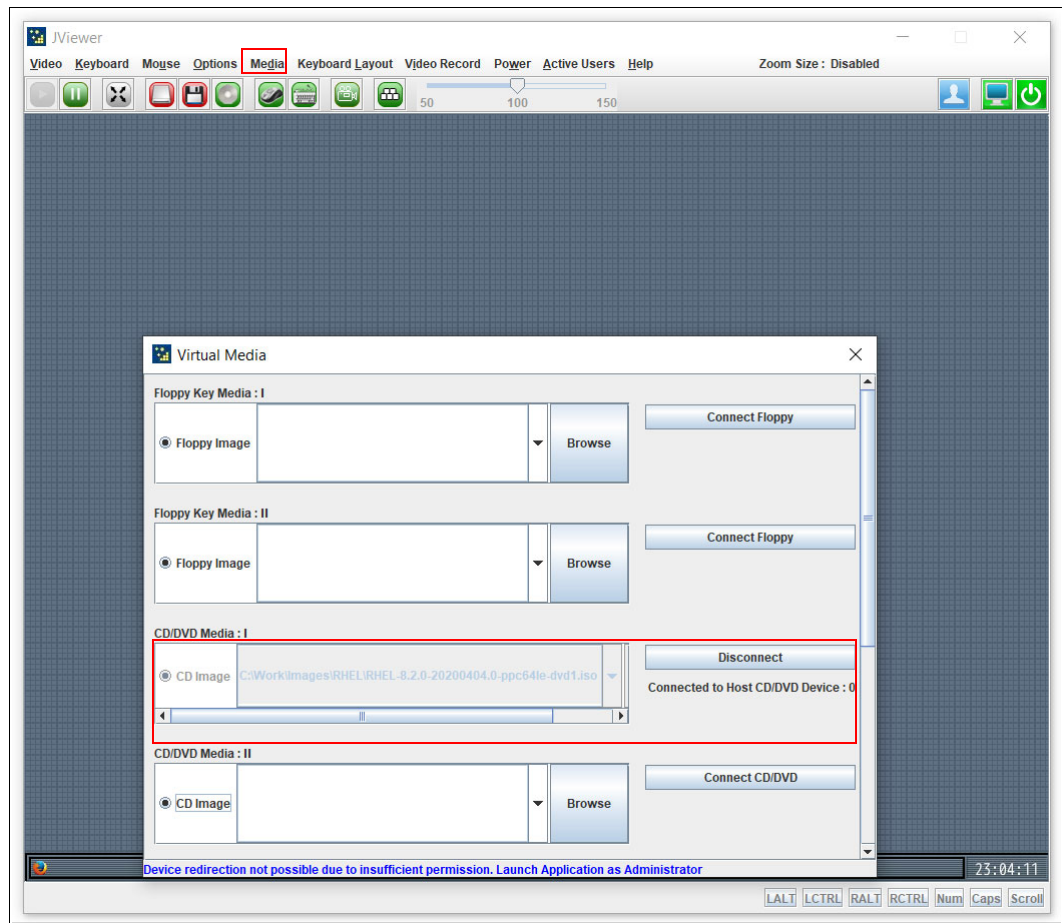


Figure 3-3 Connect install media in POWER8 BMC

3. To load the Linux installation media within the POWER9 BMC GUI click the **Server Control** tab (on the left hand side of the screen), Virtual Media view is loaded. Now click the **Choose file** button. For an example see Figure 3-4.

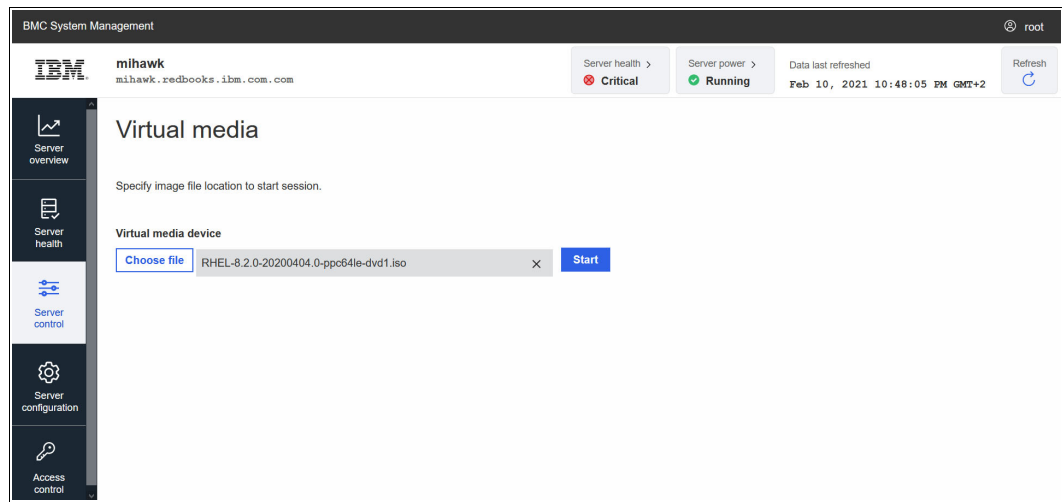


Figure 3-4 Connect install media in POWER9 BMC

4. Open **Server Control** -> **Serial over LAN console** on the desired POWER9 system and install the Linux OS as in Figure 3-5.

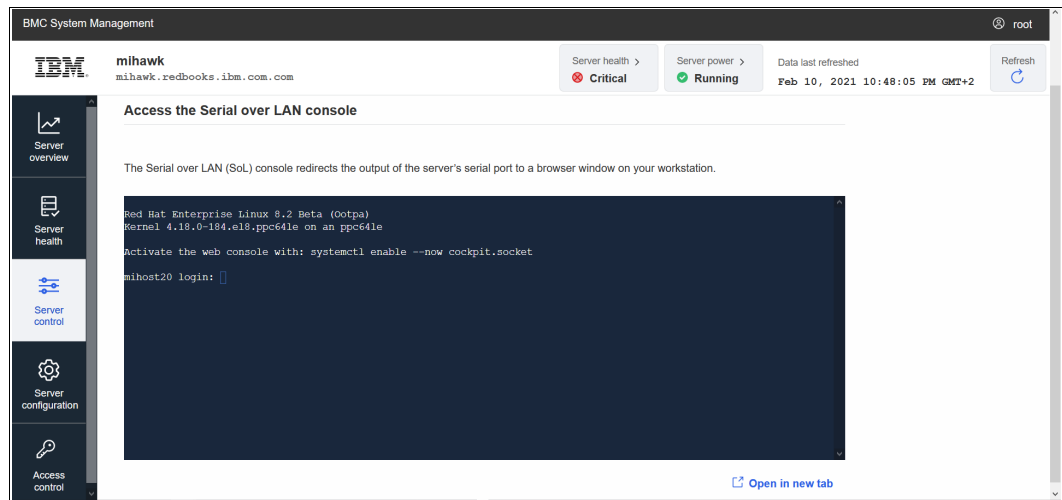


Figure 3-5 Serial over LAN console on POWER9 BMC system

Both CLI and HTML-based tools are available on OPAL bare metal systems. For more details, see the following articles:

- Linux information for IBM systems:

<https://www.ibm.com/support/knowledgecenter/en/linuxonibm/liaaa/ic-homepage.html>

- Managing OpenBMC-based systems:

https://www.ibm.com/support/knowledgecenter/POWER9/p9eih/p9eih_openbmc_kickoff.htm?pos=3

Create the VM by using Red Hat KVM or Ubuntu KVM

To create the management VM on a KVM host, you can use a command-line utility called `virsh`. The `virsh` utility is included in the KVM installation packages.

Note: PowerVC Version 2.0.0 cannot manage KVM on IBM Power Systems. Nevertheless, PowerVC Version 2.0.0 can be installed on a KVM host (ppc64le or x86_64).

Create the management VM on an x86_64 machine

To create a management VM on an x86_64 server, follow the instructions for that server.

After the VM is created, you should install the supported Linux operating system onto the management VM.

3.1.2 Downloading and installing the Linux operating system

As part of the IBM PowerVC setup, you are required to download and install the Red Hat Enterprise Linux or SUSE Linux Enterprise Server operating system. You would need a valid license and a valid copy of the PowerVC software. IBM PowerVC is not a stand-alone appliance.

Note: PowerVC is installed on top of the operating system, but it does not include the license to use Red Hat Enterprise Linux or SUSE Linux Enterprise Server.

- ▶ You obtain the software image and a valid license from the vendor's website:
 - For Red Hat, go to:
<http://www.redhat.com>
 - For SUSE, go to:
<https://www.suse.com/>
- ▶ Install Red Hat Enterprise Linux by using your preferred method. For more information, see the *Product Documentation for Red Hat Enterprise Linux 8*, that can be found at:
https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/
- ▶ Install SUSE Linux Enterprise Server SP1 or SP2 by using your preferred method. For more information, see the *SUSE Documentation for a Specific Product*, that can be found at:
 - SUSE Linux Enterprise Server 15 SP1:
<https://documentation.suse.com/sles/15-SP1/>
 - SUSE Linux Enterprise Server 15 SP2:
<https://documentation.suse.com/sles/15-SP2/>

Important: When you run the Red Hat Enterprise Linux installation wizard, on the “Software Selection” window, choose **Server** or **Server with GUI**.

- ▶ You can have Red Hat Enterprise Linux available for installation in one of the following ways:
 - The Red Hat Enterprise Linux Binary DVD ISO image is the full-installation image that contains the BaseOS and AppStream.
 - The Red Hat Enterprise Linux Boot ISO image is a minimal boot ISO image used to boot into the installation program and requires access to the BaseOS and AppStream repositories to install the software packages.
- ▶ For detailed Red Hat Enterprise Linux installation instructions, see the documentation in the Red Hat Customer Portal:
<https://access.redhat.com/>

Important: If you install Red Hat Enterprise Linux OS with **Server** option, it is required to install Python 3.0 manually. You should set the maximum file size to **unlimited** using the **ulimit** command (**ulimit -f unlimited**) on the server before you start the IBM PowerVC installation. Do not add any other packages to the server. If any other packages are needed by IBM PowerVC 2.0.00, the additional packages are obtained by the IBM PowerVC installer automatically. For more information about the packages requirements and restrictions, see IBM Knowledge Center:

https://www.ibm.com/support/knowledgecenter/SSXK2N_2.0.0/com.ibm.powervc.standa rd.help.doc/powervc_installing_rhel_mgmt_vm_hmc.html

3.1.3 Customizing the Linux operating system

Before you install PowerVC, ensure that all the hardware and software prerequisites are met, and your environment is correctly configured. Also, ensure that you prepare the management server and install the supported version of Red Hat Enterprise Linux Server or SUSE Linux Enterprise Server to meet the following IBM PowerVC requirements.

Red Hat Enterprise Linux Server

To customize your Red Hat Enterprise Linux server follow these steps:

1. Configure the network.

The first task (before you install IBM PowerVC) is to configure the network. IBM PowerVC uses the default network interface eth0. To use a different network interface, such as eth1, set the **HOST_INTERFACE** environment variable before you run the installation script. The following example shows the setting:

```
export HOST_INTERFACE=eth1
```

2. Set the Domain Name Server and host name.

Two options exist for managing name resolution. Either use Domain Name Server (DNS) or use the `/etc/hosts` file. You must pay attention to the correct setting of the name resolution of all components that are managed by IBM PowerVC.

- If you do not plan to use DNS for host name resolution, ensure that all hardware components (including virtualized components) are correctly defined in the `/etc/hosts` file.
- If you plan to use DNS for host name resolution, all hardware components must be correctly defined in your DNS. In addition, you should enable forward and reverse resolution. Host names must be consistent across the whole IBM PowerVC domain.

Important: Regardless of the host name resolution method that you use, the IBM PowerVC management host must be configured with a valid, fully qualified domain name.

3. Configure the YUM repository for the IBM PowerVC installation.

To successfully install IBM PowerVC, your system must be configured and enabled to access the Red Hat Enterprise Linux YUM repository on the host by using a mounted DVD or by being subscribed to the main server Red Hat Network (RHN) channel. You will probably need to add the two main repositories *BaseOS* and *AppStream*, required for a basic Red Hat Enterprise Linux installation. Complete the following steps:

- a. Download the correct package depending on the your system architecture:
 - Red Hat Enterprise Linux for IBM Power: ppc64le
 - Red Hat Enterprise Linux x86_64: x86_64
- b. Configure the YUM repository by using a mounted DVD as shown in Example 3-10.

Example 3-10 yum configuration

```
# ls /dev/sr*
/dev/sr0
# mount -o loop /dev/sr0 /mnt
# cp /mnt/media.repo /etc/yum.repos.d/rhel8BaseOS.repo
# cp /mnt/media.repo /etc/yum.repos.d/rhel8AppStream.repo
# sed -i '1 s/^[rhe18BaseOS] #' /etc/yum.repos.d/rhel8BaseOS.repo
# sed -i '1 s/^[rhe18AppStream] #' /etc/yum.repos.d/rhel8AppStream.repo
# echo "baseurl=file:///mnt/BaseOS/" >> /etc/yum.repos.d/rhel8BaseOS.repo
# echo "gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release" >>
/etc/yum.repos.d/rhel8BaseOS.repo
# echo "baseurl=file:///mnt/AppStream/" >> /etc/yum.repos.d/rhel8AppStream.repo
# echo "gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release" >>
/etc/yum.repos.d/rhel8AppStream.repo
```

- c. Clean, list, and verify that YUM can see the new repository file by running the following commands as in Example 3-11.

Example 3-11 Verify yum

```
# yum clean all
# subscription-manager clean
# yum --noplugins list
# yum repolist all
# yum whatprovides pstree
# subscription-manager list --installed
```

SUSE Linux Enterprise Server

To customize your SUSE Linux Enterprise Server server, follow these steps:

1. Configure the network.

The first task before you install IBM PowerVC is to configure the network. IBM PowerVC uses the default network interface eth0. To use a different network interface, such as eth1, set the **HOST_INTERFACE** environment variable before you run the installation script. The following example shows the setting:

```
export HOST_INTERFACE=eth1
```

2. Set the Domain Name Server and host name.

Two options exist for managing name resolution: Either use DNS or use the /etc/hosts file. Pay attention to the correct setting of the name resolution of all components that are managed by IBM PowerVC.

- If you do not plan to use DNS for host name resolution, ensure that all hardware components (including virtualized components) are correctly defined in the /etc/hosts file.
- If you plan to use DNS for host name resolution, all hardware components must be correctly defined in your DNS. In addition, you should enable forward and reverse resolution. Host names must be consistent across the whole IBM PowerVC domain.

Important: Regardless of the host name resolution method that you use, the IBM PowerVC management host must be configured with a valid, fully qualified domain name.

3. Configure the ZYPHER repository for the IBM PowerVC installation.

To successfully install IBM PowerVC, your system must be configured and enabled to access the SUSE Linux Enterprise Server zypper repository on the host by using a mounted DVD or by being subscribed to the SUSE provided repositories. This allows IBM PowerVC to find and install all required packages during installation. For instructions for setting up the zypper repository, see *Using Zypper*.

<https://documentation.suse.com/sles/15-SP1/html/SLES-all/book-sle-deployment.html>

See the Example 3-12 shown below to list and enable the DVD mounted repository.

Example 3-12 List and enable DVD mounted repository

```
# zypper lr
Repository priorities are without effect. All enabled repositories share the same priority.

# | Alias | Name | Enabled | GPG Check | Refresh |
-----+-----+-----+-----+-----+-----+
1 | Basesystem-Module_15.2-0 | sle-module-basesystem | No | ---- | ---- |
2 | SLES15-SP2-15.2-0 | SLES15-SP2-15.2-0 | No | ---- | ---- |
3 | Server-Applications-Module_15.2-0 | sle-module-server-applications | No | ---- | ---- |

# zypper mr -e --all
Repository 'Basesystem-Module_15.2-0' has been successfully enabled.
Repository 'SLES15-SP2-15.2-0' has been successfully enabled.
Repository 'Server-Applications-Module_15.2-0' has been successfully enabled.

# zypper sl
# | Alias | Name | Enabled | GPG Check | Refresh | Type |
-----+-----+-----+-----+-----+-----+
1 | Basesystem-Module_15.2-0 | sle-module-basesystem | Yes | ( r ) Yes | No | rpm-md |
2 | SLES15-SP2-15.2-0 | SLES15-SP2-15.2-0 | Yes | ( r ) Yes | No | rpm-md |
3 | Server-Applications-Module_15.2-0 | sle-module-server-applications | Yes | ( r ) Yes | No | rpm-md |
```

4. Download the package with the correct extension for SUSE Linux Enterprise Server on IBM Power: ppc64le.

For SUSE Linux Enterprise Server installation instructions, see the SUSE Linux Enterprise Server deployment guide:

<https://documentation.suse.com/sles/15-SP1/html/SLES-all/book-sle-deployment.html>

Installing other software

The management virtual machine should be dedicated to IBM PowerVC and the operating system on which it runs. However, PowerVC can generally coexist with other software on the same instance, assuming that there is no resource or dependency conflict between PowerVC and the other software. You need to verify and prevent any potential conflicts such as port contention, user namespace, file system capacity, firewall settings, etc, between the IBM PowerVC and the other software. Locking any RPMs is not recommended unless directed to do so by IBM support team.

Do not install other software on the host before installing PowerVC as it may cause a conflict (important packages, necessary for the successful PowerVC installation may be left out of the process).

Performance implications to IBM PowerVC and other software should be taken into consideration when installing other software on the same instance. For example, PowerVC memory and processor usage may grow and cause problems with applications that coexist with PowerVC.

Note: IBM PowerVC resource requirements (as described in section 3.1, “Setting up the IBM PowerVC environment” on page 70) are sized assuming that PowerVC is the only workload that is running on the management instance. PowerVC should be the only application running on the management system. However, if at all you plan to run other applications on the same host, plan and adjust the sizing accordingly.

3.2 Installing IBM PowerVC

This section describes how to install IBM PowerVC on your management host by using .TGZ or .ISO files that are obtained from the download site.

Before you install IBM PowerVC, ensure that all of the hardware and software prerequisites are met and that your environment is correctly configured. If you need further information on this topic, see section 2.1, “IBM PowerVC requirements” on page 14. Also ensure that you have already prepared the management host and have installed the supported version of Red Hat Enterprise Linux Server or SUSE Linux Enterprise Server on it, see section 3.1, “Setting up the IBM PowerVC environment” on page 70.

Important: The management VM is dedicated for IBM PowerVC and the operating system on which it runs. *Do not* install other software onto it.

To download and install IBM PowerVC, complete the following steps:

1. To begin the installation, open a web browser and go to the Entitled Software Support website:
<http://www.ibm.com/servers/eserver/ess/OpenServlet.wss>
2. Sign in with your IBM ID.
3. Click **Software downloads**.
4. Select either **By category** or **By product**.
5. If you choose **By product**, search and select the edition of IBM PowerVC that you purchased under IBM PowerVC 2.0.0 (5765-VS2) or IBM PowerVC for Private Cloud 2.0.0 (5765-VC2), and click **Add product**, next click on **Continue**.

6. Click **packages** next to the feature that you want to download. The available files are shown.
7. Download either the Red Hat Enterprise Linux Power Linux LE, the SUSE Linux Enterprise Server Power Linux LE or the Red Hat Enterprise Linux x86_64 Linux .TGZ or .ISO file.

Note: If your web ID is not yet registered with a customer number, click **Register** Customer ID number. If you are the first web ID to register with your customer number, that ID becomes the primary. However, if you are not the first web ID, you are forwarded to the primary contact, who must approve your web ID.

8. After you download the .TGZ file, extract it to the location that you would like to run the installation script on.
9. Change your current directory to the destination directory (where the files were extracted).

Note: You must perform these steps as root.

10. Run the prerequisite checks and confirm that all of them are met for installation, as shown in Example 3-13. Correct any errors before you proceed with the installation.

Example 3-13 Install prerequisites verification for SUSE Linux Enterprise Server ppc64le

```
# ./install -t
Checking available memory
Checking CPU speed
Power architecture detected
Checking file size limit of the file system.
Current maximum file size : unlimited
/etc directory has the user/group as root/root required for install
/opt directory has the user/group as root/root required for install
/var directory has the user/group as root/root required for install
/tmp directory has the user/group as root/root required for install
Checking wheel group usage
Checking disk space
Required disk space for installation: 6.5 GB
15
Verifying required ports are available
Required port 5000 is available.
Required port 35357 is available.
Required port 9292 is available.
Required port 9191 is available.
Required port 8774 is available.
Required port 9000 is available.
Required port 9696 is available.
Required port 8428 is available.
Required port 5470 is available.
Required port 7870 is available.
Required port 5671 is available.
Required port 4369 is available.
Required port 50110 is available.
Required port 8080 is available.
Required port 6200 is available.
Required port 6201 is available.
Required port 6202 is available.
Required port 8998 is available.
Required port 6080 is available.
Required port 8778 is available.
```

Required port 8041 is available.
 All the required ports are available
 Cleaning yum repository files.
 Configuring yum repository files.
 Repository priorities in effect: (See 'zypper lr -P' for details)
 99 (default priority) : 6 repositories
 100 (lowered priority) : 1 repository

#	Alias	Name	Enabled	GPG Check	Refresh
1	Basesystem-Module_15.2-0	sle-module-basesystem	Yes	(r) Yes	No
2	SLES15-SP2-15.2-0	SLES15-SP2-15.2-0	Yes	(r) Yes	No
3	Server-Applications-Module_15.2-0	sle-module-server-applications	Yes	(r) Yes	No
4	powervc-noarch	IBM PowerVC noarch	Yes	(p) Yes	No
5	powervc-ppc64le	IBM PowerVC ppc64le	Yes	(p) Yes	No
6	pvc-openstack-noarch	IBM PowerVC Openstack noarch	Yes	(p) Yes	No
7	pvc-openstack-ppc64le	IBM PowerVC Openstack ppc64le	Yes	(p) Yes	No

All repositories have been cleaned up.

11. Start the installation by running the installation script. Table 3-1 shows the various options to use with the IBM PowerVC `install` command.

Table 3-1 Options for the IBM PowerVC `install` command

Option	Description
<code>-c firewall</code>	Firewall configuration will be performed during installation. Unless this option is specified, admin will need to manually configure the firewall per documentation.
<code>-s <offering></code>	Run a silent installation and requires the offering value to be set to <i>standard</i> or <i>private_cloud</i> .
<code>-t</code>	Runs the prerequisite checks and exit.
<code>-u</code>	Uninstall to attempt cleanup of failed installation and exit.
<code>-f</code>	Force install to override or bypass certain checks. Used with the <code>uninstall</code> option to bypass failures during uninstall.
<code>-h</code>	Displays the help messages and exit.

12. Select the IBM PowerVC offering to install from the following two options shown in Example 3-14.

Example 3-14 IBM PowerVC 2.0.0 installation

```
#####
Starting the IBM PowerVC 2.0.0.0 installation on:
2021-02-12T15:59:45-05:00
#####

Select the edition to install:
 1 - IBM PowerVC
 2 - IBM PowerVC for Private Cloud
 9 - Exit
```

13. Accept the Licence Agreement.
14. IBM PowerVC prompts you with a list of files that would be modified during the installation.

15. Select whether the IBM PowerVC installer should configure the host firewall during the installation.

Note: If the firewall is not automatically configured during the installation of IBM PowerVC, it must be configured manually after the installation. For details on the required inbound and outbound ports, see IBM Knowledge Center:

https://www.ibm.com/support/knowledgecenter/en/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_planning_security_firewall_hmc.html

16. PowerVC prompts you with a list of files that would be modified during the installation.

17. Confirm and choose to continue with the procedure.

18. After the installation completes, you see a message similar to Example 3-15 shown below. Ensure that you download and install any fix packs that are available from Fix Central, as described in section 3.4, “Upgrading IBM PowerVC” on page 87.

Example 3-15 Installation completed

The validation of IBM PowerVC services post install was successful.

```
*****
IBM PowerVC installation successfully completed at 2021-02-12T17:34:02-05:00.
Refer to
/opt/ibm/powervc/log/powervc_install_2021-02-12-171337.log
for more details.
```

Use a web browser to access IBM PowerVC at

https://powervc_ip_address

```
*****
```

Note: The installation logs can be found under `/opt/ibm/powervc/log/`.

If the installation does not complete successfully, run the following command to remove the files that were created during the failed installation before you reinstall IBM PowerVC:

```
[IBM PowerVC_install_file_folder] ./install -u -f
```

Note: Use this command only to remove files due to a failed installation. If you need to uninstall a working instance of IBM PowerVC, run the `uninstall` command. For more information, see section 3.3, “Uninstalling IBM PowerVC” on page 86.

The installation script creates the required IBM PowerVC users and groups, as shown in Table 3-2.

Table 3-2 IBM PowerVC users

User	Description
mysql	MariaDB Server user
apache	Apache Server user
saslauth	Saslauthd user
radvd	radvd user

User	Description
unbound	Unbound DNS resolver
memcached	Memcached daemon
egoadmin	IBM EGO Daemons
epmd	Erlang Port Mapper Daemon
swift	OpenStack Swift Daemons
cinder	OpenStack Cinder Daemons
neutron	OpenStack Neutron Daemons
nova	OpenStack Nova Daemons
ceilometer	OpenStack ceilometer Daemons
glance	OpenStack Glance Daemons
keystone	OpenStack Keystone Daemon
rabbitmq	RabbitMQ messaging server
ttv-validation	IBM OpenStack Validator Daemons
clerk	IBM OpenStack SSK Daemons
mongod	Used to run the MongoDB service.
gnocchi	Use to run gnocchi service
panko	Use to run Panko service (Event, Metadata Indexing Service)
dnsmasq	Created as part of dependencies
openvswitch	Created as part of dependencies

After the installation completes, you can verify the status of IBM PowerVC services by running the **powervc-services** command, as shown in Example 3-16.

Example 3-16 The powervc-services command example

```
# powervc-services httpd status

? httpd.service - The Apache HTTP Server
  Active: active (running) since Fri 2021-02-12 09:39:43 EST; 20h ago
? memcached.service - memcached daemon
  Active: active (running) since Fri 2021-02-12 09:39:43 EST; 20h ago

# powervc-services cinder status

? openstack-cinder-api.service - OpenStack Cinder API Server
  Active: active (running) since Fri 2021-02-12 09:39:44 EST; 20h ago
? openstack-cinder-scheduler.service - OpenStack Cinder Scheduler Server
  Active: active (running) since Fri 2021-02-12 09:39:44 EST; 20h ago
? openstack-cinder-backup.service - OpenStack Cinder Backup Server
  Active: active (running) since Fri 2021-02-12 09:39:44 EST; 20h ago
? openstack-cinder-volume-9_114_181_189.service - OpenStack Cinder Volume Server
  Active: active (running) since Fri 2021-02-12 09:39:44 EST; 20h ago
```

? powervc-cinder-conductor.service - PowerVC Cinder Conductor Server
Active: active (running) since Fri 2021-02-12 09:39:44 EST; 20h ago

The IBM PowerVC commands are located in the /usr/bin and /opt/ibm/powervc/bin directories.

3.3 Uninstalling IBM PowerVC

The procedure to remove IBM PowerVC from the management host is described in this section. The uninstall process of IBM PowerVC does not remove or change anything in the environment that is managed by IBM PowerVC. Objects that were created with IBM PowerVC (VM, volumes, and so on) are left untouched by this process. Any Red Hat Enterprise Linux or SUSE Linux Enterprise Server prerequisite packages that are installed during the IBM PowerVC installation remain installed.

To uninstall IBM PowerVC, run the following command:

```
/opt/ibm/powervc/bin/powervc-uninstall
```

Example 3-17 shows the last few output lines of the uninstall process.

Example 3-17 Successful uninstall process

```
IBM PowerVC uninstallation

Are you sure you want to uninstall? [1-Yes or 2-No]1

Determining uninstall package list.

[/]

Preparing packages...
Done with cleanup actions.

Stopping all PowerVC services...
Shut down LIM on <powervc09> ..... done

Uninstalling packages.

Preparing... #####
Cleaning up / removing... #####
ibmpowervc-2.0.0-1 #####
python3-openstackclient-5.2.0-1.ibm.el#####
ibmpowervc-monitor-2.0.0-1 #####

.....

NOTE: The IBM PowerVC iptables settings have not been removed during uninstallation. The
administrator should perform the necessary steps to modify the firewall to the desired
settings.
Refer to the Knowledge Center topic for firewall ports for necessary settings to remove.

Uninstallation of IBM PowerVC completed.
#####
#####
Ending IBM PowerVC Uninstallation on:
2021-02-13T14:59:52-05:00
#####
```

```
#####  
Uninstallation was logged in /var/log/powervc-uninstall.log
```

The uninstall process writes its log to the `/var/log/powervc-uninstall.log` file.

If you encounter issues when running the `powervc-uninstall` command, you can clean up the environment by using the following command:

```
/opt/ibm/powervc/bin/powervc-uninstall -f
```

This command forces the uninstall of all components of IBM PowerVC. For the complete list of available options with the `powervc-uninstall` command, see Table 3-3.

Table 3-3 Available options for the powervc-uninstall command

Option	Description
-f	Forcefully removes IBM PowerVC.
-l	Disables uninstallation logging. Logging is enabled by default.
-y	Uninstalls without prompting.
-s	Performs IBM PowerVC backup operation and then uninstalls IBM PowerVC. The backup files are available under: <code>/var/opt/ibm/powervc/backups</code>
-h	Displays the help message and exits.

3.4 Upgrading IBM PowerVC

You can perform an upgrade to IBM PowerVC Version 2.0.0 via the backup restore procedure on Red Hat Enterprise Linux 8.2 (ppc64le, x86_64), Red Hat Enterprise Linux 8.3 (ppc64le, x86_64), SUSE Linux Enterprise Server 15 SP1 and SP2 (ppc64le) from IBM PowerVC Versions 1.4.4.0, 1.4.4.1 or 1.4.4.2 (ppc64le or x86_64) installed on your system.

Upgrades to IBM PowerVC Version 2.0.0 from earlier versions or from different host systems are described in details in 3.4.3, “Upgrading or migrating from earlier versions” on page 95.

The currently installed version and offering can be verified either via the web interface or a IBM PowerVC host command line.

To check the version via the IBM PowerVC management web interface, click **User** icon on upper left corner, as shown in Figure 3-6.

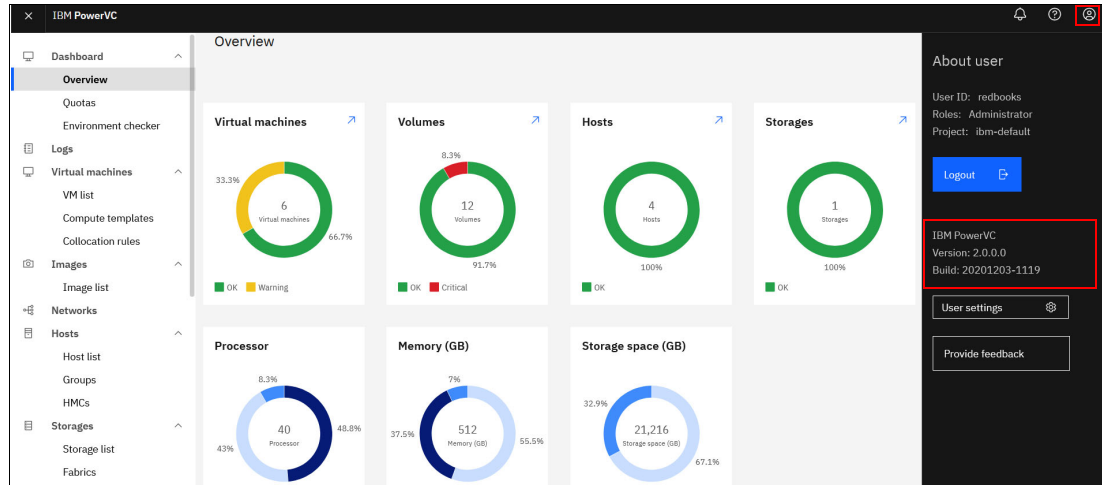


Figure 3-6 Checking the IBM PowerVC version via the graphical user interface

The installed IBM PowerVC version and offering type can be checked via the command-line interface on the IBM PowerVC host, as shown in Example 3-18.

Example 3-18 Checking the IBM PowerVC version via the command-line interface.

```
# cat /opt/ibm/powervc/version.properties
[2.0.0.0 Install]
name = IBM PowerVC
version = 2.0.0.0
build = 20201203-1119
oem = no
install-date = 2021-02-09
offering = standard
cloud_enabled = yes
```

3.4.1 Before you begin

Before you begin the software upgrade, complete the following tasks:

- ▶ Review the hardware and software requirements for IBM PowerVC Version 2.0.0, for more details see section “IBM PowerVC requirements” on page 14.
- ▶ Ensure that all compute and storage hosts are running before you start the upgrade.
- ▶ Verify your environment before you start the backup and restore process to ensure that the resources are properly configured and free from environmental issues, for more details refer to “Environment verification” on page 156.
- ▶ Ensure that no tasks (such as resizing, migration, or deployment) are running on the VM when you start the upgrade. Any tasks that are running on the VM during the upgrade might cause the VM to enter an error state after the upgrade is complete.
- ▶ If CA signed certificate is used, ensure that you manually copy any customized `powervc.crt` and `powervc.key` files from the previous version of IBM PowerVC to IBM PowerVC Version 2.0.0.

- ▶ Any operating system users from the Administrator, Deployer, or Viewer groups on the previous version must be added again to the groups on the host operating system.
- ▶ If PowerVM NovaLink is used, upgrade the host operating system (Red Hat Enterprise Linux or Ubuntu) to a supported version, next upgrade the pvm-novalink packages on the PowerVM NovaLink hosts to Version 2.0.0, for additional details see the link below:

https://www.ibm.com/support/knowledgecenter/POWER9/p9eig/p9eig_upgrade_novalink.htm

3.4.2 Performing an in-place upgrade

An IBM PowerVC host (ppc64le or x86_64) running IBM PowerVC Versions 1.4.4.0, 1.4.4.1 or 1.4.4.2 or later can be upgraded to IBM PowerVC Version 2.0.0 via backup restore procedure. The upgrade process is handled with a custom **backup_n_uninstall_144.sh** script available with IBM PowerVC 2.0.0 (.TGZ or .ISO) installation files or can be downloaded from IBM cloud repo:

<http://public.dhe.ibm.com/systems/virtualization/powervc/>

The custom script performs backup and uninstall, in addition the smooth upgrade process is ensured by the following additional operations on the backup system:

- ▶ A backup of the configuration files on the NovaLink and the HMC managed hosts is created. The configuration files are restored on PowerVC 2.0.0 system.
- ▶ All PowerVC packages that block leapp upgrade are removed, if you plan to perform an in-place upgrade on existing PowerVC instance to Red Hat Enterprise Linux 8.2 or Red Hat Enterprise Linux 8.3 (ppc64le or x86_64). The detailed backup procedure is explained below for each PowerVC version.

Note: Do not use the regular powervc-backup script for the upgrade process.

Backup procedure for PowerVC Version 1.4.4.x

Before you begin the backup process of PowerVC Versions 1.4.4.x, complete the following tasks:

1. Download all files available in custom_backup_and_uninstall_script_for_PowerVC144 folder from the IBM cloud repository location given below, and place them in your 1.4.4.0, 1.4.4.1, or 1.4.4.2 installed PowerVC management system (ppc64le or x86_64):

<http://public.dhe.ibm.com/systems/virtualization/powervc/>

Note: Custom **backup_n_uninstall_144.sh** script is available with IBM PowerVC 2.0.0 (.TGZ or .ISO), installation files. The custom script is in the `/powervc-2.0.0.0/lib/` directory.

2. Create a backup of all existing scripts **backup_n_uninstall_144.sh** and **backup_144.py** to a directory on 1.4.4.x management system. To ensure smooth upgrade, you must use custom scripts provided for upgrade instead of using old backup scripts.
3. Make sure that all compute nodes and storage hosts are operational before running the custom **backup_n_uninstall_144.sh** script.
4. Remove the network node from backup system (PowerVC Version 1.4.4.x) as PowerVC Version 2.0.0 does not support network nodes.

5. The newly created backup file must be transferred to the PowerVC 2.0.0 for the restore. The filename will have the PowerVC version and timestamp, for example, `powervc_backup_1.4.4.x_<timestamp>.tar.gz`.
6. The use of a custom backup and uninstall script will expedite your work.

Check the file permissions and make sure that the script is executable, next run the backup and uninstall custom script. Table 3-4 shows the available options to use with the IBM PowerVC custom `backup_n_uninstall_144.sh` script.

Table 3-4 Options for the custom backup script

Option	Description
<code>backup_n_uninstall_144.sh</code>	Run backup first then uninstall.
<code>backup_n_uninstall_144.sh -s</code>	Backup only and not to uninstall PowerVC.
<code>backup_n_uninstall_144.sh -f</code>	Run only uninstall of PowerVC.

Run the script with `-s` option to do the backup only, without uninstalling the IBM PowerVC, as shown in Example 3-19.

Example 3-19 Creating a backup without uninstalling the PowerVC instance

```
# ./backup_n_uninstall_144.sh -s

##### The purpose of this script to support IBM PowerVC upgrade from version
144x to 2.0 #####

Attempting to run backup on IBM PowerVC.

Continuing with this operation will restart some PowerVC services. Do you want
to continue? (y/N):y
Stopping services: panko, swift, gnocchi, validator, clerk, health, ceilometer,
nova, neutron, ego, cinder, glance ...
Stopping services: bumblebee ...
Stopping services: httpd ...
Stopping services: db ...
Backing up the databases and data files...
Database and file backup completed. Backup data is in archive
/var/opt/ibm/powervc/backups/20210216004647239370/powervc_backup.tar.gz
Starting services: db ...
Starting services: httpd ...
Starting services: bumblebee ...
Starting services: glance, cinder, ego, neutron, nova, ceilometer, health,
clerk, validator, gnocchi, swift, panko ...
PowerVC backup completed successfully.
INFO: Validating remote host information
INFO: No novalink host registered, skipping novalink conf collection

+-----+
| Summary of the opts you specified |
+-----+
Archive: powervc_backup_1.4.4.2_210216-004854.tar.gz has novalinkinfo dir added
which has registered novalink backup confs
Re archived powervc_backup.tar.gz to
powervc_backup_1.4.4.2_210216-004854.tar.gz at location
/var/opt/ibm/powervc/backups/20210216004647239370
```

```
Output directory      :  
/var/opt/ibm/powervc/backups/20210216004647239370  
Archive filename     : powervc_backup_1.4.4.2_210216-004854.tar.gz  
Interactive mode     : True  
Backup completed to the following location:  
/var/opt/ibm/powervc/backups/20210216004647239370/
```

PowerVC Version 2.0.0 restore procedure

PowerVC 2.0.0 must be installed (based on the instructions provided in 3.2, “Installing IBM PowerVC” on page 81) before you start the restore process. To restore the taken backup from IBM PowerVC Version 1.4.4.x instance to the IBM PowerVC Version 2.0.0 instance use the **powervc-restore** command.

If you use a local OS registry and not the LDAP server as a back-end authentication on PowerVC Version 1.4.4.x instance, make sure all local OS users are re-created on the new IBM PowerVC Version 2.0.0 instance as well (before initiating the restore procedure). To prevent any user login issues after the restore, verify that all users are members of **powervc-filter** group.

Before starting with the restore procedure, ensure that the target PowerVC 2.0.0 system is properly configured and all PowerVC services are in a running state. Perform the following steps to restore the backup from PowerVC Version 1.4.4.x to PowerVC Version 2.0.0:

- ▶ Copy the backup file to the target PowerVC Version 2.0.0 instance.
- ▶ Run **powervc-restore** command to restore PowerVC from backup. By default, the **restore** command looks for backup archives in `/var/opt/powervc/backups` directory. You can specify a different directory with `--targetdir` option if a backup is available at different path, see command below:

```
# powervc-restore --targetdir=</tmp/backups>
```

This procedure updates only the managed nodes, if you have PowerVM NovaLink hosts managed by IBM PowerVC make sure to update the NovaLink hosts to a supported version (PowerVM NovaLink 2.0.0), see section 2.2, “IBM PowerVM NovaLink requirements” on page 19.

To upgrade IBM PowerVC to Version 2.0.0 use one of the two available methods:

- ▶ Create a backup on PowerVC Version 1.4.4.x and restore it in version 2.0.0. Use this method to take the backup in a PowerVC Version 1.4.4 system with the custom script **backup_n_uninstall_144.sh** and then restore it on a PowerVC Version 2.0.0 system. This method is supported on SUSE Linux Enterprise Server and Red Hat Enterprise Linux operating systems.
- ▶ In-place upgrade is available only on Red Hat Enterprise Linux operating system. To perform an in-place upgrade, complete the following steps:
 - a. Download and check the file permission of the custom script. Ensure the script is executable:
 - i. Run the **backup_n_uninstall_144.sh** script to removes all packages that could lead to failure of the *leapp* upgrade and to create a backup on PowerVC Version 1.4.4.x.

Note: You can choose to run the custom script in steps, first with `-s` option to perform only a backup and next with `-f` option to uninstall PowerVC and all packages that could lead to failure of the *leapp* upgrade. See the previous section for usage and available options of **backup_n_uninstall_144.sh** script.

- b. Upgrade the operating system to Red Hat Enterprise Linux 8.2 or 8.3 (ppc64le or x86_64, depending on your system architecture).
- c. Once the in-place upgrade completes, you can obtain the IBM PowerVC Version 2.0.0 installation package for the current architecture.
- d. Extract the installation package to the local file system of the IBM PowerVC host.
- e. Go to the directory that contains the extracted installation script:
./install
- f. Start the upgrade process, as shown in Example 3-20.

Note: During the upgrade to IBM PowerVC Version 2.0.0, the IBM PowerVC installation can be upgraded to IBM PowerVC for Private Cloud offering or kept as the IBM PowerVC offering.

Example 3-20 IBM PowerVC upgrade to Version 2.0.0

```
# ./install
#####
Starting the IBM PowerVC 2.0.0.0 installation on:
2021-02-16T09:02:45-05:00
#####

Select the edition to install:
  1 - IBM PowerVC
  2 - IBM PowerVC for Private Cloud
  9 - Exit
2
LOG file is /opt/ibm/powervc/log/powervc_install_2021-02-16-090245.log

Checking disk space
Checking file size limit of the file system.
Current maximum file size : unlimited

International Program License Agreement

Part 1 - General Terms

BY DOWNLOADING, INSTALLING, COPYING, ACCESSING, CLICKING ON
AN "ACCEPT" BUTTON, OR OTHERWISE USING THE PROGRAM,
LICENSEE AGREES TO THE TERMS OF THIS AGREEMENT. IF YOU ARE
ACCEPTING THESE TERMS ON BEHALF OF LICENSEE, YOU REPRESENT
AND WARRANT THAT YOU HAVE FULL AUTHORITY TO BIND LICENSEE
TO THESE TERMS. IF YOU DO NOT AGREE TO THESE TERMS,

* DO NOT DOWNLOAD, INSTALL, COPY, ACCESS, CLICK ON AN
"ACCEPT" BUTTON, OR USE THE PROGRAM; AND

* PROMPTLY RETURN THE UNUSED MEDIA, DOCUMENTATION, AND

Press Enter to continue viewing the license agreement, or
enter "1" to accept the agreement, "2" to decline it, "3"
to print it, "4" to read non-IBM terms, or "99" to go back
to the previous screen.
1

Do you want the IBM PowerVC setup to configure the firewall? 1-Yes or 2-No?
1
```

NOTE: The following list of system configurations will be modified as part of installation.
/etc/sudoers
/etc/httpd/conf.d/ssl.conf
/etc/sysctl.conf
/etc/httpd/conf/httpd.conf

Continue with the installation: 1-Yes or 2-No?
1

Checking system prerequisites.
Installation task 1 of 7

Checking available memory
Checking CPU speed
Checking file size limit of the file system.
Current maximum file size : unlimited
/etc directory has the user/group as root/root required for install
/opt directory has the user/group as root/root required for install
/var directory has the user/group as root/root required for install
/tmp directory has the user/group as root/root required for install
Checking wheel group usage
Checking disk space
Required disk space for installation: 6.5 GB
Verifying required ports are available
.....
Installing IBM PowerVC for Private Cloud
Installation task 6 of 7

Stopping all PowerVC services...
Shut down LIM on <powervc09> done
Starting all PowerVC services...

Cleaning up temporary and support files.
Installation task 7 of 7

Done with cleanup actions.

The validation of IBM PowerVC services post install was successful.

IBM PowerVC installation successfully completed at 2021-02-16T09:03:05-05:00 Refer to
/opt/ibm/powervc/log/powervc_install_2021-02-16-090245.log
for more details.

Use a web browser to access IBM PowerVC at

https://powervc__ip_address

-
7. During the upgrade to IBM PowerVC Version 2.0.0, if a selection was made to enable the IBM PowerVC for Private Cloud edition, the offering is changed from IBM PowerVC to IBM PowerVC for Private Cloud Edition. This option is available only for IBM PowerVC editions managing PowerVM hosts.
 8. After the upgrade completes successfully, run the **powervc-restore** command to automatically picks the backed-up tar.gz file generated before uninstalling PowerVC.

Updating PowerVC on the NovaLink host

After the restore process completes successfully on the PowerVC 2.0.0 management system, perform the following steps to upgrade the PowerVC version on the NovaLink hosts:

1. From PowerVC Version 2.0.0, place all PowerVM hosts in maintenance mode, see section 4.5.1, “Planning for host management - NovaLink” on page 121. The hosts will be unavailable for deployments and migrations. Virtual machines on this host will have limited functions.
2. On the PowerVM NovaLink hosts, upgrade or re-install the operating system to a supported version (Red Hat Enterprise Linux 8.2, Red Hat Enterprise Linux 8.3, or Ubuntu 18.04), see section 2.2.1, “PowerVM NovaLink system requirements” on page 19. If you have changed the NovaLink credentials during OS re-installation process, or if the host is registered to PowerVC via ssh key based authentication prior to re-installation, then you must update the username and password in Edit Host Connection section of the Host details page.
3. Ensure NovaLink 2.0.0 is properly installed and all services are in active state.
4. Select and exit the NovaLink host from maintenance mode (one at a time), as shown in Figure 3-7.

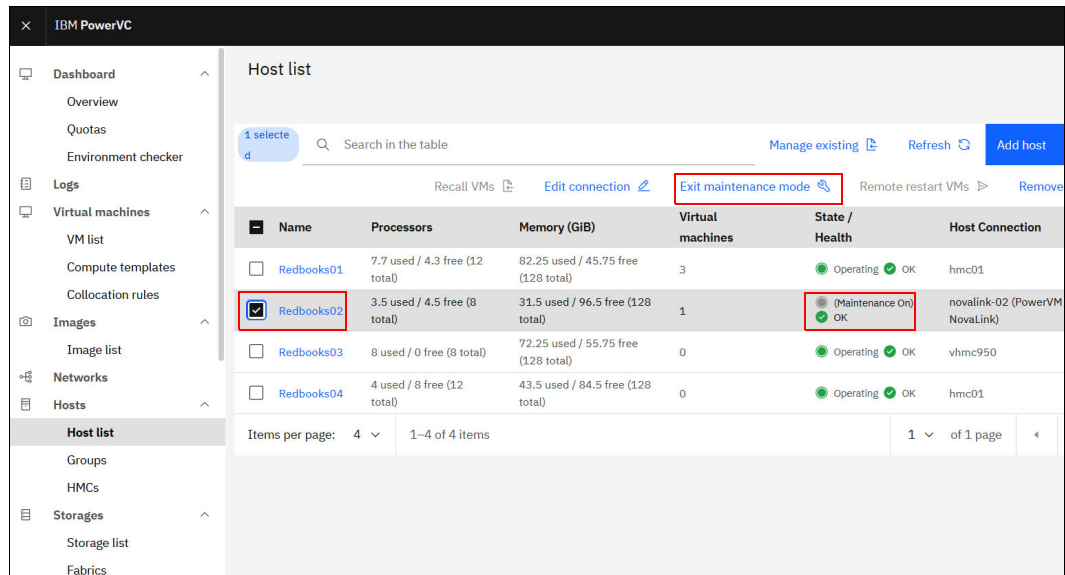


Figure 3-7 Exiting maintenance mode

5. Login to PowerVC 2.0.0 GUI, start the NovaLink update and wait until completes successfully, see Figure 3-8 on page 95. You can choose one of the following methods to update the NovaLink:
 - a. From Overview page, click **Update NovaLink** to start the update process, monitor notifications pane to display a message that NovaLink is successfully updated.
 - b. Navigate to the Hosts page, from the submenu go to Host list, choose the NovaLink managed system and click **Update NovaLink**.

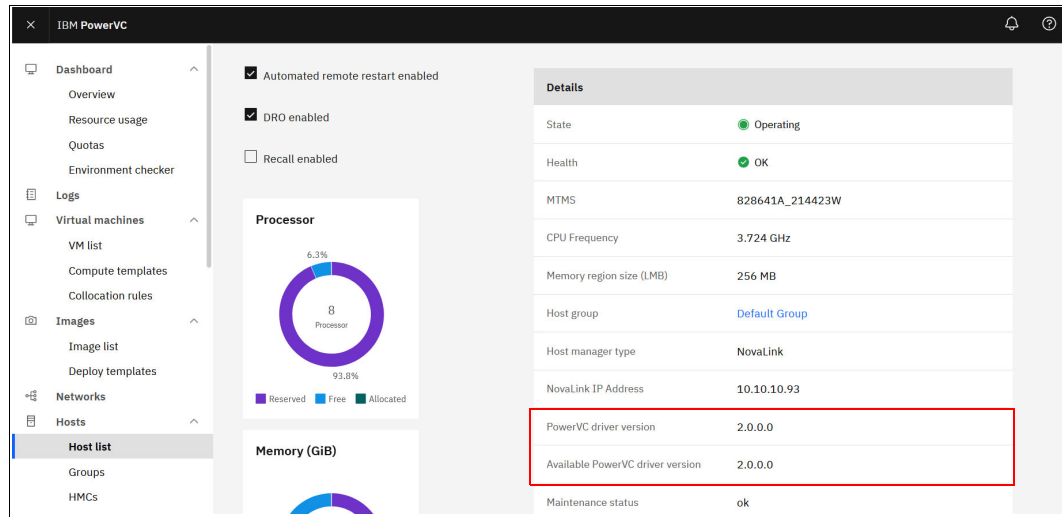


Figure 3-8 NovaLink successfully updated.

3.4.3 Upgrading or migrating from earlier versions

This section provides the methods to upgrade or migrate from versions of IBM PowerVC which are earlier than PowerVC 1.4.4.

Upgrading from earlier versions on ppc64le and x86_64 platforms

To upgrade IBM PowerVC to Version 2.0.0 from versions earlier than Version 1.4.4.x (where x is 0 or later) on the ppc64le or x86_64 platforms, the PowerVC software must be first upgraded to 1.4.4.x. The upgrade to version 1.4.4 will have to be a multi-step process if the current installed PowerVC version is 1.4.1 or lower.

Note: On the Red Hat Enterprise Linux platform, the PowerVC Version 2.0.0 only supports Red Hat Enterprise Linux versions 8.2 and 8.3. However, the Red Hat Enterprise Linux version 8.x is not supported for PowerVC Version 1.4.4. An operating system upgrade will be needed hence, when upgrading from PowerVC 1.4.4 to version 2.0.0 on the same system.

The following steps should be performed for upgrading from PowerVC versions 1.4.2 or 1.4.3 to Version 1.4.4:

1. Review the hardware and software requirements for IBM PowerVC Version 1.4.4 and PowerVC Version 2.0.0.
2. Obtain the IBM PowerVC Version 1.4.4.x (where x is 1 or later) installation package for current architecture.
3. Ensure that all the compute and storage hosts are up and running.
4. Ensure no tasks such as resizing, migrating, or deploying, are running on the virtual machine when you start the upgrade.
5. Upgrade the current IBM PowerVC installation to Version 1.4.4.x on the current system. PowerVC versions 1.4.2.x and 1.4.3.x on ppc64le or x86 platforms can be directed upgraded to PowerVC Version 1.4.4. To start the upgrade process, go to the directory which contains the extracted install script and begin the upgrade process as shown:

```
# cd <local directory>/powervc-1.4.4.x (where x is 1 or later)
# ./install
```

6. Upgrade the operating system to versions supported by PowerVC Version 2.0.0.
7. Perform an upgrade of IBM PowerVC to Version 2.0.0, as described in “Performing an in-place upgrade” on page 89.

Note: Prior to the upgrade, use the **powervc-backup** command to backup the PowerVC data. It is a best practice to always have latest backups available, especially prior to any software upgrade operation.

3.4.4 Migrating PowerVC from ppc64 to ppc64le or x86_64

Support for the PowerVC installation on ppc64 hardware is being withdrawn from PowerVC Version 1.4.4. Older PowerVC versions 1.4.2.x and 1.4.3.x on ppc64 cannot be upgraded directly to PowerVC 1.4.4. You need to take a backup of these older versions of PowerVC on ppc64 and restore on ppc64le or x86 system that has the newly installed PowerVC Version 1.4.4. The migration steps are as follows:

1. Go to the `/opt/ibm/powervc/bin` directory.
2. Use the **powervc-backup** command to take a backup of PowerVC data on ppc64. An example command is shown below, where the backup file will be written to the `/tmp/backups` directory:

```
# powervc-backup --noprompt --targetdir=/tmp/backups
```

3. Install PowerVC 1.4.4 on the ppc64le system.
4. Copy the backup file from the ppc64 to the ppc64le system.
5. Use the **powervc-restore** command from the `/opt/ibm/powervc/bin` directory, to recover the backed up PowerVC data on ppc64le. An example command is shown below:

```
# powervc-restore --noprompt --targetdir=/tmp/backups
```

The `/tmp/backups` is the directory where the backup file is located.

3.4.5 Upgrading to PowerVC for Private Cloud from PowerVC Standard Edition

The IBM PowerVC product offering is available in two editions:

- ▶ PowerVC
- ▶ PowerVC for Private Cloud

The edition is usually selected during the product installation. Refer to 3.2, “Installing IBM PowerVC” on page 81 where the detailed steps for a fresh installation are described.

You can also upgrade an existing PowerVC installation to the PowerVC for Private Cloud offering. Complete the following steps:

1. Log in to the IBM PowerVC host as the root user.
2. Go to the following folder:

```
/opt/ibm/powervc/images/cloud
```

3. Start the installation by running the `./install-cloud` command. Example 3-21 shows the output of the command.

Example 3-21 Upgrading from PowerVC to PowerVC for Private Cloud

```

./install-cloud
#####
Starting the IBM PowerVC installation on:
2021-02-12T13:58:14-05:00
#####
LOG file is
/opt/ibm/powervc/log/powervc_cloudmanager_install_2021-02-12-135814.log
....
Output truncated
....
Installed:
  ibmpowervc-clerk-2.0.0.0-1.noarch  python3-ibmpowervc-clerk-2.0.0.0-1.noarch

Complete!
Stopping services: noncore ...
Shut down LIM on <powervc04> ..... done
Starting services: noncore ...
#####
IBM PowerVC installation successfully completed at 2021-02-12T14:03:37-05:00.
Refer to
  /opt/ibm/powervc/log/powervc_cloudmanager_install_2021-02-12-135814.log
  for more details.
#####

```

3.5 Updating IBM PowerVC

IBM PowerVC updates are published on the IBM Fix Central repository. Log in with your IBM ID at the following website to get the update package:

<http://www.ibm.com/support/fixcentral>

The update can be an interim fix or fix pack. An interim fix is released to address a single product defect and hence made available as and when the identified defect is addressed. A fix pack is a collection of fixes, consisting of multiple fixes since the last release. The fix packs are generally released at an interval of 4-6 months.

Interim Fix installation

Refer to the README.txt for the Interim fix installation information. The steps for the fix pack installation may vary for each Interim fix. The README.txt is available on Fix Central along with the Interim fix download file.

Fix Pack Installation

Complete the following steps for fix pack installation:

1. Login to Fix Central and download the fix pack.

Note: Ensure the name of the maintenance file is not changed during the download process, either intentionally or inadvertently.

2. Extract the file, and run the **update** command. To extract the file, run this command:

```
tar -zxvf [location_path]/powervc-update-ppcle-rhel-version.tgz
```

This command extracts the package in the current directory and creates a directory that is named `powervc-version`.
3. Run the **update** script by running the following command:

```
/[location_path]/powervc-[version]/update
```
4. The status of the update operation is displayed on the screen upon completion. The detailed information can be viewed in the log file at the following location:

```
/opt/ibm/power/log/powervc_update_[timestamp].log
```

3.5.1 IBM Notifications

IBM provides the My Notifications facility, using which you can sign up and stay informed about IBM support updates for PowerVC and other IBM products. You can subscribe to My Notifications at the following website:

<https://www.ibm.com/support/pages/stay-date-my-notifications>

3.6 IBM PowerVC backup and recovery

This section shows the PowerVC backup capability, using which you can always save your data regularly as part of a broader system backup and recovery strategy. You may use the operating system scheduling tool to perform regular backups or any other automation tool.

Backup and recovery tasks can be performed only by using the command-line interface (CLI).

3.6.1 Backing up IBM PowerVC

Use the **powervc-backup** command to back up your essential PowerVC data. You can then restore to a working state in the event of data corruption or a disaster.

The **powervc-backup** command is located in the `/opt/ibm/powervc/bin` directory. See the command syntax and options in Table 3-5.

Table 3-5 Options for powervc-backup command

Flag	Meaning
-h, --help	Displays help information about the command.
--noprmt	If specified, no user intervention is required during execution of the backup process.
--targetdir LOCATION	Target location in which to save the backup archive. Default value is <code>/var/opt/ibm/powervc/backups</code> .
--active	If specified, the services will not be stopped and the backup will be performed as active. The backup may fail if operations are in-progress so minimize activity during the backup interval and retry the backup as necessary.

When running **powervc-backup** command, the following data is backed up:

- ▶ PowerVC databases, such as the nova database where information about your registered hosts is stored
- ▶ PowerVC configuration data, such as `/etc/nova`
- ▶ SSH private keys that are provided by the administrator
- ▶ Glance image repositories: Glance is the database name for the image repository.

Note: During the backup operation, most PowerVC services are stopped and all other users are logged off from PowerVC until the operation completes.

Steps to back up PowerVC data

To back up PowerVC, perform the following steps.

1. Ensure that the user has sufficient privileges to perform execution of this command and read access to the files.
2. Open a command-line interface to the operating system on the virtual machine on which IBM PowerVC is installed.
3. Navigate to the `/opt/ibm/powervc/bin/` directory if the **powervc-backup** command does not have a symbolic link to `/usr/bin/powervc-backup`.
4. Run the **powervc-backup** command with any needed options. If prompts are not suppressed, respond to them as needed.
5. When the backup operation completes, new file `powervc_backup.tar.gz` is placed in a new timestamped subdirectory of the target directory. You can then leave the archive there or move or copy it to another directory as needed.

Example 3-22 uses the **powervc-backup** command with no prompts and a non default mounted file system target directory.

Example 3-22 IBM PowerVC backup

```
# powervc-backup --targetdir /powervc/backups
Continuing with this operation will restart some PowerVC services. Do you want to
continue? (y/N):y
Stopping services: ui-server, gnocchi, health, ceilometer, nova, neutron, cinder, glance,
ego, panko, swift, validator ...
Stopping services: bumblebee ...
Stopping services: httpd ...
Stopping services: db ...
Backing up the databases and data files...
Database and file backup completed. Backup data is in archive
/powervc/backups/20210209141441181472/powervc_backup.tar.gz
Starting all PowerVC services...
PowerVC backup completed successfully.

# cd /powervc/backups/20210209141441181472

# ls -l
-rw----- 1 root root 24734606 Feb  9 14:16 powervc_backup.tar.gz
```

If you experienced an error while running **powervc-backup**, check for errors in the `powervc-backup` logs in the `/opt/ibm/powervc/log` directory.

Active backup of PowerVC

You can use `powervc-backup` with the `--active` option to take backup of PowerVC management server even when operations are in progress. You can consider regularly, for example on a daily basis, backing up your PowerVC data as part of a broader system backup and recovery strategy. If you want to regularly back up, use your operating system's scheduling tool to set the `powervc-backup` command to run on a schedule.

Example 3-23 demonstrates an active backup of PowerVC.

Example 3-23 Active backup of PowerVC without an interruption to services running

```
# powervc-backup --targetdir /powervc/backups --active
Backing up the databases and data files in active state...
Completed backup of database files...
Database and file backup completed.
Backup data is in archive /powervc/backups/20210209143122194270/powervc_backup.tar.gz
PowerVC backup completed successfully.
```

3.6.2 Recovering IBM PowerVC data

Use `powervc-restore` command to recover IBM PowerVC data that was previously backed up so that you can restore a working state after a data corruption situation or disaster.

The `powervc-restore` command is located in `/opt/ibm/powervc/bin` directory. See the command syntax and options in Table 3-6.

Table 3-6 Options for powervc-restore command

Flag	Meaning
<code>-h, --help</code>	Show this help message and exit.
<code>--noprompt</code>	If specified, no user intervention is required during execution of the restore process.
<code>--targetdir LOCATION</code>	Target location where the backup archive is located. Default value is <code>/var/opt/ibm/powervc/backups/<most recent></code> .

Note: During the recovery operation, most PowerVC services are stopped and all other users are logged off from PowerVC until the operation completes.

Complete steps to recover PowerVC data

Complete the following steps to recover PowerVC data:

1. Ensure that the user has sufficient sudo privileges configured to perform execution of this command and read access to the files.
2. Open a command-line interface to the operating system on the virtual machine on which IBM PowerVC is installed.
3. Navigate to the `/opt/ibm/powervc/bin/` directory if the `powervc-restore` command is not having a symbolic link to `/usr/bin/powervc-restore`.
4. Run the `powervc-restore` command with any needed options. If prompts are not suppressed, respond to them as needed.

When the restore operation completes, PowerVC runs with all of the data from the targeted backup file.

Example 3-24 provides a sample of the **powervc-restore** command with no prompts and a **--targetdir**.

Example 3-24 IBM PowerVC restore

```
# powervc-restore --targetdir /powervc/backups/20210209141441181472
Continuing with this operation will stop some PowerVC services and overwrite critical
PowerVC data in both the database and the file system. Do you want to continue? (y/N): y
Using archive /powervc/backups/20210209141441181472/powervc_backup.tar.gz for the
restore.
Stopping PowerVC services...
Restoring the databases and data files...
Starting PowerVC services...
PowerVC restore completed successfully.
```

If you experienced an error while running **powervc-restore**, check for errors in the `/opt/ibm/powervc/log` directory.

3.6.3 Status messages during backup and recovery

In a backup operation, you can use **--active** option as shown previously in Example 3-23 on page 100. However during the normal PowerVC backup (without **--active** option) or a PowerVC restore, all IBM PowerVC processes and databases are shutdown. Any user that is working with IBM PowerVC receives the maintenance message as shown in Example 3-25.

Example 3-25 Maintenance message for logged-in users

```
The system is undergoing maintenance. All services will be shutdown during this process.
You will be logged out and redirected to a status page. This application will remain
unavailable until this process completes.
```

3.6.4 Considerations about backup and recovery

There are some considerations to note about IBM PowerVC backup and recovery:

- ▶ IBM PowerVC backup and recovery task must be part of a backup plan for the organization infrastructure.
- ▶ The IBM PowerVC backup and recovery commands save only information that relates to IBM PowerVC. As a preferred practice, save the management station operating system by using the tool that you prefer, at the same time that you back up IBM PowerVC.
- ▶ The backup archive can only be restored to a system that is running the same level of PowerVC as the system from which the backup was taken. However it can be done to a system with a different hardware architecture than the backup.
- ▶ The new environment should duplicate the original user environment, such as users and groups that are assigned roles in PowerVC. Otherwise, PowerVC users will not be available after the restore.
- ▶ When restoring PowerVC backup on another PowerVC system, make sure the system does not have any existing resources.
- ▶ Any interruption while running **powervc-restore** might cause the system to go in bad state. If that happens then reinstall PowerVC to restore the system.

3.7 IBM PowerVC command-line interface

IBM PowerVC offers a CLI to perform tasks outside of the GUI. The CLI is used mainly for maintenance and for troubleshooting problems.

Table 3-7 lists the IBM PowerVC commands that are available.

Table 3-7 IBM PowerVC available commands

Command	Description
powervc-audit-export	Allows you to export previously collected audit data.
powervc-backup	Backs up essential PowerVC data so that you can restore to a working state in the event of data corruption or a disaster. The powervc-backup --active option can be used to take a backup of PowerVC even when operations are in progress.
powervc-cloud-config	Allows an administrator to work with projects, policies, and VMs.
powervc-config	Configures the PowerVC management server. It is the only supported method for changing the PowerVC config files.
powervc-diaq	Collects diagnostic data from your IBM PowerVC installation.
powervc-encrypt	Prompts the user for a string, then encrypts the string and returns it. Use the command to encrypt passwords, tokens, and strings that are stored by IBM PowerVC.
powervc-log	Enables or disables debug log level for PowerVC. Run this command without any option to view the current settings. Use this command to enable debug level without having to restart services.
powervc-log-management	Provides a utility to view/modify the settings for log management.
powervc-register	Registers an OpenStack supported storage provider or fabric.
powervc-restore	Recovers IBM PowerVC data that was previously backed up.
powervc-services	Starts, stops, restarts, and views the status of IBM PowerVC services.
powervc-uninstall	Uninstalls IBM PowerVC from your management server or host.
powervc-image	List, export, and import images in PowerVC. You can export to an OVA package and import from several sources.
powervc-volume-image-import	Creates a deployable image by using one or more volumes.
powervc-tape-reg	Use this command to register a Fibre Channel tape driver in PowerVC. This driver enables zoning of a tape library to a virtual machine that uses NPIV during deployment.
powervc-edit-scg	Use this command to update the storage connectivity group assigned to a virtual machine.

Command	Description
<code>powervc-validate</code>	Validates that your environment meets certain hardware and software requirements.
<code>openstack</code>	Facilitates communication with OpenStack services.

3.7.1 Exporting audit data

IBM PowerVC provides auditing support for the OpenStack services. Use the `powervc-audit-export` command to export audit data to a specified file.

An *audit record* is a recording of characteristics, including user ID, time stamp, activity, and location, of each request that is made by IBM PowerVC.

Reviewing audit records is helpful when you are trying to solve problems or resolve errors. For example, if a host was deleted and you must determine the user who deleted it, the audit records show that information.

The `powervc-audit-export` command is in the `/usr/bin` directory. The syntax and options are shown in Example 3-26.

Example 3-26 Using the powervc-audit command

```
powervc-audit-export [-h] [-u <user name>] [-n <number of records>] [-o <output file>] [-f <filter file>] [-x {json,csv}]
```

Table 3-8 explains the `powervc-audit-export` command options.

Table 3-8 Options for the powervc-audit-export command

Option	Description
<code>-h, --help</code>	Displays help information about the command.
<code>-u <user name>, --user_name <user name></code>	Sets the user that requests audit data. This flag is optional. The default is the logged-in user.
<code>-p <project name>, --project_name <project name></code>	Sets the name of the project for which to authenticate and request audit data. This is an optional flag. The default is <code>ibm-default</code> .
<code>-n <number of records>, --top_n <number of records></code>	Sets the upper limit for the number of audit records to return. The request and response audit records are returned in pairs. This flag is optional.
<code>-o <output file>, --output <output file></code>	Sets the file that contains the exported audit data. This flag is optional. The default file is <code>export_audit.json</code> or <code>export_audit.csv</code> , depending on the specified output format.
<code>-f <filter file>, --filter <filter file></code>	Sets the file that contains the filter records. The format of the records is JSON. For examples of filter records, see the IBM PowerVC IBM Knowledge Center. This flag is optional.
<code>-x {text,csv}, --output_format {text,csv}</code>	The format of the exported audit data. This flag is optional. The formats are <code>text</code> (JSON format) and <code>csv</code> . If not specified, the default is <code>json</code> .

3.8 Virtual machines that are managed by IBM PowerVC

This section focuses on the best practices for the operating system on the managed VMs.

3.8.1 IBM AIX virtual machines

To install VMs when your system runs on the IBM AIX operating system, no additional setup is necessary. After the IP address is configured, an RMC connection is automatically created.

The IBM PowerVC, PowerVM, and the HMC rely on the active RMC services. When these services are down, most of the concurrent and dynamic tasks cannot be run. Make sure the RMC status is active every time you change the VM dynamically.

Note: By default, AIX does not contain SSH or SSL tools unless you do a custom install (change installation setting menu) during the AIX installation and select the SSH package.

As a preferred practice, install these tools if you want to access a managed machine other than the telnet protocol.

See below link for details about `cloud-init` on AIX:

https://www.ibm.com/support/knowledgecenter/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_install_cloudinit_aix_hmc.html

3.8.2 IBM i virtual machines

IBM PowerVC can also manage the IBM i VMs. After you add the Power Systems hosts, import the IBM i VMs. No unique requirements exist for IBM i, AIX, or Linux on Power VMs.

If you need to capture the IBM i virtual machine image later, then you must install and configure the `cloud-init` on IBM i.

Before you capture a virtual machine, install the `cloud-init` initialization package using the below procedure:

- ▶ Download and install the required license programs at:
 - <http://www-304.ibm.com/servers/eserver/ess/index.wss>
 - 5770DG1 with *ALL option
 - 5770SS1 with Option 30 and 33
 - 5770SC1 *BASE and option 1
- ▶ Use `yum` to install below required rpm packages by following the steps in:
 - <https://bitbucket.org/ibmi/opensource/src/master/docs/yum/>
 - `python3-ibm_db`
 - `python3-six`
 - `python3.ppc64`
 - `cloud-init-1.3-1.ppc64`
- ▶ Install all required PTFs, as shown in the technical document below:
 - <https://www.ibm.com/support/pages/node/1119129>

- ▶ Capture the virtual machine. See details in the below link:

https://www.ibm.com/support/knowledgecenter/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_install_cloudinit_ibmi_hmc.html

3.8.3 Linux on Power virtual machines

If you plan to use Live Partition Mobility (LPM) or Dynamic Logical Partitioning with your Linux VM, you must install the IBM Installation Toolkit, especially the Reliable Scalable Cluster Technology (RSCT) utilities and RSCT core tools. Run the following command to start the IBM Installation Toolkit installation process as shown in Example 3-27.

Example 3-27 'IBM Installation Toolkit' installation

```
# ./install
Do you want to copy the repository of IBM packages to your machine? [y/n]
y
Do you want to configure your machine to receive updates of IBM packages? [y/n]
n
IBMIT needs the ports 4234 and 8080 to be accessed remotely. Would you like to
open those ports? [y/n]
y
The licenses BSD, GPL, ILAN and MIT must be accepted. You can read their text
using the options below and then accept or decline them.
1) Read license: BSD
2) Read license: GPL
3) Read license: ILAN
4) Read license: MIT
5) I have read and accept all the licenses
6) I do not accept any of the licenses
#? 5
Configuring an installation repository for your Linux distribution
Where is the installation media to be used?
1) DVD
2) Network (HTTP or FTP)
3) Directory
4) I already have a repository configured. Skip.
5) I don't know
#? 1
Insert the DVD in the drive
Press Enter to continue
Verifying if there is a repository on DVD
Available DVD devices: /dev/sr1 /dev/sr0
Checking /dev/sr1
Adding repository configuration to repository manager
Repository successfully configured
Package ibmit4linux was successfully installed
```

After the Installation Toolkit, install `ibm-power-managed-rhel.ppc64` by running:

```
# yum install -y ibm-power-managed-rhel7.ppc64
```

Once the installation completes, check the Resource Monitoring Control (RMC) status by running `lssrc -a` command as shown in Example 3-28.

Example 3-28 Resource Monitoring Control status

```
# lssrc -a
Subsystem      Group      PID      Status
ctrmc          rsct       2279     active
```

IBM.HostRM	rsct_rm	2495	active
IBM.MgmtDomainRM	rsct_rm	2498	active
IBM.DRM	rsct_rm	2553	active
IBM.ServiceRM	rsct_rm	2618	active
ctcas	rsct		inoperative
IBM.ERRM	rsct_rm		inoperative
IBM.AuditRM	rsct_rm		inoperative
IBM.SensorRM	rsct_rm		inoperative

See the link below for details about IBM Installation Toolkit for Linux on Power:

<https://www14.software.ibm.com/support/customer/sas/f/lopdiags/installtools/home.html>

See also the link below for setting up an RMC connection:

https://www.ibm.com/support/knowledgecenter/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_configuring_vlans_hmc.html

If you plan to capture a linux virtual machine, you need to install the cloud-init package.

See the link below for installing cloud-init on Red Enterprise Linux:

https://www.ibm.com/support/knowledgecenter/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_install_cloudinit_rhel_hmc.html

See the link below for installing cloud-init on SUSE Linux Enterprise Server:

https://www.ibm.com/support/knowledgecenter/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_install_cloudinit_sles_hmc.html

See the link below for installing cloud-init on Ubuntu:

https://www.ibm.com/support/knowledgecenter/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_install_cloudinit_ubuntu_hmc.html



IBM PowerVC for managing IBM PowerVM

This chapter describes the general setup of IBM Power Virtualization Center for managing IBM PowerVM. There are two IBM PowerVC Version 2.0.0 offerings that can be installed:

- ▶ IBM PowerVC
- ▶ IBM PowerVC for Private Cloud

This chapter also describes the new graphical user interface, a key difference between the IBM PowerVC and IBM PowerVC for Private Cloud version 2.0.0 offerings compared to its predecessors.

The following sections explain the discovery or configuration of the managed objects. They describe the verification of the environment and the operations that can be performed on virtual machines (VMs) and images:

- ▶ 4.1, “IBM PowerVC graphical user interface” on page 108
- ▶ 4.2, “Introduction to the IBM PowerVC setup” on page 118
- ▶ 4.3, “Managing resources outside of IBM PowerVC” on page 118
- ▶ 4.4, “Connecting to IBM PowerVC” on page 120
- ▶ 4.5, “Host setup” on page 121
- ▶ 4.6, “Hardware Management Console management” on page 131
- ▶ 4.7, “Storage and SAN fabric setup” on page 133
- ▶ 4.8, “VIOS port tags setup” on page 136
- ▶ 4.9, “Storage connectivity group setup” on page 138
- ▶ 4.10, “Storage template setup” on page 140
- ▶ 4.11, “Storage volume setup” on page 143
- ▶ 4.12, “Network setup” on page 144
- ▶ 4.13, “Compute template setup” on page 154
- ▶ 4.14, “Environment verification” on page 156
- ▶ 4.15, “Management of virtual machines” on page 157
- ▶ 4.16, “Images” on page 210
- ▶ 4.17, “Users, groups, and roles setup” on page 247

4.1 IBM PowerVC graphical user interface

This section introduces the IBM Power Virtualization Center Version 2.0.0 new user interface that has a design based on Carbon framework. Carbon is IBM's open source design system for products and digital experiences. With the IBM Design Language as its foundation, the system consists of working code, design tools and resources, human interface guidelines, and a vibrant community of contributors. For more information on Carbon Design System, see the link below:

<https://www.carbondesignsystem.com/>

The PowerVC Version 2.0.0 GUI is developed over the MERN stack, a web development framework that stands for MongoDB, Express, ReactJS, and NodeJS. At a high level, MongoDB is used as a caching layer, which is used by NodeJS or Express to feed data to the GUI. The data loading on the GUI is much faster than any of the previous releases. However, for all the other actions excluding GET are triggered on PowerVC through PowerVC REST API through the NodeJS server. The results of the actions are synchronized on NodeJS through OpenStack generated events.

Depending on the installed IBM PowerVC offering the available functions vary. The IBM PowerVC Edition 2.0.0 offering Dashboard window is shown in Figure 4-1. The hot links are highlighted in red in the illustration.

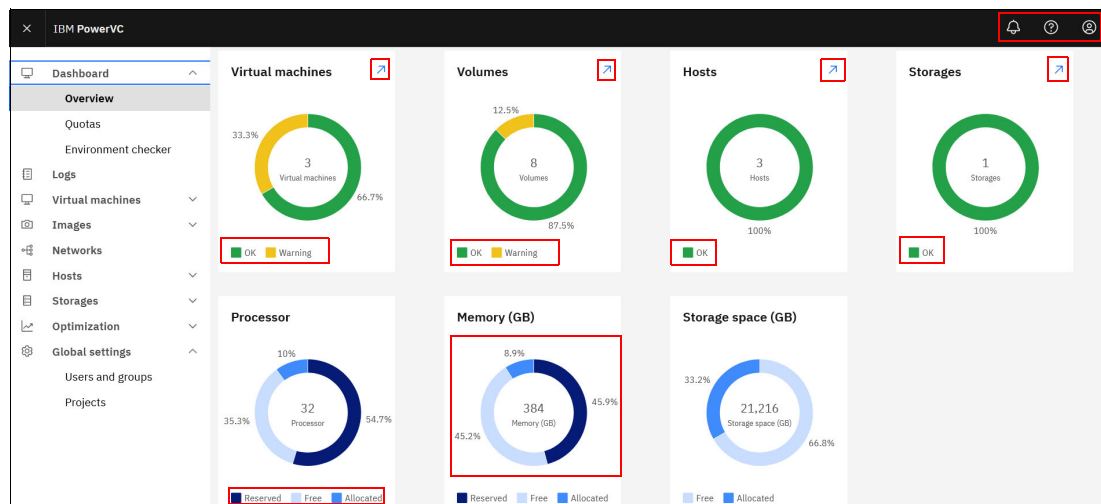


Figure 4-1 IBM PowerVC Dashboard Overview page

The Dashboard window from IBM PowerVC for Private Cloud Version 2.0.0 is shown in Figure 4-2. IBM PowerVC for Private Cloud Version 2.0.0 contains all functions available on IBM PowerVC Edition including the exclusive self-service portal (for example Quotas, Approvals, Enterprise policies, Email Server, Email template) highlighted in Figure 4-2.

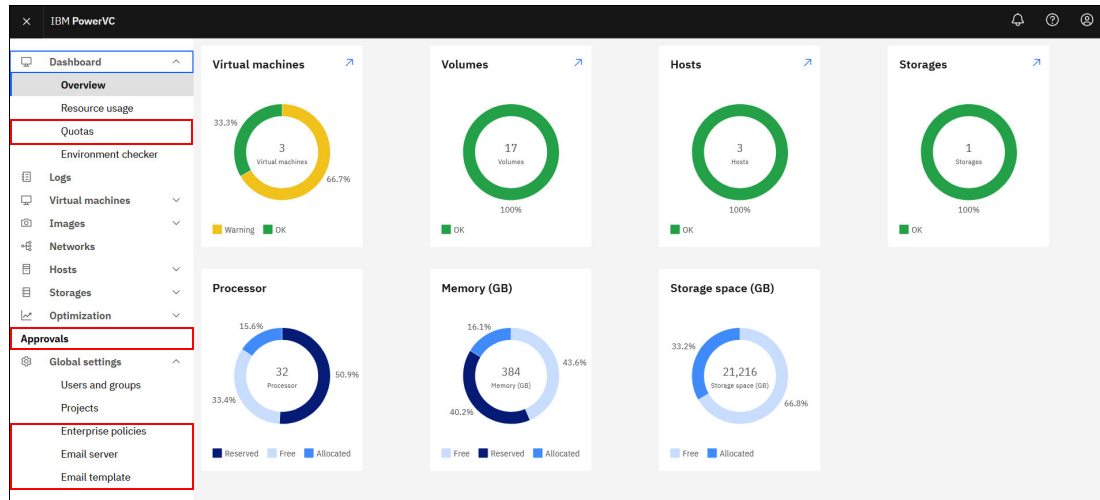


Figure 4-2 Additional functionalities with IBM PowerVC for Private Cloud 2.0.0

Functional highlights

The new PowerVC Version 2.0.0 user interface provides an ease-of-work in terms of the behavioral patterns. The following is a glimpse of the functional highlights:

- ▶ The new user interface is developed asynchronous in most places. When you trigger any resource action, the UI notifies that the action is being processed as shown in Figure 4-3. When an action is being processed, you can access other features without getting blocked.
 - a. When you initiate an Add host or Remove host operation, you never get blocked by the spinning wheel (indicating progress of action). The action gets triggered on PowerVC and the task progress can be seen on the Alerts pane. The Alerts pane is a new feature on the UI that displays top 20 log messages in a chronological order.

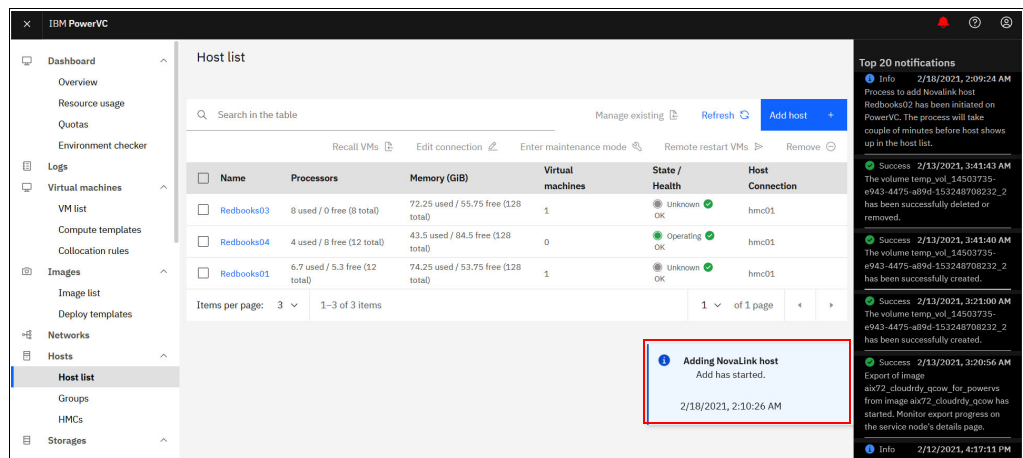


Figure 4-3 PowerVC 2.0.0 GUI notifications

- b. Additionally, for any action on a virtual machine, the VM list displays the running task *state* on the VM list page (apart from getting alerts on the Alerts pane), see Figure 4-4. This behavior is same as in the previous version of the UI.

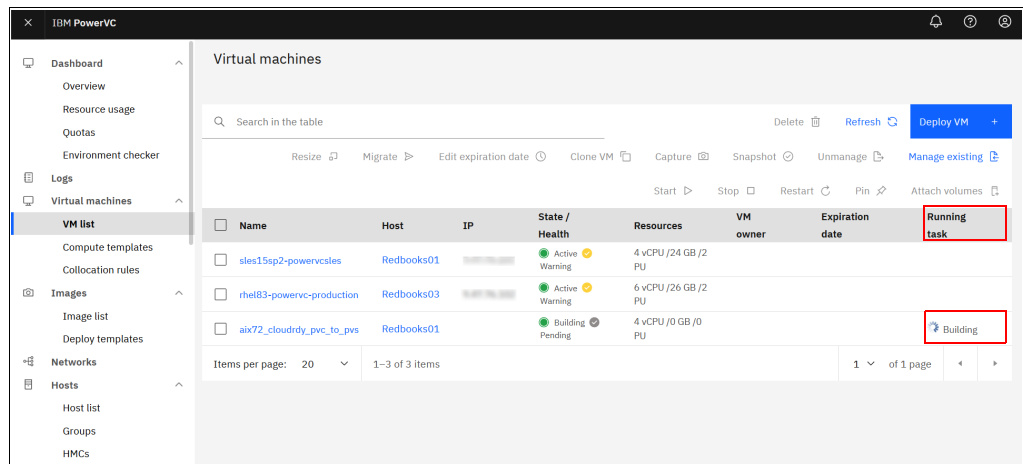


Figure 4-4 VM list displays the running task

- c. At certain places, the actions are made synchronous. For example, when you add a storage provider, it is necessary to connect to the storage provider to get a list of storage pools. So, adding a storage provider is synchronous as its 2-step process (Figure 4-5).

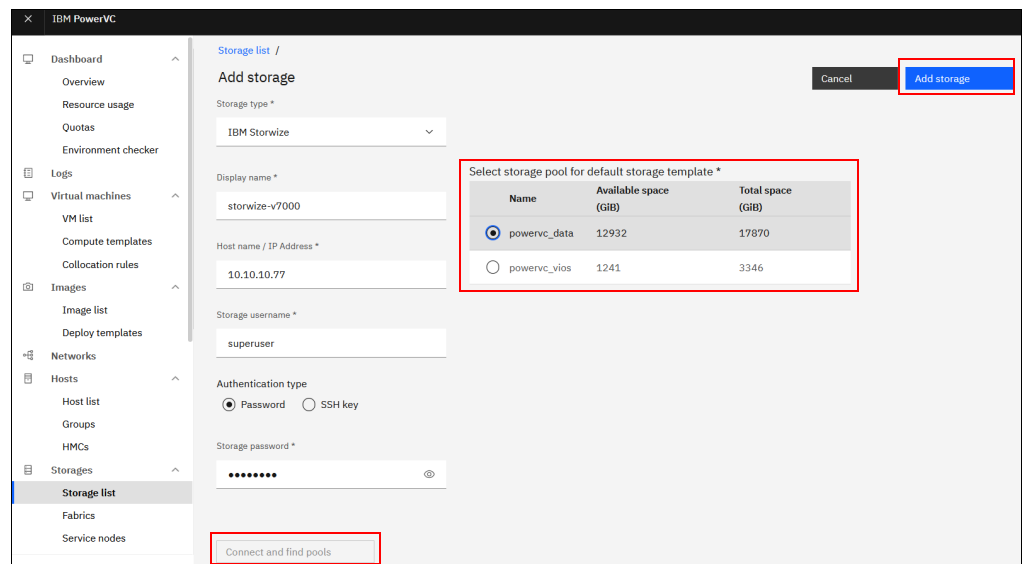


Figure 4-5 2-step add storage dialog

- ▶ On the UI, now you can do both, view details and edit attributes of a resource. For example, when you click a VM name, not only can you view the details, but also edit the details. The editable attributes are segregated from non-editable attributes. Figure 4-14 on page 115 shows the change of multiple attributes at the same time with a single click on the **Save** button.
- ▶ All edit pages of a resource are mostly a form. Instead of waiting for the auto-refresh, you can now Save the details or perform another action. See Figure 4-14 on page 115.

- ▶ The edit page displays tabs based on the resource type such as VM or Host. The tab displays the details of the associated resources. For example, a compute host edit page has a tab for showing VMs that are running on the host. On the tab, you can Start or Stop a VM. However, since the tabs are part of the edit page as a whole, the PowerVC Version 2.0.0 auto-refresh is not provided for the tabs. However, click **Refresh** to view the updated information on the tab (Figure 4-6).

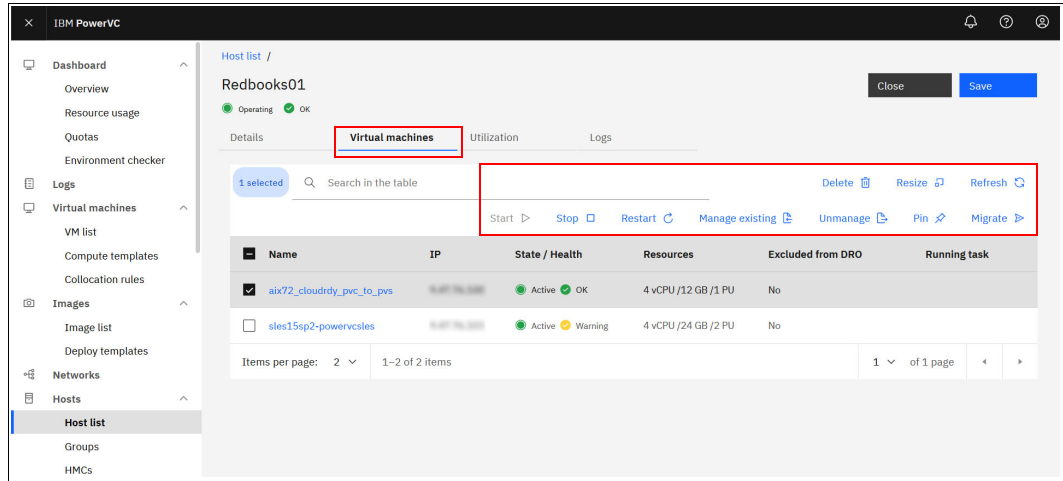


Figure 4-6 Virtual machines tab at the Host edit page

Prominent features

PowerVC Version 2.0.0 has these new major features:

- ▶ All menu items are organized under the left navigation. The menu items such as Compute templates, Storage templates, DRO events, which were on the top menu in previous PowerVC releases, can now be found under either the main menu or submenu in the left navigation bar. Additionally, the left navigation bar can be expanded as shown in Figure 4-7 or collapsed as shown in Figure 4-8 on page 112 from the UI shell.

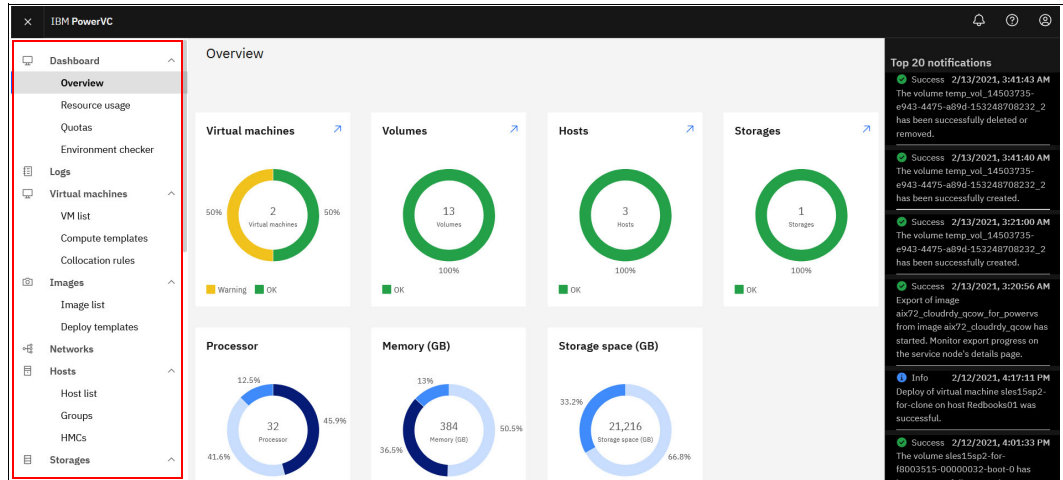


Figure 4-7 Left navigation bar - expanded

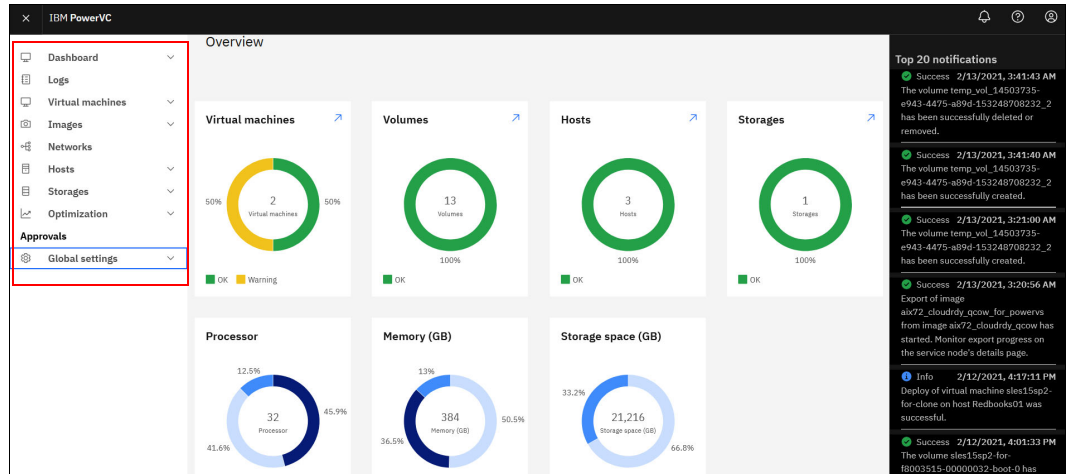


Figure 4-8 Left navigation bar - collapsed

- ▶ On the Logs page, you have the option to download the messages' log as a .csv file by clicking the **Download** button (Figure 4-9).

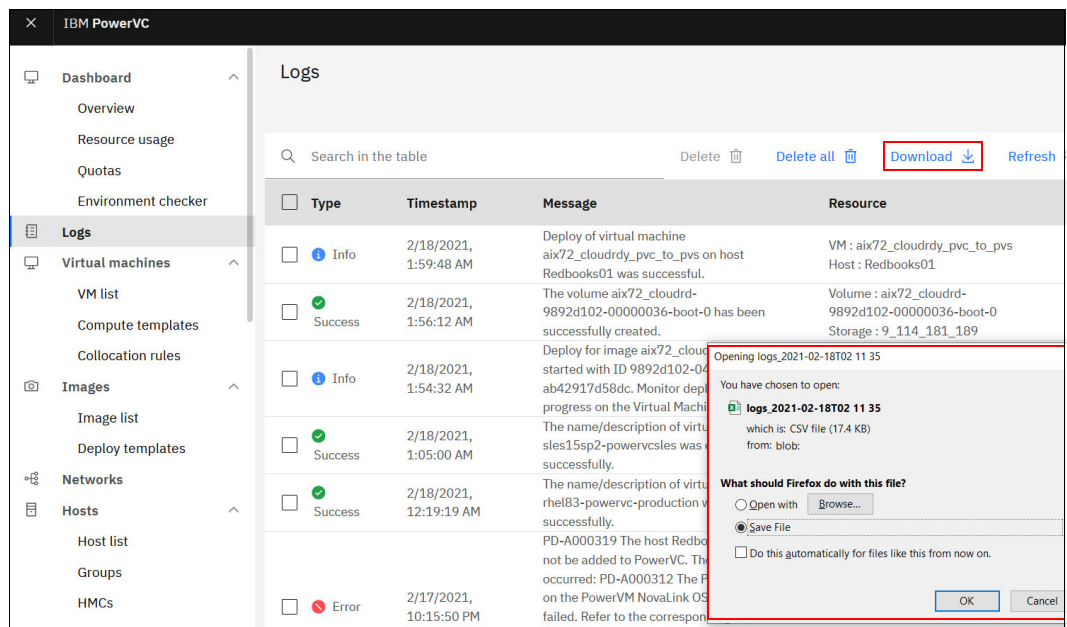


Figure 4-9 Download and export the logged messages.

- ▶ The Alerts pane is identified by a Bell icon on the upper right-hand corner of the UI. Click this icon to view the latest 20 alerts, as shown in Figure 4-10 on page 113. The Alerts pane can be kept open to view alerts pulled from PowerVC server. When the Alerts pane is closed and there are newly reported alerts, then the bell icon turns red. You can click the bell icon and review the new alerts.

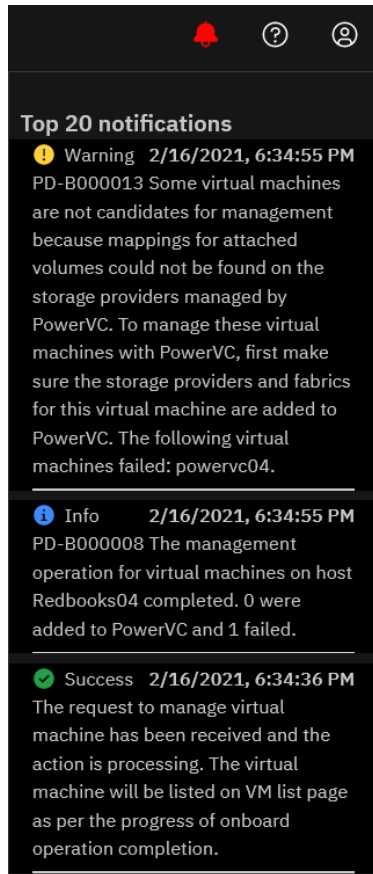


Figure 4-10 Alerts pane

- ▶ As soon as you log into the PowerVC UI, you are navigated to the Overview sub-menu. It shows utilization statistics such as virtual machines, volumes, hosts, storage, processor, to name a few. Also, the UI displays additional graphs with respect to processor, memory, and storage space as shown in Figure 4-11.

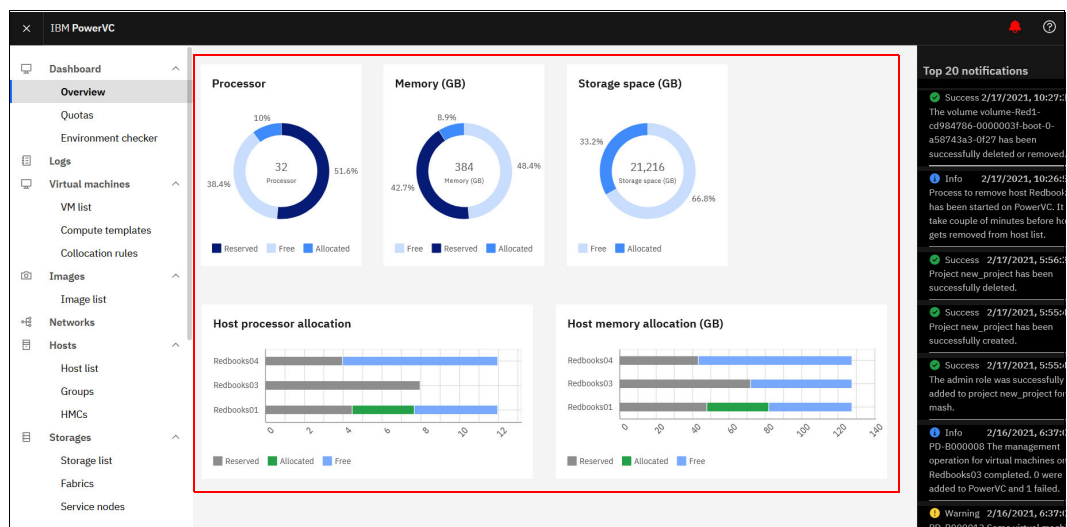


Figure 4-11 PowerVC Version 2.0.0.0 Dashboard

- ▶ An Inline help pane is identified by a Question mark icon on upper middle side of the UI as shown in Figure 4-12. Click the icon to view the context-sensitive help with respect to different pages of UI. Three separate UI places can be used to show various help topics:
 - The right pane, see Figure 4-12
 - The pop-up level help shown in Figure 4-13
 - Tooltip and messages are shown in Figure 4-12

In addition to these, the inline help pane also has links to PowerVC IBM Knowledge Center and YouTube channel. If preferred, you can also open the API References topics.

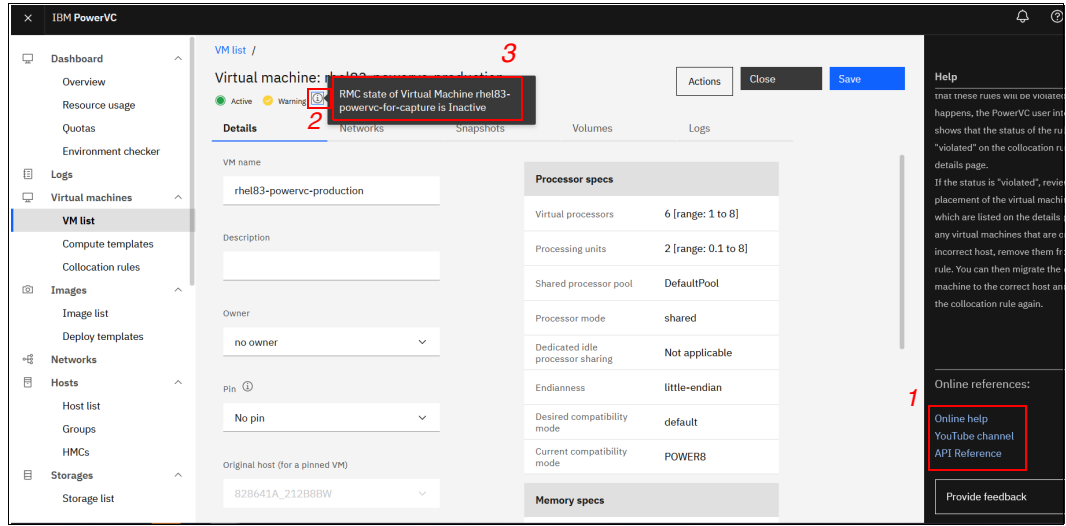


Figure 4-12 Inline help pane and tooltip

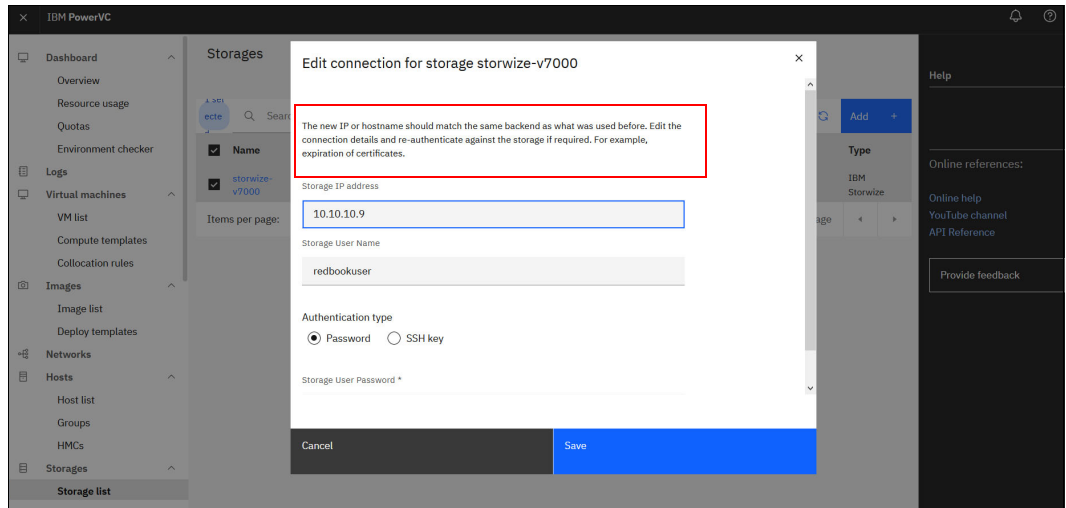


Figure 4-13 Pop-up level help

- ▶ From the resource edit page, you can change multiple attributes at the same time with a single click on the **Save** button, all changes take effect instantly, see Figure 4-14 on page 115.

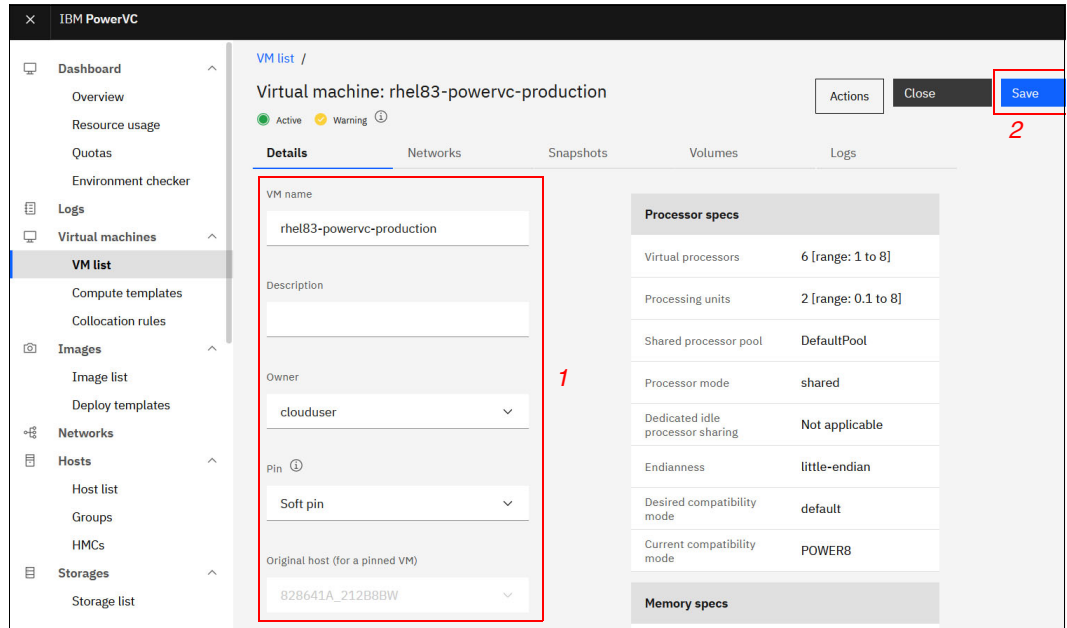


Figure 4-14 Changing multiple attributes at the same time

- In PowerVC Version 2.0.0, you can choose to deploy a VM either from the VM list page or from the Image list page. See Figure 4-15.

Note: In the previous versions of PowerVC, you could only deploy a VM from the Images page.

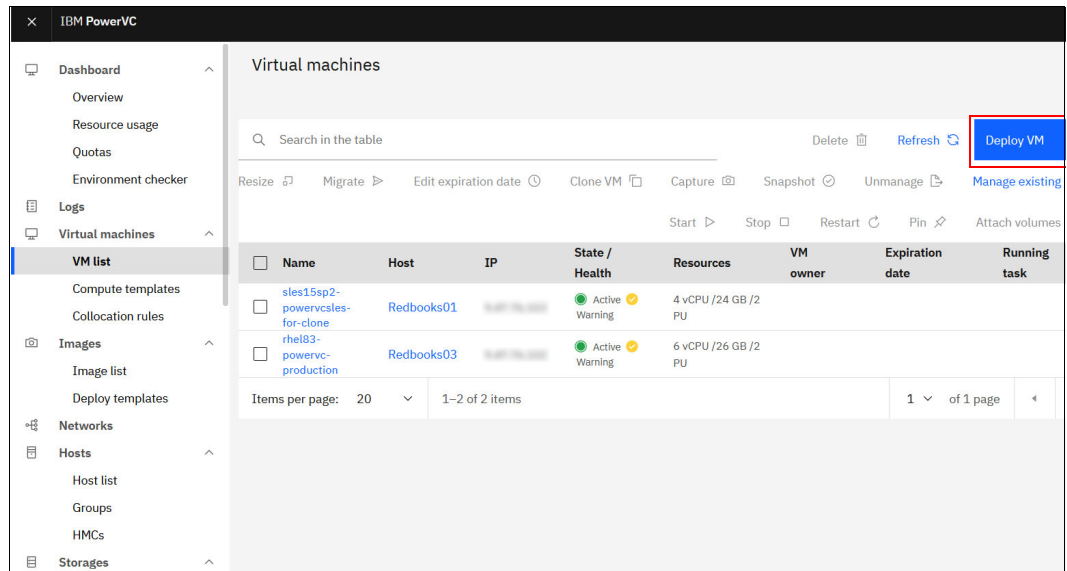


Figure 4-15 VM deployment using the VM list page

- ▶ New or modified columns in list pages are as follows, as shown in Figure 4-16.
 - The resources column on the VM list page displays processor units, virtual processor, and memory information for each VM.
 - Health and Status columns are combined into one.
 - Network list page shows you the free IP addresses in the network.

Name	Host	IP	State / Health	Resources	VM owner	Expiration date	Running task
sles15sp2-powervcsles	Redbooks01	[REDACTED]	Active Warning	4 vCPU / 24 GB / 2 PU			
rhel83-powervc-production	Redbooks03	[REDACTED]	Active Warning	6 vCPU / 26 GB / 2 PU			

Figure 4-16 Resources column on the VM list page

- ▶ Each resource edit page has a Logs tab named Context-sensitive logs which shows log messages specific to the particular resource. See Figure 4-17 Context-sensitive logs on VM edit page.

Type	Timestamp	Message	Resource	User
Info	2/12/2021, 1:25:52 PM	Stop of virtual machine rhel83-powervc-for-capture on host Redbooks03 was successful.	VM : rhel83-powervc-for-capture Host : Redbooks03	system
Info	2/12/2021, 1:05:52 AM	Capture of virtual machine rhel83-powervc-for-capture on host Redbooks03 was successful.	VM : rhel83-powervc-for-capture Host : Redbooks03	root
Info	2/11/2021, 11:26:35 PM	Stop of virtual machine rhel83-powervc-for-capture on host Redbooks03 was successful.	VM : rhel83-powervc-for-capture Host : Redbooks03	system
Info	2/11/2021, 11:16:11 PM	Deployment of virtual machine rhel83-powervc-for-capture on host Redbooks03 was successful.	VM : rhel83-powervc-for-capture Host : Redbooks03	root

Figure 4-17 Context-sensitive logs in VM edit page

- ▶ Navigating back and forth within the UI is now easier. The PowerVC Version 2.0.0 shows the breadcrumb navigation bar as you navigate through the tabs or features of a menu item as shown in Figure 4-18 on page 117.

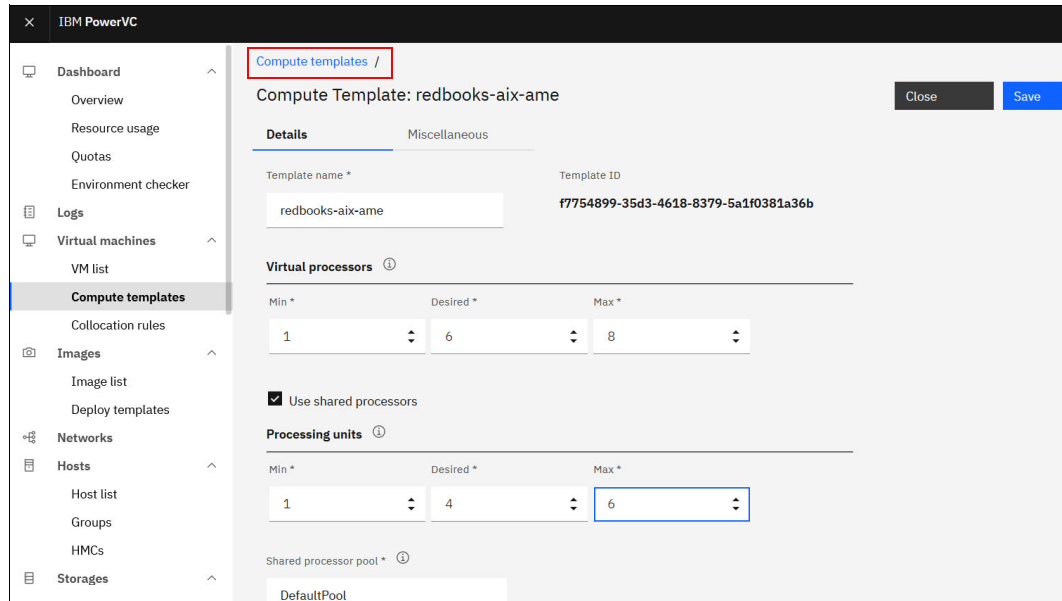


Figure 4-18 Breadcrumb of the trail

- ▶ The eye icon on the UI login screen helps you to verify if the entered password is correct, see Figure 4-19.

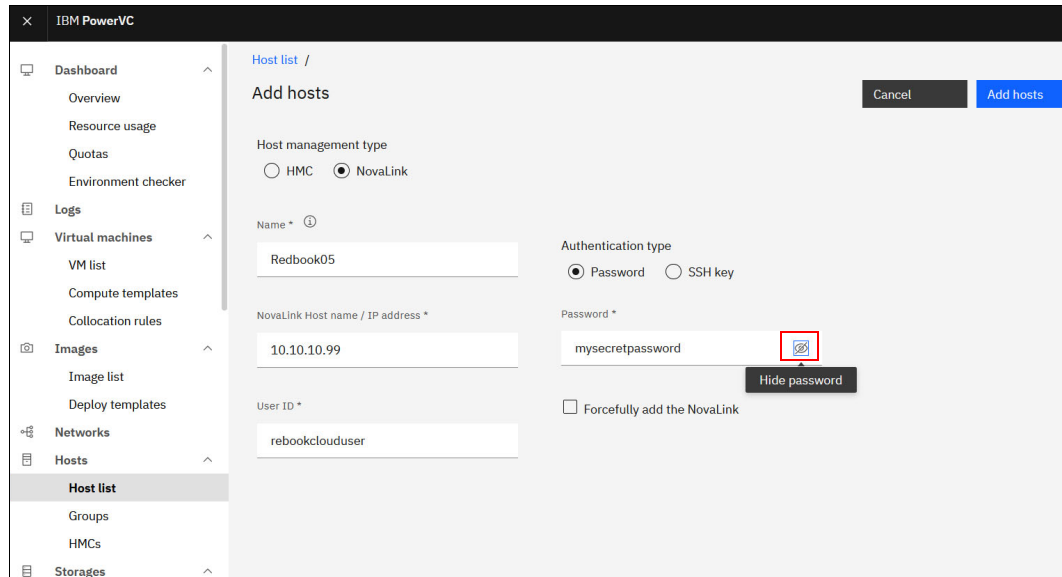


Figure 4-19 The eye icon

- ▶ PowerVC Version 2.0.0 comes with enhanced performance that you experience as you work with the UI features.
- ▶ PowerVC Version 2.0.0 UI is AVT compliant with Google Chrome's ChromeVox extension.
- ▶ You might come across a few issues or limitations while working with PowerVC Version 2.0.0. However, these issues do not hinder any of the core functionalities of the product. These issues will be fixed during future releases. You can find the list of current issues on the link below:

<https://ibm.biz/Bdfr9g>

4.2 Introduction to the IBM PowerVC setup

Before you can perform tasks in IBM PowerVC, you must discover and register the resources that you want to manage. You can register storage systems and hosts, and you can create networks to use when you deploy images. When you register resources with IBM PowerVC, you make them available to the management functions of IBM PowerVC (such as deploying a VM on a discovered host or storing images of captured VMs).

This discovery or registration mechanism is the key to the smooth deployment of IBM PowerVC in an existing environment. For example, a host can consist of several partitions while you deploy IBM PowerVC. You first register the host without registering any of the hosted partitions. All IBM PowerVC functions that relate to host management are available to you, but no objects exist where you can apply the functions for managing partitions. You can then decide whether you want to manage all of the existing partitions with IBM PowerVC. If you prefer a progressive adoption plan instead, start by managing only a subset of these partitions.

Ensure that the following preliminary steps are complete before you proceed to 4.4, “Connecting to IBM PowerVC” on page 120.

1. Secure the environment with the various PowerVC security features. Restart PowerVC services if you chose to enable SELinux enforcing mode after the PowerVC installation.
2. Configure the IBM Power Systems environment to be managed through the Hardware Management Console (HMC) or PowerVM NovaLink.
3. Set up the user accounts with an administrator role on IBM PowerVC, as described in 2.8, “Planning users and groups” on page 61. LDAP can be optionally configured with PowerVC to manage users and groups.
4. Set up the host name, IP address, and an operator user ID for the HMC or PowerVM NovaLink connection.

4.3 Managing resources outside of IBM PowerVC

When you switch from using IBM PowerVC to manage your resources to accessing the managed resource directly, you might see unexpected or adverse results in IBM PowerVC.

As you manage resources in IBM PowerVC, you might want to perform certain operations directly on the resource. For example, you are managing a VM with IBM PowerVC, but you use another method to stop the VM. In this case, IBM PowerVC might not be notified that the VM was stopped and might not immediately reflect the updated status of the VM. IBM PowerVC typically polls a resource to obtain an update the next time that you use IBM PowerVC to manage the resource. As a result, operations from IBM PowerVC might fail until the state of the VM in IBM PowerVC is the same as the state on the VM itself.

When you perform a management operation outside of IBM PowerVC, such as adding or deleting resources, the action can adversely affect the operation of IBM PowerVC and the data center. For example, you might delete a VM by using a Platform Manager for the VM. The VM goes into Error state and you must take the additional step of deleting the VM in IBM PowerVC. Similar results can occur when you remove a disk or network device for a VM that IBM PowerVC is managing.

For appropriate results, use IBM PowerVC to perform management tasks on the VMs and associated resources in your IBM PowerVC environment.

4.3.1 Changes of the LPAR in the HMC

Although it is not recommended to do changes outside of PowerVC, changes in the LPAR are possible and discovered by PowerVC after a short amount of time. Examples are LPAR name, enable or disable Performance Information Collection, activate/deactivate Remote Restart, the processor compatibility mode, shared or dedicated processors, or the minimum /desired /maximum values of the processors and memory.

If you want to do these changes outside PowerVC the recommended procedure is to unmanage the VM, do the changes outside, and then onboard the VM again into PowerVC.

4.3.2 Changes in the storage or zoning

Although some changes are recognized by PowerVC, for example the change of a volume name, this may lead to issues later. For example if you try to delete a VM with volumes where the name changed, you get errors, the VM will be deleted, but not the volumes.

There may be good reasons for doing storage or zoning related changes outside of PowerVC, for example to comply to naming standards. If you have the need to do such changes, unmanage the VM, do the changes and onboard the VM again.

If you use automation tools like Ansible you may want to create a Playbook that executes all necessary steps for you. Have a look into 6.2, “Ansible automation” on page 283.

4.4 Connecting to IBM PowerVC

After IBM PowerVC is installed and started on a Linux partition, you can connect to the IBM PowerVC management GUI by completing the following steps:

1. Open a web browser on your workstation and point it to the IBM PowerVC address:

`https://<ipaddress or hostname>/`

2. Log in to IBM PowerVC as an administrative user (Figure 4-20). The first time that you use IBM PowerVC, this administrative user is root. As a preferred practice, after the initial setup of IBM PowerVC, define other user IDs and passwords rather than using the root user. For more information about how to add, modify, or remove users, see section 2.8.1, “User management planning” on page 61.




Figure 4-20 IBM PowerVC Login window

3. The IBM PowerVC Overview screen as shown in Figure 4-21 is displayed.

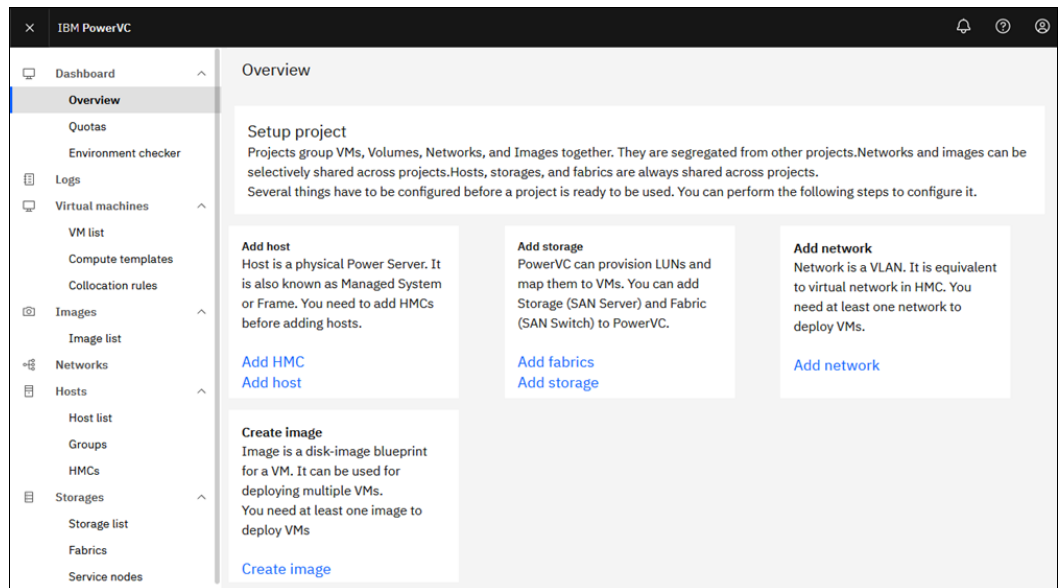


Figure 4-21 First time login

As shown in Figure 4-21 on page 120, no resources are yet configured and the screen recommends to add hosts, storage, networks, and images. These are the topics covered in the next sections.

Before you start, make sure your environment meets all of the hardware and software requirements and be configured correctly. Have a look into IBM Knowledge Center for details on the requirements:

https://www.ibm.com/support/knowledgecenter/SSVSPA_2.0.0/com.ibm.powervc.cloud.help.doc/powervc_hwandsw_reqs_cloud.html

4.5 Host setup

In this section we will be discussing adding hosts to IBM PowerVC. IBM PowerVC can either use PowerVM NovaLink or the hardware management console (HMC) to manage PowerVM hosts. IBM PowerVC can manage any combination of HMC managed and PowerVM NovaLink managed hosts. Planning on which management method to use is vital when considering the amount of hosts systems to be managed. PowerVM NovaLink can manage more hosts than the HMC.

4.5.1 Planning for host management - NovaLink

PowerVM NovaLink enables a highly scalable cloud management and deployment for critical enterprise workloads. Managing a host through PowerVM NovaLink has its advantages such as, more hosts can be managed and more simultaneous VM deployments can take place when using PowerVM NovaLink relative to the HMC. Also, the console of VMs can be viewed from IBM PowerVC GUI and DRO operations can be performed taking into consideration memory utilization (memory capacity on demand) for VMs managed by PowerVM NovaLink.

4.5.2 Planning for host management - HMC

When managing hosts managed by the HMC you can utilize all the great features of IBM PowerVC, excluding the features exclusive to PowerVM NovaLink managed hosts that were mentioned in the previous sub section. With hosts managed by the HMC, administrators can configure primary and secondary HMCs. The switch to the secondary HMC in the case the primary HMC has been taken down for maintenance is carried out by IBM PowerVC.

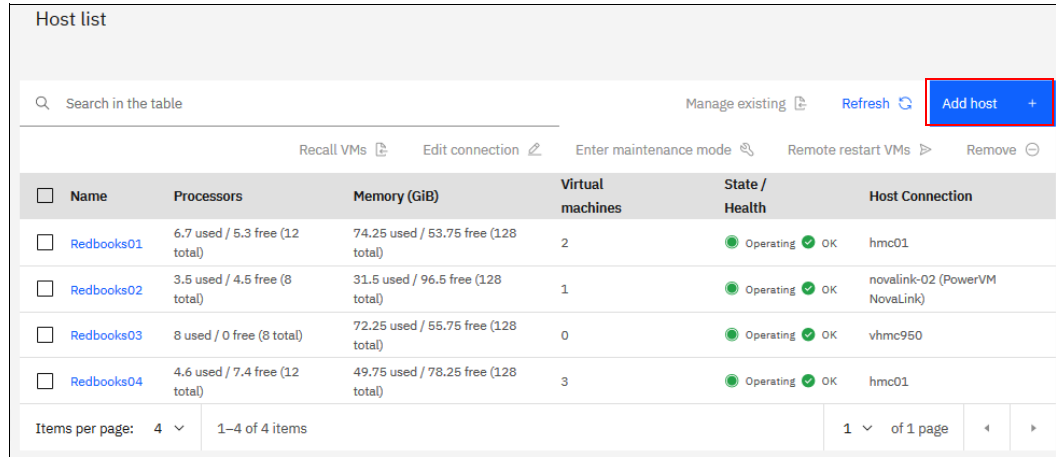
4.5.3 Adding a host

In order to add a host to IBM PowerVC, there must be network communication between IBM PowerVC and the HMC or PowerVM NovaLink.

To add hosts for IBM PowerVC management, complete the following steps:

1. On the left side of the IBM PowerVC window, click on Host list.

2. In the Host list page, click on **Add Hosts** on the top left corner as shown in Figure 4-22.



Host list

Search in the table

Manage existing Refresh Add host +

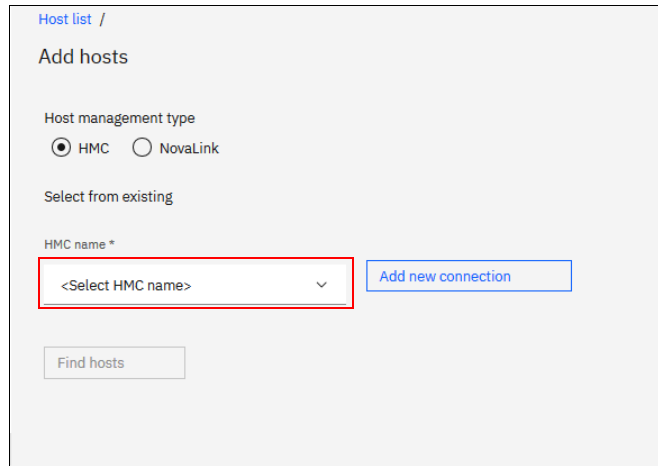
Recall VMs Edit connection Enter maintenance mode Remote restart VMs Remove

<input type="checkbox"/>	Name	Processors	Memory (GiB)	Virtual machines	State / Health	Host Connection
<input type="checkbox"/>	Redbooks01	6,7 used / 5,3 free (12 total)	74,25 used / 53,75 free (128 total)	2	Operating OK	hmc01
<input type="checkbox"/>	Redbooks02	3,5 used / 4,5 free (8 total)	31,5 used / 96,5 free (128 total)	1	Operating OK	novalink-02 (PowerVM NovaLink)
<input type="checkbox"/>	Redbooks03	8 used / 0 free (8 total)	72,25 used / 55,75 free (128 total)	0	Operating OK	vhmc950
<input type="checkbox"/>	Redbooks04	4,6 used / 7,4 free (12 total)	49,75 used / 78,25 free (128 total)	3	Operating OK	hmc01

Items per page: 4 1-4 of 4 items 1 of 1 page

Figure 4-22 Adding a host in IBM PowerVC

3. If you are adding a host managed by the HMC see Figure 4-23. You can choose to add a host from an already managed HMC or you can add a new connection. If you are adding from a managed HMC, select the HMC from the drop down box highlighted in red in Figure 4-23 to choose the hosts to manage with IBM PowerVC, select it using the check box. You can select several host names simultaneously.



Host list /

Add hosts

Host management type

HMC NovaLink

Select from existing

HMC name *

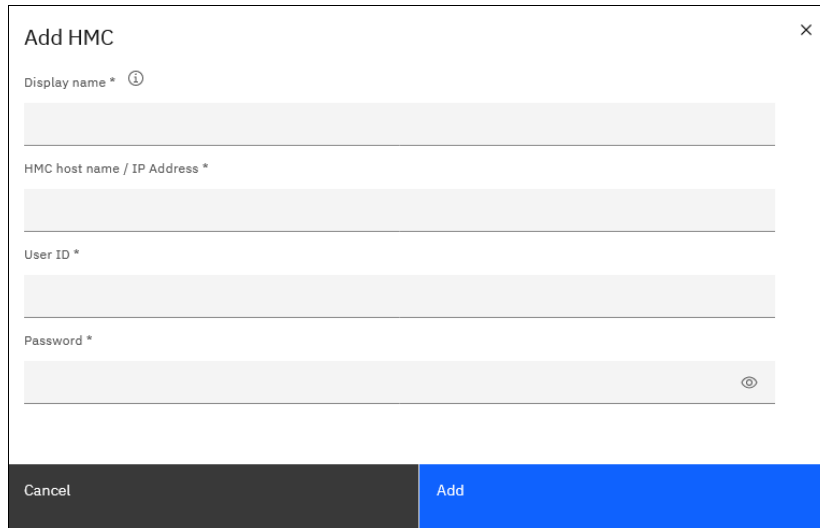
<Select HMC name>

Add new connection

Find hosts

Figure 4-23 Adding a HMC managed host through existing HMCs or a new connection

When adding a new connection, see Figure 4-24, use the credentials of a user in the HMC that has the role `hmcsuperadmin`. In the Display name field, enter the string that will be used by IBM PowerVC to refer to this HMC. The entered string is local to IBM PowerVC and does not alter any settings of the HMC. You may be prompted to accept a the X.509 certificate. Review the message details to determine whether you are willing to override this warning. If you are willing to trust the certificate, click **Connect** to continue.

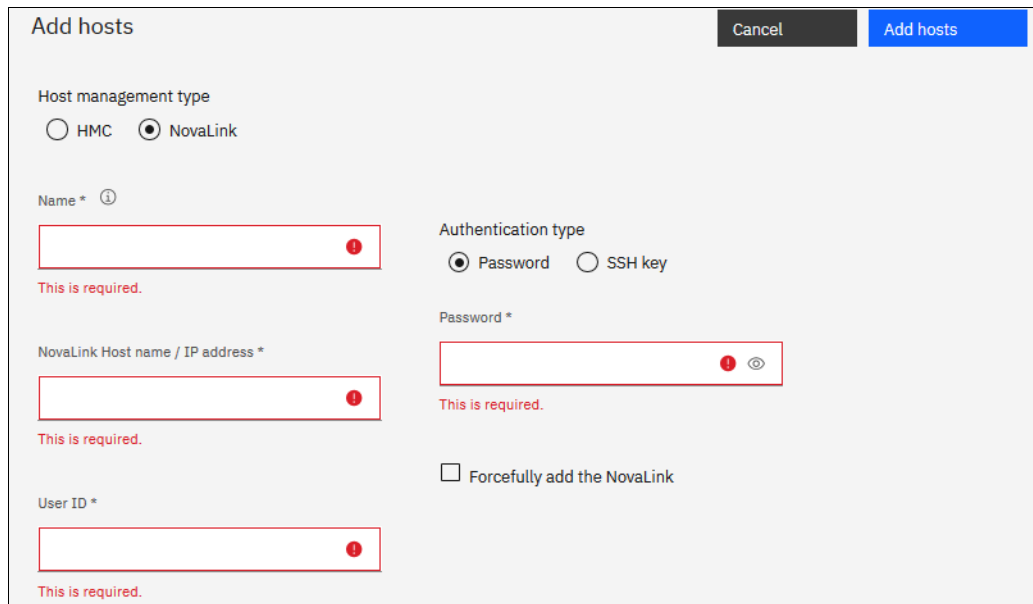


The 'Add HMC' dialog box contains the following fields and controls:

- Display name ***: A text input field with an information icon.
- HMC host name / IP Address ***: A text input field.
- User ID ***: A text input field.
- Password ***: A password input field with a visibility toggle icon.
- Buttons**: 'Cancel' (black) and 'Add' (blue).

Figure 4-24 Add HMC dialog

4. If you are adding a PowerVM NovaLink managed host, follow Figure 4-25. Provide the name and credentials for the NovaLink partition. Specify a user ID with root privileges or one that is part of the `pvm_admin` group and has `sudo` capabilities. You might be prompted to accept the SSH Fingerprint. Review the message details to determine whether you are willing to override this warning. If you are willing to accept the SSH key, click **Connect** to continue.



The 'Add hosts' dialog box contains the following fields and controls:

- Host management type**: Radio buttons for HMC and NovaLink.
- Name ***: A text input field with a red border and error message 'This is required.' and an information icon.
- Authentication type**: Radio buttons for Password and SSH key.
- NovaLink Host name / IP address ***: A text input field with a red border and error message 'This is required.'
- Password ***: A password input field with a red border, error message 'This is required.', and a visibility toggle icon.
- User ID ***: A text input field with a red border and error message 'This is required.'
- Forcefully add the NovaLink**: A checkbox.
- Buttons**: 'Cancel' (black) and 'Add hosts' (blue).

Figure 4-25 Adding a NovaLink managed host to IBM PowerVC

After managing hosts, the hosts will take a couple of minutes to become available. The progress of the hosts can be viewed in the host list window as shown in Figure 4-26.

<input type="checkbox"/>	Name	Processors	Memory (GiB)	Virtual machines	State / Health	Host Connection
<input type="checkbox"/>	Redbooks01	5.2 used / 6.8 free (12 total)	61.75 used / 66.25 free (128 total)	2	Operating OK	vhmc950
<input type="checkbox"/>	Redbooks02	4 used / 4 free (8 total)	35.5 used / 92.5 free (128 total)	1	Operating OK	novalink-02 (PowerVM NovaLink)
<input type="checkbox"/>	Redbooks03	7.5 used / 0.5 free (8 total)	106.75 used / 21.25 free (128 total)	0	Operating OK	vhmc950
<input type="checkbox"/>	Redbooks04	3 used / 9 free (12 total)	31.5 used / 96.5 free (128 total)	0	Operating OK	vhmc950

Figure 4-26 Host list window on IBM PowerVC

4.5.4 Host maintenance mode

In every data center there may be a need to carry out maintenance on systems. This can range from replacing of hardware to firmware updates. Whatever the case may be, the show must go on as down time remains unacceptable. IBM PowerVC has this covered with its maintenance mode feature.

Maintenance mode is sometimes referred to as one-click evacuation. When a host is put in maintenance mode, based on the request, the virtual machines on the host are migrated to a different host. If there are any errors during migration, the maintenance operation fails and enters the error state until an administrator resolves the problem. By default, PowerVC migrates the virtual machines that are allocated the most memory first. In a case where two virtual machines have the same memory, then IBM PowerVC checks for the Availability Priority setting. For example, a virtual machine that has the priority set to 50 is remote restarted after virtual machine with the priority value set to 100.

After the maintenance mode is requested, the hosts maintenance state is *Entering Maintenance* while the VMs react according to the maintenance option selected. The status changes to *Maintenance On* after the maintenance request is completed, and the host is fully in the maintenance state. To remove a host from maintenance mode, select the host and select Exit Maintenance Mode.

When a host is ready for use, exit the maintenance mode from the IBM PowerVC user interface. When the host is available in IBM PowerVC again, you must manually move VMs back to the host if they were not set with the feature to automatically return.

The following requirements has to be met before moving the host into maintenance mode:

- ▶ The host must be in an operating state. If it is not, a virtual machine migrations might fail.
- ▶ If all VMs on the host going into maintenance are to be migrated to other hosts, the following must be true:
 - The hypervisor must be licensed for live partition mobility.
 - The state of the VMs cannot be in error, paused or building.

- All VMs must be in an OK state, with active RMC connections.
- All requirements for live migration must be met
- ▶ If the request was made to migrate active or all virtual machines (including virtual machines in shut off state) when entering maintenance mode, the following cannot also be true, or the request will fail:
 - There is a virtual machine on the host that is a member of a collocation rule that specifies affinity and has multiple members.
 - The collocation rule has a member that is already undergoing a migration or is being remote restarted.

Figure 4-27, “Putting a host in maintenance mode” on page 125 shows how to put a host in maintenance mode. Navigate to the host list page, select the host and click **maintenance mode**.

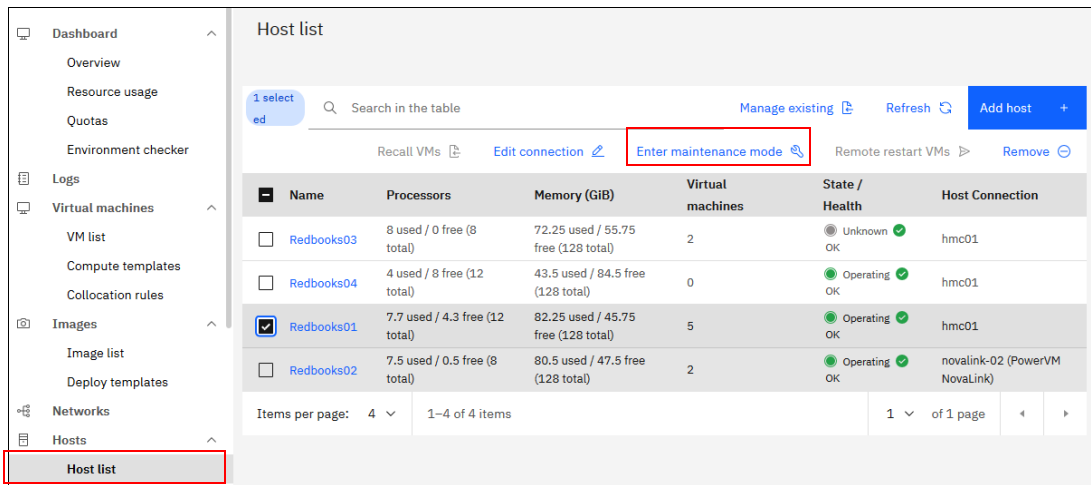


Figure 4-27 Putting a host in maintenance mode

After clicking on the maintenance mode link, the next action is to select what happens when the host goes into maintenance mode. Three self explanatory options are provided as shown in Figure 4-28. You can choose to specify the host the VMs are migrated to. Leave the default option, this will migrate the VMs, based on the host group placement policy, if Migrate active virtual machines to another host or Migrate all virtual machines to another host are selected.

Figure 4-28 Selecting the action to carry out when setting a host to maintenance mode

After the maintenance mode is requested, the hosts maintenance state is Entering while the VMs react according to the maintenance option selected. The status changes to On after the maintenance request is completed and the host is fully in the maintenance state. To remove a host from maintenance mode, select the host and select **Exit Maintenance Mode**.

After the host exits the maintenance mode (after 5 minutes) or when the host becomes active, any pinned virtual machines that were moved from the host or if the host has Recall enabled option is set to Yes, PowerVC automatically recalls the virtual machines to the host.

Note: You can edit the period after which the migration operation times out and the maintenance mode enters an error state by running the following commands:

```
/usr/bin/powervc-config compute maintenance_migration_timeout --set
<duration_in_seconds> --restart
```

The <duration_in_seconds> is the timeout period in seconds. For example, to set the timeout for 1 hour, run:

```
/usr/bin/powervc-config compute maintenance_migration_timeout --set 3600
--restart.
```

4.5.5 Host groups

Host groups are logical groupings of hosts, regardless of what they might have in common. Key rules guiding host groups will paint a better picture as to why they are required:

- ▶ Every host must be in a host group.
- ▶ VMs are kept within the host group when migrations are initiated.
- ▶ Placement policies guides placement of deployed VM amongst members of the same host group.

There are several reasons clients will want to utilize host groups a few are:

- ▶ To maintain cloud service level agreements (SLA). For example, to ensure if a VM is migrated, it is migrated to another POWER9 or later server.
- ▶ To employ multiple placement policies.
- ▶ To separate hosts based on business requirements.

IBM PowerVC comes with a default host group called *Default Group* as shown in Figure 4-29, you can choose to use this host group or you can create more host groups. The default host group cannot be deleted. A host can be a member of one host group only.

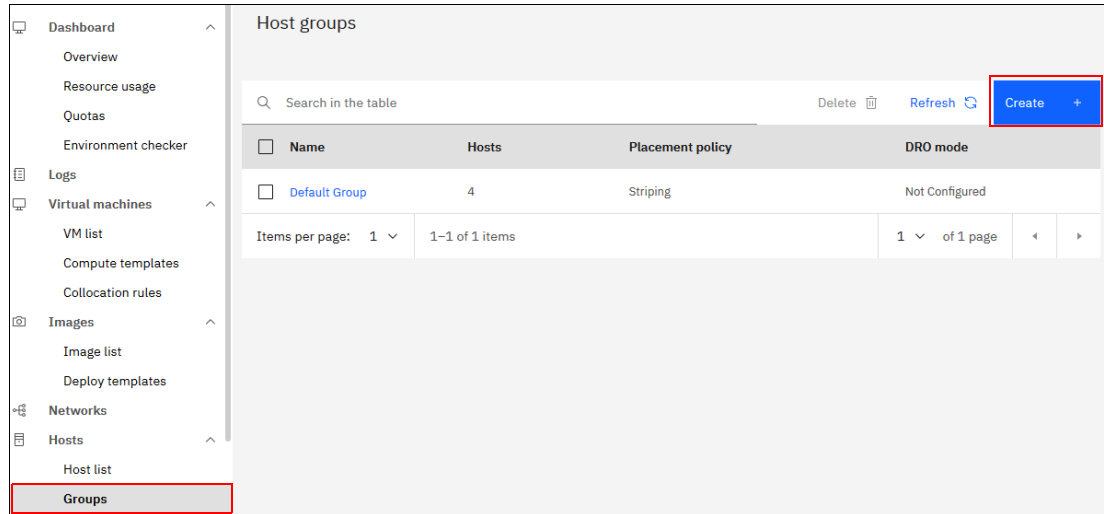


Figure 4-29 Creating a host group

When creating a host group, there are decisions to be made in order to achieve the goal of the client. The two major decisions are:

Placement policy PowerVC uses placement policies to determine which host to deploy a VM. The placement policies are associated with host groups are as follows:

Striping

The striping placement policy is the default placement policy. It distributes VMs evenly across all host members. For each deployment, PowerVC determines which hosts have enough processing units and memory to meet the requirements of the virtual machine. Other factors for determining eligible hosts include the storage and network connectivity that are required by the virtual machine. From the group of eligible hosts, PowerVC chooses the host that contains the fewest number of virtual machines and places the virtual machine on that host.

Packing

This placement policy is quite opposite relative to striping, the packing placement policy places virtual machines on a single host until its resources are completely consumed and then it moves on to the next host. For each VM deployment, IBM PowerVC determines which hosts have enough processing units and memory to meet the requirements of the virtual machine.

Other factors for determining eligible hosts include the storage and network connectivity that are required by the VM. From the group of eligible hosts, PowerVC chooses the host that contains the most virtual machines and places the virtual machine on that host.

CPU utilization balanced

This placement policy takes into consideration the CPU utilization computed as a running average over the last 15 minutes. Newly deployed VMs are placed on the host that has the lowest CPU utilization.

CPU allocation balanced

This placement policy places VMs on the host that has the lowest percentage of its CPU allocated after the deploy or relocation. For example if a VM is to be migrated and it has 2 CPUs, the two available host, host A and host B have 8 of 12 CPUs left and 2 of 6 CPUs left respectively. For host A, that will be 10 of 12 used after migration, for host B, that will be 4 of 6 left after migration. The migrated VM will be placed on host B.

Memory utilization balanced

This placement policy places VMs on the host that has the lowest memory utilization in the host group. The memory utilization is computed as a running average over the last 15 minutes.

Memory allocation balanced

This placement policy places VMs on the host that has the lowest percentage of its memory that is allocated after a VM deploy or relocation. This is similar to CPU allocation balanced, but memory is the resource taken into consideration.

Note: HMC managed hosts do not accurately report their memory utilization. Therefore, host groups that use this policy should not contain HMC managed hosts.

Dynamic resource optimizer (DRO)

IBM PowerVC DRO continuously monitors the resource utilization of physical hosts and virtual machines. IBM PowerVC DRO optimizes the infrastructure environment when usage imbalances are detected, by making adjustments for workloads whose characteristics change over time. DRO utilizes the powerful Capacity on Demand (CoD) tool from IBM Power Systems to activate unused cores or mobile memory for a host, giving it more computing power or memory instantly. On eligible systems, the PowerVC DRO works with CoD to relieve overburdened systems. For systems that are not Enterprise Pool members, DRO migrates virtual machines. DRO can monitor CPU or memory and use thresholds to initiate actions or advise the administrator on actions to take. Specific hosts or VMs can be excluded for DRO actions, this is to ensure that mission-critical VMs are not migrated and to ensure licensing control.

Note: DRO ignores hosts that are in maintenance mode. Such hosts are not monitored and are not targets for migrations.

DRO can migrate virtual machines or activate mobile memory to correct imbalances. This is only available for PowerVM NovaLink managed host.

With placement policy and DRO explained, we can proceed with an example showing how to create a host group. Figure 4-30 shows the creation of a host group. When creating a host group, the process has three pages:

- ▶ The detail page as shown in Figure 4-30
- ▶ The host page where members are added
- ▶ The DRO page, where DRO policies are set

The screenshot shows the 'Create host group' interface with the 'Details' tab selected. It includes a 'Name' input field, a 'Placement policy' dropdown menu set to 'Striping', an unchecked checkbox for 'Automated remote restart enabled', a 'Run interval' dropdown set to '1', and a 'Stabilization' dropdown set to '5'. 'Cancel' and 'Create' buttons are visible in the top right corner.

Figure 4-30 Details page when creating a host group

The name placement policy and the option to enable automated remote restart are indicated on the details page when creating a host group. With automated restart enabled, DRO will search hosts for failure, if a failure is detected, IBM PowerVC will automatically remote restart the failed VMs on other hosts within the same host group, taking into consideration the selected placement policy.

Figure 4-31 shows the selection of hosts page. All hosts that are managed by IBM PowerVC are shown and the administrator can select which ones will be part of the host group.

The screenshot shows the 'Create host group' interface with the 'Hosts' tab selected. A table lists available hosts with their properties and DRO status.

Name	Virtual Machines	Processors	Memory (GB)	Management type	DRO threshold status
<input checked="" type="checkbox"/> Redbooks01	3	3.2 Used (12Total)	34	hmc	OK
<input type="checkbox"/> Redbooks03	0	0 Used (8Total)	0 Used (128Total)	hmc	OK
<input checked="" type="checkbox"/> Redbooks04	0	0 Used (12Total)	0 Used (128Total)	hmc	OK

Figure 4-31 Selecting hosts when creating a host group

Up next is the DRO page where DRO policies are defined. When creating a host group, you can choose to configure or not configure DRO. The default host group, has DRO disabled by default. Figure 4-32 shows the configuration of DRO.

Figure 4-32 Configuration of DRO

Based on Figure 4-32, it is evident that the configuration of DRO can be very flexible. A host being over burdened can mean different things depending on the threshold values set when configuring DRO.

For example, assume the administrator configured DRO to monitor the CPU utilization in host group A with the following values:

- ▶ Resource to monitor= CPU
- ▶ Utilization threshold = 70%
- ▶ Run interval = 5 minutes
- ▶ Stabilization = 3 times

Every 5 minutes, DRO checks the utilization of every member of the host group. If any host has a utilization greater than 70%, its threshold is exceeded. If the host's utilization exceeds the utilization threshold three times in a row, the appropriate DRO action is triggered based on operation selected by the administrator.

Settings to customize DRO

The following settings can be used to customize the DRO:

Mode DRO can run in Active or Advise Only mode, virtual machines are not moved, and neither cores nor memory are activated. Instead, optimization events are logged on the DRO Events page so the administrator can review the recommended optimization actions and act.

Resource to monitor This indicates if the CPU, Memory, CPU and Memory or Shared processor pool is monitored in host systems.

Operations

Migrate virtual machines

Migrate virtual machines to other host in the host group.

Activate mobile cores or memory on the host

Activate an unused mobile core or unused mobile memory on the host. This operation requires the host to be a member of an Enterprise Pool.

Note: If both operations are selected, DRO tries to perform the CoD operation (activate a dormant core or available memory) before it migrates a virtual machine.

Schedule DRO can be scheduled to be effective at all times or during specific times, defined by the administrator.

4.6 Hardware Management Console management

With IBM PowerVC users can always add redundant HMCs for Power Systems servers. If one HMC fails, the user can change the HMC to one of the redundant HMCs.

4.6.1 Adding an HMC

You can add redundant HMCs for Power System servers. To add an HMC, in the HMC Connections section in the Hosts window, do the following:

1. Open the Hosts window.
2. Select the HMCs tab.
3. Click **Add** button.
4. Enter the HMC host name or IP address, display name, user ID, and password.
5. Click **Add** button
6. The new HMC is added, and you can easily click **Remove HMC** to remove an HMC.

4.6.2 Changing the HMC credentials

If you want to change the credentials that are used by IBM PowerVC to access the HMC, do the following:

1. Open the Hosts window.
2. Select the HMCs tab.

3. Click on the HMC name that you want to work with.
4. A window opens where you can specify another user ID, which must already be defined on the HMC with the hscsuperadmin role. Then, click **Save**.

4.6.3 Changing the HMC

IBM PowerVC uses only one HMC for one server. If one HMC fails, you must change the management console to another HMC. You can change it using the following procedure:

1. On the Hosts window, select all of the servers that you want to change.
2. Click **Edit Host Connection** and select the HMC you want.
3. Click **OK**.

4.6.4 Change management connection, HMC, and PowerVM NovaLink

POWER8 processor-based servers and later POWER models can be co-managed by PowerVM NovaLink and an HMC that is at version 8.4.0, or later. An HMC that co-manages a POWER8 processor-based server can also manage POWER6 and POWER7 processor-based servers. However you cannot use IBM PowerVC to change your hosts to be PowerVM NovaLink managed if you are using shared processor pools or shared storage pools (SSPs).

When a managed system is co-managed by the HMC and PowerVM NovaLink, you set either PowerVM NovaLink or the HMC to be in the master mode. Certain system management, partition management, and Capacity on Demand (CoD) tasks can be performed only from the interface that is in master mode. For example, if PowerVM NovaLink is in the master mode, you can run partition change operations only by using PowerVM NovaLink. If you want to run partition change operations by using the HMC, you must set the HMC to master mode.

Firmware updates for a co-managed system can be performed only from the HMC. The HMC must be set to the master mode to update the firmware.

- ▶ To set the HMC to the master mode, run the following command from the HMC command line, where **managed_system** is the name of the managed system:

```
chcomgmt -m managed_system -o setmaster -t norm
```

- ▶ Once you have completed your tasks on the HMC, run the following command from the HMC command line to make PowerVM NovaLink the master again:

```
chcomgmt -m managed_system -o relmaster
```

If you are using the DRO with Capacity on Demand (CoD), you can switch to PowerVM NovaLink. DRO continues to monitor and adjust workloads, but cannot take advantage of CoD.

If you want to prepare your system for PowerVM NovaLink management, follow the below steps:

1. Install the PowerVM NovaLink software on a system that is managed by an HMC.
2. From the IBM PowerVC user interface, open the Hosts list window.
3. Select the host that you want to update, and then select **Edit Connection**.
4. For the Host management type, select **NovaLink**. Enter the appropriate information and click **Save**. IBM PowerVC will recognize the host as PowerVM NovaLink managed and the host comes out of maintenance mode.

4.7 Storage and SAN fabric setup

The host has been successfully added to IBM PowerVC in the previous sections now its time to add the storage infrastructure. This entails first configuring the SAN fabric and storage controller before adding to IBM PowerVC.

For the SAN fabric, connect the designated Fibre Channel cables from the VIOS and storage controller Fibre Channel adapters to the SAN fabric:

1. Initialize the SAN fabric
2. Set up the host names, IP addresses, and administrator user ID and password
3. Create a configuration file

For the storage controller:

1. Initialize the storage
2. Set up IP addresses, user authentication and create pools.

After all of the above is done, the SAN fabric and storage controller can be added to IBM PowerVC.

4.7.1 Adding a storage controller to IBM PowerVC

To add a storage controller to IBM PowerVC, go to the storage list page and click **Add**, as shown in Figure 4-33.



Figure 4-33 Adding a storage controller

1. Input the storage details as shown in Figure 4-34. This page may vary depending on the storage system selected.

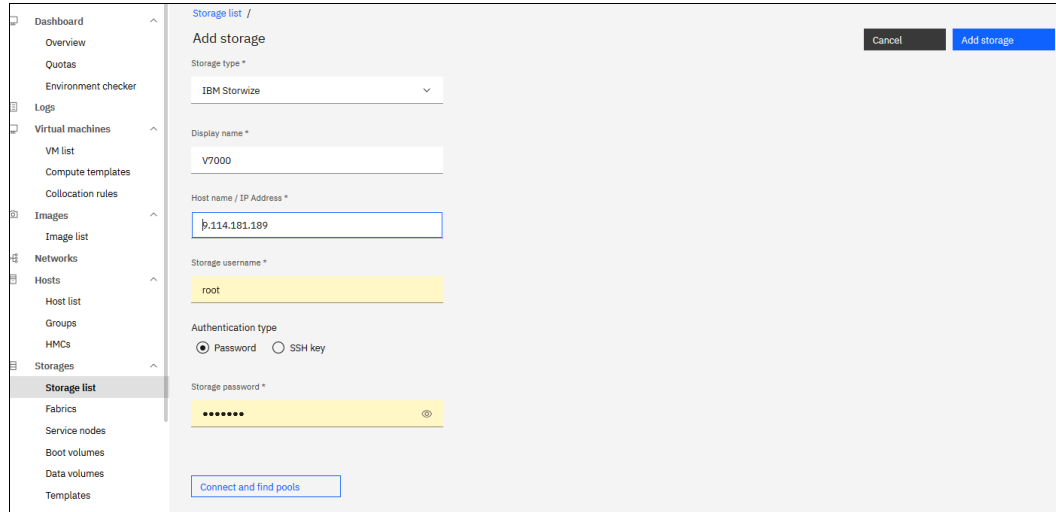


Figure 4-34 Adding an IBM Storwize storage controller

2. After adding the storage controller user authentication and IP address details, click on **Connect** and find pool. Select the pool to be added to IBM PowerVC and finally click on **Add storage**. The storage will be added as part of the storages on the storage list page. This is shown in Figure 4-35.

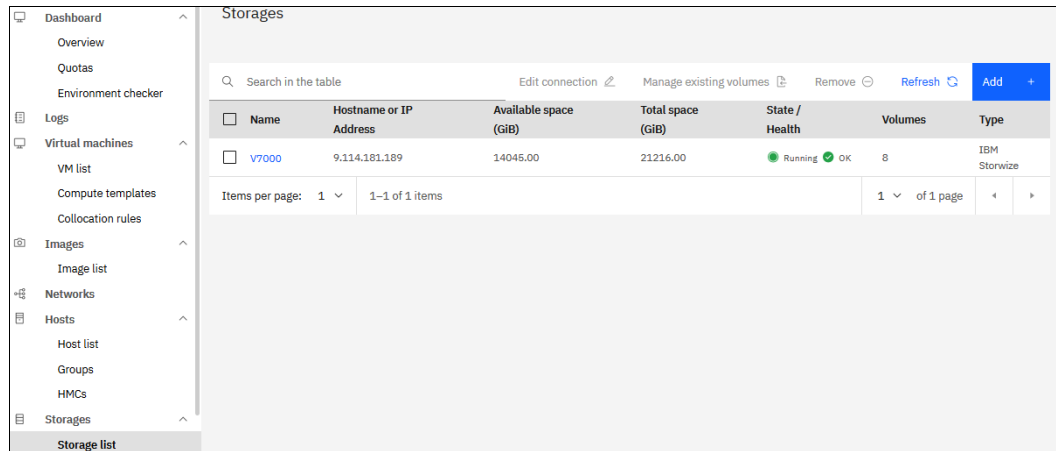


Figure 4-35 List of added storage controllers

4.7.2 Adding a SAN fabric to IBM PowerVC

Before adding a SAN fabric to IBM PowerVC, there is a key decision to be made by the administrative team. What zoning policy will be used? When adding a SAN fabric, there are three zoning policies to choose from namely:

Initiator zoning Creates a one-to-many relationship between the host and the target storage controller node.

initiator-target zoning Creates a one-to-one relationship between the host and the storage controller.

initiator-vfc

Creates a zone with the active and migration initiators, along with the targets for the storage provider.

Based on enterprise policy or choice the administrator can select what zoning policy best suits their purpose. To add a SAN fabric navigate to the fabrics page and click on **Add** as shown in Figure 4-36.

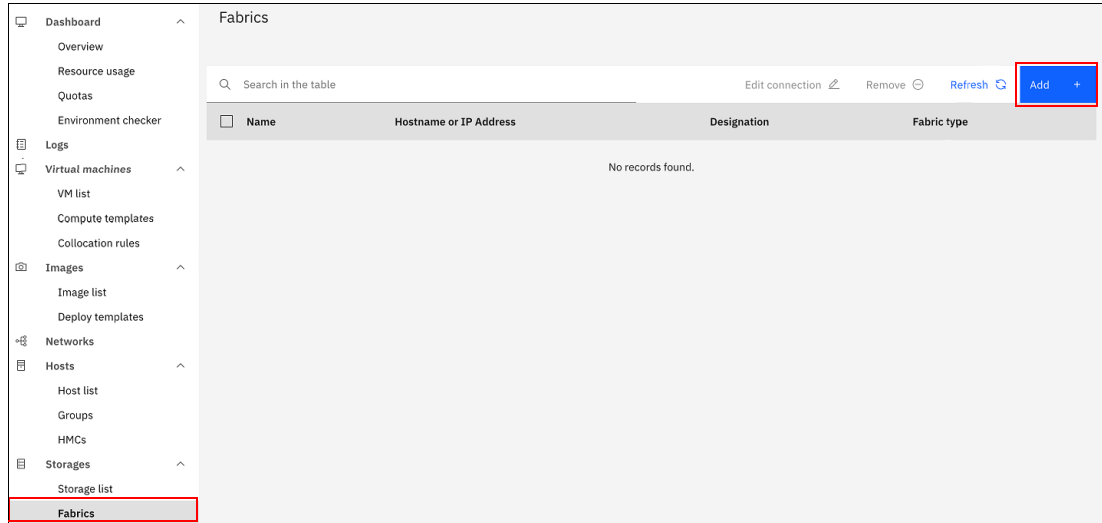


Figure 4-36 Adding a SAN fabric to IBM PowerVC

The type of SAN fabric, user authentication, zoning policy and virtual information is added, then the **Add** button is clicked afterwards as shown in Figure 4-37.

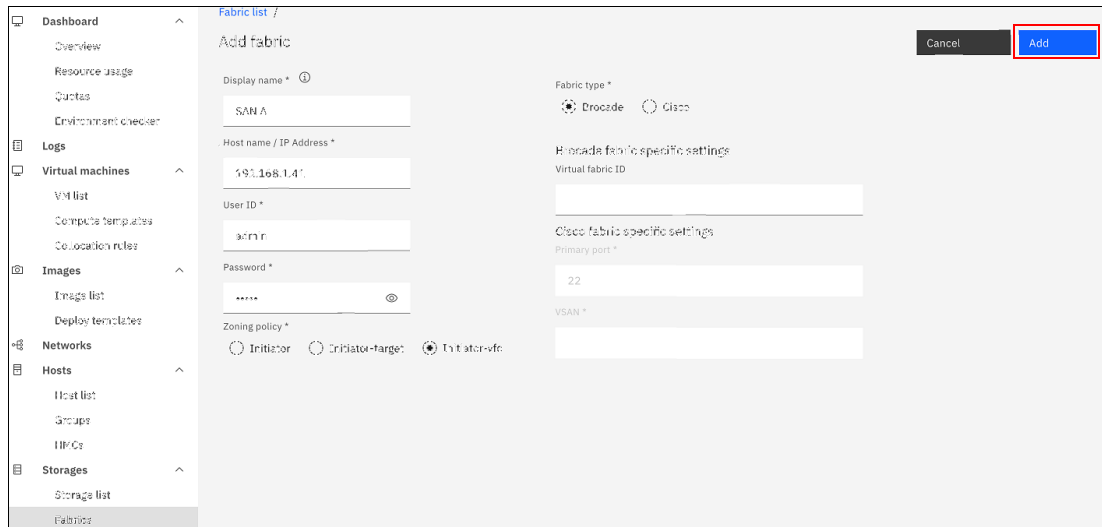


Figure 4-37 Adding SAN fabric details

The SAN fabric can be viewed in the fabric list window after being added to IBM PowerVC as shown in Figure 4-38.

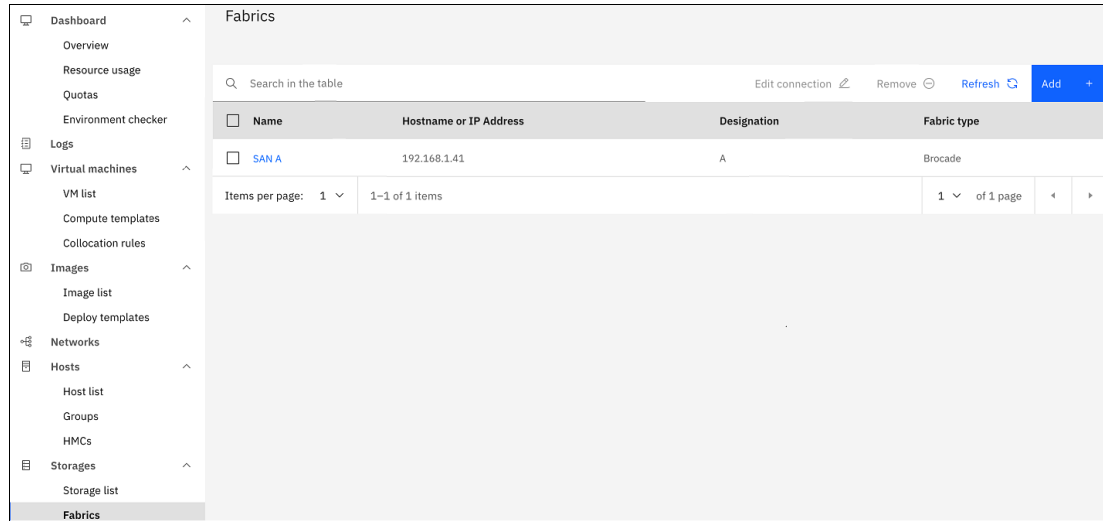


Figure 4-38 Viewing the added SAN fabric

4.8 VIOS port tags setup

Individual FC ports in the VIOSs that are managed by IBM PowerVC can be tagged with named labels. For more information about IBM PowerVC tags and storage connectivity groups, see section 2.6.3, “Storage connectivity group and tags” on page 53.

Note: Tagging is optional. It is needed only when you want to partition the I/O traffic and restrict certain traffic to use a subset of the available FC ports.

To configure VIOS Fibre Channel ports tagging, complete the following steps:

1. On the left side of the IBM PowerVC window, click on **Storages** → **FC ports**. For each FC adapter in all VIOSs that is managed by IBM PowerVC, you can enter a Tag (arbitrary name). You can define your own tag. Select the port you want to configure (IBM PowerVC automatically recognizes the fabric that the port is connected to) and click on **Set port tag** as shown in Figure 4-39.

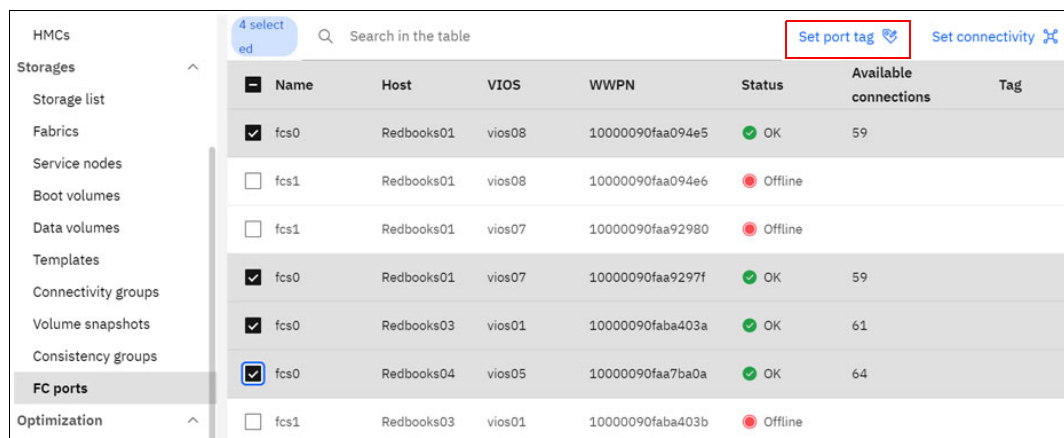


Figure 4-39 Fibre Channel ports selection

2. Enter the production tag and click **Set** as shown in Figure 4-40.

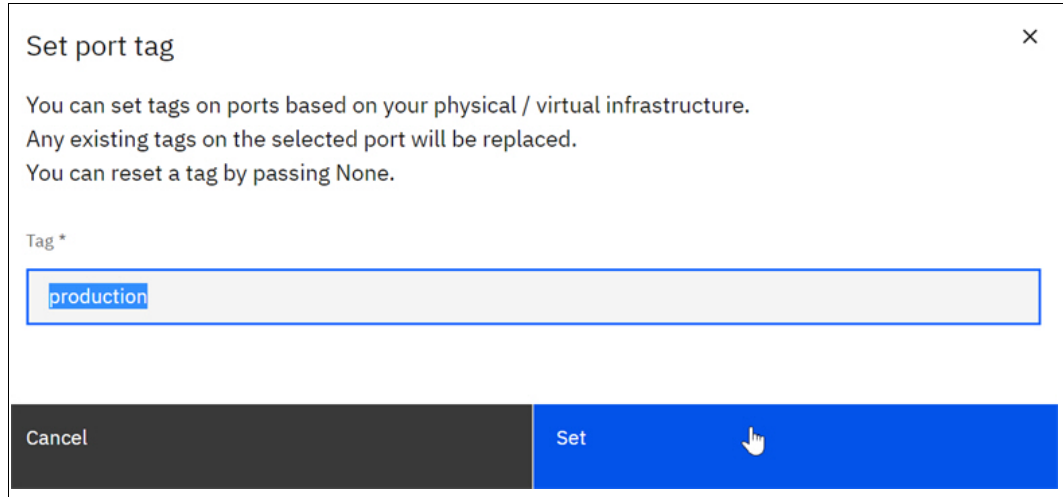


Figure 4-40 Enter tag

You can also select N_Port ID Virtualization (NPIV) or virtual SCSI (vSCSI) for the Connectivity field to restrict the port to special SAN access. To disable the FC port in IBM PowerVC, set the connectivity to None. By default, all ports are set to Any connectivity, which allows all connectivity methods.

3. Click on **Set connectivity** and select **NPIV** as shown in Figure 4-41.

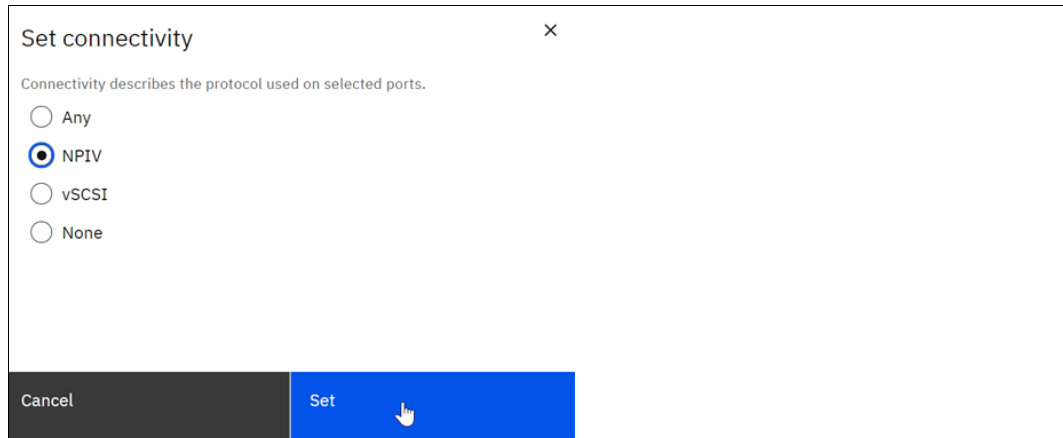


Figure 4-41 Set connectivity

4. Do not forget to click **Save** to validate your port settings as shown in Figure 4-42.

Name	Host	VIOS	WWPN	Status	Available connections	Tag	Connectivity	Fabric
<input checked="" type="checkbox"/> fcs0	Redbooks01	vios08	10000090faa094e5	OK	59	production	NPIV	B
<input checked="" type="checkbox"/> fcs0	Redbooks01	vios07	10000090faa9297f	OK	59	production	NPIV	A
<input checked="" type="checkbox"/> fcs0	Redbooks04	vios05	10000090faa7ba0a	OK	64	production	NPIV	A
<input checked="" type="checkbox"/> fcs0	Redbooks03	vios01	10000090faba403a	OK	61	production	NPIV	A

Figure 4-42 Fibre Channel port configuration

4.9 Storage connectivity group setup

A *storage connectivity group* is a set of VIOSs with access to the same storage controllers. The storage connectivity group also controls the boot volumes and data volumes to use NPIV or vSCSI storage access. For more information about IBM PowerVC tags and storage connectivity groups, see section 2.6.3, “Storage connectivity group and tags” on page 53.

IBM PowerVC creates a default storage connectivity group named *Any Host, All VIOS* during the initial setup and adds all hosts and VIO servers into that group with NPIV connection by default. Defining additional storage connectivity groups is an optional task that can be performed any time after the initial installation. All VMs that are deployed in IBM PowerVC require a storage connectivity group to be assigned during the deployment.

To configure a Storage Connectivity group, complete the following steps:

1. On the left side of the IBM PowerVC window, click on **Storages** → **Connectivity groups**. You can then create your own storage connectivity group. Click **Create**. In the next window, enter information or select predefined options for the new storage connectivity group in the **Details** tab as shown in Figure 4-43 on page 139.

- Name of the storage connectivity group
- Boot and Data volume connectivity types: NPIV or vSCSI
- Automatically add applicable VIOSs from newly registered hosts to this storage connectivity group

If checked, from now on, newly VIOSs are added to this group if they can access the same storage (fabrics and tags) as the other members of the group.

- Fibre channel ports tag settings

This setting is optional. If you use tags, you can select a specific tag. VMs that are deployed using this storage connectivity group (with a selected tag) can access storage only through FC ports with the specified tag.

Create a storage connectivity group

Details * VIOS redundancy * Fabrics redundancy IPG

Name *

vSCSI

Volumes connectivity

Boot *

vSCSI

Data *

vSCSI

Fibre channel ports tag settings

Deploy VMs with tag restrictions

production

Automatically add applicable Virtual I/O Servers from newly registered hosts to this storage connectivity group.

Allow deployments using this storage connectivity group (enabled).

Figure 4-43 Storage connectivity group details

2. Enter information in the **VIOS redundancy** tab as shown in Figure 4-44:

- At least** Select this option to define the minimum number of VIOSs per host that can connect to the volume. From the drop-down menu, select 1, 2, 3, or 4 VIOSs.
- Exactly** Select this option to define the exact number of VIOSs per host that can connect to the volume. From the drop-down menu select 1, 2, 3, or 4 VIOSs.

3. Select the VIOSs to add to the storage connectivity group.

Create a storage connectivity group

Cancel Create

Details * **VIOS redundancy *** Fabrics redundancy IPG

Storage path redundancy avoids connectivity loss during VIOS updates / down time. At least 2 VIOSes are recommended for production environments.

At least Exactly

2 VIOS

2 selected Search in the table Refresh

VIOS	Host	State	RMC state	Total FC ports	FC ports ready	Connectivity type
<input checked="" type="checkbox"/> vios07	Redbooks01	Running	Active	2	NPIV:1.vSCSI:1	vSCSI,NPIV,ISCSI
<input checked="" type="checkbox"/> vios08	Redbooks01	Running	Active	2	NPIV:1.vSCSI:1	vSCSI,NPIV,ISCSI
<input type="checkbox"/> vios01	Redbooks03	Running	Active	2	NPIV:1.vSCSI:1	vSCSI,NPIV,ISCSI
<input type="checkbox"/> vios06	Redbooks04	Running	Active	2	NPIV:1.vSCSI:2	vSCSI,NPIV,ISCSI

Figure 4-44 Storage connectivity group VIOS redundancy

- Then, click **Create**, and the storage connectivity group is created and available for VM deployment.
- A storage connectivity group can be disabled to prevent the deployment of VMs in this group. To disable a group, you must clear the check box for **Enable VM deployment** on the Details tab as shown in Figure 4-45.

Storage connectivity group: vSCSI-NPIV

Details * VIOS members * Fabrics IPG Member health

Name *
vSCSI-NPIV

Boot volume connectivity type *
vSCSI

Data volume connectivity type
NPIV

Match FC Port tag
No restriction

You may create a new port tag [here](#).

Include untagged ports

Auto-add VIOS from new hosts ⓘ

Enable VM deployment ⓘ

Details	
ID	9ed21f0f-7399-458c-8a88-5e888b83d189
System generated ⓘ	No
Boot health	<input checked="" type="checkbox"/> Inactive
Data health	<input checked="" type="checkbox"/> SCG_Disabled
Fabric access requirements met	Yes
Data VIOS redundancy requirements met	Yes
Boot VIOS redundancy requirements met	No
Ready members	Boot: 0, Data: 1
Fibre channel ports ready	NPIV:2, vSCSI:0

Figure 4-45 Disable a storage connectivity group

4.10 Storage template setup

Storage templates provide a predefined storage configuration to use when you create a disk. You must define different information on the storage template for different types of storage. For more information about IBM PowerVC tags and storage connectivity groups, see 2.6.2, “Storage templates” on page 49.

A default storage template is automatically created by IBM PowerVC for each storage provider. However, if the storage contains several storage pools, create a storage template for each storage pool that you want to use. For IBM Storwize storage, you also must create a storage template for each I/O group that you want to use, and each volume mirroring pool pair that you want to use.

To configure a Storage template for a **hyperswap** configuration (vdisk mirroring), complete the following steps:

- On the left side of the IBM PowerVC window, click on **Storages** → **Templates**. You can then create your own storage template, click **Create**, in the next window, enter information or select the predefined options for the new storage template in the **Details** tab as shown in Figure 4-46 on page 141:
 - Name of the storage template.
 - Select a Storage provider.
 - Select a Template type:

Thin

Additional options are available:

- I/O group
- Real capacity % of virtual storage
- Warning threshold
- Thin-provisioned grain size
- Use all available worldwide port names (WWPNs) for attachment
- Flash Copy Rate

For more information about how these settings affect IBM PowerVC disk allocation, see section 2.6.2, “Storage templates” on page 49.

Compressed

Additional options are available:

- I/O group
- Real capacity % of virtual storage
- Enable threshold check
- Auto expand
- Use all available worldwide port names (WWPNs) for attachment
- Flash Copy Rate

Generic

Additional options are available:

- I/O group
- Use all available worldwide port names (WWPNs) for attachment
- Flash Copy Rate

The screenshot shows the 'Create storage template' dialog box with the 'Details' tab selected. The 'General properties' section includes fields for 'Template name' (hyperswap), 'Storage provider' (SVC), and 'Template type' (Thin). The 'Thin/Compressed provisioning settings' section includes 'Thin provisioning grain size' (256), 'Real capacity (% of virtual capacity)' (2), 'Enable threshold check' (checked), and 'Warning threshold (%)' (80). At the bottom, there are two checkboxes: 'Use this storage template as default' and 'Use all available WWPNs'.

Figure 4-46 Hyperswap Storage template details

2. Enter information in the **Storage pools** tab as shown in Figure 4-47.
3. Select a Pool in site 1.
4. Check the box **Enable mirroring**.
5. Select a Pool in site 2.

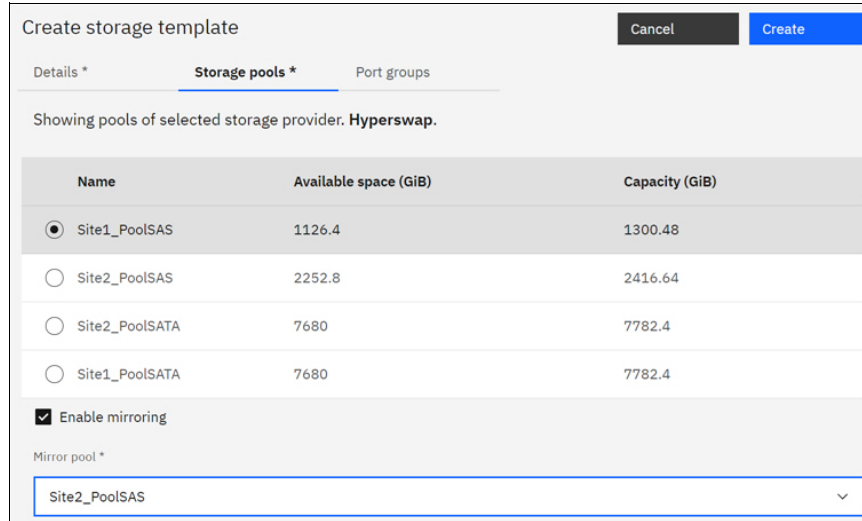


Figure 4-47 Hyperswap Storage template Storage pools

6. Select the Storage provider ports you want to use in the **Port groups** tab as shown in Figure 4-48.

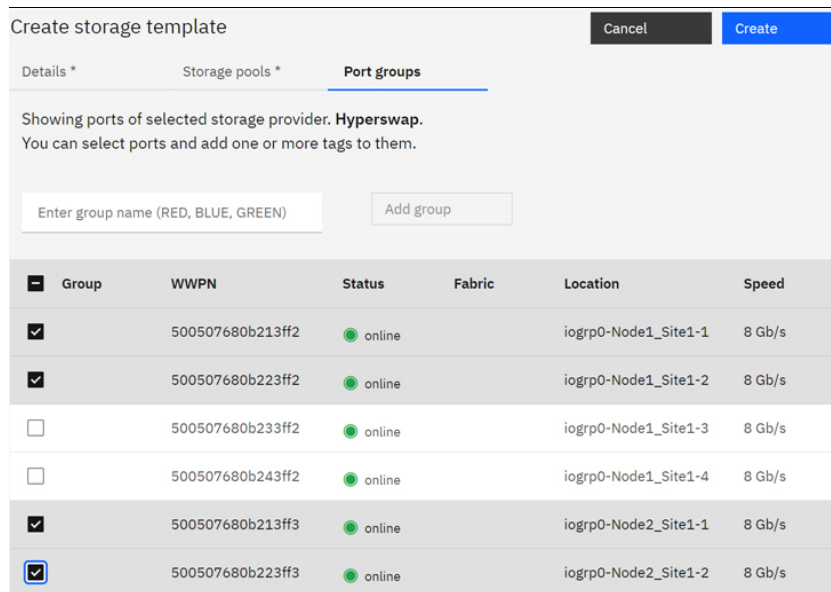


Figure 4-48 Hyperswap Storage template Port groups

7. Click **Create**. The Storage template is created and it is available for use when you create storage volumes.

4.11 Storage volume setup

When you create a volume, you must select a storage template that determines where (which storage controller and pool) and what the parameters are (thin or thick provisioning, grain size, and so on) for the volume to create.

When you create a volume, you must select these elements:

- ▶ A storage template
- ▶ The new volume name
- ▶ A short description of the volume (optional)
- ▶ The volume size (GB)
- ▶ The Volume count (1 or more)
- ▶ Can be shared between VMs or not

If this option is selected, the volume can be attached to multiple VMs. This option is for PowerHA or similar solutions.

To create a hyperswap volume, complete the following steps:

1. On the left side of the IBM PowerVC window, click on **Storages** → **Data volumes**. Click **Create**, in the next window, enter information or select predefined options for the new data volume as shown in Figure 4-49.

The screenshot displays the 'Create data volume' window with the following details:

- Volume name ***: hyperswap1
- Description**: hyperswap volume
- Storage template ***: Hyperswap
- Size (GiB) ***: 10
- Volume count ***: 1
- Can be shared between VMs

Projected storage usage:

- Total space (GiB)**: 7,782
- Free**: 98.7%
- Used**: 1.3%
- Required**: 0%

Required space: < 1 GiB
Free space: 7680 GiB
Used space: 102 GiB

Figure 4-49 Create Hyperswap Storage volume

- After you click **Create**, the volume is created. A list of existing volumes is displayed, as shown in Figure 4-50. This figure shows that the provisioned disk is in the available state.

Name	Size (GiB)	State / Health	VMs	Storage template	Storage provider
hyperswap1	10	Available OK	0	Hyperswap	Hyperswap

Figure 4-50 List available storage volume

- On the SVC, there is one volume created with one copy in each pool as shown in Figure 4-51.

Name	State	Synchroniz...	Pool	Host Mappings	Capacity
volume-hyperswap1-767c51d4-76ab	Online		Site1_PoolSA...	No	10.00 GiB
Copy 0*	Online	Yes	Site1_PoolSA...	No	10.00 GiB
Copy 1	Online	Yes	Site2_PoolSA...	No	10.00 GiB

Figure 4-51 Show Hyperswap volume on SVC

4.12 Network setup

When you create a VM, you must select a network. PowerVC need information to make your new VM reachable over the network and to make sure that no IP address conflicts occur. Therefore all networks that should be used need to be defined in PowerVC. For planning of your network configuration in IBM PowerVC, see section 2.7, “Network management planning” on page 56 for some useful information. The details how to configure networks in Power will be discussed in the following sections.

4.12.1 Basic network setup

Initially, IBM PowerVC contains no network definition, so you must create at least one network definition.

To create a network definition in IBM PowerVC, click on **Networks** in the navigation area on the left to see the networks overview as shown in Figure 4-52 on page 145.

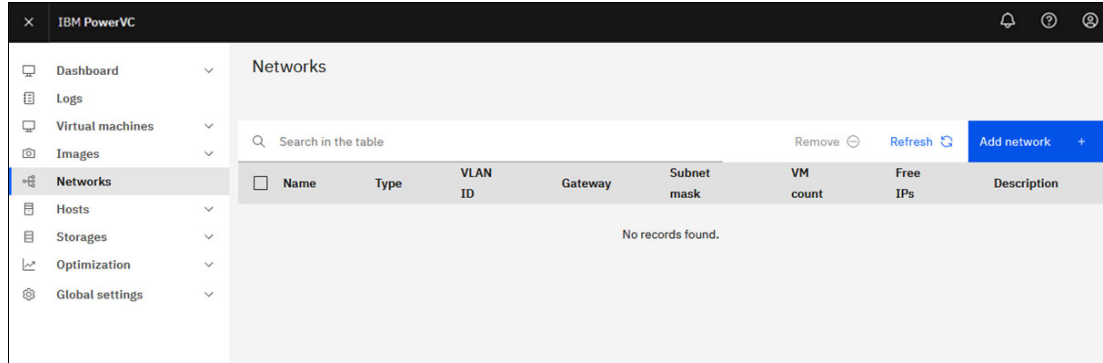


Figure 4-52 Networks in PowerVC

1. To add a network, click **Add network**.

Figure 4-53 shows a basic network definition with static IP addresses using configured SEA adapters.

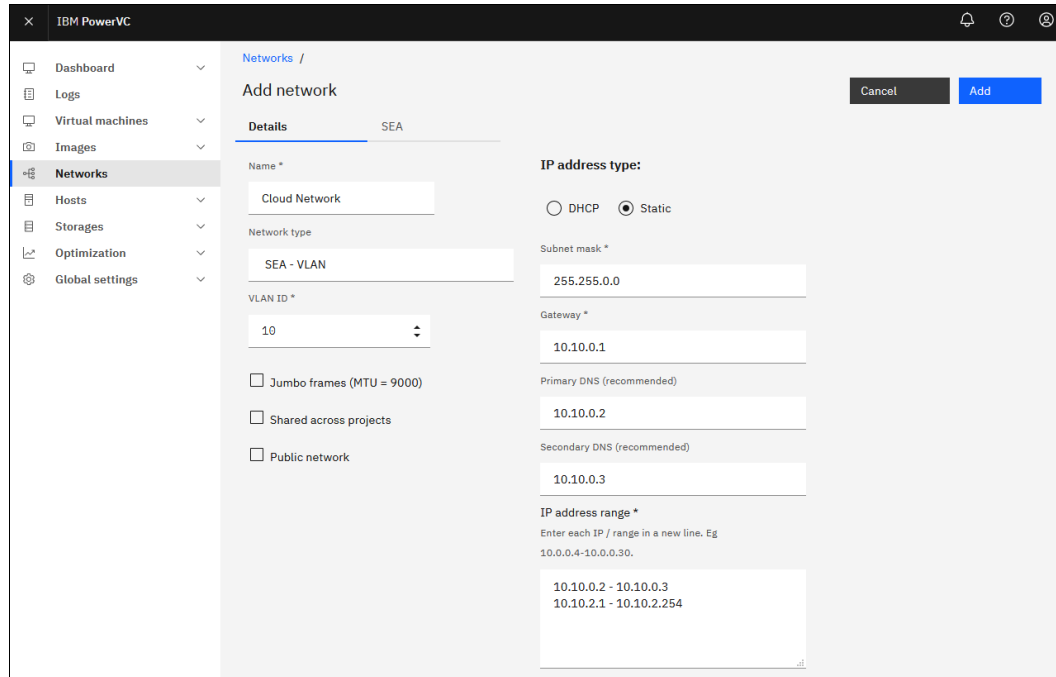


Figure 4-53 Definition of a basic network

2. You can provide the following data on the Details tab when you create a network:

- Network name
- Virtual LAN (VLAN) ID
- Jumbo frames (MTU = 9000)

Jumbo frames have a Maximum Transmission Unit (MTU) size of 9000, which is much more than the standard of 1500. This means a higher junk of data could be transmitted in one package. For a big amount of data this can increase the bandwidth usage. Jumbo frames are often used in backup networks.

If you select Jumbo Frames, make sure that the whole network path is capable of this MTU size, otherwise the network splits the packages, which means usually even more overhead.

– Shared across projects

If checked this network will be available in all projects of PowerVC, otherwise it will be private to the current project.

– Public network

Select this option to mark that the network is public. Then in VM deployment page, you can choose to mark this network as primary network.

– For IP address type, select **DHCP** or **Static**.

– The following attributes are only visible for a network with static IP addresses (for DHCP these value are provided by the DHCP server):

- Subnet mask
- Gateway
- Primary/Secondary DNS (This fields are optional if you do not use DNS.)
- IP address range(s)

PowerVC supports multiple ranges for a network, Figure 4-53 on page 145 shows an example.

Hint: If you want to add a single IP address, specify the IP address as range, for example 10.10.3.33 - 10.10.3.33. Just the IP address will not be stored.

The IP address ranges can be changed or extended after defining the network, which was not the case in previous PowerVC versions.

On the SEA tab you find all hosts with a configured SEA with the following details:

- VIOS
- The VIOS providing the shared network
- SEA
- The device name of the SEA
- Status
- Redundant adapter
- VIOS and device name of a redundant SEA (if available)
- vSwitch

vSwitch used for that SEA. You can have multiple vSwitches with multiple SEAs. Make sure to only select the SEA that should be part of the defined network.

• Primary VLAN

The primary VLAN shows the Port Virtual Ethernet ID defined in the virtual adapter used for the SEA.

Important: The VLAN number that you specify does *not* need to match the primary VLAN.

3. Select all SEAs that should be part of the network. By default all that are capable to support the network are selected.

PowerVC deploys the VLAN that is associated with the virtual machine to a host system when the VLAN is not present on the Shared Ethernet Adapter. If every virtual machine that requires a specific network is deleted from a host, the associated VLAN is removed from the VIOS. However, the VLAN is never removed from the Shared Ethernet Adapter's primary virtual Ethernet adapter.

Hint: If a SEA should handle multiple VLANs, configure the adapter as IEEE 802.1q compatible adapter to be able to support these VLANs.

4. When finished, click **Add** to create the network. Figure 4-54 shows the new network.

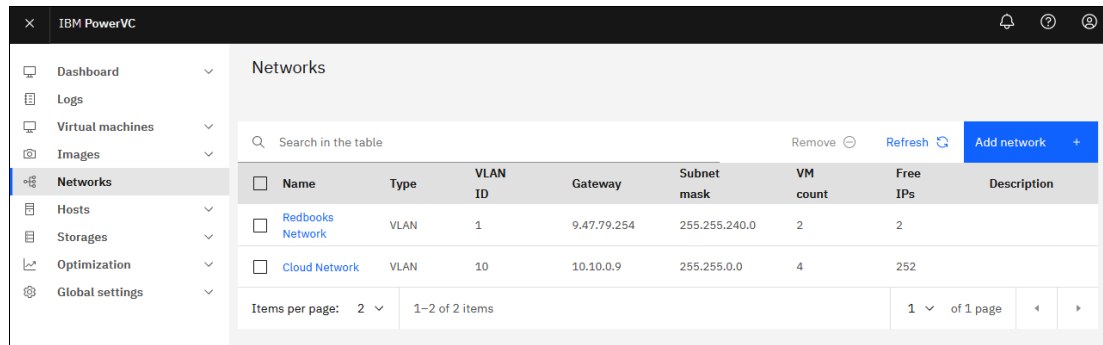


Figure 4-54 Networks in PowerVC

To edit a network, click the network name in the Networks tab. Some values can be changed after creation of a network, for example the IP address range, or if a network should be shared across projects.

In the edit mode, it is also possible to see the used IP addresses of the network in IP Address tab as shown in Figure 4-55.

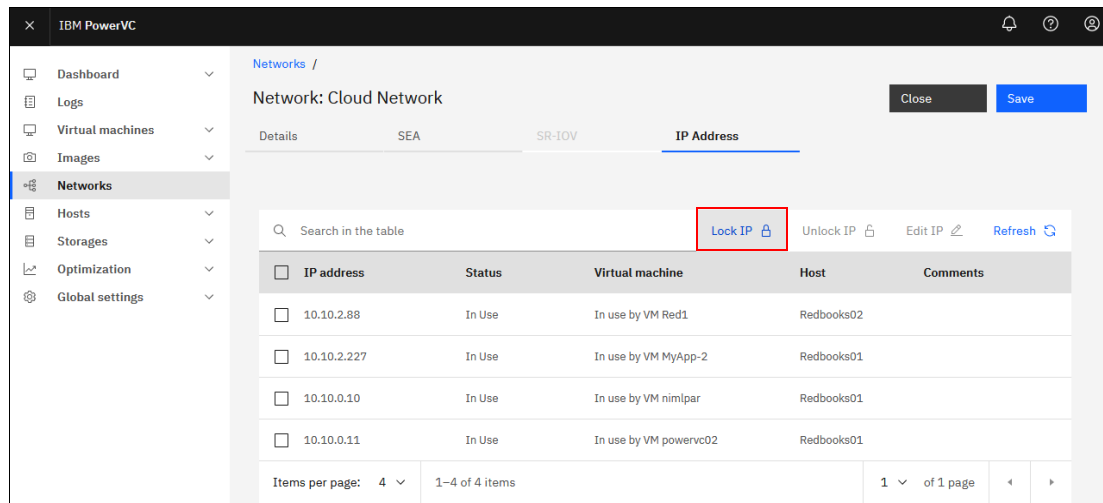
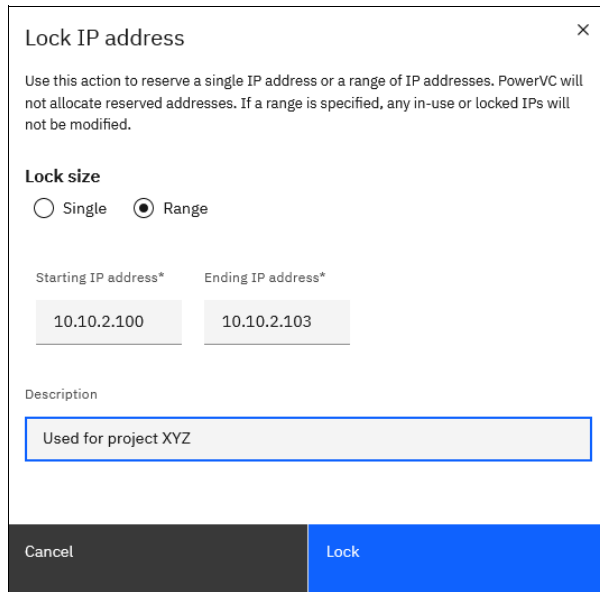


Figure 4-55 Used IP addresses of a network

PowerVC also offers to lock IP addresses. This means PowerVC will prevent deployments of new VMs to use the locked IP addresses. To do that, click on **Lock IP** and the dialog shown in Figure 4-56 opens.



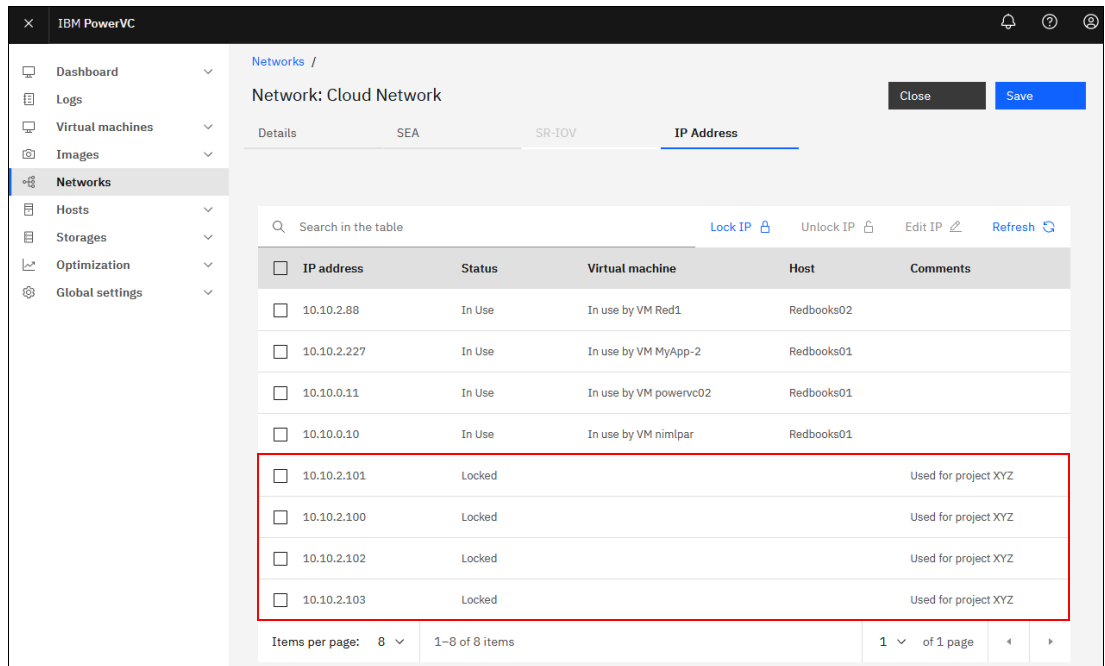
The dialog box titled "Lock IP address" contains the following elements:

- Instructions:** "Use this action to reserve a single IP address or a range of IP addresses. PowerVC will not allocate reserved addresses. If a range is specified, any in-use or locked IPs will not be modified."
- Lock size:** Two radio buttons, "Single" (unselected) and "Range" (selected).
- Starting IP address*:** A text input field containing "10.10.2.100".
- Ending IP address*:** A text input field containing "10.10.2.103".
- Description:** A text input field containing "Used for project XYZ".
- Buttons:** "Cancel" (grey) and "Lock" (blue).

Figure 4-56 Lock IP addresses

Here you can select a single IP address or a range of IP addresses that should be locked. Enter the IP addresses and a Description if you want. After that click **Lock**.

Back to the IP Address tab, click **Refresh** and you will see the locked IP addresses as shown in Figure 4-57.



The screenshot shows the "IP Address" tab in the IBM PowerVC interface. A table lists IP addresses with their status and associated virtual machines. The last four rows are highlighted with a red border, indicating they are locked.

IP address	Status	Virtual machine	Host	Comments
<input type="checkbox"/> 10.10.2.88	In Use	In use by VM Red1	Redbooks02	
<input type="checkbox"/> 10.10.2.227	In Use	In use by VM MyApp-2	Redbooks01	
<input type="checkbox"/> 10.10.0.11	In Use	In use by VM powervc02	Redbooks01	
<input type="checkbox"/> 10.10.0.10	In Use	In use by VM nimlpar	Redbooks01	
<input type="checkbox"/> 10.10.2.101	Locked			Used for project XYZ
<input type="checkbox"/> 10.10.2.100	Locked			Used for project XYZ
<input type="checkbox"/> 10.10.2.102	Locked			Used for project XYZ
<input type="checkbox"/> 10.10.2.103	Locked			Used for project XYZ

Figure 4-57 Locked IP addresses in a network

Note: When you unmanage a VM, the IP address will switch from In Use to Locked in the IP Address page to avoid double usage of the IP address. In the Comments field you see the hint Unmanaged VM <Name>. If you delete the unmanaged VM you may want to unlock the IP address.

If you want to unlock IP addresses, simply select the IP addresses and click **Unlock IP**. You then will be asked for confirmation and after a refresh the IP address disappears. You can unlock IP addresses in the state In Use, only Locked.

To delete a whole network enter the Networks page as shown in Figure 4-54 on page 147. There select the network to delete and click **Remove**. You will be asked for a confirmation. After clicking the red **Remove** button, the network will be deleted.

Note: Networks can only be deleted if no virtual machines are using it.

4.12.2 SR-IOV network setup

You can deploy virtual machines that leverage Single Root Input/Output Virtualization (SR-IOV). SR-IOV supports pass-through of Ethernet data from guest virtual machines directly to the hardware. This improves performance by allowing data to pass directly from guest virtual machines to physical adapters with minimal processing between, allowing a guest virtual machine to achieve near wire-speed Ethernet performance. SR-IOV also supports some additional configuration options, such as quality of service (QoS) for enforcing bandwidth allocations to guest virtual machines.

A given SR-IOV adapter can have multiple physical ports, connected to external switches. Each physical port is divided into logical ports, so called Virtual Functions (VFs). These logical ports are connected to a virtual machine for network connectivity. These logical ports allow a single physical hardware device to appear as multiple devices to guest virtual machines.

VFs could be assigned directly to a VM. In this case the VM sees a physical adapter similar to a complete network port. Such a configuration is not supported with PowerVC as this is against the principal that all resources must be virtual. Therefore PowerVC uses the vNIC technology, which abstracts the hardware layer. In this case the VF is assigned to a VIOS that connects it to a vNIC server. The client VM has a vNIC client that connects to the vNIC server. This is a similar approach as NPIV for the storage. Additionally a vNIC client can attach to multiple vNIC servers to provide redundancy.

Figure 4-58 illustrates the different ways that SR-IOV can be used to connect physical ports to virtual machines.

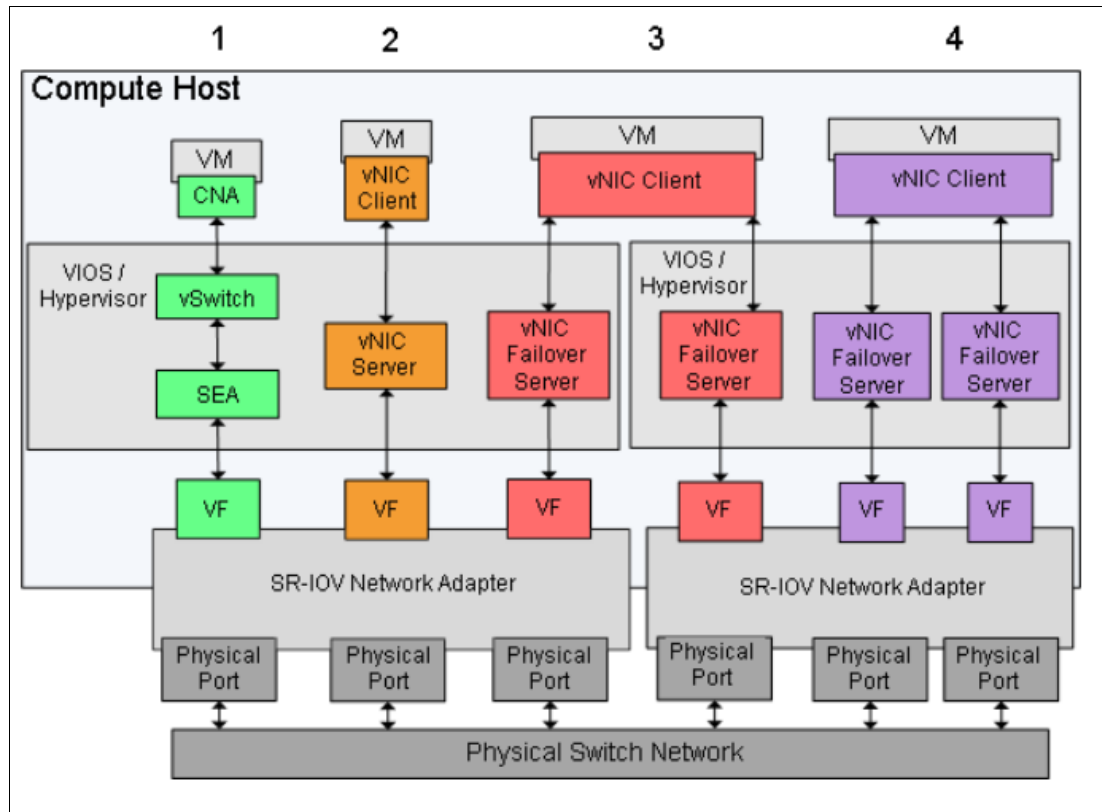


Figure 4-58 SR-IOV connectivity options

PowerVC can connect a virtual machine that uses an SR-IOV adapter by using options 2, 3 or 4, depending on the specified redundancy levels. Users can implement option 1 outside of PowerVC by carving a virtual function out of a physical port and assigning it as a backing networking device for a Shared Ethernet Adapter (SEA). PowerVC treats that SEA the same as any other SEA, regardless of the backing device.

Note: Previous versions of PowerVC supported SR-IOV with vNIC only for NovaLink-connected hosts. This is no longer the case. Also HMC-connected systems are supported for SR-IOV with vNIC usage.

To configure an SR-IOV network, complete the following steps:

1. Enter the host list under Hosts. There click on the hostname with SR-IOV ports. Inside the host details, click on **SR-IOV ports** to see the available SR-IOV of the host as shown in Figure 4-59. If there is no SR-IOV ports tab, then most likely this host does not have adapters configured in SR-IOV mode.

Physical port	Status	Configured Port Speed	Available capacity	Available logical ports	Total logical ports	Physical network name
<input type="checkbox"/> U78D2.001.WZ500AL-P1-C11-T2	True	10Gbps	100	16	16	default
<input type="checkbox"/> U78D2.001.WZ500AL-P1-C11-T1	True	10Gbps	100	11	11	default
<input type="checkbox"/> U78D2.001.WZ500AL-P1-C11-T3	False	Unknown	100	16	16	default
<input type="checkbox"/> U78D2.001.WZ500AL-P1-C11-T4	False	Unknown	100	16	16	default
<input type="checkbox"/> U78D2.001.WZ500AL-P1-C8-T2	True	10Gbps	90	15	16	default
<input type="checkbox"/> U78D2.001.WZ500AL-P1-C8-T1	True	10Gbps	88	14	16	default
<input type="checkbox"/> U78D2.001.WZ500AL-P1-C8-T4	True	10Gbps	100	16	16	default
<input type="checkbox"/> U78D2.001.WZ500AL-P1-C8-T3	True	10Gbps	100	16	16	default

Figure 4-59 SR-IOV ports of a host

In the example shown in Figure 4-59, you see the host has two SR-IOV capable adapters (P1-C8 and P1-C11) with four ports each. Six ports have a status of True and can be used. You also can see:

- Configured Port Speed
The network speed at which the physical port operates.
- Available capacity
The percentage of bandwidth capacity that is not yet allocated to any logical port.
- Available logical ports
Remaining available logical ports.
- Total logical ports
The total number of logical ports to use.
- Physical network name
Current name that is associated to the physical port. Specifying a physical network name allows you to configure which ports are attached to which physical network in your environment. You can choose an existing physical network name or give this network a new name. This will be the basis for a network definition in PowerVC.
To define a physical network with redundancy give at least two ports the same physical network name.

Note: For adapter redundancy select ports from different adapters. For optimal redundancy connect the ports to different switches.

To change the Physical network name of a port or multiple ports, select them and click **Edit physical network** as shown in Figure 4-60.

The screenshot shows the SR-IOV ports configuration page for host Redbooks05. A table lists 8 ports with columns for Physical port, Status, Configured Port Speed, Available capacity, Available logical ports, Total logical ports, and Physical network name. Two ports are selected (checked): U78D2.001.WZS00AL-P1-C11-T1 and U78D2.001.WZS00AL-P1-C8-T4. The 'Edit physical network' button is highlighted with a red box.

Physical port	Status	Configured Port Speed	Available capacity	Available logical ports	Total logical ports	Physical network name
<input checked="" type="checkbox"/> U78D2.001.WZS00AL-P1-C11-T1	True	10Gbps	100	11	11	default
<input type="checkbox"/> U78D2.001.WZS00AL-P1-C11-T3	False	Unknown	100	16	16	default
<input type="checkbox"/> U78D2.001.WZS00AL-P1-C11-T2	True	10Gbps	100	16	16	default
<input type="checkbox"/> U78D2.001.WZS00AL-P1-C8-T2	True	10Gbps	90	15	16	default
<input type="checkbox"/> U78D2.001.WZS00AL-P1-C8-T1	True	10Gbps	88	14	16	default
<input type="checkbox"/> U78D2.001.WZS00AL-P1-C11-T4	False	Unknown	100	16	16	default
<input type="checkbox"/> U78D2.001.WZS00AL-P1-C8-T3	True	10Gbps	100	16	16	default
<input checked="" type="checkbox"/> U78D2.001.WZS00AL-P1-C8-T4	True	10Gbps	100	16	16	default

Figure 4-60 Available SR-IOV ports

In this example we choose two redundant adapters and set the name to MySRIOVNet as shown in Figure 4-61.

The dialog box titled 'Edit physical network name' contains the following fields and options:

- Instructions: Use this form to edit the physical network name for the selected physical ports.
- Radio buttons: Add a new name, Select from existing names
- Name *:
- Physical network name *:
- Comment:
- Buttons: Cancel, Save

Figure 4-61 Set the name of a physical network

2. Back in the SR-IOV ports tab, click **Refresh** to see the result as shown in Figure 4-62.

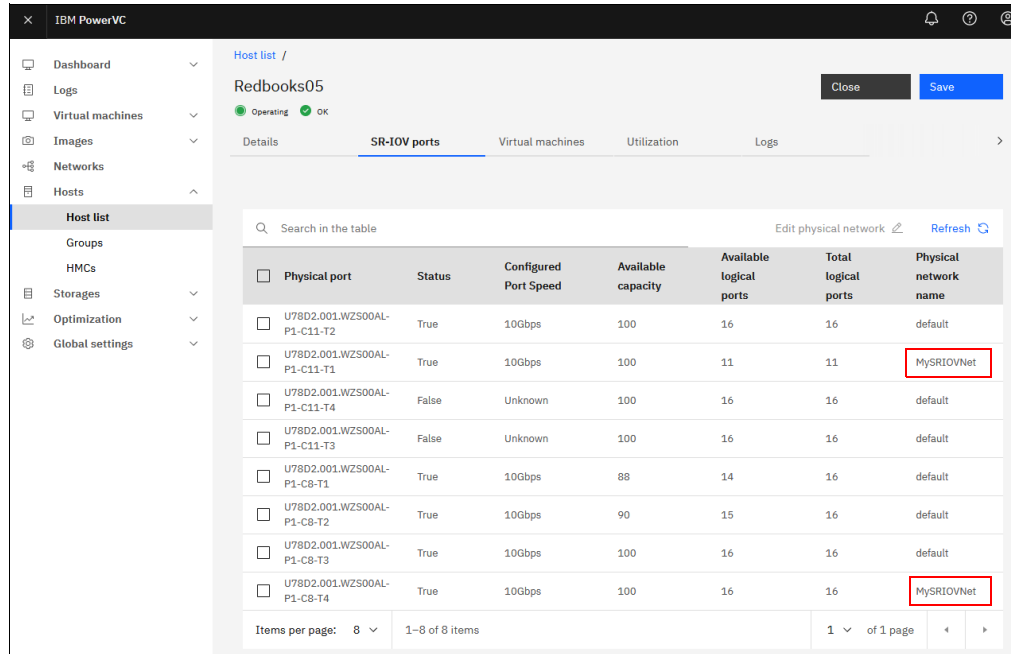


Figure 4-62 Physical SR-IOV network in PowerVC

The network name in PowerVC corresponds to the Label of the port in the HMC or NovaLink as shown in Figure 4-63.

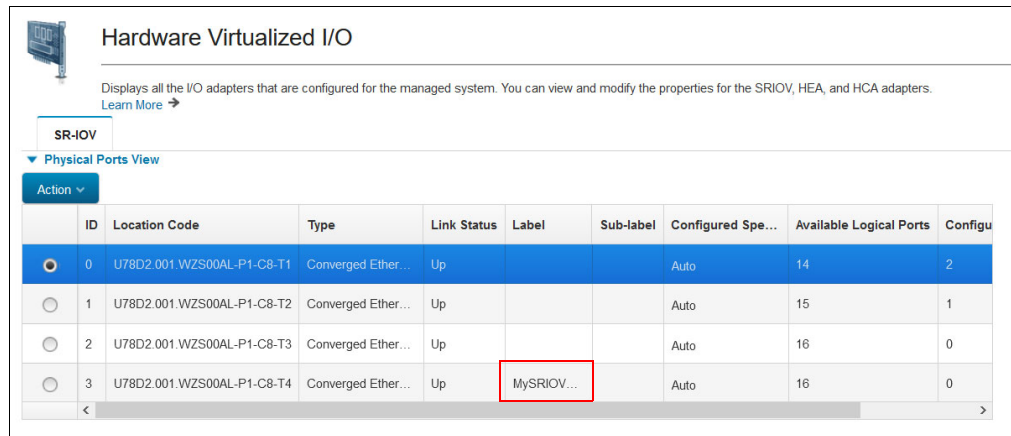


Figure 4-63 SR-IOV ports in the HMC

Notes:

You cannot rename a physical network that is in use by other VMs. Also if these VMs are not managed by PowerVC.

The port label is updated to match the physical network name, but the port *sublabel* is not changed and remains available for your use.

3. Define the PowerVC network.

To define a network in PowerVC that uses SR-IOV with vNIC first follow the steps described in 4.12.1, “Basic network setup” on page 144. After filling the attributes in the Details tab click on the **SR-IOV tab**. Now select the Physical network name from the drop down list that you want to use for that network as shown in Figure 4-64.

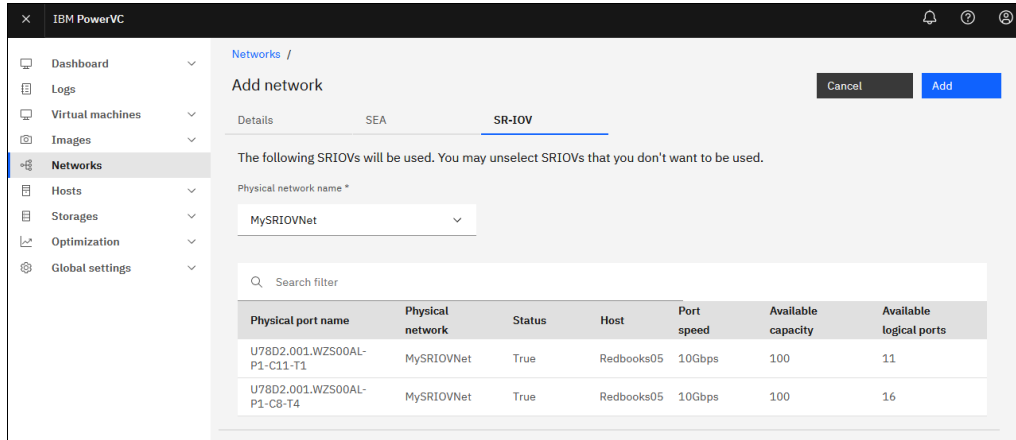


Figure 4-64 Physical SR-IOV networks available for the PowerVC network

After clicking **Add**, the PowerVC network with SR-IOV ports will be created.

The usage of a network with SR-IOV is described in 4.16.7, “Deploying a new virtual machine” on page 239.

4.13 Compute template setup

A *compute template* provides a predefined compute configuration to use when you deploy a new VM. By default, IBM PowerVC provides six compute templates with different sizes, from tiny to xlarge. You select a compute template when you deploy a VM. You can change the values that are set in the compute template that is associated with a VM to resize it. You can also create your own compute templates.

For the full description about compute templates, see section 2.4.4, “Information that is required for compute template planning” on page 34.

To create a new Compute template, complete the following steps:

1. On the left side of the IBM PowerVC window, click on **Virtual machines** → **Compute templates**. and click **Create**. In the next window, enter information for the new Compute template in the **Details** tab as shown in Figure 4-65 on page 155:
 - Name of the template
 - Virtual processors
 - Processing units
 - Shared processor pool
 - Memory

Create compute template Cancel Create

Details Miscellaneous

Template name *

Virtual processors ⓘ

Min * Desired * Max *

Use shared processors

Processing units ⓘ

Min * Desired * Max *

Shared processor pool * ⓘ

Uncapped

Weight (0-255)

Dedicated idle processor sharing

Memory (GB) ⓘ

Min * Desired * Max *

Figure 4-65 Create Compute template

- After you click **Create**, the list of available compute templates is displayed as shown in Figure 4-66.

Compute templates

Search in the table Delete Refresh Create

<input type="checkbox"/>	Name	Processors	Processing units	Memory (GB)
<input type="checkbox"/>	1VP4RAM	1	1	4
<input type="checkbox"/>	Dedicated-1	1		2
<input type="checkbox"/>	Small (Remote Restartable)	2	0.2	2

Figure 4-66 Compute template list

4.14 Environment verification

After you added the hosts, storage providers, networks, and templates, as a preferred practice, verify your IBM PowerVC environment before you try to capture, deploy, or onboard VMs.

Virtualization management function failures might occur when dependencies and prerequisite configurations are not met.

IBM PowerVC reduces the complexity of virtualization and cloud management. It can check for almost all required dependencies and prerequisite configurations and clearly communicate the failures. It can also accurately pinpoint validation failures and remediation actions when possible.

To check an environment, complete the following steps:

1. On the left side of the IBM PowerVC window, click on **Dashboard** → **Environment checker** and click **Check now** as shown in Figure 4-67.

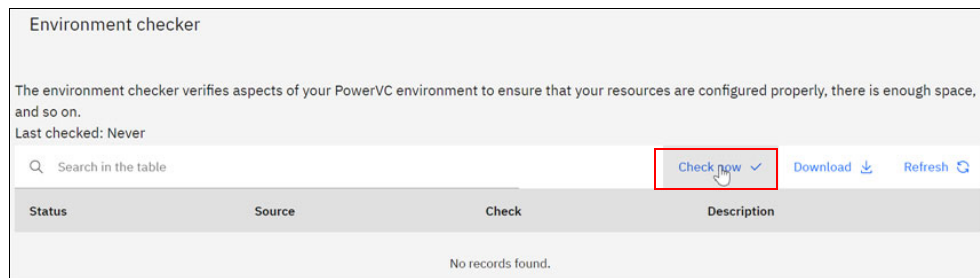


Figure 4-67 Check environment

2. After the environment check finish, verify the results that can be, passed, warning, failed or info as shown in Figure 4-67.

The screenshot shows the 'Environment checker' interface after a check. It displays the same description and search bar as Figure 4-67. The 'Last checked' time is now '2/23/2021, 9:55:04 PM'. Below this, it shows the results: 'Environment Checker Results: Passed(162),Warning(11), Failed(2), Info(0)'. The 'Check now' button is still present. The table below shows the following data:

Status	Source	Check	Description
Passed	9.114.181.33	Brocade Fabric validations	The switch's firmware level is acceptable.
Passed	9.114.181.33	Brocade Fabric validations	Zoning is enabled on the fabric managed by this switch.
Passed	9.114.181.34	Brocade Fabric validations	The switch's firmware level is acceptable.
Passed	9.114.181.34	Brocade Fabric validations	Zoning is enabled on the fabric managed by this switch.
Passed	powervc02	Certificate expiration validations	PowerVC web certificate will not expire in the next six months.
Passed	powervc02	Certificate expiration validations	PowerVC messaging CA certificate will not expire in the next six months.

Figure 4-68 Check environment result

4.15 Management of virtual machines

The following sections describe the operations that can be performed on VMs by using the IBM PowerVC management host:

- ▶ 4.15.1, “Virtual machine onboarding” on page 158
- ▶ 4.15.2, “Start, stop, and restart of virtual machines” on page 167
- ▶ 4.15.3, “Opening a console window” on page 169
- ▶ 4.15.4, “Resizing the virtual machine” on page 171
- ▶ 4.15.5, “Attaching and detaching of network adapters” on page 174
- ▶ 4.15.6, “Attaching, detaching, and cloning of storage volumes” on page 176
- ▶ 4.15.7, “Migration of virtual machines” on page 181
- ▶ 4.15.8, “VM snapshots” on page 186
- ▶ 4.15.9, “Cloning a virtual machine” on page 190
- ▶ 4.15.10, “Host maintenance mode” on page 192
- ▶ 4.15.11, “Restarting virtual machines remotely from a failed host” on page 196
- ▶ 4.15.12, “Resetting the state of a virtual machine” on page 205

Most of these operations can be performed from the **Virtual machines** window, as shown in Figure 4-69.

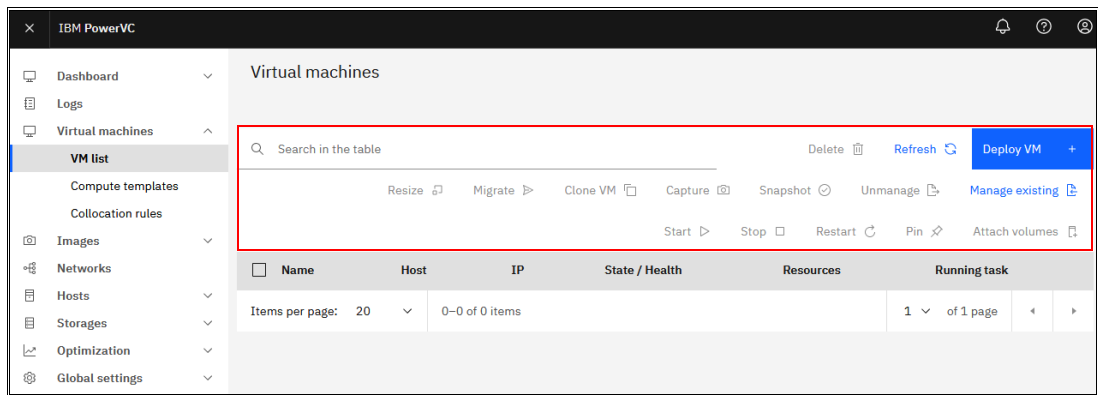


Figure 4-69 VM management operations in the VM list

However, remote restarting of VMs or putting a host into maintenance mode, that may trigger Live Partition Mobility operations, are initiated from the **Host list** as shown in Figure 4-70.

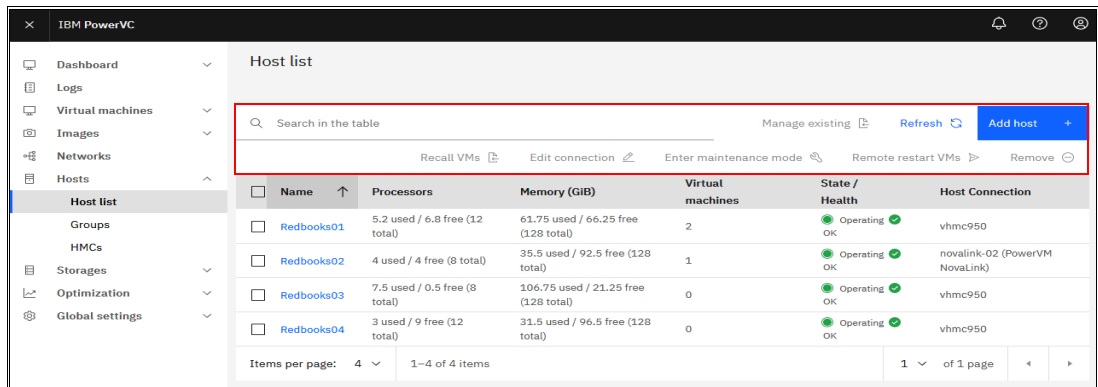


Figure 4-70 VM management operations in the host list

4.15.1 Virtual machine onboarding

IBM PowerVC can manage VMs that were not deployed by IBM PowerVC, such as VMs that were created manually or by other tools. To add an existing VM, complete the following steps:

1. In the **VM list** under **Virtual machines** as shown in Figure 4-69 on page 157 select the **Manage existing** button.
2. In the following dialog, select the host where the VM is running. After that you may select all virtual machines that are fully supported by PowerVC, or select specific virtual machines. By selecting **Select specific virtual machines** you can select the VMs you want to import under the Supported tab as shown in Figure 4-71. You can also find VMs that are not supported under the Not supported tab. Most likely virtual machines are not supported if they are not fully virtualized, meaning they have hardware components in the LPAR configuration, such as PCIe adapters or SR-IOV ports (without vNIC).

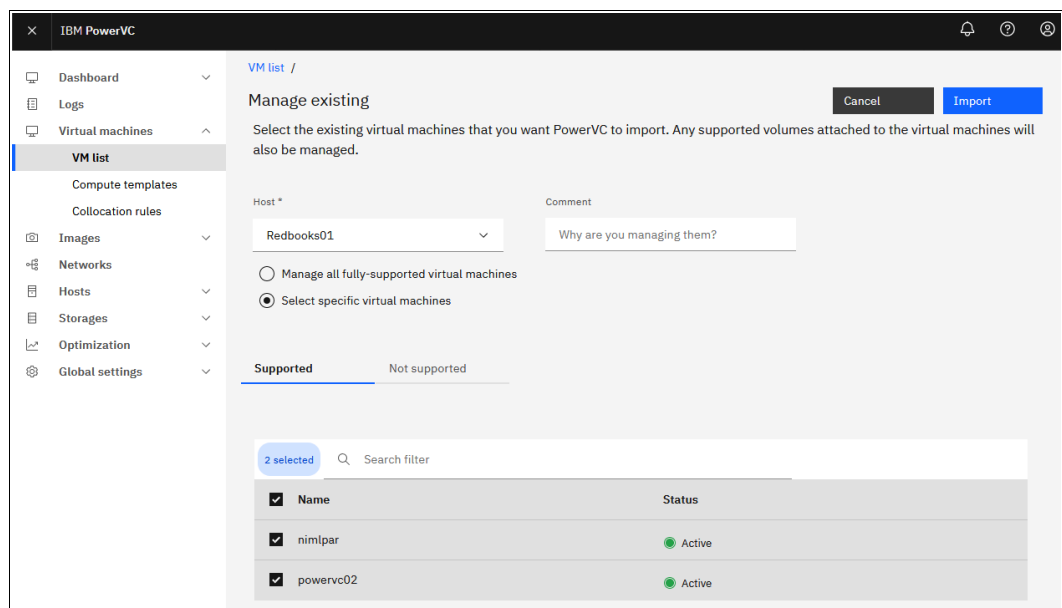


Figure 4-71 Onboard a VM

3. After you selected the VMs that you want to onboard, click the **Import** button.

Note: The detailed eligibility requirements to add a VM into an IBM PowerVC managed PowerVM host are available in IBM Knowledge Center:

https://www.ibm.com/support/knowledgecenter/en/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_capture_reqs_hmc.html

After you click **Import**, it may take a while until the virtual machines show up in the list, because the import is done in the background.

PowerVC does many tasks in the background to give the user the ability to continue working with PowerVC while the task is running. This may lead sometime to confusion if the task is really executed or not. This behavior is much distinct with PowerVC Version 2.0.

When you are unsure if the task was really started, have a look into the messages column on the right. You can show it, by clicking the bell symbol in the upper right corner as shown in Figure 4-72.

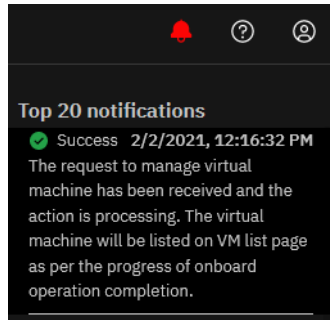


Figure 4-72 Onboard message

You may click the **Refresh** button to see it before the auto update of the page takes place to see the onboarded VM.

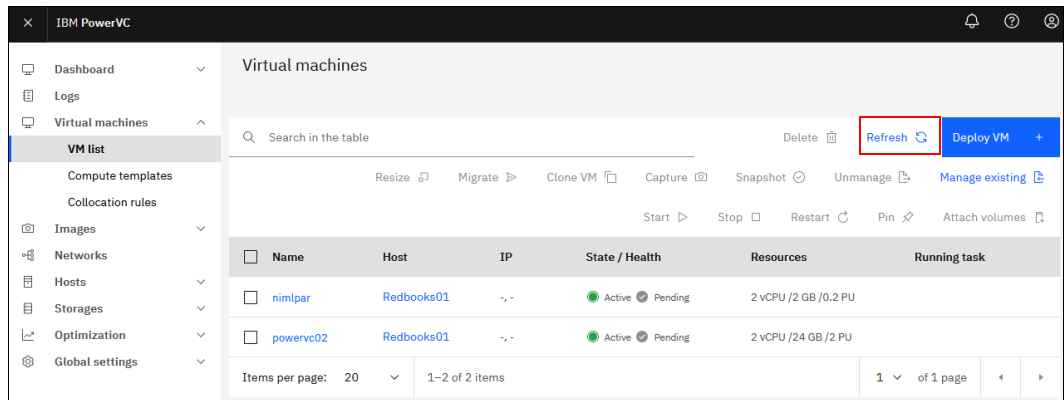


Figure 4-73 Onboarded VMs

Note: On many IBM PowerVC windows, you can see a Refresh button, as shown for example in Figure 4-73. Most windows update asynchronously through long polling in the background. By clicking the **Refresh** icon, a Representational State Transfer (REST) call is made to the IBM PowerVC server to get the latest data that is available from IBM PowerVC.

4. Manually add additional information for onboarded VMs.

There is some information that PowerVC is not able to discover automatically that should be added manually. These will be discussed now.

By clicking on the **VM name** you can enter the VM details as shown in Figure 4-74.

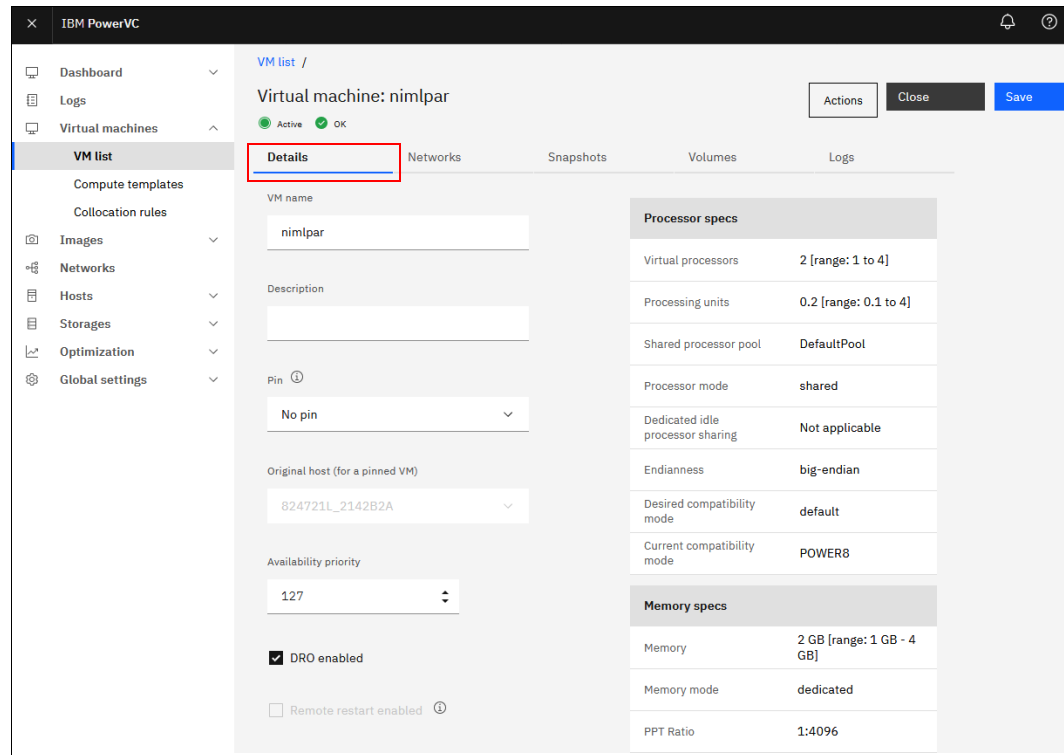


Figure 4-74 VM details

In the VM details you can find a lot of information such as the Processor specs or Memory specs that have been imported. In the details there are also a lot of attributes that can be changed (Table 4-1).

Table 4-1 VM attributes that can be changed

Field	Description
Name	The name of the VM. If a VM was deployed with PowerVC a unique ID will be added to the VM name. In the HMC you can see complete name with ID, in PowerVC the ID is removed from the name. When you import a VM that was formerly deployed by a PowerVC instance, the ID may also show up. In this case you can edit then name and remove the ID. This does not remove the ID from the LPAR in HMC, only in PowerVC.
Description	A user editable field for providing descriptive information.
Pin	When a VM is hard pinned, a partition migration (Live Partition Mobility) or Remote Restart is not permitted in PowerVC. Although such an operation is still permitted from the HMC or in NovaLink, as this is a PowerVC feature. Soft pinned VMs may move using partition migration in a system evacuation (maintenance mode). If the Recall option in the host is enabled, the VM will also migrated back, when the maintenance mode is over. Below the Pin option the original host of the VM is displayed. When a recall is executed after a maintenance mode, this is the host where the VM will be moved to.

Field	Description
Availability priority	The Availability priority is an attribute of a Power LPAR. It specifies the importance of a VM. In case a processor needs to be deconfigured from the system and the entitled capacity is fully assigned to VMs, then VM(s) need to be stopped to be able to not have more assigned entitled capacity as available processors after the deconfiguration. Remember: the entitled capacity are guaranteed processor resources to the VM, therefore you can not assign more than the number of available processors.
DRO enabled	Shall the VM be part of Dynamic Resource Optimizer (DRO) operations? DRO is covered in section 4.5.5, “Host groups” on page 126.
Enforce affinity check during LPM	After deployment, the virtual machine is assigned an affinity score based on processor and memory resource placement within the host. If affinity score checking is on, the virtual machine can only be migrated to a host if the resulting affinity score will be at least as high as the current score. This setting is supported on POWER9 hosts only.

In Figure 4-73 on page 159 you can see that the IP addresses for the network adapters of the VM are not imported. You can add them under the Networks tab for each network as shown in Figure 4-75. A prerequisite is that the IP address is available in the network IP address ranges. In case the you imported a VM with an IP address that is not in the IP address ranges, then first add the IP address. For more information about adding an IP address range to a network refer to section 4.12.1, “Basic network setup” on page 144.

When the IP address is added to the network adapter in the VM, this also makes sure that the IP address is flagged in use in the network to prevent duplicated IP addresses.

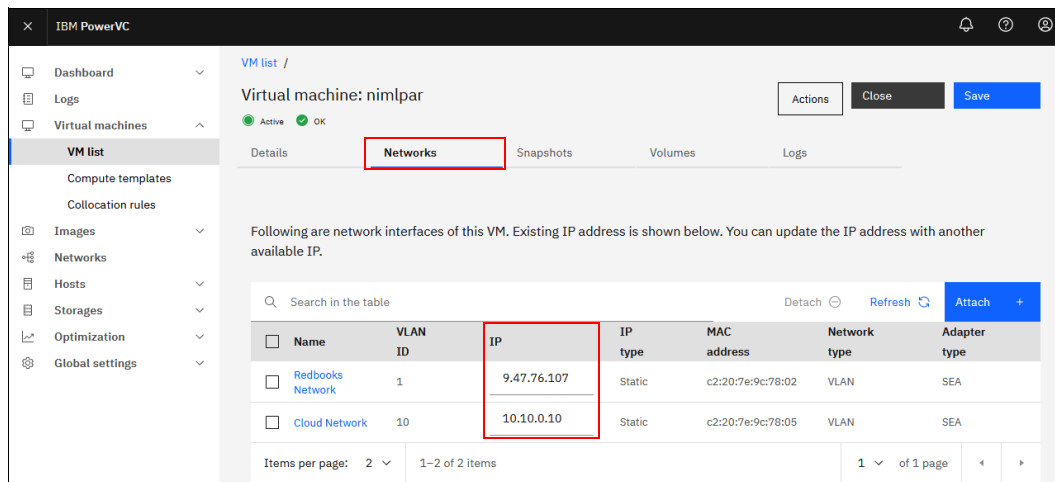


Figure 4-75 Enter IP addresses for onboarded VMs

Also SR-IOV (with vNIC) connected networks are imported and shown in the Networks tab. In this case PowerVC also shows additional SR-VIO related information, such as Redundancy, which is the number of redundant path to the network, and also the vNIC capacity. Figure 4-76 shows an example of such a network port.

<input type="checkbox"/>	Name	VLAN ID	IP	IP type	MAC address	Network type	Adapter type	Redundancy	vNIC capacity
<input type="checkbox"/>	Cloud Network	100	10.10.2.2	Static	fa:16:3e:d0:f7:f2	VLAN	SRIOV	2	2 %

Items per page: 1 of 1 items

Figure 4-76 SR-IOV connected network adapter

PowerVC does not set the Operating system attribute of the VM. To set the Operating system attribute, you can click **Actions** in the VM details view as shown in Figure 4-74 on page 160. In the following screen shown in Figure 4-77 you can select the Operating system from the dropdown list. You also should select the Endianness, if it is not set automatically.

In the list of volumes, the boot disk has 0 as the first disk in the Boot Order. If you have mirrored boot disks, you should define the second boot disk with a 1 in the Boot Order attribute of the volume.

Edit details ✕

Fill in the details for the virtual machine and click save.

Operating system

AIX ▼

Endianness

Big Endian Little Endian

Please specify a number indicating boot order for boot volume. Boot order must start with 0. Please specify "Data" for data volume.

<input type="checkbox"/>	Name	Size	Boot Order
<input checked="" type="checkbox"/>	volume-nimvg- cea7bdf7-c04c	256	Data
<input checked="" type="checkbox"/>	volume-nimlpar- 8e817a65-00000003- boot-0-60b79fc3-3836	64	0

Cancel
Save

Figure 4-77 VM details for an imported VM

Figure 4-78 on page 163 shows the Volumes tab. There you find all LUNs that are connected to the VM. You can see that the imported volumes are not assigned to a Storage Template. It is also recommended to assign a storage template for the imported volumes in the Storages section of PowerVC, as described in the next step. But for that you need to leave the Virtual Machine details and therefore don't forget to click the **Save** button for your unsaved changes.

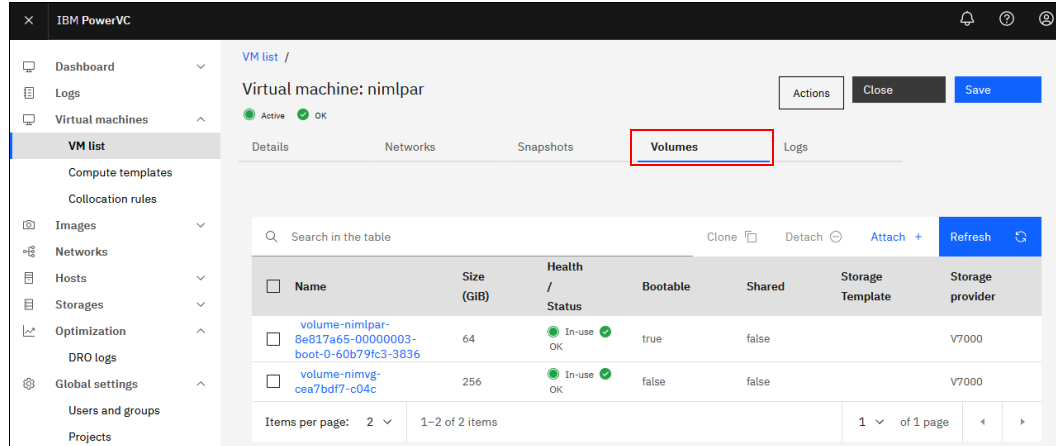


Figure 4-78 Volumes of a virtual machine

To assign a storage template for the boot volumes, enter **Boot volumes** under **Storages** in the navigation area and select the volume to change as shown in Figure 4-79. Then select the **Set storage template** button.

Note: You can only select one volume to set the storage template after another.

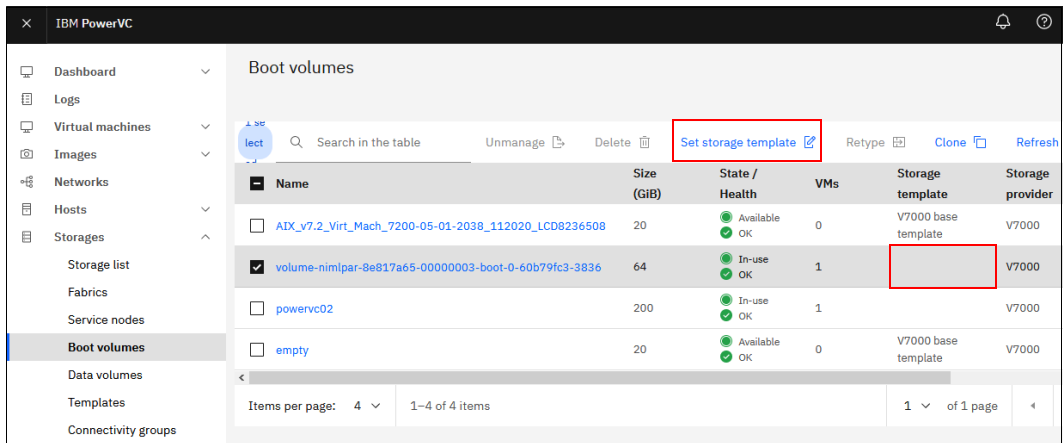


Figure 4-79 Boot volumes

In the next screen shown in Figure 4-80 select the storage template from the drop-down list and click the **Save changes** button.

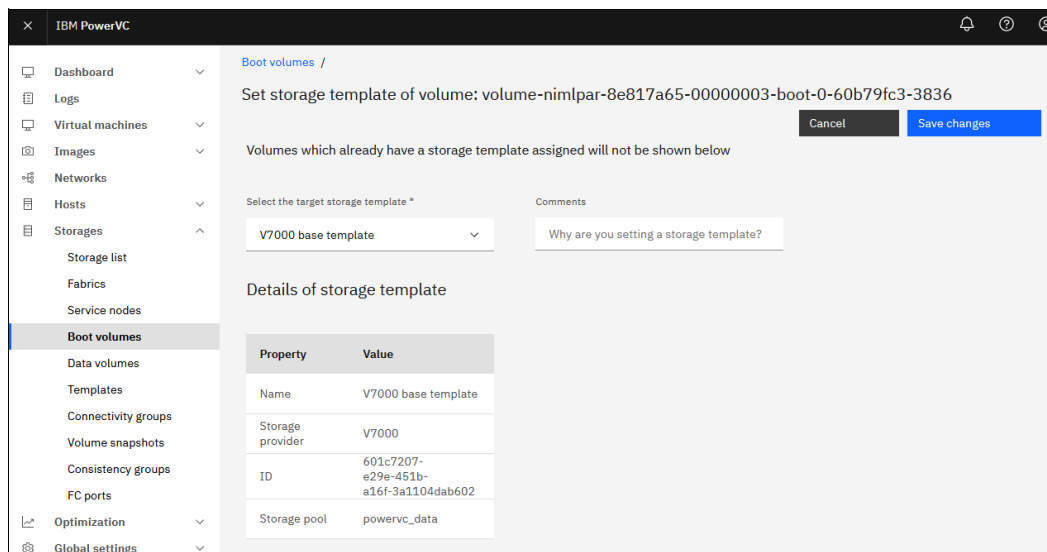


Figure 4-80 Set storage template

To set the storage template for the data disks attached to the VM, do the same procedure, but in the **Data volumes** list under **Storages**.

Useful information about a virtual machine in PowerVC

When you enter a VM in the **VM list** under **Virtual machines** as shown in Figure 4-74 on page 160 you find a lot of useful information. Table 4-2 shows informations about status icons.

Table 4-2 Information about status icons

Field	Description
State icon	The actual state for the VM (Active, Stopped)
Health Icon	The actual health status for the VM: <ul style="list-style-type: none"> ▶ OK: The target resource, all related resources, and the IBM PowerVC management services for the resources report zero problems. ▶ Warning: The target resource or a related resource requires user attention. Nova or Cinder host services that manage the resources report problems and require user attention. ▶ Critical: The target resource or a related resource is in an error state. ▶ Unknown: IBM PowerVC cannot determine the health status of the resource.

In Table 4-3 on page 165 you find information about the Processor specs of the virtual machine.

Table 4-3 Information about the processor specs

Field	Description
Virtual processors	Number of virtual processors assigned to the VM. In brackets you find the minimum and maximum number of VPs for that VM. You can dynamically resize the VPs in that range without restarting the VM.
Processing units	The entitled capacity (EC) for the VM. In brackets the minimum and maximum values.
Shared processor pool	The shared processor pool for the VM (none for VMs with dedicated processors)
Processor mode	Shared or dedicated processor allocation.
Dedicated idle processor sharing	If the VM is in dedicated processor mode, spare cycles can be donated to the shared processor pool. The values can be: <ul style="list-style-type: none"> ▶ Allow when virtual machine is inactive: The resources are donated when the VM is inactive. ▶ Allow when virtual machine is active: The resources are donated when the VM is active. ▶ Allow always: Donating is always allowed. ▶ Never: Donating is never allowed.
Endianness	Does the VM run in big- or little-endian mode.
Desired compatibility mode	Desired processor compatibility mode when the VM starts.
Current compatibility mode	Currently running processor compatibility mode.

The specs for memory are listed in Table 4-4.

Table 4-4 Information about the memory specs

Field	Description
Memory	Amount of memory that is assigned to the VM. In brackets you find the minimum and maximum amount of memory for that VM. You can dynamically resize the memory in that range without restarting the VM.
Memory mode	Shared or dedicated memory usage.
PPT Ratio	The PPT ration defines the size of the translation table from virtual memory addresses to the physical memory addresses. The parameter affects the size of an addressable memory chunk. This is important for Live Partition Mobility. If the chunks are to big in system with a high memory change rate, an LPM operation may not come to an end. On the other side, if the size is small more memory is needed for the translation table. More information is available in the following blog: https://community.ibm.com/community/user/power/blogs/hariganesh-muralidharan1/2020/06/08/hmc-v8-r870-sp1-enhancements

The last part in the Details tab of the VM are the Details specs as shown in Table 4-5.

Table 4-5 Information about the memory specs

Field	Description
ID	The unique ID of the virtual machine. Each host, network, VM, and any other resource that is created in the IBM PowerVC management host has its own ID number. This ID uniquely identifies each resource to the IBM PowerVC management host.
Partition name	The name of the LPAR as shown in the Hardware Management Console.
Host	Host of the VM.
Deployed image	Deployed image or No Image (Onboarded VM).
Last update	Date of the last change of the VM in PowerVC. This is not necessarily a user change. A simple state change of the VM is sufficient to change the date.
Remote restart state	State of Remote Restart (Invalid, Local_Storage_Update_Failed, Remote_Restartable). The state of Invalid usually means the VM is not enabled for remote restart, usually because the attribute in the VM is not set. You may change the attribute in PowerVC, if the VM is stopped. A different status usually means that there may be something wrong.
Power state	Running, Stopped.
Task state	State of the task that is running. States may be for example stopping, starting or resizing.
Secure boot	Secure boot activated for this VM?
Performance data collection enabled	Is performance data collection enabled?
Storage connectivity group	Storage connectivity group that was used for the deployment. This attribute will not be visible, when the VM was imported.
Operating system	Version of the operating system. For AIX the technology level as well as the service pack will be displayed. For Linux it shows the name of the operating system and the running kernel version. For IBM i it shows the version and the Technology Refresh (the last number on the right). If the operating system can't be discovered, this attribute is "Unknown".

At the bottom of the details tab you can also find processor and memory usage information if these are available for the VM.

Note: To see performance information, the performance data collection must be turned on for the host in the HMC.

Memory usage information will be only displayed for NovaLink connected hosts.

All the details existing Collocation rules are displayed for the VM. For more information about collocation rules, please refer to 4.15.1, "Virtual machine onboarding" on page 158.

Under the Networks tab as shown in Figure 4-75 on page 161, you find all the information about the networks, such as VLAN ID, IP address, IP address type (static or DHCP), the MAC address, the network type and the adapter type (SEA or SR-IOV).

Under the Snapshots tab you can manage snapshots for a VM. This topic is covered in 4.15.8, “VM snapshots” on page 186.

On the Volumes tab all volumes for this VM are displayed. You find also information about these volumes, such as the size, health status, bootable, shared, the Storage Template and the Storage provider. Shared means that this volume can be connected to multiple VMs, for example to be used in an high availability cluster.

The Logs tab on the right shows all log entries that are related to this VM. This is very helpful, because you do not need to search for VM related entries in the general log with all other entries. The general log is available in the left hand navigation area.

4.15.2 Start, stop, and restart of virtual machines

The following sections guide you through the major operations of a virtual machine.

Start virtual machines

To start a virtual machine enter the VM List under Virtual machines as shown in Figure 4-81.

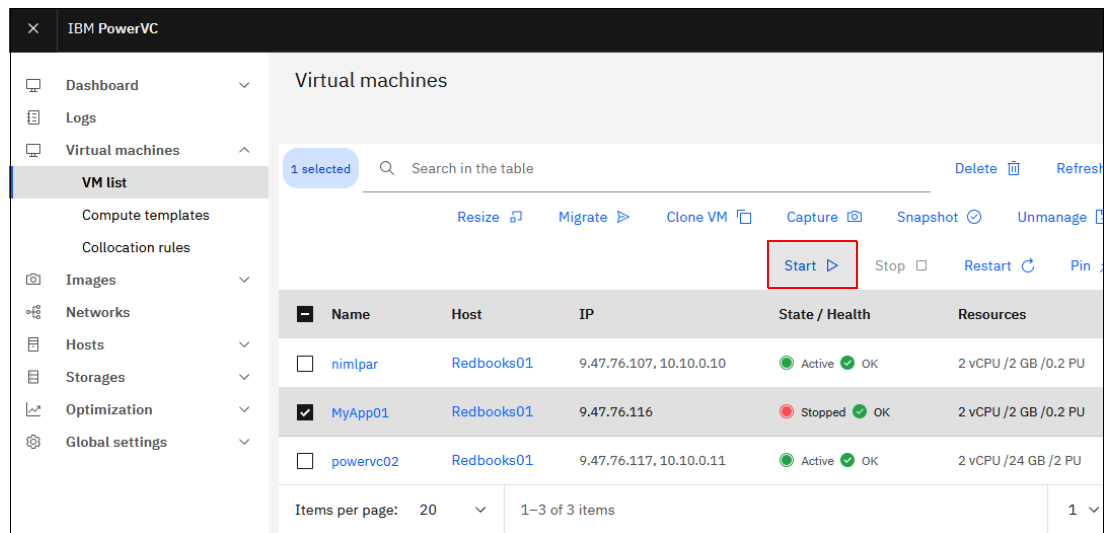


Figure 4-81 Start a virtual machine

Note: Since PowerVC 2.0 it is possible to select multiple VMs to be started, stopped, or restarted.

After the VM finishes the start process, it is available for operations that are performed through the IBM PowerVC management host. The process takes more time than the boot process of the operating system. IBM PowerVC waits until the RMC service is available to communicate with the VM. Even though the State field is Active (because the VM is powered on), the Health field may be still in Warning. Wait for a few minutes for the Health field to display a status of OK before you manage the VM from IBM PowerVC.

Note: When a VM is stopped, but was started with PowerVC and had an Health state of OK, the state remains when it is stopped. When you start this VM, the Health state may change to Warning, because after the start the RMC connection is not immediately available. Usually the state should turn to OK again, after some minutes.

Stop virtual machines

To stop one or multiple virtual machines, select the VMs to be stopped and click the **Stop** button. After clicking **Stop**, PowerVC asks for confirmation as shown in Figure 4-82. PowerVC triggers an operating system shutdown. If a VM is in a hang condition you may want to also select **Shut down immediately** to force the shutdown.

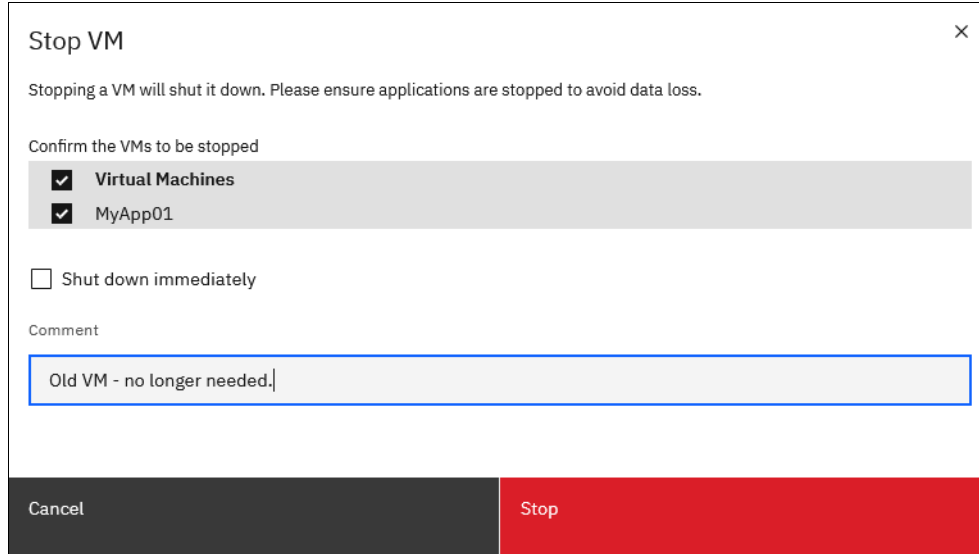


Figure 4-82 Stop a VM

When the VM completes the shutdown process, the state changes to Stopped. This takes a few minutes to complete.

Note: If an active RMC connection exists between IBM PowerVC and the target VM, a shutdown of the operating system is triggered. If no active RMC connection exists, the VM is shut down without shutting down the operating system.

Restart virtual machines

To restart a VM you can select the VMs to be restarted and then click the **Restart** button. In the dialog shown in Figure 4-83 you can select if you want to do a soft or hard restart. Soft means the operating system will be stopped before the stop of the VM.

Restart VM

Restarting a VM will do a OS reboot. Please ensure applications are stopped to avoid data loss.
Soft restart - Shuts down the virtual machine normally.
Hard restart - Shuts down the virtual machine immediately without shutting down the operating system first. A hard restart might result in loss of data.

Confirm the VMs to be restarted

- Virtual Machines
- MyApp01

Soft Hard

Comment

Restart after OS update

Cancel Restart

Figure 4-83 Restart a VM

4.15.3 Opening a console window

IBM PowerVC supports access to a virtual console for NovaLink-managed VMs from the user interface. This allows convenient command-line access without leaving the PowerVC web interface. Additionally, it provides access to a VM even if the VM's networking is not functioning correctly.

To open a virtual console enter the details of a virtual machine by clicking on the name of the VM under the Virtual machines list. If the VM runs on a NovaLink-managed system there is a separate tab Console visible. After you clicked on the **Console** tab, the console opens automatically as shown in Figure 4-84.

Note: For HMC-connected systems this tab is not shown.

For a console no RMC connection is necessary, it will be available immediately after a deployment, when the VM was created on the host.

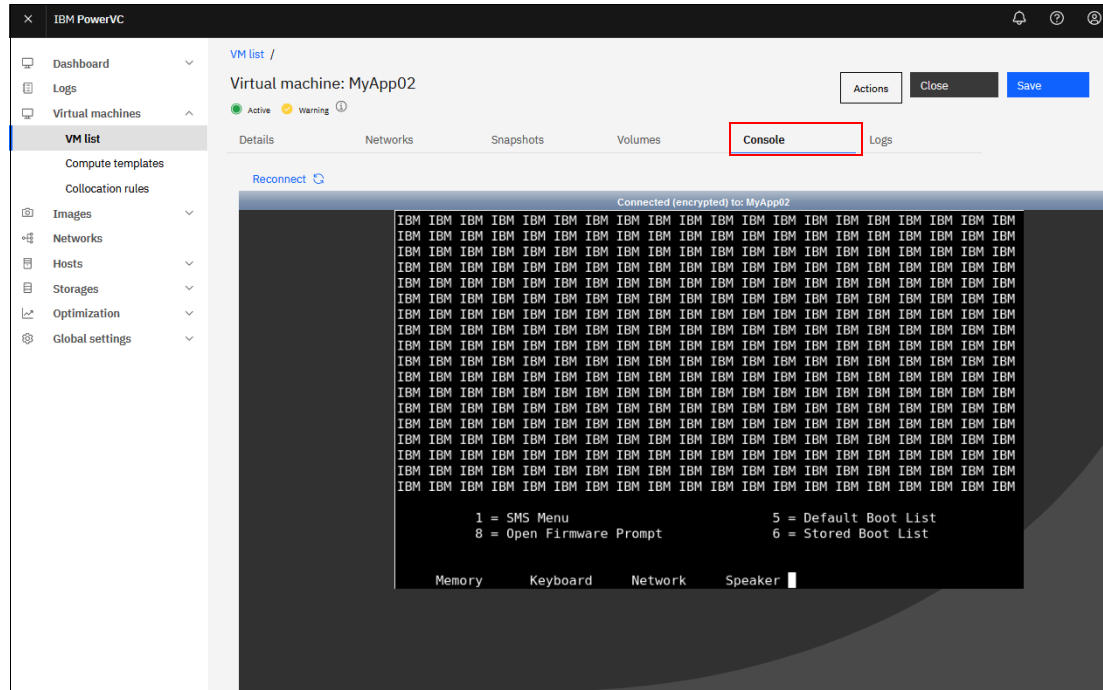


Figure 4-84 NovaLink console

When the console does not open, you may debug the issue in the NovaLink command line. There try to open a console by entering the following command:

```
# mkvterm --id <ID of the VM>
```

Note: To find the ID you may use the `pvmtl vm list` command.

If the command does not work, you get some errors to debug further.

One issue could be the error “There are no DR capable slots on this system”. In this case the VM may has no more virtual adapters. This could have happened, when you installed NovaLink as a normal VM and did not increase the number of Maximum Virtual Adapters from 10 to something higher, maybe 200. When you have an HMC connected to the NovaLink system in co-management mode, you can change the number of virtual adapters as follows:

1. Stop the NovaLink VM.
2. Let the HMC take over the control of the system with the following command:

```
# chcomgmt -m <system name> -o setmaster -t norm
```

You can verify if the HMC took over the control by looking at the lock on the server symbol in the HMC. If it changed from red locked to grey unlocked, the HMC has control.

3. In the General VM information in the HMC under Advanced Settings you find the Maximum Virtual Adapters. Change the value to a higher number.
4. Activate the NovaLink VM.
5. Give the control back to NovaLink with the following command:

```
# chcomgmt -m <system name> -o relmaster
```

With that the open grey lock in the HMC turns back to a closed red lock.

For HMC connected systems it is not possible to open a console window out of PowerVC. In this case either open a console from the HMC GUI, or login to the HMC using SSH and use the `vtmenu` command to open a menu-based console program as shown in Example 4-1.

Example 4-1 Output example of vtmenu

```
-----
Managed Systems:
-----
```

- 1) Redbooks01
- 2) Redbooks02
- 3) Redbooks03
- 4) Redbooks04

```
Enter Number of Managed System. (q to quit):
-----
```

From there select the number of the host where the VM runs and in the next screen select the number of the virtual machine and the console will open.

Note: To end a console you can use `~.` to disconnect.

4.15.4 Resizing the virtual machine

The IBM PowerVC management host can resize the managed VMs dynamically. Complete the following steps:

1. In the **VM list** under **Virtual machines** select the VM that you want to change and then click **Resize** on the upper bar on the window, as shown in Figure 4-85.

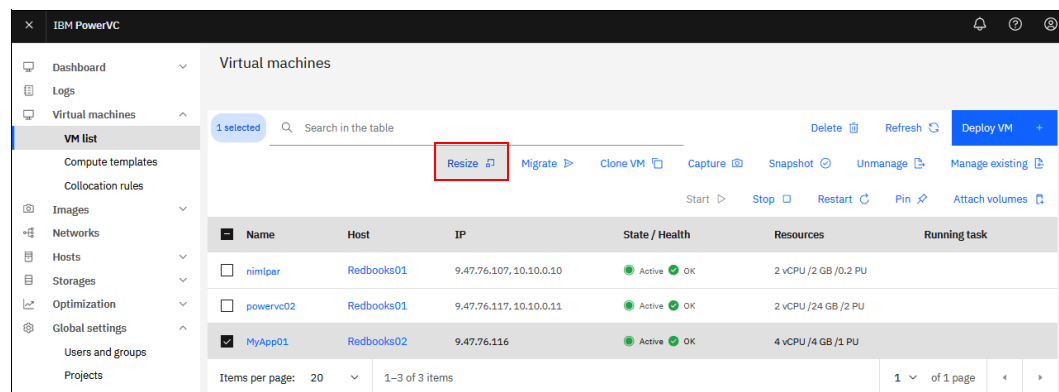


Figure 4-85 Resizing a VM

Make sure that the Health state of the VM is OK, which usually means that an active RMC connection to the VM is established. If this is not the case, the resizing will not work.

- In the next window (Figure 4-86), enter the new values for the resources or choose an existing compute template. In the brackets on the right of the values you can see the minimum and maximum values of the LPAR. You must enter a value that is in the interval that is defined in the brackets to ensure a concurrent resizing (Dynamic LPARing) while the VM is running.

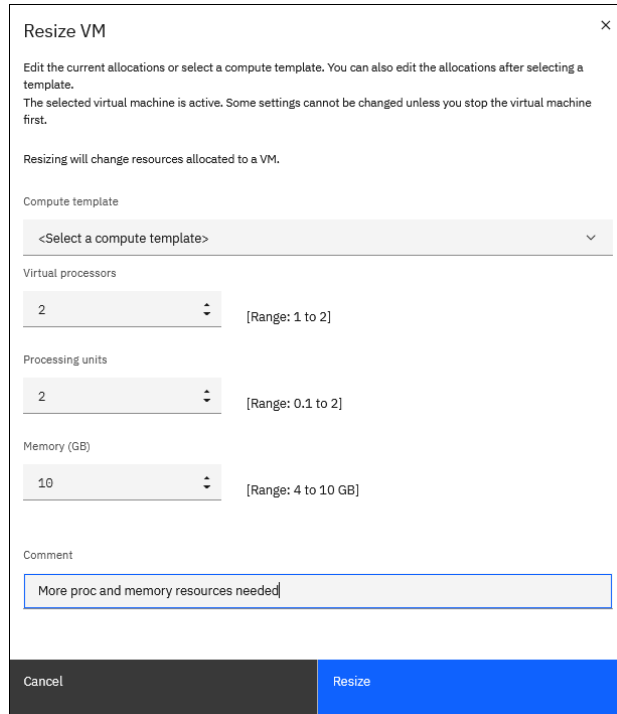


Figure 4-86 Values for VM resizing

- After you complete the information that is required in this window, click **Resize** to start the resizing process. You may see resize messages in the column Running tasks for the selected virtual machine. Also the state changes for a short time to resize. The change of the resources will be visible in the Resources column after the resize operation completed.

To change the maximum boundaries for processors and memory of a virtual machine you need to stop it first. When the VM is Stopped, you can select it and click **Resize**, as described above and enter values that exceed the maximum values. After clicking **Resize**, you will see that the resizing starts for the stopped VM as shown in Figure 4-87.

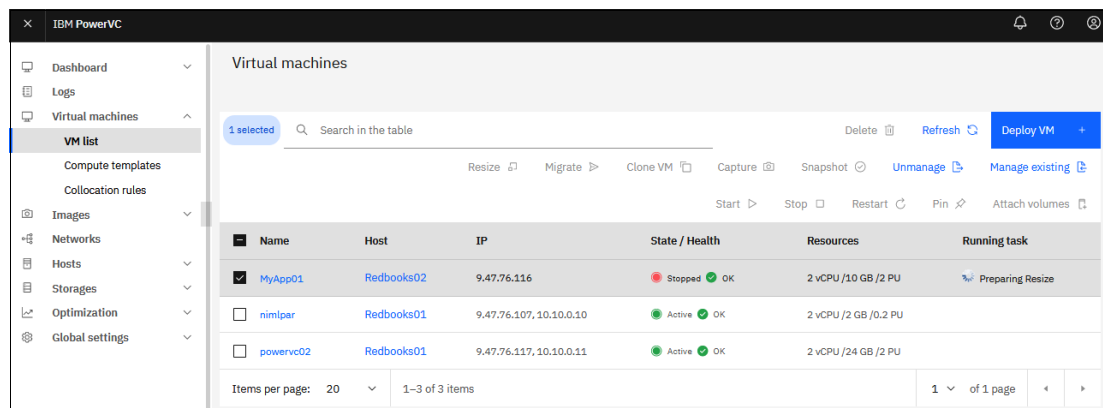


Figure 4-87 Resizing a stopped VM

If you want to change the minimum values or want to lower the maximum value, then do the changes in the Hardware Management Console for the stopped LPAR as shown in Figure 4-88 for memory and in Figure 4-89 for the virtual processors and the entitled capacity.

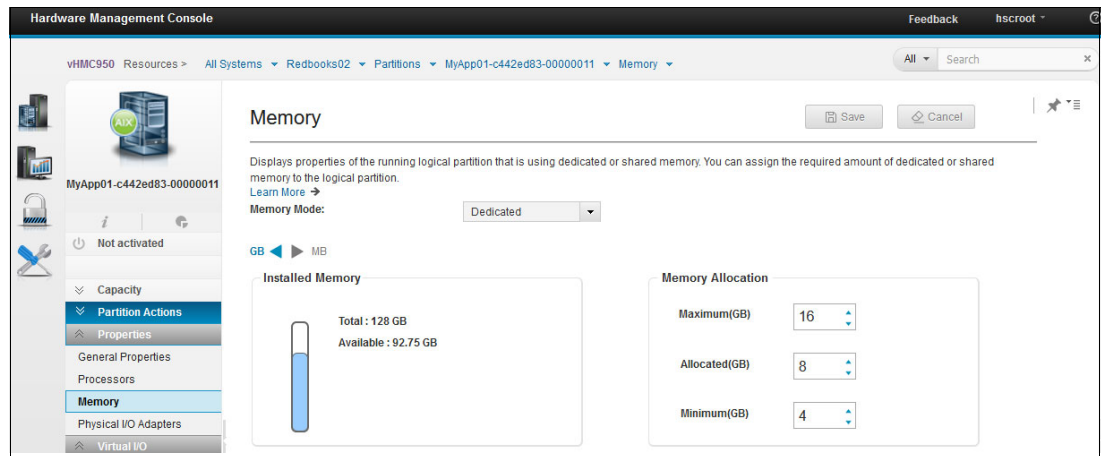


Figure 4-88 Resize the memory boundaries for a VM in the HMC

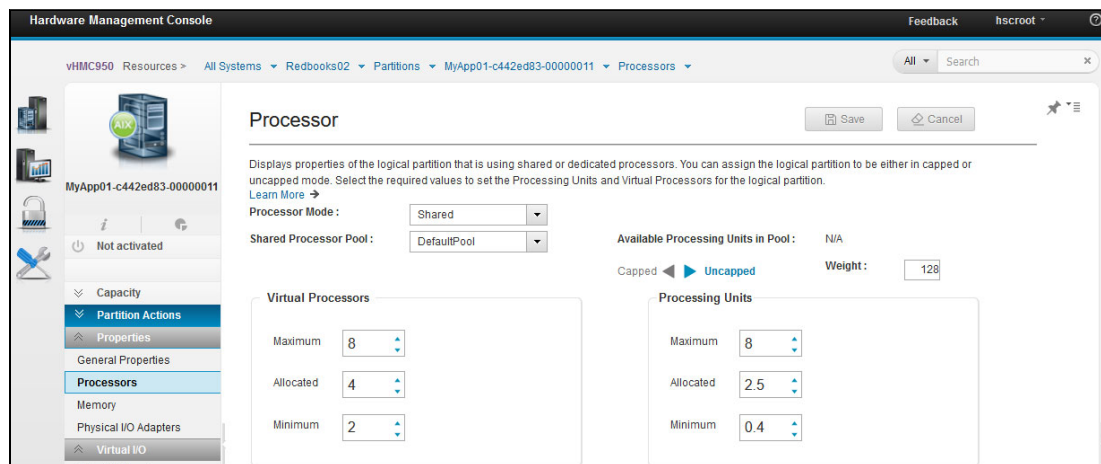


Figure 4-89 Resize the processor boundaries for a VM in the HMC

When your system is NovaLink-connected you can use the following commands in the NovaLink partition to change the values as shown in the following Example 4-2.

Example 4-2 How to change memory and processor resources for a VM in NovaLink

#Find the ID of your VM:

pvmtl lpar list

Logical Partitions

Name	ID	State	RMC	Env	Ref Code	Mem	CPU	Ent
NovaLink	1	running	----	AIX/Linux	Linux ppc64le	8192	2	1.0
MyApp01-c442ed83-00000011	4	not activated	inactive	AIX/Linux	00000000	10240	4	2.5

#List and change the memory values:

```
# pvmtl vm list -i id=4 --display-fields PartitionMemoryConfiguration.min
PartitionMemoryConfiguration.desired PartitionMemoryConfiguration.max
min=4096,desired=10240,max=10240
```

```
# pvmtl lpar update -i id=4 -s PartitionMemoryConfiguration.desired=8192
PartitionMemoryConfiguration.min=4096 PartitionMemoryConfiguration.max=16384
```

```

# pvmtcl vm list -i id=4 --display-fields PartitionMemoryConfiguration.min
PartitionMemoryConfiguration.desired PartitionMemoryConfiguration.max
min=4096,desired=8192,max=16384
#List and change the number of virtual processors (VPs):
# pvmtcl vm list -i id=4 --display-fields SharedProcessorConfiguration.min_virtual
SharedProcessorConfiguration.desired_virtual SharedProcessorConfiguration.max_virtual
min_virtual=1,desired_virtual=4,max_virtual=4

# pvmtcl lpar update -i id=4 -s SharedProcessorConfiguration.min_virtual=2
SharedProcessorConfiguration.desired_virtual=4 SharedProcessorConfiguration.max_virtual=8

# pvmtcl vm list -i id=4 --display-fields SharedProcessorConfiguration.min_virtual
SharedProcessorConfiguration.desired_virtual SharedProcessorConfiguration.max_virtual
min_virtual=2,desired_virtual=4,max_virtual=8
#List and change the number of entitled capacity (EC):
# pvmtcl vm list -i id=4 --display-fields SharedProcessorConfiguration.min_units
SharedProcessorConfiguration.desired_units SharedProcessorConfiguration.max_units
min_units=0.1,desired_units=2.5,max_units=2.5

# pvmtcl lpar update -i id=4 -s SharedProcessorConfiguration.min_units=0.4
SharedProcessorConfiguration.desired_units=2.5 SharedProcessorConfiguration.max_units=8.0

# pvmtcl vm list -i id=4 --display-fields SharedProcessorConfiguration.min_units
SharedProcessorConfiguration.desired_units SharedProcessorConfiguration.max_units
min_units=0.4,desired_units=2.5,max_units=8.0

```

PowerVC will discover the new values after some minutes automatically. Now you can start the VM in PowerVC with the new values.

4.15.5 Attaching and detaching of network adapters

To attach or detach a network adapter in PowerVC for a running VM, enter the Networks tab of the VM as shown in Figure 4-90. Attaching and detaching of networks only works when the health state of the VM is OK, because an RMC connection to the VM is needed to do that concurrently.

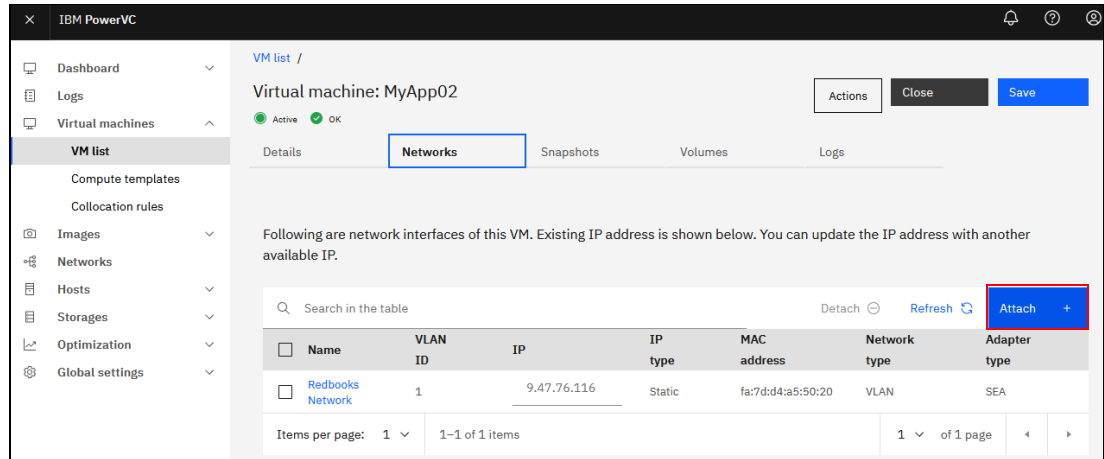


Figure 4-90 Attach network adapter

After clicking **Attach**, a window with all available networks opens as shown in Figure 4-91.

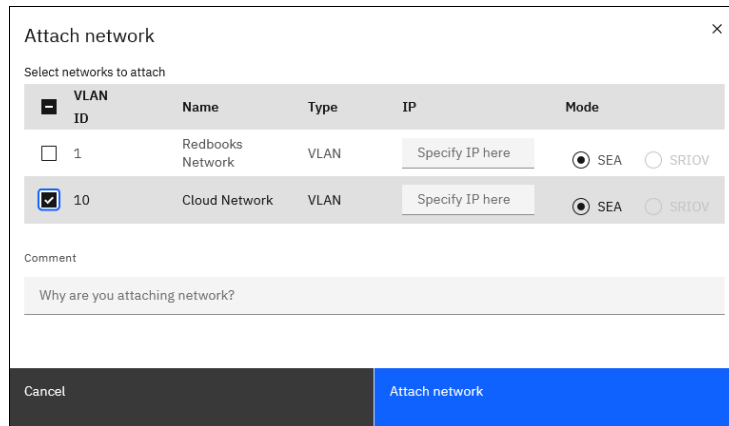


Figure 4-91 Select network to be attached

Select the network you want to attach. You can enter a specific IP address or let PowerVC pick an address out of the IP address pool for that network. If the network is available through SEA and SR-IOV, you can select the connection method. In case of SR-IOV, you also must specify the percentage of the bandwidth you want to assign. Now click **Attach network**.

It takes a view seconds until the adapter is connected. You can click the **Refresh** button, to update the network view. In Figure 4-92 you see that the adapter was attached and an IP address was assigned.

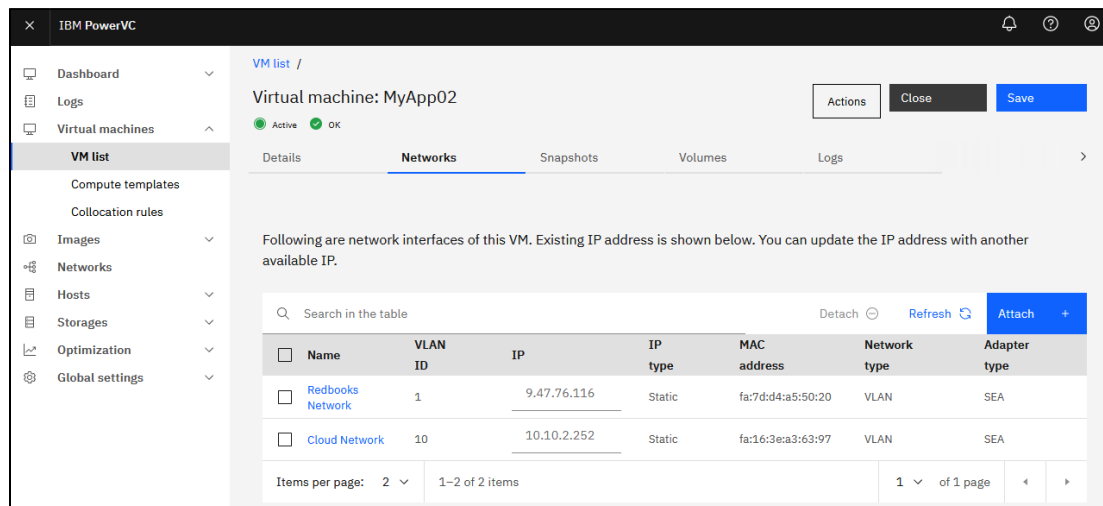


Figure 4-92 Network successfully attached

You can now leave the VM networks tab. When you click **Close** the network is nevertheless created and stays in the VM. Press **Save** if you did some other changes of the VM that needs to be saved.

Note: You can use this method to attach multiple adapters of the same network if needed.

PowerVC does not configure the operating system. Therefore make the adapter visible in the operating system (in AIX for example use **cfgmgr**). After that configure the adapter with the IP address assigned from PowerVC.

To detach a network adapter, first unconfigure the adapter in the operating system. Remove the adapter (in AIX for example use the `rmdev` command). Now select the network to be removed in PowerVC and click **Detach** as shown in Figure 4-93.

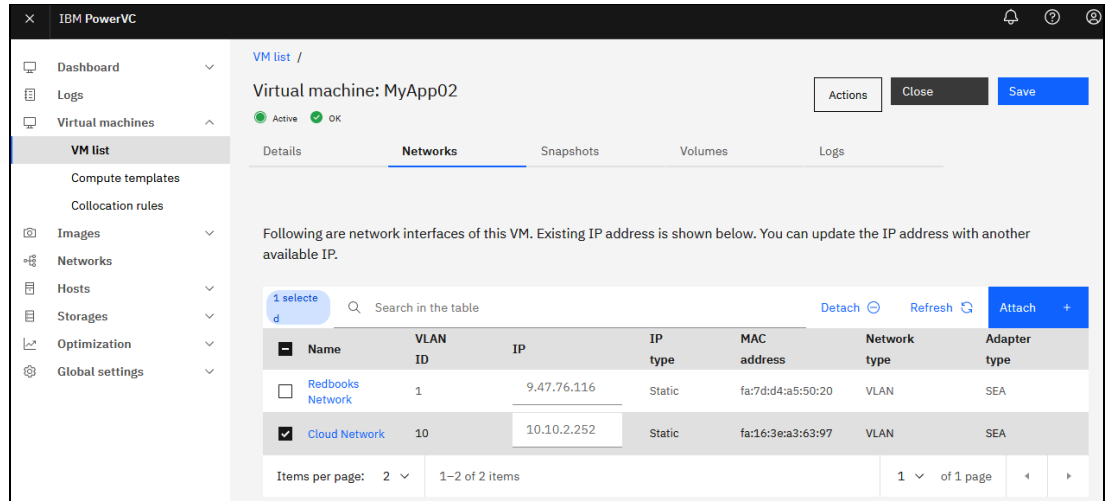


Figure 4-93 Detach a network

Note: If the adapter was not removed from the operating system you may get an error and the detach will not work.

4.15.6 Attaching, detaching, and cloning of storage volumes

In this section we discuss how to attach, detach, and clone a storage volume.

Attach volumes

The IBM PowerVC management server allows the concurrent attachment or detachment of volumes.

To attach a volume complete the following steps:

1. In the **VM list** under **Virtual machines** select the VM that you want to attach volumes to and then click the **Attach volumes** button as shown in Figure 4-94.

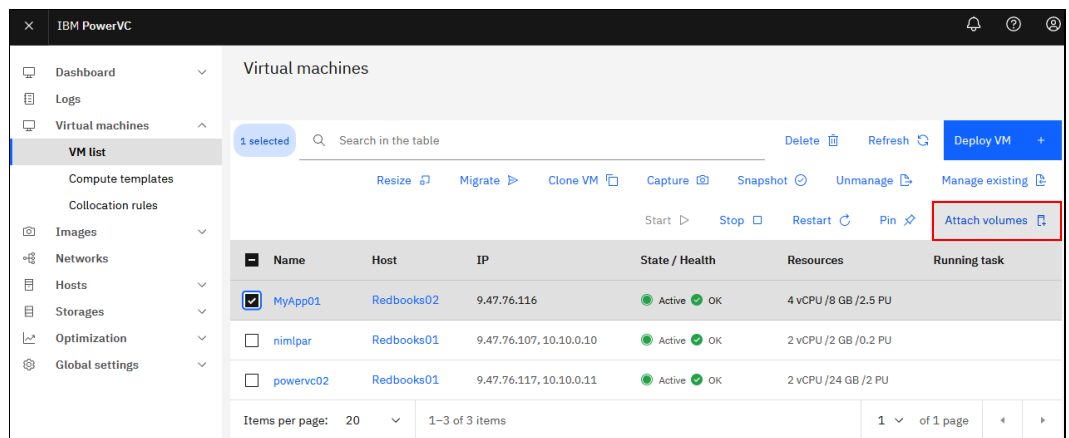


Figure 4-94 Attach a volume to a VM

- In the next screen shown in Figure 4-95 existing volumes that may be attached are provided. In this example the Software repository volume will be attached to the VM after you clicked on the **Attach volumes** button.

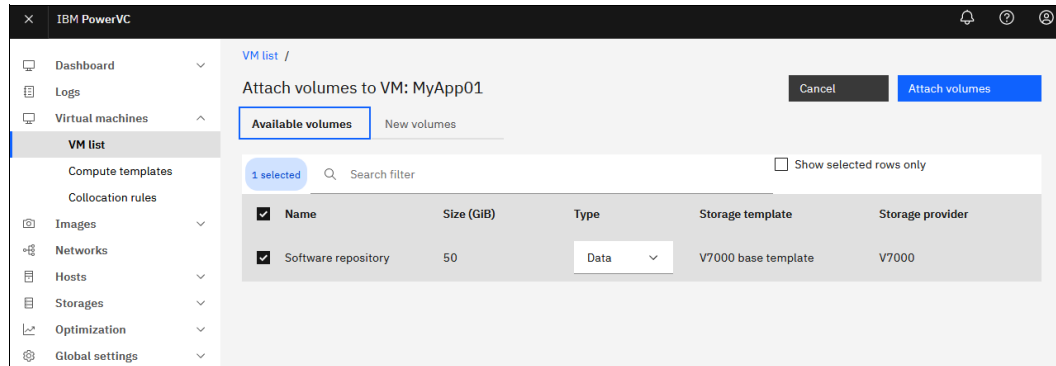


Figure 4-95 Attach an existing volume

- If you want to attach new volumes, first select the **New volumes** tab. In the example shown in Figure 4-96 two volumes with 80 GB each will be attached to the VM. The volumes are flagged Shared, which means PowerVC allows these volumes to be attached to two VMs (for example in an High Availability cluster).

Note: IBM i does not support shared volumes.

If you have multiple Storage templates that may refer to different storages or different storage pools in the storage, you can select the correct template. For more information on storage templates refer to 4.10, “Storage template setup” on page 140.

For multiple volumes a <dash><number> (ex. -1, -2, ...) will be added to the volume name.

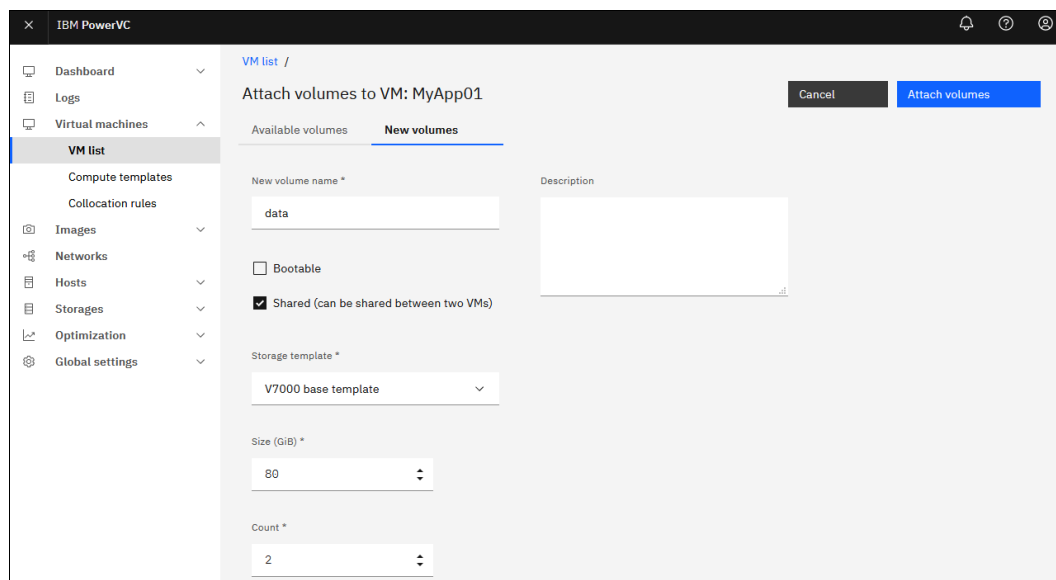


Figure 4-96 Attaching new volumes to a VM

After you clicked **Attach volume**, the selected or the new to be created volumes will be attached. While the attachment is running, you can see Attaching volume in the Running task column in the Virtual machines list.

Note: Attaching of existing volumes and new volumes with one task using the **Attach volumes** button is not supported. To do both you need to click the button on each tab in two separate tasks.

To verify the attached volumes in PowerVC you can enter the VM details under **VM List** in **Virtual machines**. There click on the name of the virtual machine that you want to see. Then select the Volumes tab as shown in Figure 4-97.

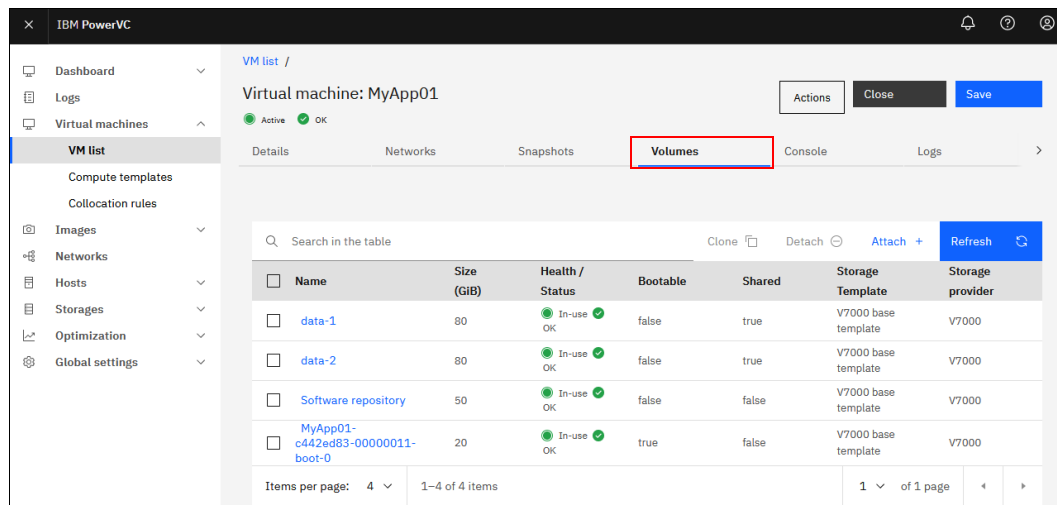


Figure 4-97 Attached volumes in a VM

Tip: As you can see in Figure 4-97, in the VM volumes information there is also a **Attach** button that could be used to attach additional volumes.

4. To complete the process, you must run the correct command in the operating system of the VM:

- For IBM AIX operating systems, run this command as root:

```
# cfmgr
```

- For Linux operating systems, run the following commands as root:

```
# rescan-scsi-bus.sh -a
```

or

```
# echo "- -" > /sys/class/scsi_host/host_N/scan
```

host_N is the controller that manages the disks in the VM.

Detaching volumes

To detach a volume from the VM, you must first remove it from the operating system.

Note: As a preferred practice, cleanly unmount all file systems from the disk, remove the logical volume, and remove the disk from the operating system before you detach the disk from IBM PowerVC.

For the IBM AIX operating system, run this command as root, where *hdisk_N* is the disk that you want to remove:

```
# rmdev -d1 hdisk_N
```

For the Linux operating system, restart after you detach the volume.

Now you can detach the volumes in the VM details: under **VM List** in **Virtual machines** click on the name of the VM where you want to detach volumes. Now click the **Volumes** tab and there select the volumes to be detached as shown in Figure 4-98.

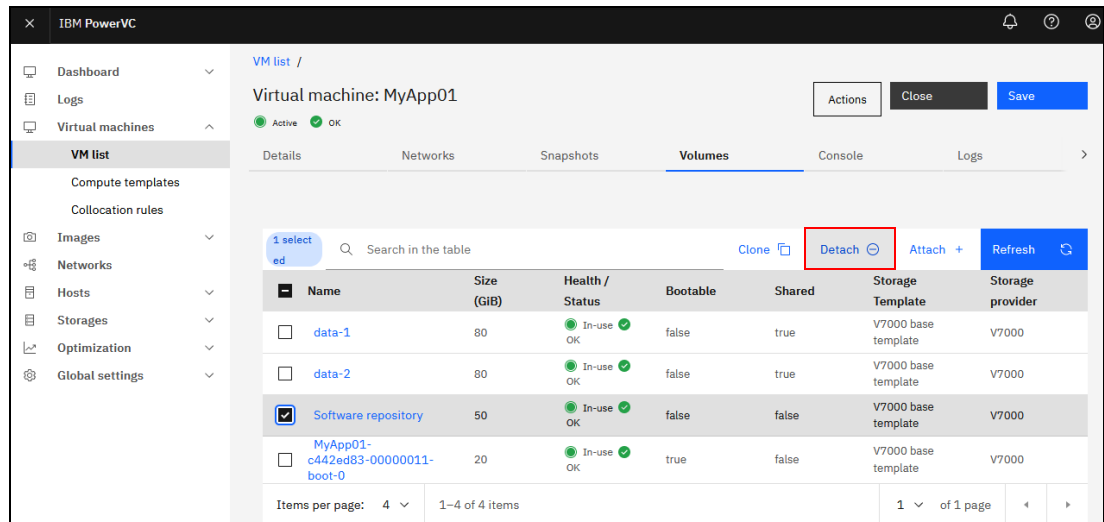


Figure 4-98 Detach a volume from a VM

After you clicked the **Detach** button, a confirmation dialog opens as shown in Figure 4-99.

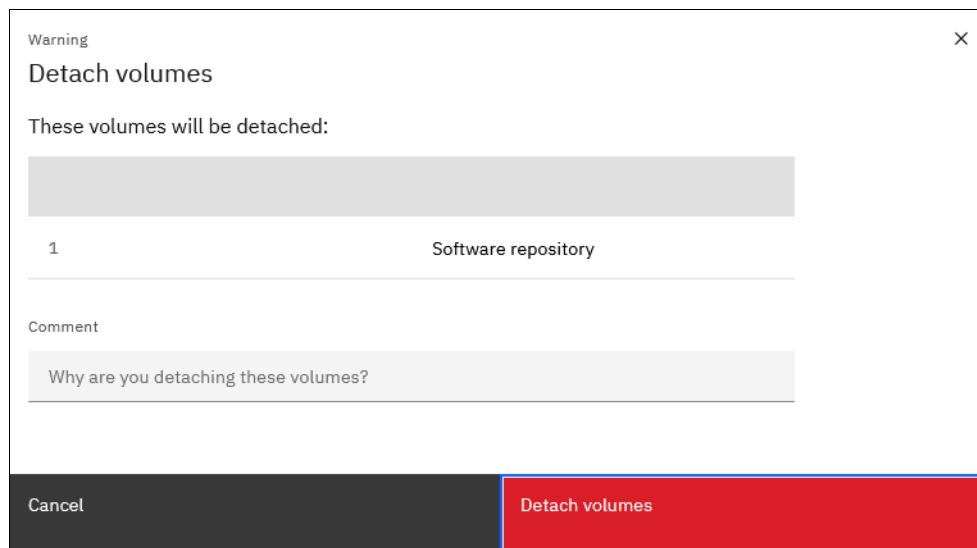


Figure 4-99 Detach a volume from a VM - confirmation

To verify that the detachment is done, have a look at the messages by clicking on the bell symbol in the right, upper corner or by clicking the **Refresh** button in the Volumes tab of the VM details as shown in Figure 4-98. When the detachment was successful, the volume disappeared.

The detached volumes are still managed by the IBM PowerVC management host. You can see the volume from in the Storages section.

Cloning of volumes

A feature that is available since the version 2.0 of PowerVC is cloning of volumes.

Before you clone volumes make sure that these volumes are in a consistent state. Best is to unmount all filesystems and volume groups on these volumes.

Now you can clone them directly in the Volumes tab of a VM. There select the disks that you want to clone and click the **Clone** button as shown in Figure 4-100.

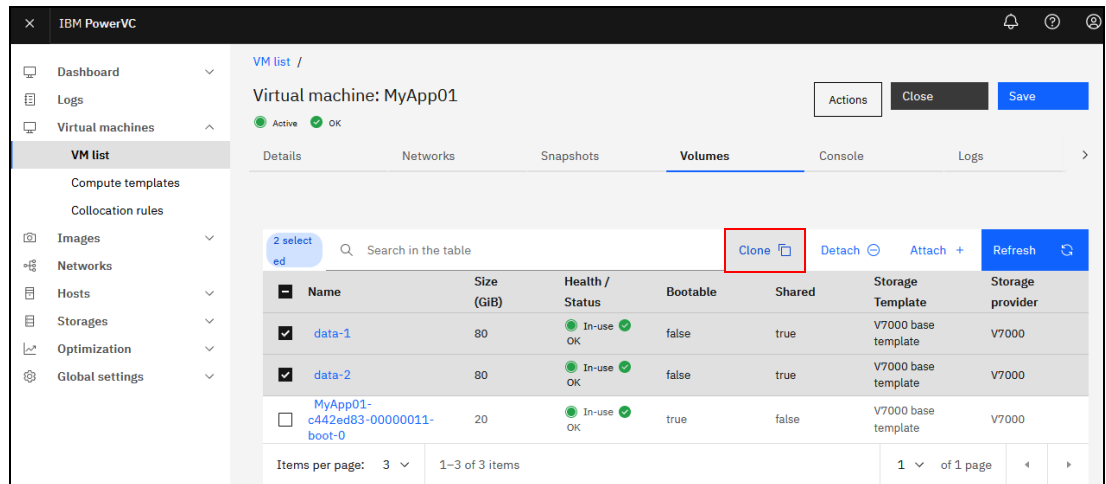


Figure 4-100 Clone volumes in the VM volumes tab

Now another dialog opens to confirm the cloning and also to assign a suffix that will be attached to the volumes clone. In the example shown in Figure 4-101 the two volumes data-1 and data-2 will be cloned with the attached suffix test.

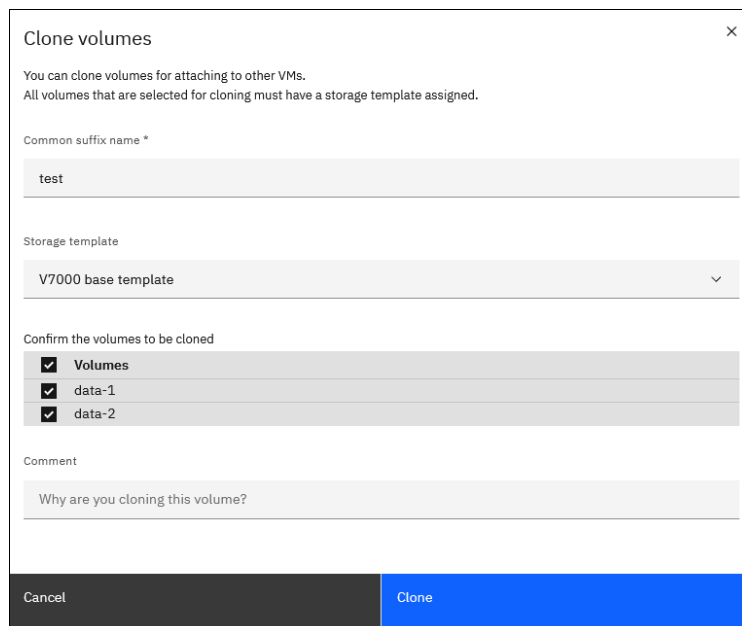


Figure 4-101 Clone volume suffix

The name of the cloned volumes will be clone-<suffix>-<number of volume>. As the cloned volumes will be probably used in a different VM, the clones will not be attached to the origin VM. You can find the cloned VMs under Data volumes in Storages as shown in Figure 4-102.

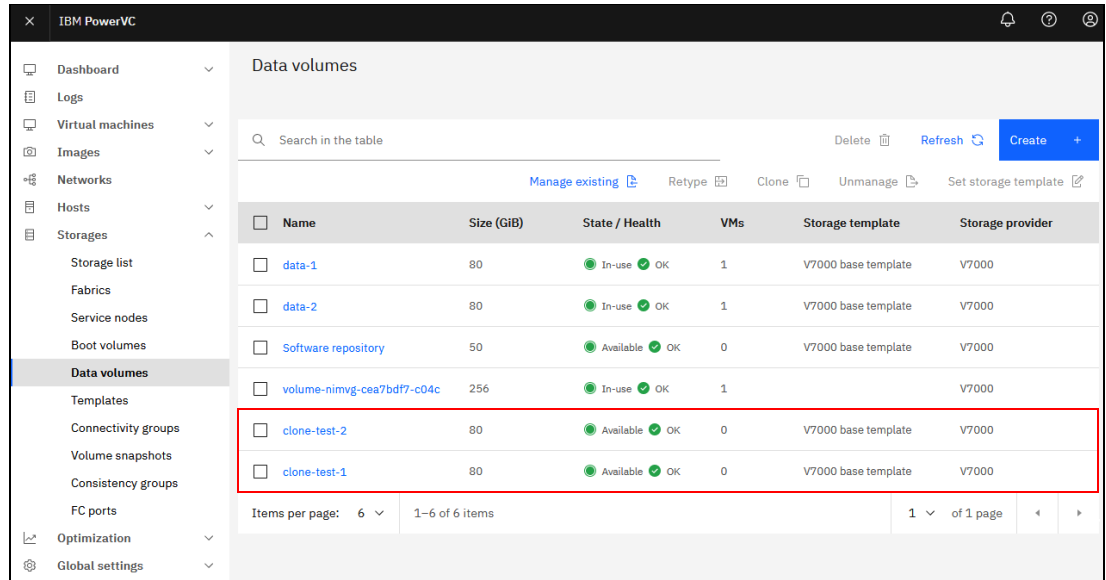


Figure 4-102 Cloned volumes

4.15.7 Migration of virtual machines

IBM PowerVC can manage the Live Partition Mobility (LPM) feature. Use the LPM feature to migrate VMs from one host to another host.

Migration requirements

To migrate VMs by using the IBM PowerVC management server, ensure that the source and destination hosts and the VMs are configured correctly.

To migrate a VM, the following requirements must be met:

- ▶ The VM is in Active or Stopped status in the IBM PowerVC management host.
- ▶ The PowerVM Enterprise Edition or PowerVM for IBM Linux on Power hardware feature is activated on your hosts. This feature enables the use of the LPM feature.
- ▶ The SEA networks for both source and target hosts must use the same virtual Ethernet switch.
- ▶ As a preferred practice, set the maximum number of virtual resources (virtual adapters) to at least 200 on all of the hosts in your environment. This value ensures that you can create enough VMs on your hosts.
- ▶ The logical memory block size on the source host and the destination host must be the same.
- ▶ Both the source and destination hosts must have an equivalent configuration of VIOS that belong to the same storage connectivity group.

Note: If the source host has two VIOSs and the target host has only one VIOS, it is not possible to migrate a partition by accessing its storage through both VIOSs on the source. However, if a partition on the source host is using only one VIOS to access its storage, it can be migrated (assuming that other requirements, such as port tagging, are met).

- ▶ The processor compatibility mode on the VM that you want to migrate must be supported by the destination host.
- ▶ The AIX/Linux VM must have an enabled RMC connection. IBM i does not require RMC.
- ▶ To migrate a VM with a vSCSI attachment, the destination VIOS must be zoned to the backing storage. If the disks are not yet mapped to the destination VIOS, PowerVC will take care about that and map the LUNs to the VIOS.
- ▶ No virtual optical device must be mapped to the VM.
- ▶ At least one pair of VIOS VMs must be storage-ready and members of the storage connectivity group. Each of these VIOS VMs must have at least two physical FC ports ready.
- ▶ Each of the two physical FC ports must be connected to a distinct fabric, and the fabric must be set correctly on the FC ports' Configuration pages.

The following restrictions apply when you migrate a VM:

- ▶ You cannot migrate a VM to a host that is a member of a different host group.
- ▶ If the VM is running a Little Endian guest, the target host must support Little Endian guests.
- ▶ If the VM was created as remote restart-capable, the target host must support remote restart.
- ▶ Certain IBM Power System servers can run only Linux workloads. When you migrate an AIX or IBM i VM, these hosts are not considered for placement.
- ▶ You cannot exceed the maximum number of simultaneous migrations that are designated for the source and destination hosts. The maximum number of simultaneous migrations depends on the number of migrations that are supported by the VIOS that are associated with each host.
- ▶ A source host in a migration operation cannot serve concurrently as a destination host in a separate migration operation.
- ▶ PowerVM does not support the migration of a VM whose attachment type changes its multipathing solution between the source and destination VIOS. For example, a VM on a path-control module (PCM)-attached VIOS can be successfully migrated only to a PCM-attached VIOS. However, PowerVM does not enforce this requirement. To avoid unsupported migrations, create separate storage connectivity groups for PCM and PowerPath multipathing solutions.
- ▶ Collocation rules are enforced during migration:
 - If the VM is a member of a collocation rule that specifies affinity and multiple VMs are in that collocation rule, you cannot migrate it. Otherwise, the affinity rule is broken. To migrate a VM in this case, remove it from the collocation rule and then add it to the correct group after the migration.

- If the VM is a member of a collocation rule that specifies anti-affinity, you cannot migrate it to a host that has a VM that is a member of the same collocation rule. For example, assume the following scenario:
 - VM A is on Host A.
 - VM B is on Host B.
 - VM A and VM B are in a collocation rule that specifies anti-affinity.
 Therefore, VM A cannot be migrated to Host B.
- Only one migration or remote restart at a time is allowed for VMs in the same collocation rule. Therefore, if you try to migrate a VM or restart a VM remotely and any other VMs in the same collocation rule are being migrated or restarted remotely, that request fails.

Migrating the virtual machine

To migrate a VM, complete the following steps:

1. Open the Virtual Machines window by clicking **VM list** under **Virtual machines** and select the VM(s) that you want to migrate.
2. Click **Migrate**, as shown in Figure 4-103.

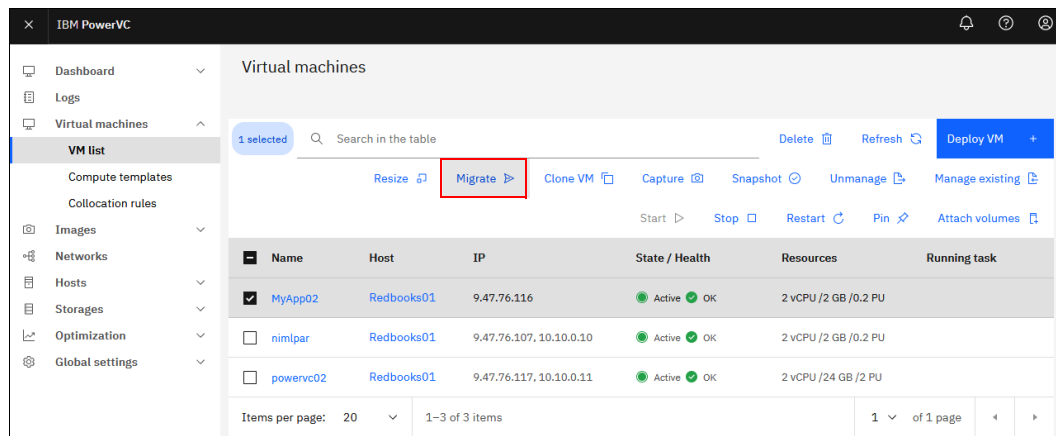


Figure 4-103 Migrate selected virtual machines

- You can select the target host or the placement policy can determine the target, as shown in Figure 4-104.

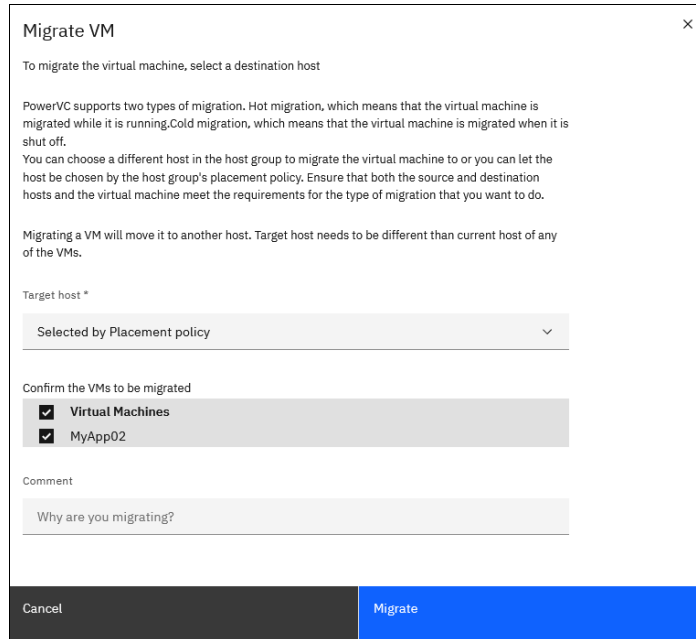


Figure 4-104 Migration destination

- Figure 4-105 shows in the Running task column of the selected VM that the migration running.

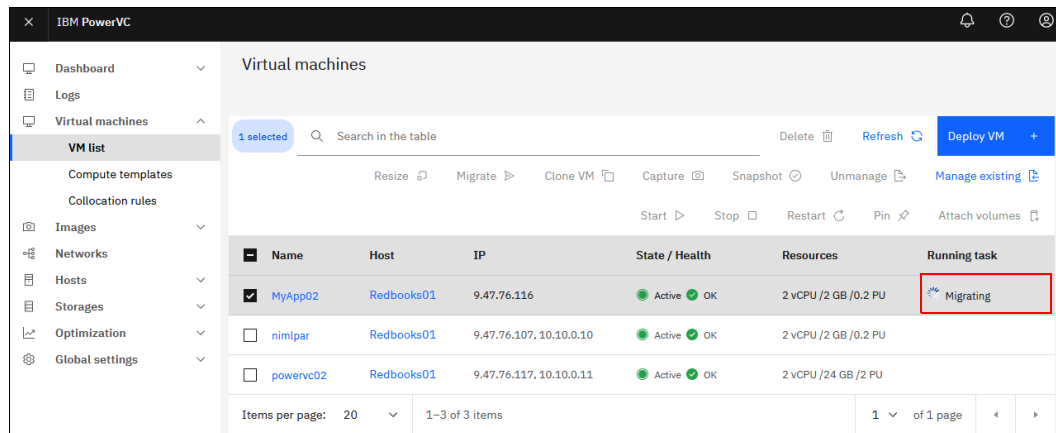


Figure 4-105 Virtual machine migration in progress

- After the migration completes, you can check the Virtual Machines window to verify that the partition is now hosted on the target host, as shown in Figure 4-106.

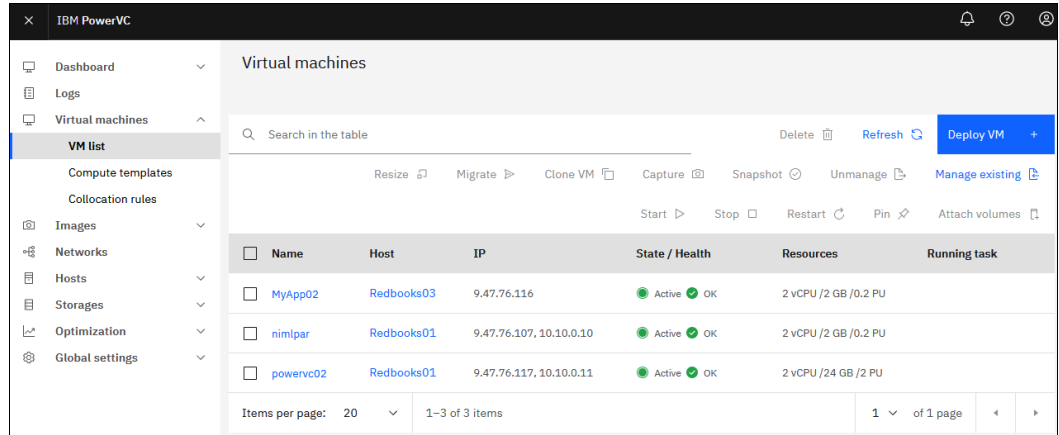


Figure 4-106 Virtual machine migration finished

Note: A warning message in the Health column is normal. It takes a few minutes to change to OK.

Pinning

In the VM details you can specify, if this VM should be pinned to the host where it is running, by selecting Soft or Hard pin as shown in Figure 4-107.

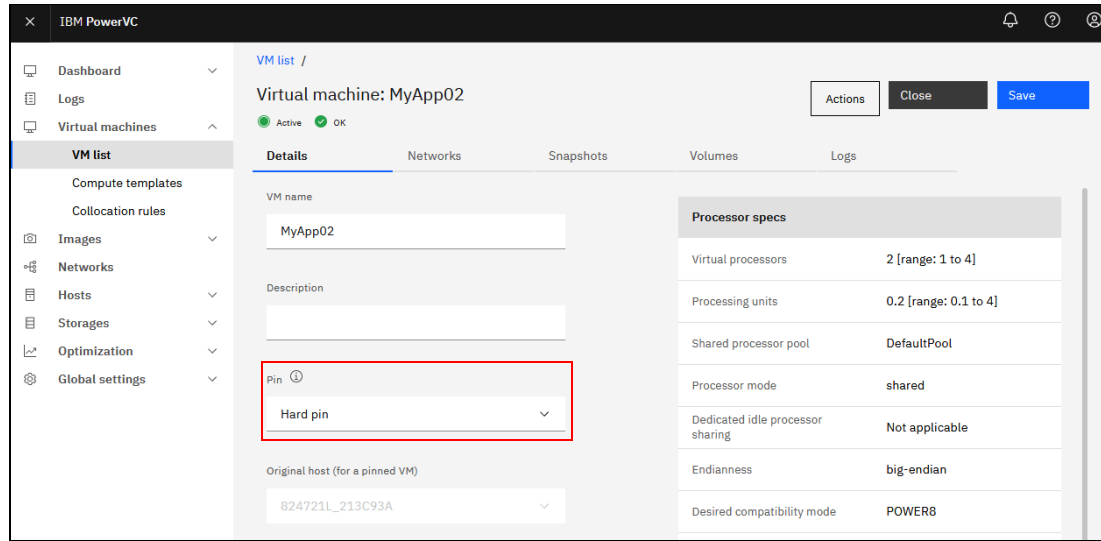


Figure 4-107 Hard pinned VM

If you try to migrate a pinned VM, this will not work and PowerVC will show you a message similar to the example in Figure 4-108. Soft pinned VMs can be moved with remote restart or when you put a host into maintenance mode. Hard pinned VM can neither be migrated in a DRO, or remote restart operation, nor when you put a host into maintenance mode.

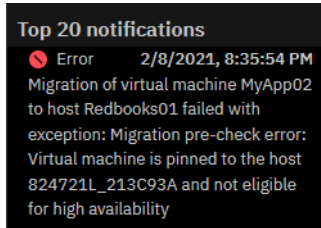


Figure 4-108 Error when migrating a hard pinned VM

Migration performance considerations

In PowerVC it is possible to define which VIOS / MSP (Mover Service Partition) IP addresses should be used for migration. This makes it possible to select for example a fast LPM/Backup network instead of a slower client network.

Since PowerVC 1.4.3 it is also possible to specify multiple MSPs to increase the bandwidth and to get redundancy.

This can be done from the command line only using the **powervc-config** command:

```
# powervc-config compute mover-service-par --host <Managed System> --set <LPM IP VI01 [,LPM IP VI02]>
```

The IP addresses that you specify with **powervc-config** must be a dedicated physical Ethernet port or an IP address that is configured on top of a SEA adapter. Just a virtual adapter for a specific VLAN with an IP address will not work.

4.15.8 VM snapshots

VM snapshots are available since PowerVC Version 2.0. This feature lets you create a snapshot of a VM before you do some work that may be disruptive for the VM, or just as a backup.

Virtual Machine snapshot is only supported for IBM Storwize, PowerMAX, and Hitachi storage providers with non-replicated volumes:

1. To create a snapshot, enter the **VM list** under **Virtual machines**. There you can select the VM that you want to snapshot and click **Snapshot** as shown in Figure 4-109.

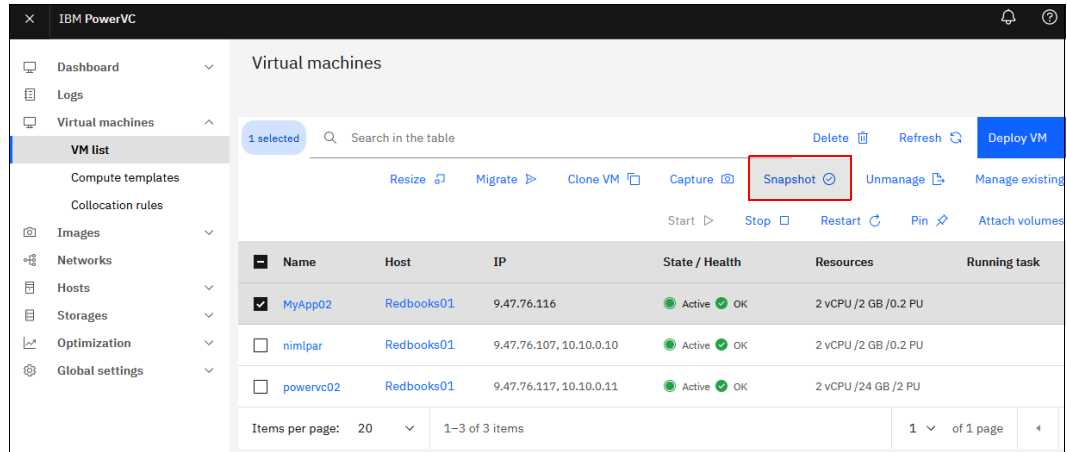


Figure 4-109 Create snapshot from a VM

Note: To select multiple VMs for a snapshot is not possible, only one at a time.

2. In the next screen you can give the snapshot a meaningful name and maybe some description. It is also possible to select which disks should be included into the snapshot. Possible selections are All volumes, Boot set only, or Select specific volumes as shown in Figure 4-110.

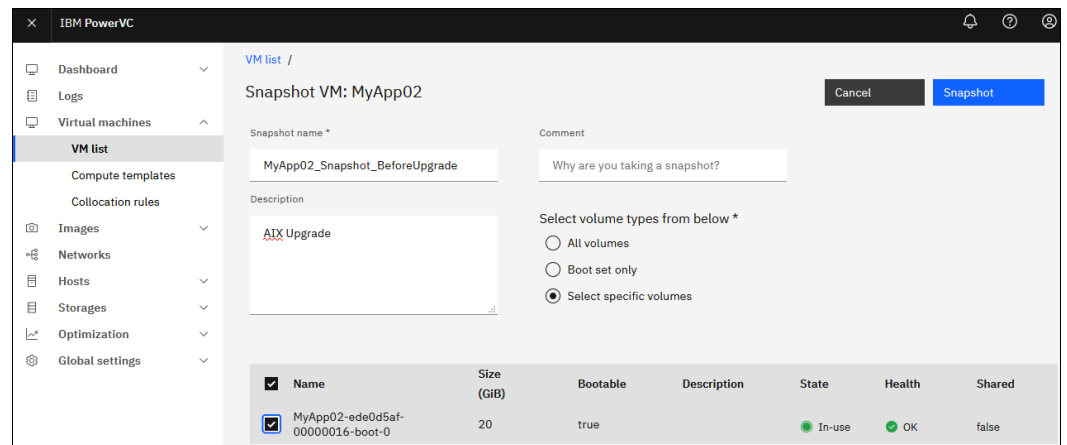


Figure 4-110 Details of the snapshot

The snapshot may be taken while the virtual machine is running. But in this case make sure that the disks that are included into the snapshot are in a consistent state.

When the snapshot has finished you get a message in the messages bar on right when you click on the bell symbol in the upper, right corner of the screen.

- To view the snapshot, enter the VM details, by clicking on the VM name in the Virtual machines list and then select the **Snapshots** tab as shown in Figure 4-111.

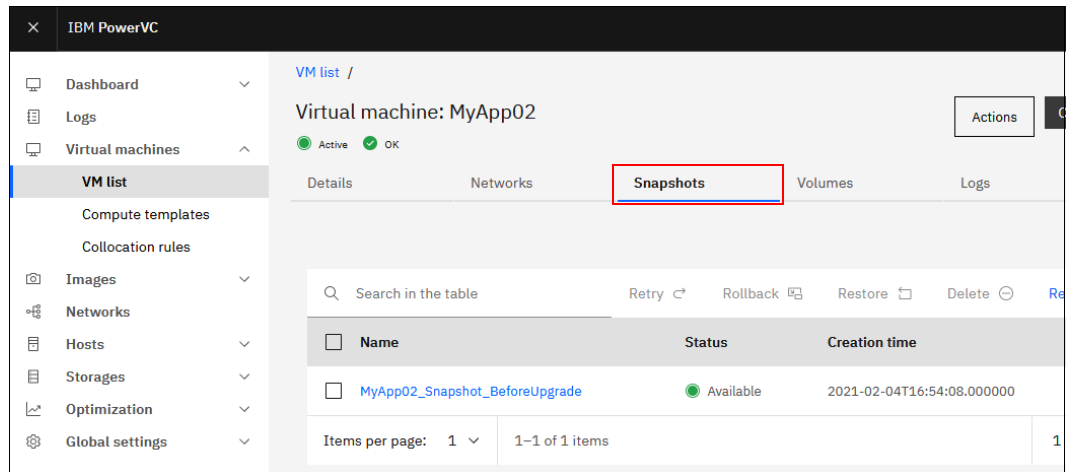


Figure 4-111 Snapshots of a VM

- In case you want to restore a snapshot, select the snapshot on the **Snapshots** tab in the VM and click **Restore** as shown in Figure 4-112.

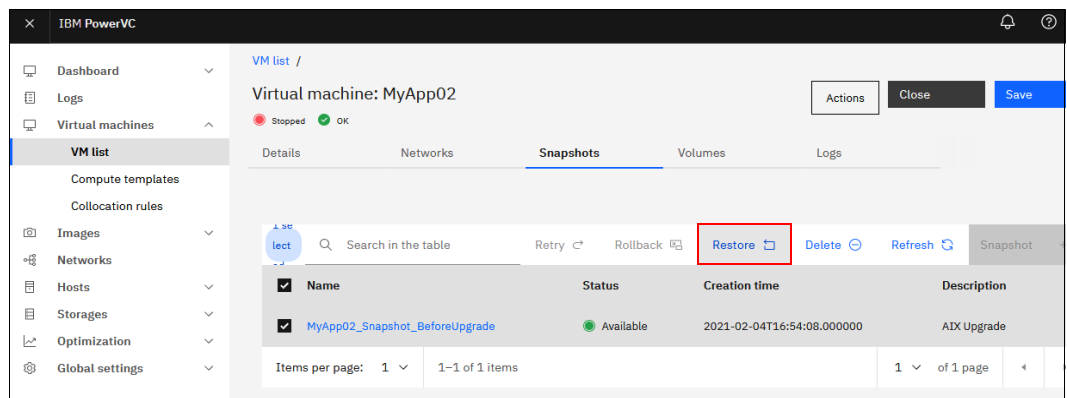


Figure 4-112 Restore a snapshot

- On the next screen you will be asked for confirmation. When you try to restore a snapshot while the VM is running, you get a Warning as shown in Figure 4-113. If you restore the bootset of a VM, stop the VM first, otherwise you may get inconsistencies in the operating system. If you restore only data volumes, like a volume group that is varyoff, you may restore such a snapshot while the VM is running.

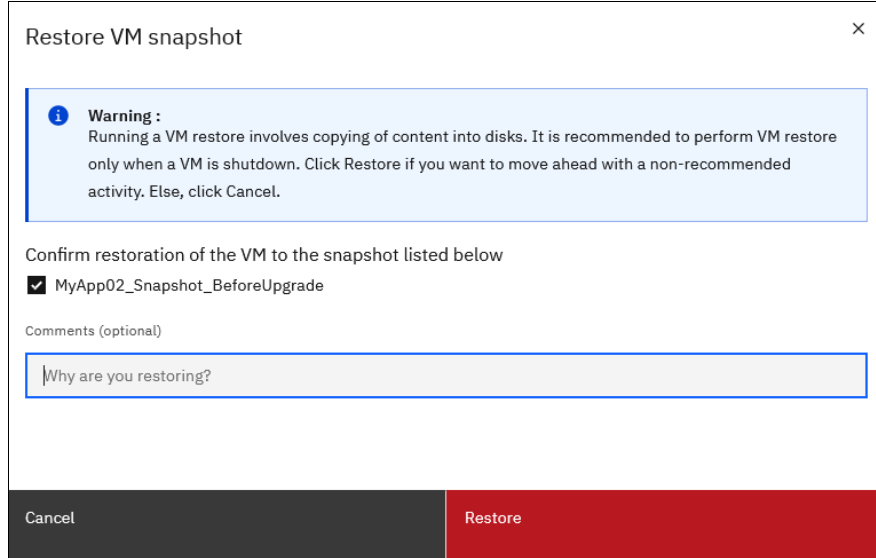


Figure 4-113 Confirmation to restore the snapshot

While the restore is running, you can see restoring in the Running task column of the VM. When the restore has finished, start the VM if it was stopped.

To delete a snapshot, simply click **Delete** on the Snapshots tab in the VM details as shown in Figure 4-114.

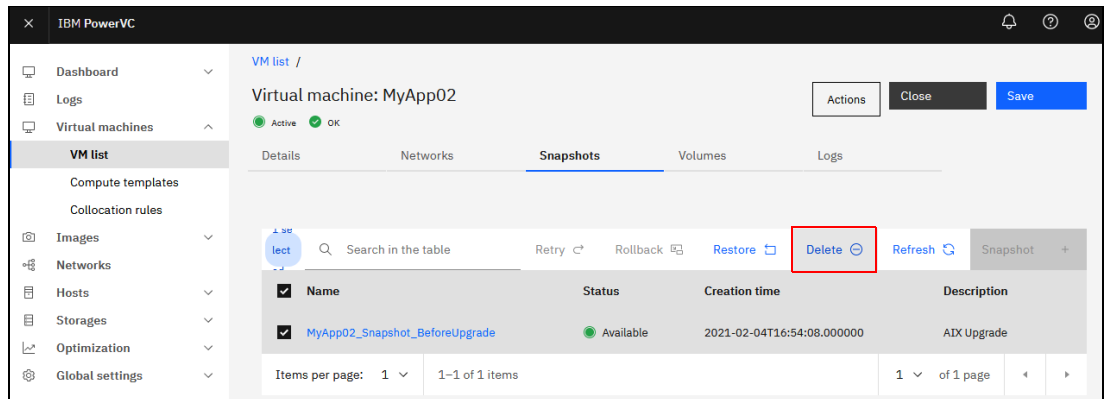


Figure 4-114 Delete a snapshot

Also before the snapshot will be deleted, you must confirm the deletion as shown in Figure 4-115.

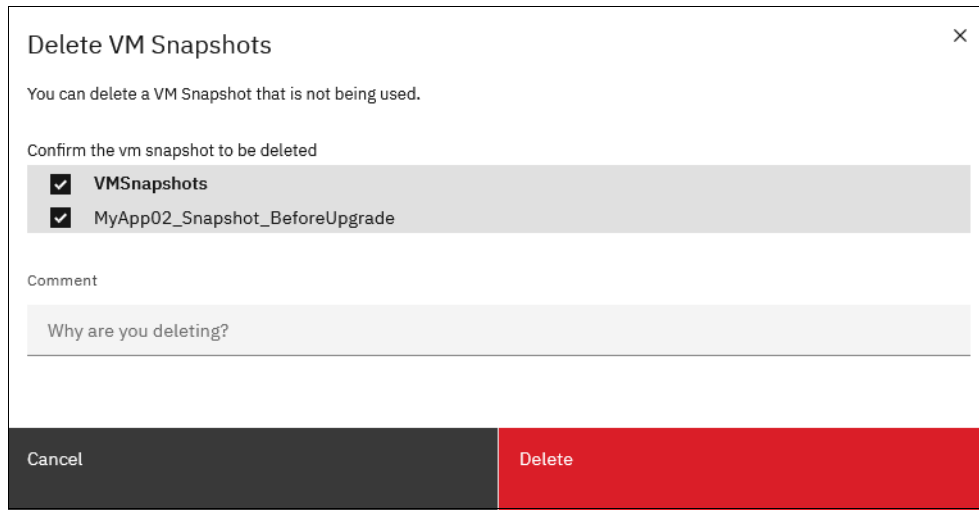


Figure 4-115 Confirmation to delete a snapshot

Before PowerVC performs a virtual machine restore, it takes a backup snapshot for all the volumes to be restored. In case the restore of the snapshot fails, you can click **Rollback** in the Snapshots tab under the VM details to restore to the original volume state. There is also a **Retry** button to try the restore again.

4.15.9 Cloning a virtual machine

Cloning of a VM is available since PowerVC 2.0. Before this version you could capture a VM and then re-deploy it. In PowerVC 2.0 these two steps are summarized for you convenience.

1. To clone a VM, enter the **VM list** under **Virtual machines**. There select the VM that you want to clone. Only one VM could be selected for cloning. Now click **Clone VM** as shown in Figure 4-116.

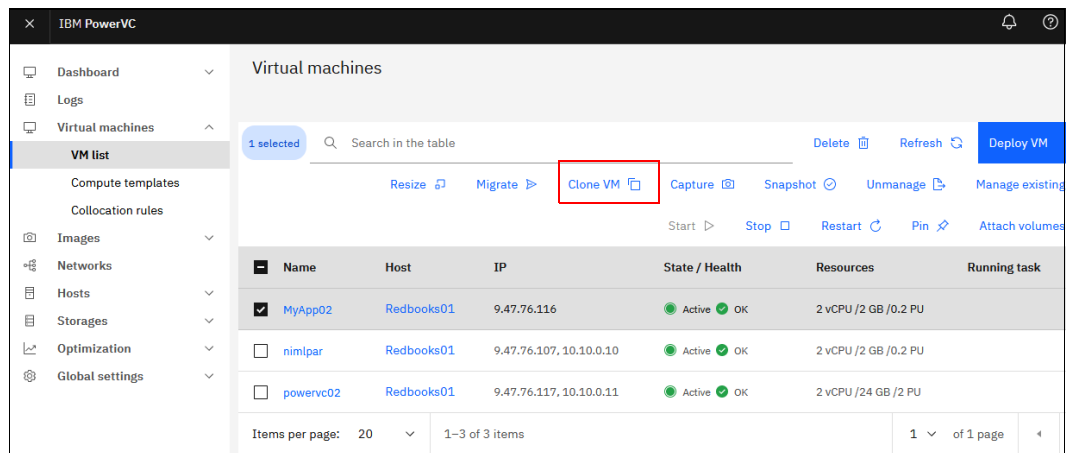


Figure 4-116 Clone a VM

- The next screens are similar to a VM deployment. In the first screen you can modify attributes of the clone as shown in Figure 4-117.

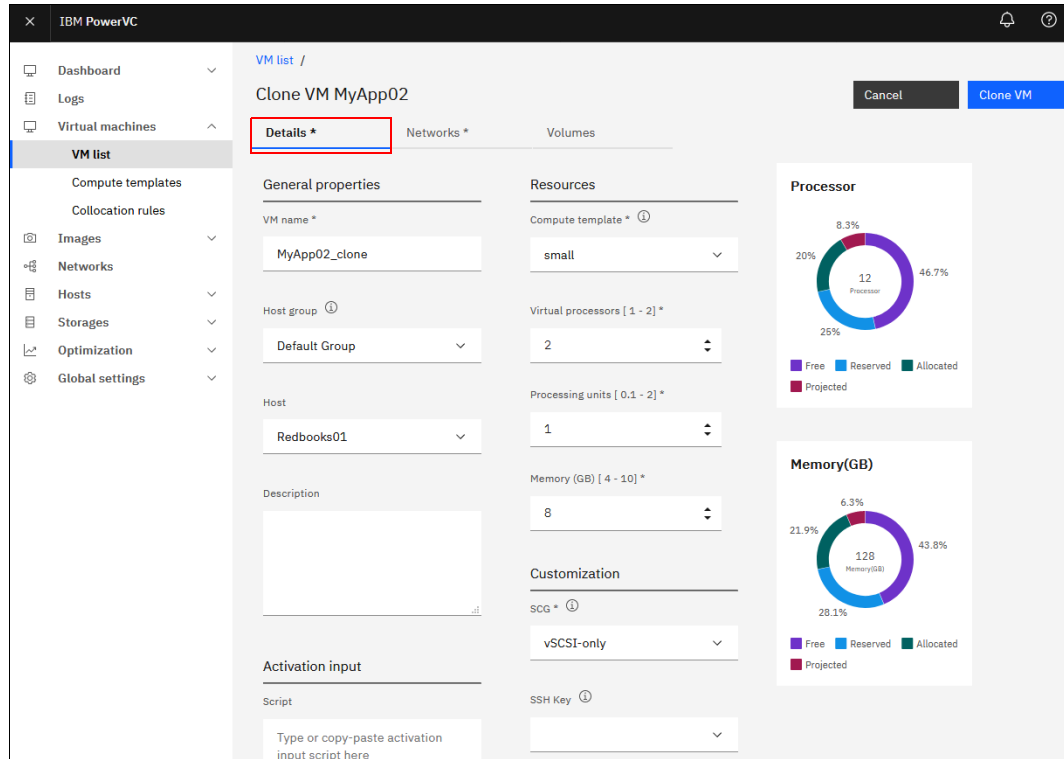


Figure 4-117 VM details for the VM clone

Here you can give the clone a name, specify the processor and memory resources by selecting a Compute template, and change the processor and memory values. You define the Host group and potentially a specific host for the clone. You also select the Storage Connectivity Group (SCG) that defines how the volumes should be attached to the clone and you also may enter a description, an activation script or select SSH keys for injection.

As this screen is similar to a VM deployment, please refer to 4.16.7, “Deploying a new virtual machine” on page 239 for more details on these attributes.

- Click on **Networks** to specify the network details for the clone as shown in Figure 4-118.

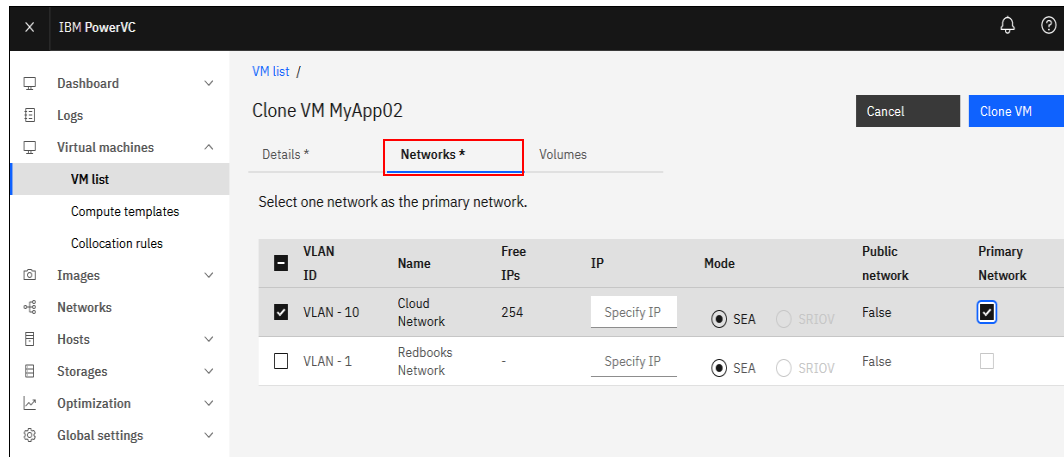


Figure 4-118 Network details for the clone

4. Select the networks and the primary network.

The last tab **Volumes** shows the volumes of the VM (Figure 4-119).

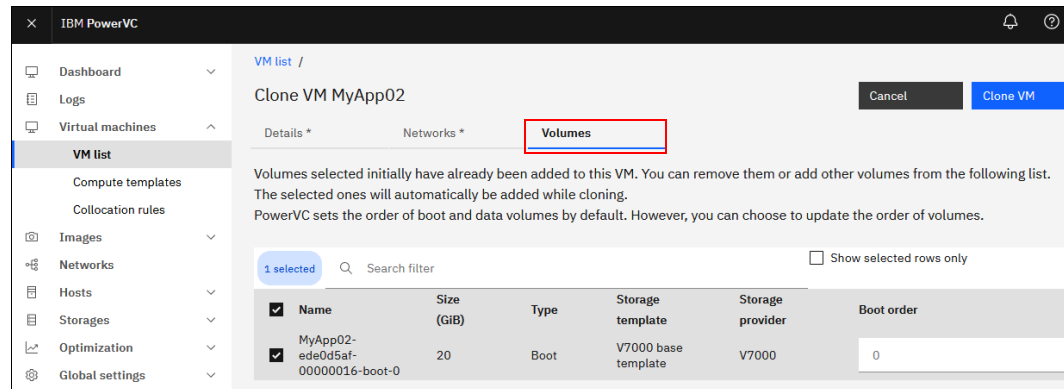


Figure 4-119 Volumes for the clone

5. Select all volumes that should be included into the clone and click **Clone VM** to start cloning.

Note: Additional volumes cannot be created or attached at the time of a clone virtual machine operation.

In the list of virtual machines you can see first Cloning in the task column of the source VM and later Building in the new create clone VM. During the cloning, PowerVC also creates an image called `Image_Do_not_use_to_deploy`, that will be automatically deleted when the cloning has finished. When the cloning has finished, the clone shows up, as in Figure 4-120.

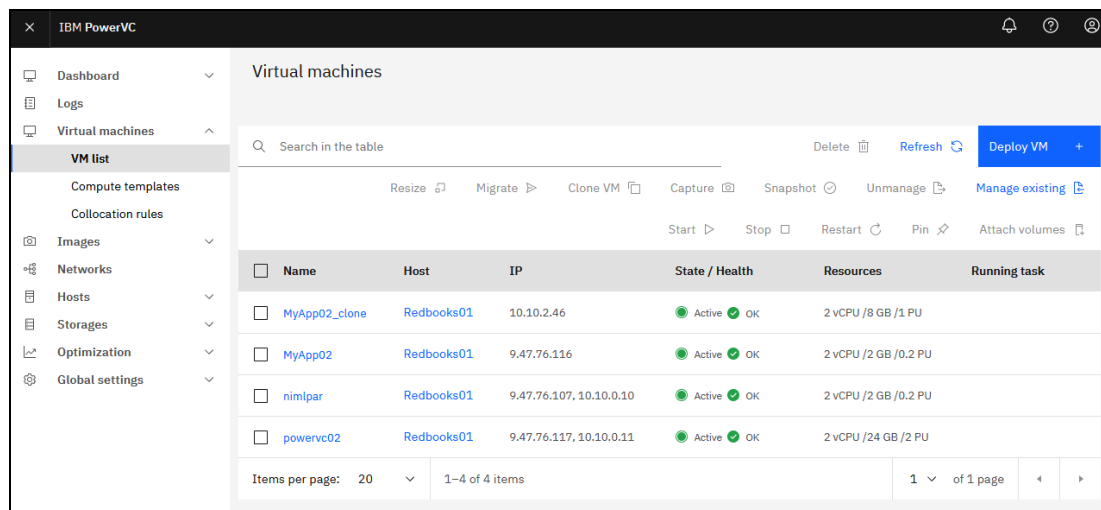


Figure 4-120 VM Clone available

4.15.10 Host maintenance mode

You put a host to maintenance mode to perform maintenance activities on a host, such as updating firmware or replacing hardware. When you enter the maintenance mode for a host, the VMs can be migrated using live partition migration. When a host is in maintenance mode, PowerVC also prevents that deployments of new VMs will be done on this host.

Maintenance mode requirements

Before you put the host into maintenance mode, check whether the following requirements are met:

- ▶ If the request was made to migrate active VMs when the host entered maintenance mode, the following conditions must also be true:
 - The hypervisor must be licensed for LPM.
 - The VMs on the host cannot be in the Error, Paused, or Building states.
 - On all active VMs, the health must be OK and the RMC connections of the AIX/Linux VMs must be Active.
 - All requirements for live migration must be met, as described in “Migration requirements” on page 181.
- ▶ The host’s hypervisor state must be Active. If it is not, VM migrations might fail.
- ▶ If the request was made to migrate active VMs when the host entered maintenance mode, the following conditions cannot also be true, or the request fails:
 - A VM on the host is a member of a collocation rule that specifies affinity and has multiple members.
 - The collocation rule has a member that is already undergoing a migration or is being restarted remotely.

Putting the host in maintenance mode

If all of the requirements are met, you can put a host in maintenance mode by following these steps:

1. On the Hosts window, select the host that you want to put into maintenance mode, and click **Enter Maintenance Mode**, as shown in Figure 4-121.

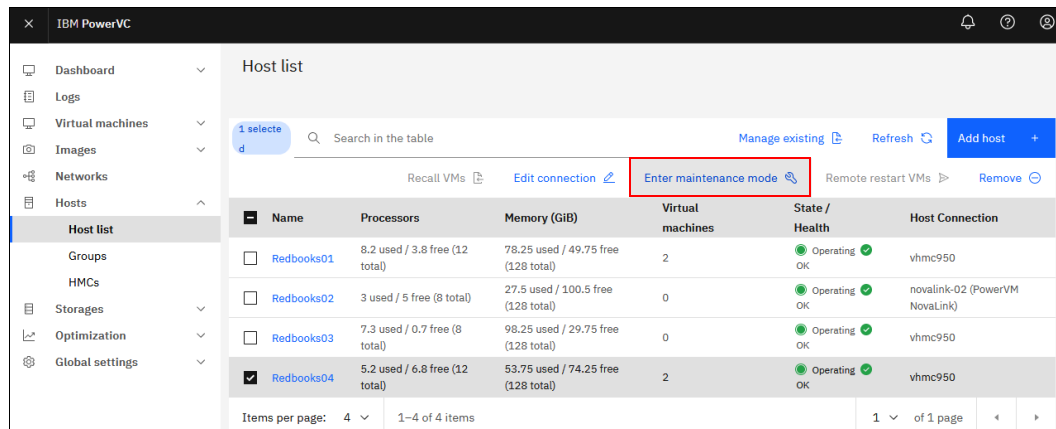


Figure 4-121 Put host in maintenance mode

- In the confirmation window you can select if you want to migrate all virtual machines to another host, or just active VMs or none. You also can select if you want the VMs to be moved to a specific host or if PowerVC shall select the hosts based on the placement policy of the host group where this host belongs to. Figure 4-122 shows an example.

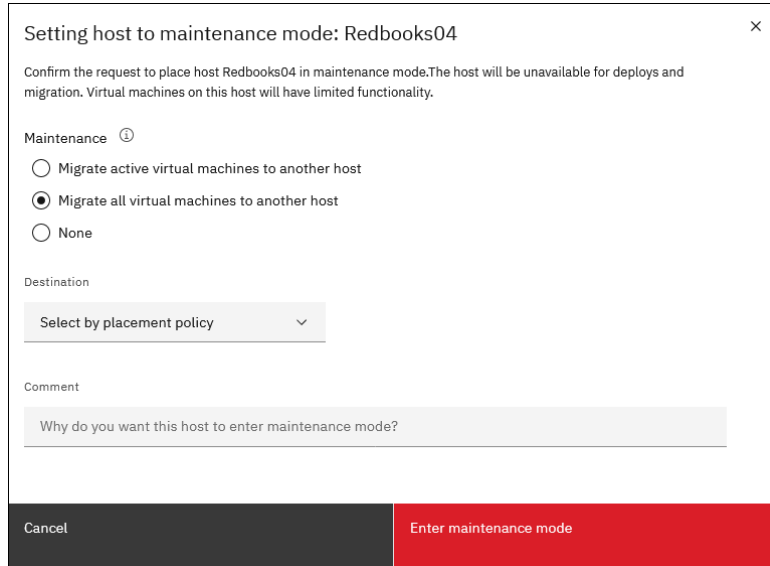


Figure 4-122 VMs to be migrated in maintenance mode

After clicking on **Enter maintenance mode** PowerVC starts the migration of VMs as selected in the confirmation window. During the migration the host shows Entering Maintenance in the state column of the host list.

When the maintenance mode is reached you can see Maintenance on in the state column of the host list as shown in Figure 4-123.

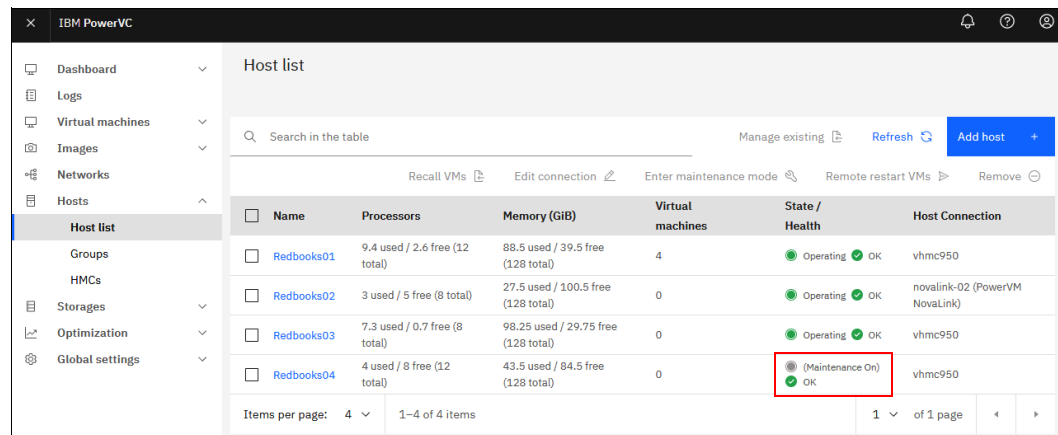


Figure 4-123 Host in maintenance mode

Exit the maintenance mode

To remove a host from maintenance mode, select the host and select **Exit Maintenance Mode** as shown in Figure 4-124.

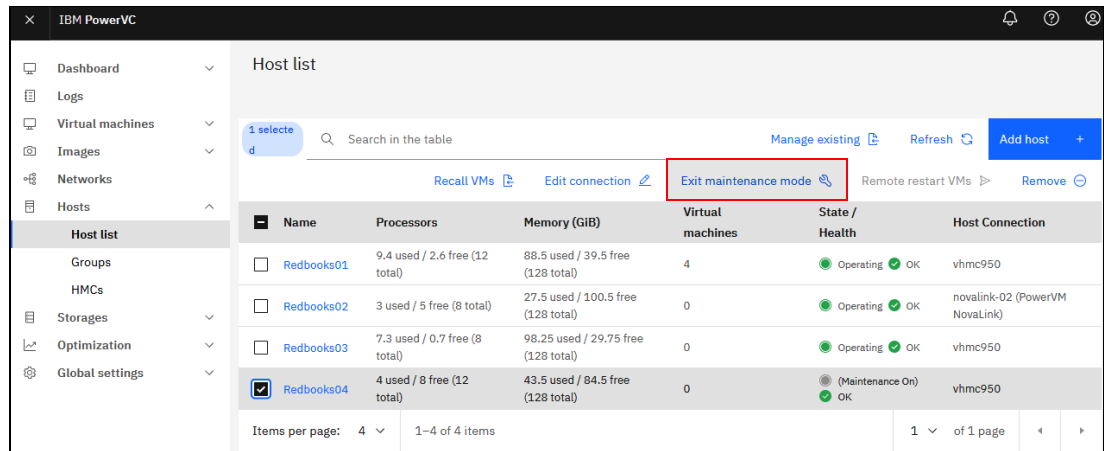


Figure 4-124 Exit maintenance mode

You will be asked for confirmation in the next screen. After you clicked **Exit** maintenance mode as shown in Figure 4-125 PowerVC will bring the host back into Operating state.

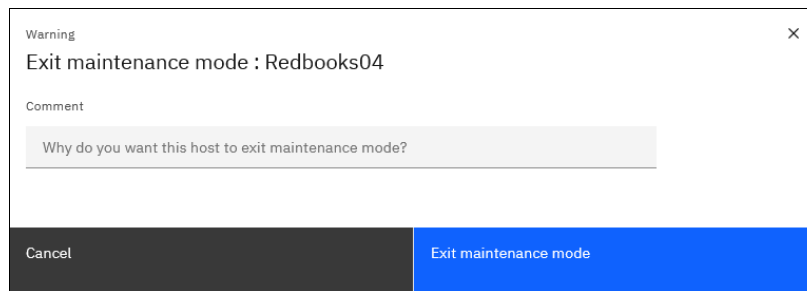


Figure 4-125 Exit maintenance mode confirmation

You can add VMs again to the host after it is brought out of maintenance mode.

Tip: You can edit the period after which the migration operation times out and the maintenance mode enters an error state by running the following commands:

```
/usr/bin/openstack-config --set /etc/nova/nova.conf DEFAULT prs_ha_timeout_seconds <duration_in_seconds>
```

For example, to set the timeout for two hours, run this command:

```
/usr/bin/openstack-config --set /etc/nova/nova.conf DEFAULT prs_ha_timeout_seconds 7200
```

Then, restart openstack-nova-ibm-ego-ha-service:

```
service openstack-nova-ibm-ego-ha-service restart
```

VM recall

When a host with the recall option enabled, as shown in Figure 4-126, exits the maintenance mode, all VMs that have been migrated to other hosts will be called back.

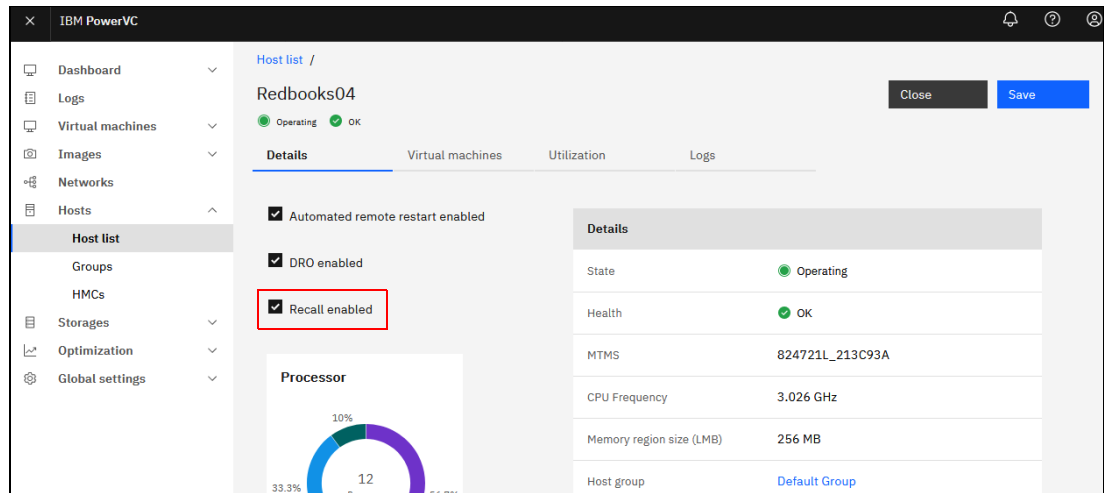


Figure 4-126 Recall option

After the host has exited the maintenance mode it takes a while until the VM migrations start.

4.15.11 Restarting virtual machines remotely from a failed host

IBM PowerVC can restart VMs remotely from a failed host to another host. To successfully restart VMs remotely by using IBM PowerVC, you must ensure that the source host and destination host are configured correctly.

Remote restart requirements

To restart a VM remotely, the following requirements must be met:

- ▶ PowerVC remote restart capability is required at the host and virtual machine level.
- ▶ Remote restart must be enabled on the virtual machine. You can set this option while deploying or resizing the virtual machine.
- ▶ Remote restart must be enabled on the host. You can verify support on the host details page. You cannot change this setting by using PowerVC.
- ▶ The source and destination hosts must have the same management type (PowerVM NovaLink or Hardware Management Console (HMC)).
- ▶ The new and old hosts must be members of the same host group.
- ▶ The Virtual I/O Servers on the target host must be members of the storage connectivity group specified on the virtual machine.
- ▶ PowerVC must be managing two or more hosts.
- ▶ The source and destination hosts must have all of the appropriate virtual switches that are required by networks on the virtual machine.
- ▶ The virtual machine state must be in Active, Error, or Shutoff with no task state. When you select to remote restart all virtual machines, PowerVC resets their state to Active and task state to None if their current state does not meet the remote restart requirements.
- ▶ The hosts and virtual machines must be capable of simplified remote restart capability.

- ▶ Only soft pinned virtual machines can be remote restarted. You cannot remote restart the hard pinned virtual machines. When soft pinned virtual machines are moved during host maintenance or remote restart, PowerVC automatically migrates the virtual machines back to the original host once the host is back to operating state.
- ▶ If soft pinned virtual machines is moved to a different host and if the host to which it is initially pinned is operating, then you cannot migrate the virtual machine to a new host.
- ▶ Remote Restart is supported with POWER8 or later servers. Consider the following:
 - The hosts must be running firmware 820 or later.
 - The HMC must be running with HMC 820 Service Pack (SP) 1 or later, with the latest program temporary fix (PTF).
 - Remote restart for virtual machines with shared storage pools must use HMC 840 and VIOS 2.2.4.0.
- ▶ The virtual machine remote restart state must be Remote restartable before the source host is down. For details on remote restart states of a virtual machine, see the Simplified Remote Restart documentation.
- ▶ The source host compute service must be down or the host must be in Error, Power off, Error - dump in progress, or FSP unreachable state. If the host is in FSP unreachable state, verify that the host is actually down before initiating the remote restart process.
- ▶ Only virtual machines that are using dedicated processors or assigned to the default processor pool can be remote restarted.

Before you can restart a VM on PowerVM remotely, you must deploy or configure the VM with remote restart capability. This can be done in two ways:

1. Create a compute template with the enabled virtual remote restart capability. You can find this attribute under the Miscellaneous tab of the compute template as shown in Figure 4-127. A new deployed VM with that compute template, gets the Remote Restart enabled flag.

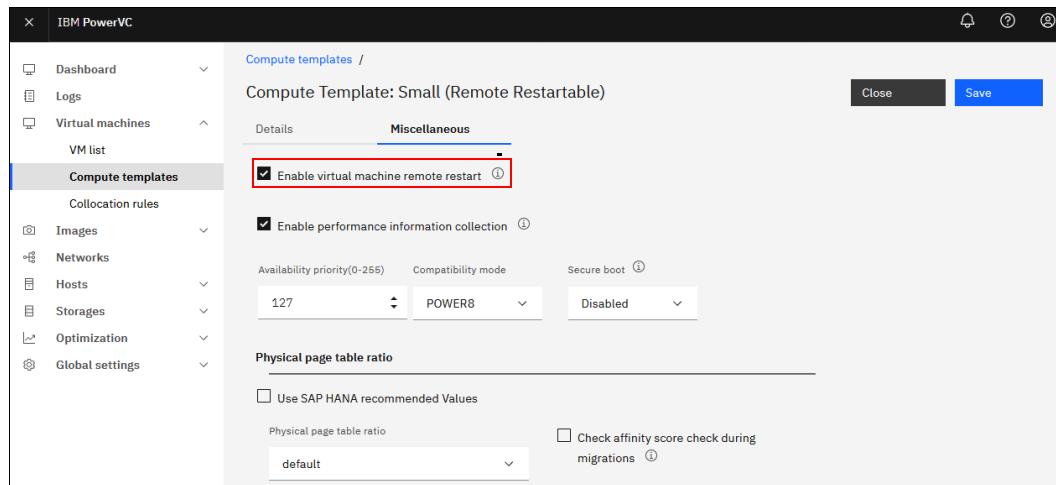


Figure 4-127 Remote Restart enabled Compute Template

2. Modify the remote restart property after the VM is deployed. For that, the VM must be stopped, because this attribute is not changeable while the VM is running. Now enter the VM details by clicking the VM name under **Virtual machines, VM list** as shown in Figure 4-128. There check the **Remote restart enabled** checkbox.

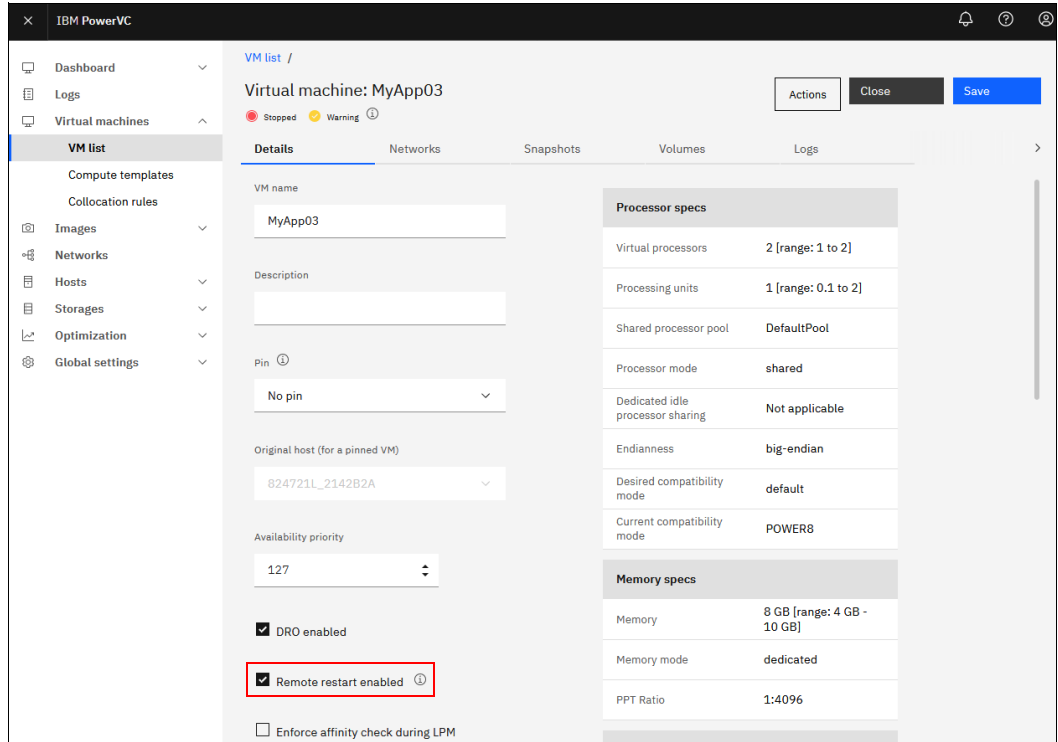


Figure 4-128 Remote restart attribute in the VM details

Note: When the VM is running, the Remote Restart attribute is disabled and cannot be changed.

To check, if a VM is remote restartable you can verify the state on the details tab of a VM as shown in Figure 4-129.

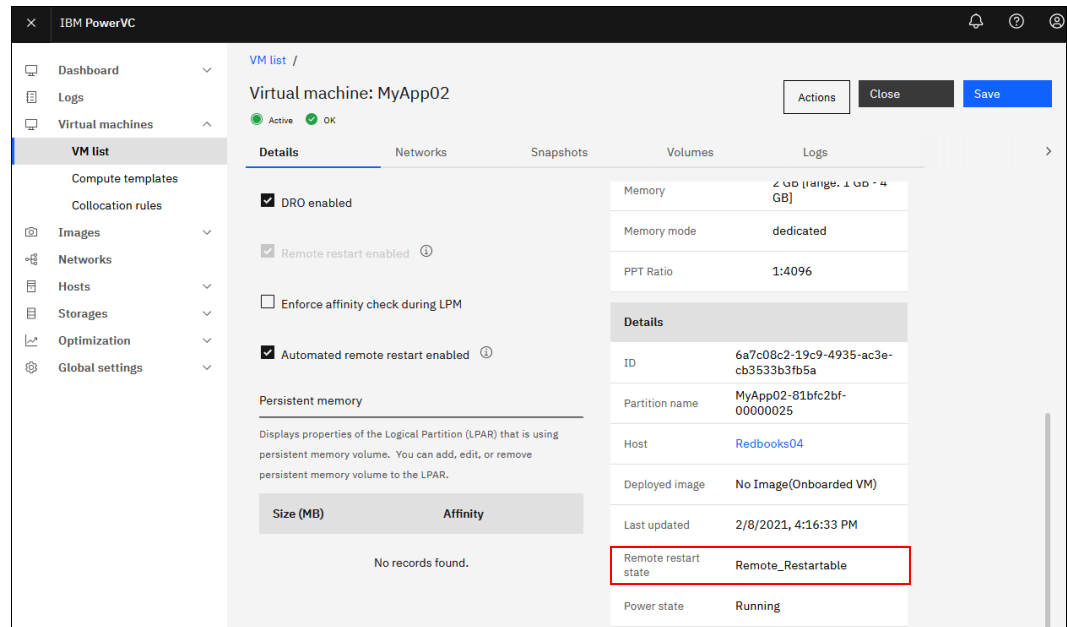


Figure 4-129 Remote Restart state of a VM

A VM can be restarted remotely in PowerVM only if the Remote Restart state is Remote restartable. When a VM is deployed initially, the HMC must collect partition and resource configuration information. The remote restart state changes from Invalid to different states. When it changes to Remote restartable, IBM PowerVC can initiate the remote restart operation for that VM.

For a new deployed remote restart enabled VM the status is usually Local_Storage_Update_Failed. This is normal, because PowerVC adds in one of the Virtual I/O servers a little virtual optical disk to the VM. This disk is used to pass information such as hostname or IP address to the VM. As this disk is not available on a potential remote restart destination, the remote restart state is not Remote_Restartable. It may take more than an hour until the disk will be automatically removed from the VIOS by PowerVC. When it was removed the state for remote restart changes automatically (if there are no other reasons that prevent it).

Tip: If you do not want to wait, remove the virtual optical device in the VIOS and click **Refresh Remote Restart Data** button (the two blue curved arrows) in the advanced settings under the general properties of the host in the HMC.

Restarting virtual machines remotely

You can start a remote restart task, when the host is in Power off, Error or Error - dump condition.

1. To start it, select the host and click **Remote restart VMs** as shown in Figure 4-130.

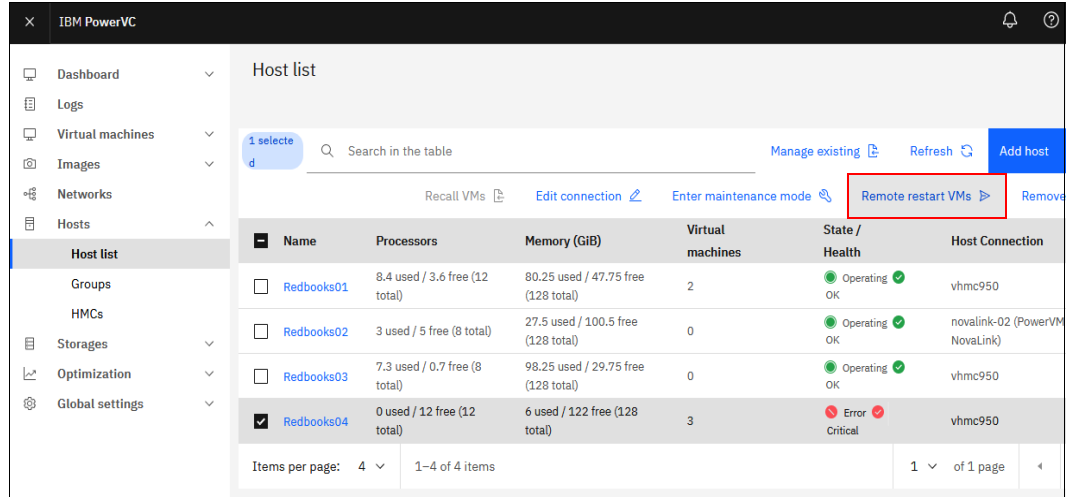


Figure 4-130 Start remote restart

2. In the next screen you can select the remote restart enabled VMs that should be remote restarted and it is possible to select a destination host for the VMs. Alternatively, let PowerVC select the host(s) by the placement policy linked to the host group of the selected host as shown in Figure 4-131.

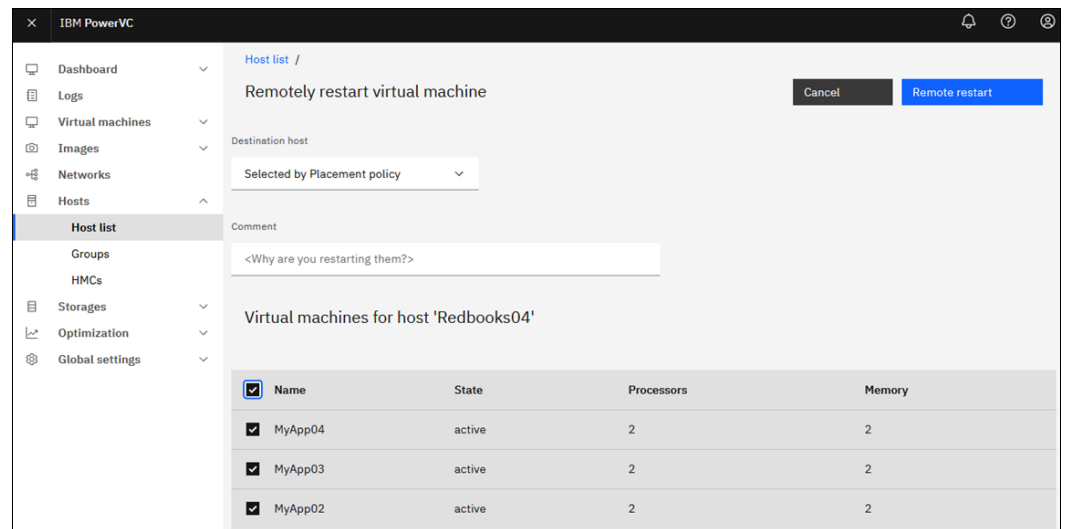


Figure 4-131 Select VMs to be remote restarted

In this example we made MyApp03 soft-pinned and MyApp04 hard-pinned. MyApp02 was not pinned. More information about pinning can be found in section “Pinning” on page 185.

- After you clicked on **Remote restart**, the host will enter the Remote Restart Rebuilding state as shown in Figure 4-132.

The screenshot shows the IBM PowerVC interface with the 'Host list' tab selected. The table displays the following data:

Name	Processors	Memory (GiB)	Virtual machines	State / Health	Host Connection
Redbooks01	6.7 used / 5.3 free (12 total)	74.25 used / 53.75 free (128 total)	2	Operating OK	hmc01
Redbooks02	3.5 used / 4.5 free (8 total)	31.5 used / 96.5 free (128 total)	1	Operating OK	novalink-02 (PowerVM NovaLink)
Redbooks03	8 used / 0 free (8 total)	72.25 used / 55.75 free (128 total)	0	Operating OK	vhmc950
Redbooks04	0 used / 12 free (12 total)	6 used / 122 free (128 total)	3	Remote Restart Rebuilding OK	hmc01

Figure 4-132 Host in Remote Restart Rebuilding

Note: In case Remote Restart Error appears, an error occurred while one or more VMs were moved to the destination host. You can check the reasons for the failure in the corresponding compute log file in the `/var/log/nova` directory.

During the Remote Restart Building state you also can see Rebuilding or rebuilding_spawning in the VM list under Virtual machines. In the state rebuilding_spawning the VMs show up in the HMC under the new host(s). Figure 4-133 shows our example:

The screenshot shows the IBM PowerVC interface with the 'Virtual machines' tab selected. The table displays the following data:

Name	Host	IP	State / Health	Resources	Running task
Red1	Redbooks02	10.10.2.88	Active Warning	1 vCPU / 4 GB / 0.5 PU	
MyApp04	Redbooks04	9.47.76.116	Active OK	2 vCPU / 2 GB / 0.2 PU	
MyApp03	Redbooks04	10.10.2.213	Active Warning	2 vCPU / 2 GB / 0.2 PU	rebuild_spawning
MyApp02	Redbooks04	9.47.76.114	Active Warning	2 vCPU / 2 GB / 0.2 PU	rebuild_spawning
nimlpar	Redbooks01	9.47.76.107, 10.10.0.10	Active OK	2 vCPU / 2 GB / 0.2 PU	
powervc02	Redbooks01	9.47.76.117, 10.10.0.11	Active OK	2 vCPU / 24 GB / 2 PU	

Figure 4-133 VMs during Remote Restart

Figure 4-133 also shows that in our example MyApp04 will be not rebuilt, although we selected it in Figure 4-131 on page 200. The reason is, that this VM was hard-pinned.

Note: Pinning is a PowerVC feature. The HMC does not know about pinned VMs. Therefore you can force a remote restart of the VM on the HMC using the command:

```
$ rrestartlpar -o restart -m <source system> -p <partition name> -t <Destination system>
```

Keep in mind that the partition name in PowerVC does not show the unique suffix, therefore make sure to use the complete VM name shown in the details of the VM in PowerVC or in the partition list in the HMC.

Recovery after Remote Restart

When the host that was remotely restarted is back for operation, an automatic cleanup is initiated (from the HMC or NovaLink) that removes the remote restarted VMs from the failed host to prevent data access from both VMs, the failed and the remotely restarted one.

The hard-pinned VM MyApp04 stayed on the host Redbook04 and remained stopped.

You can see that information in the notifications. To show them click on the bell symbol in the right upper corner of the screen as shown Figure 4-134.

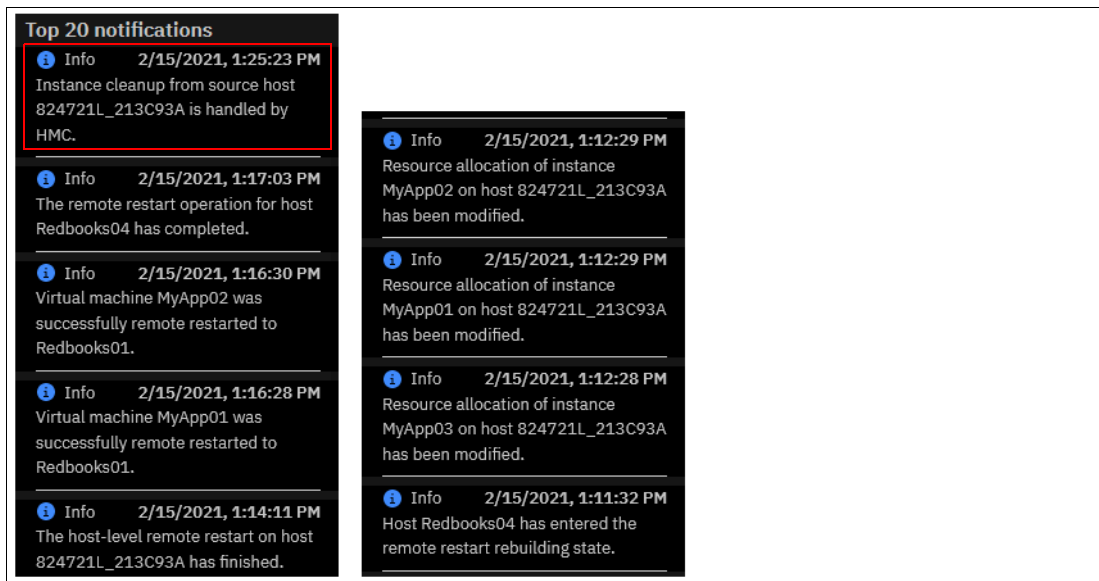


Figure 4-134 Remote restart related notifications

Manual recall

VMs that are soft pinned save the original host, where they have been pinned. When the failed host is back for operation, you can move the soft pinned VMs back with **Migrate**, or simply start a **Recall VMs** from the host list as shown in Figure 4-135.

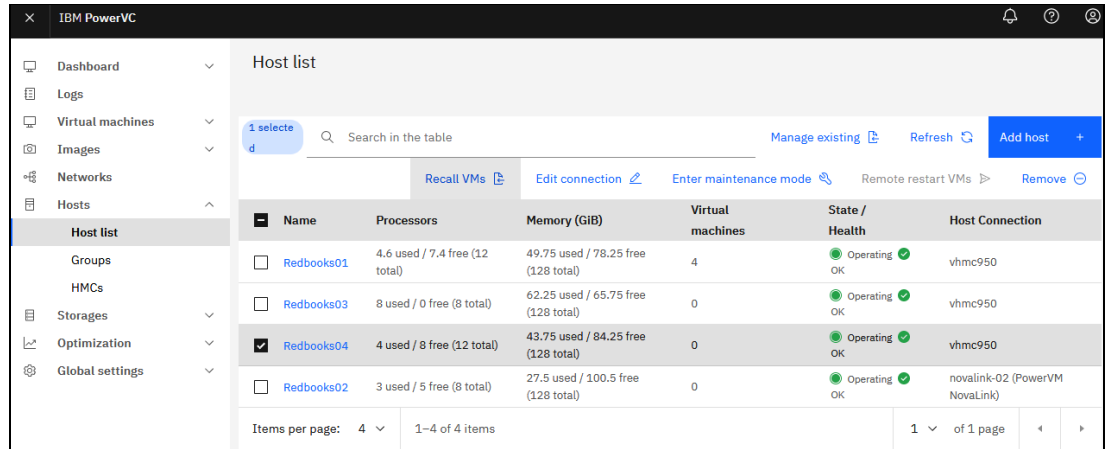


Figure 4-135 Recall VMs

Before the recall starts you must confirm the operation as show in Figure 4-136.

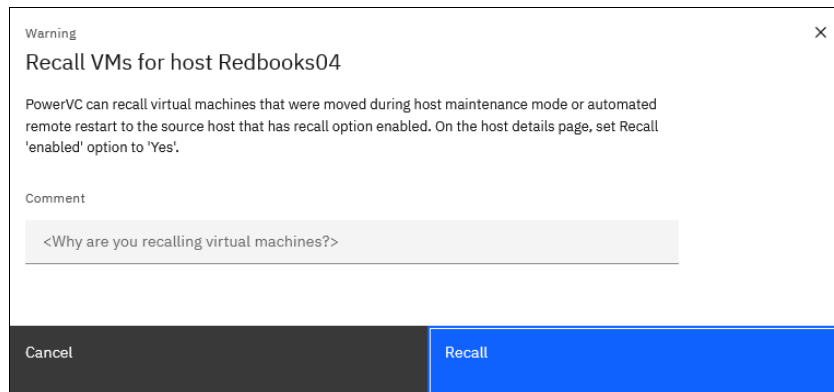


Figure 4-136 Recall VMs - confirmation

Now all VMs that were located on the failed host will be migrated back to its source host.

Automated remote restart

PowerVC is also able to do an automated remote restart. This is not possible from an Hardware Management Console without PowerVC, because the HMC only has a host view. PowerVC also has the ability to check if there is traffic on the storage volumes to prevent split situations where the host is not reachable anymore, but the VMs on it are still running and accessing data.

Automated remote restart is enabled on a host group level as shown in Figure 4-137. If activated you can set the following attributes:

- ▶ **Run interval**
The frequency that the state of the host is checked.
- ▶ **Stabilization**
The number of consecutive run intervals that the host must be down before an automated remote restart operation is initiated.

In the example shown in Figure 4-137, the failed host will be checked five times in a interval of 1 minute.

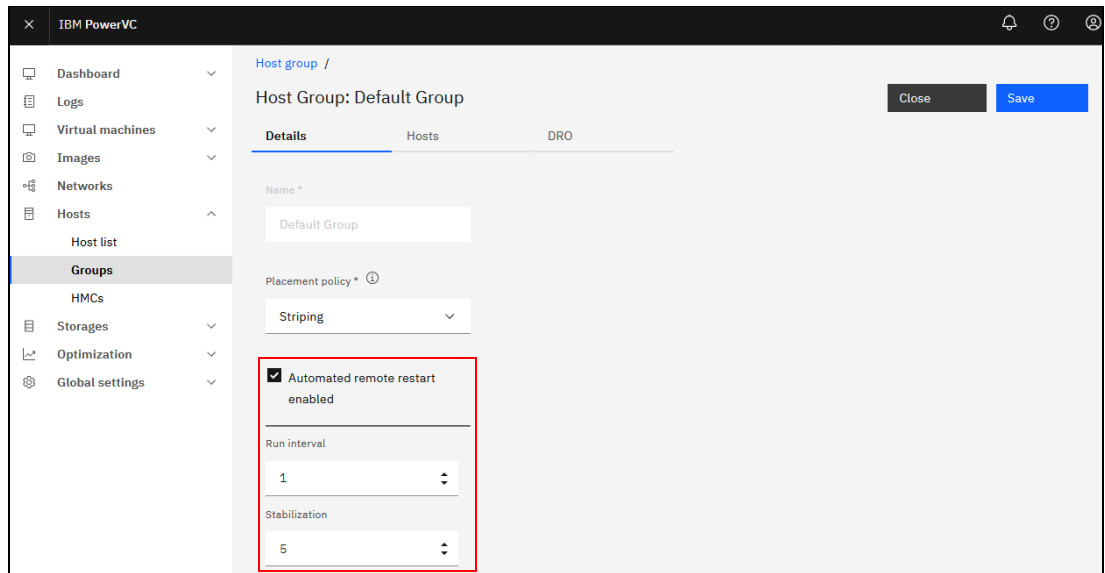


Figure 4-137 Automated remote restart in a host group

After you enabled automated remote restart in general, you can select if a host should be included in a remote restart task by checking the Automated remote restart enabled attribute in the details tab of a host as shown in Figure 4-138. This is the default for a host.

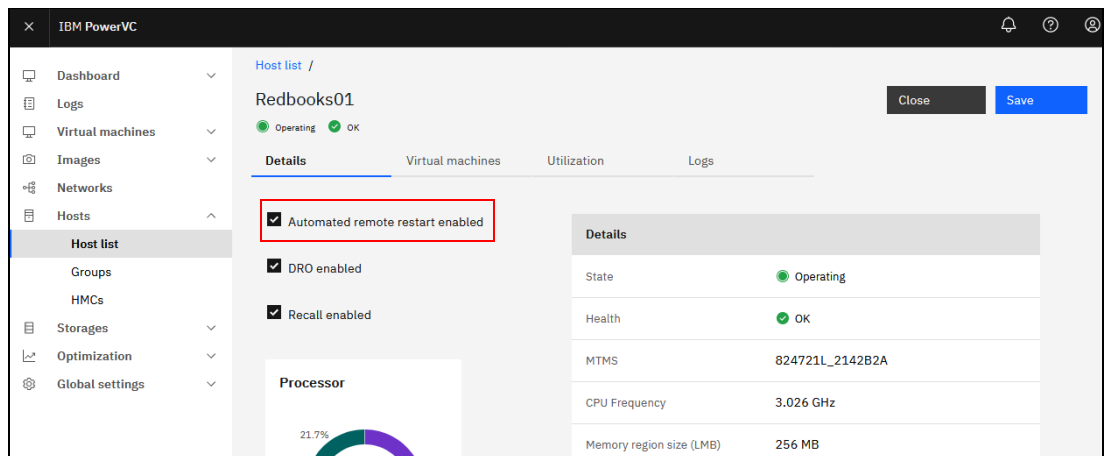


Figure 4-138 Host included in an automated remote restart task

It is also possible to exclude a virtual machine from an automated remote restart task. To do that, enter the VM details tab of a VM and uncheck the Automated remote restart enabled attribute as shown in Figure 4-139.

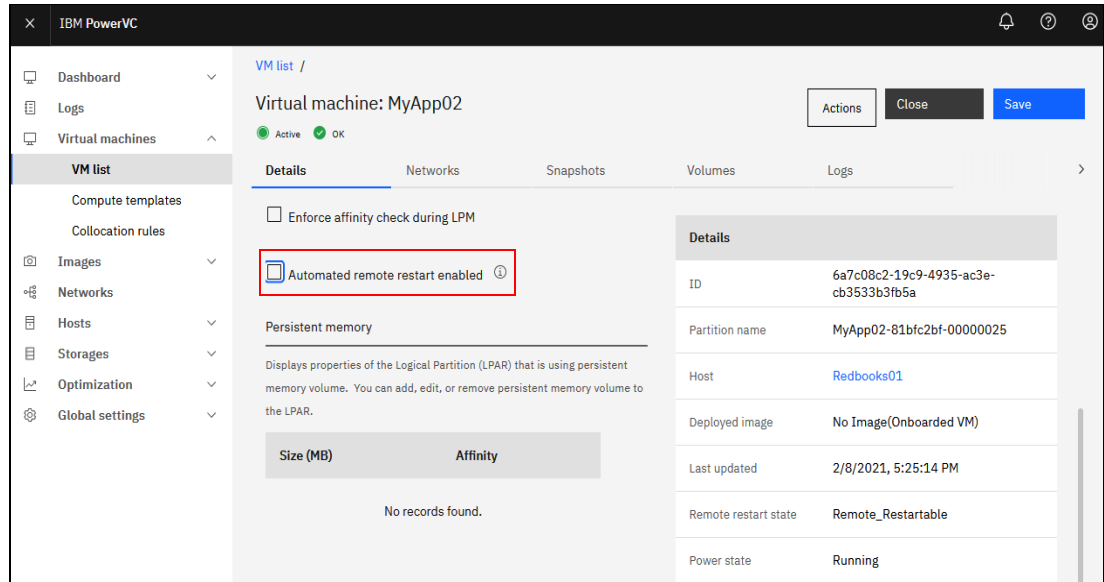


Figure 4-139 Exclude a VM from automated remote restart

Recall from automated remote restart

When a host is back for operation after an automated remote restart, the HMC will do the remote restart cleanup as mentioned above. PowerVC will also automatically recall the VMs to the original host using a migration task (live partition mobility). This is the same functionality as with putting a host into maintenance mode. Refer to 4.15.11, “Restarting virtual machines remotely from a failed host” on page 196 for details.

4.15.12 Resetting the state of a virtual machine

In certain situations, a VM becomes unavailable or it is in an unrecognized state for the IBM PowerVC management server. When these situations occur, you can run a Reset VM state procedure. This process sets the machine back to an consistent state.

One example for such a situation is when you try to start a VM on a NovaLink-managed system, when NovaLink is not in control of the system. In such a situation the VM is constantly in a Starting condition and needs a reset of the state.

Figure 4-140 shows this example. When you click on the **information icon** you be presented additional information. In this case the host is maintenance mode as NovaLink does not have control over the system.

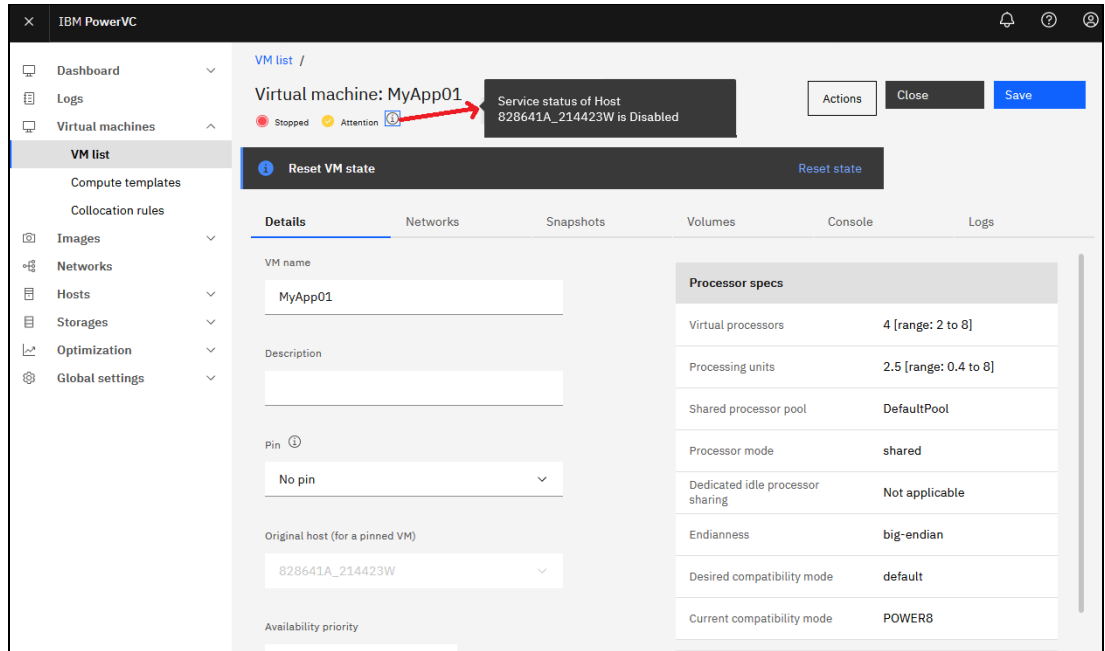


Figure 4-140 Reset the virtual machines state

Click **Reset State** to start the reset process. This process can take a few minutes to complete.

Note: No changes are made to the connection or database.

4.15.13 Unmanaging a virtual machine

The Unmanage function is used to discontinue the management of a VM from IBM PowerVC. After a VM becomes unmanaged, the VM is no longer listed in the Virtual Machines window, but the VM still exists. The VM and its resources remain configured on the host. The VM remains running and can still be managed from the HMC.

1. To unmanage a VM, open the Virtual Machines window, and select the VM(s) that you want to remove from IBM PowerVC. Click **Unmanage** to remove this VM or multiple VMs from the IBM PowerVC environment.

Figure 4-141 shows the Unmanage option to unmanage a VM.

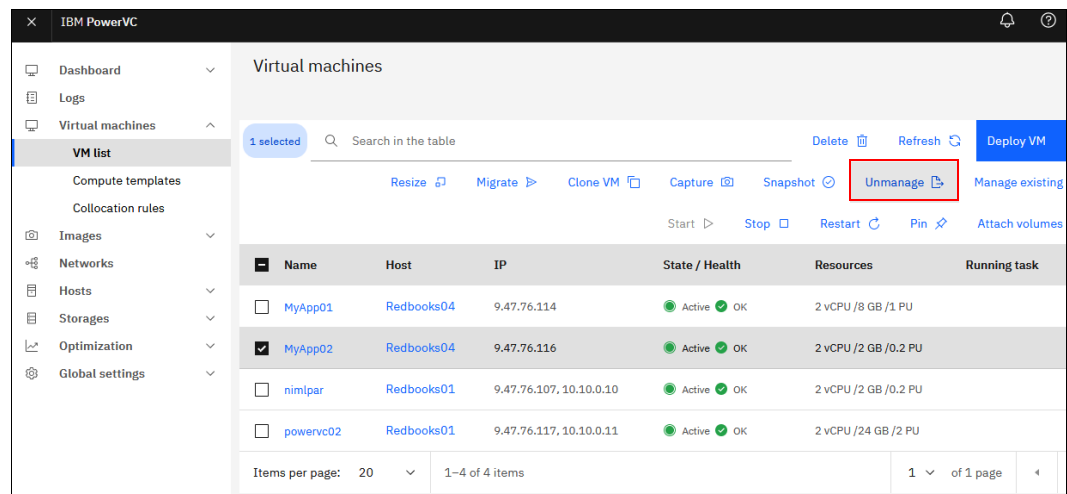


Figure 4-141 Unmanage a VM

2. After clicking **Unmanage**, you will be asked to confirm the operation as shown in Figure 4-142.

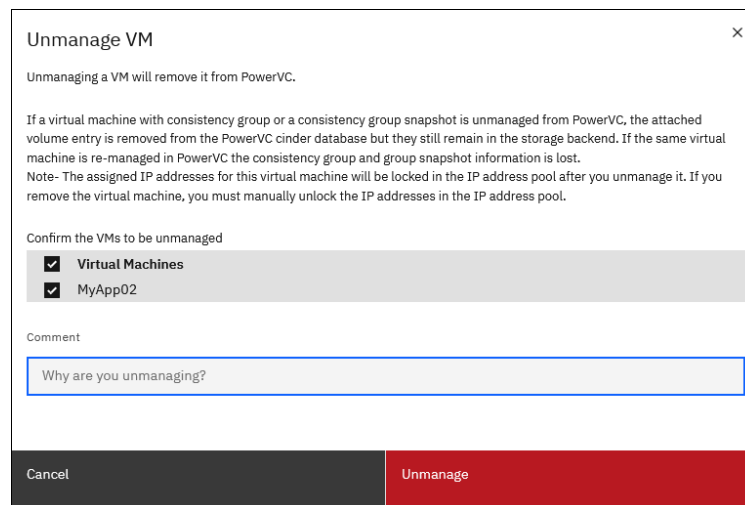


Figure 4-142 Unmanage a VM - confirmation

After a short amount of time, the VM disappears in PowerVC, but is still visible in the HMC or Novalink.

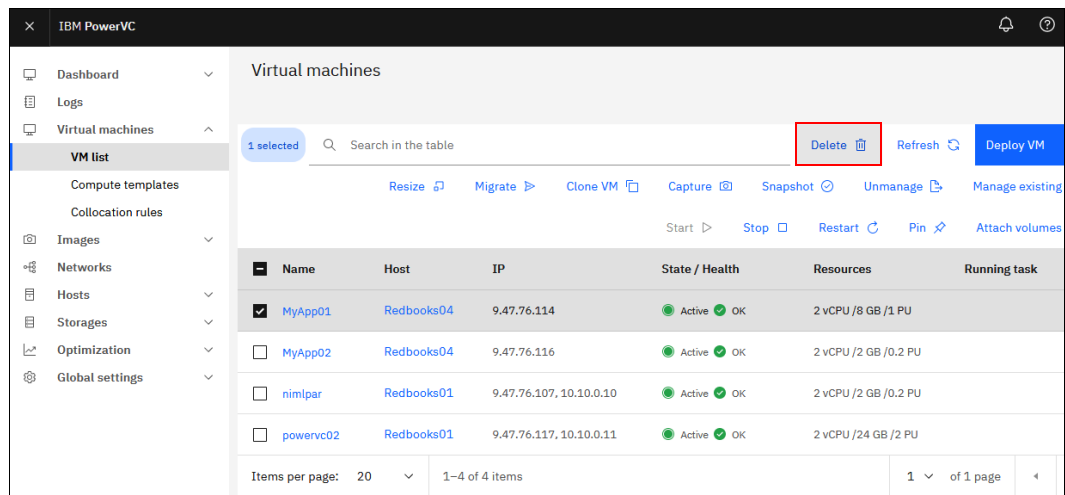
4.15.14 Deleting a virtual machine

IBM PowerVC can delete VMs completely from your systems.

Important: By deleting a VM, you completely remove the VM from the host system and from the HMC or PowerVM NovaLink configuration, and IBM PowerVC no longer manages it. Volumes that are used by VM can be deleted with granularity.

1. To remove a VM, open the Virtual Machines window and select the VM that you want to remove. See Figure 4-143 for an example. It is also possible to select and delete multiple VMs.

Note: You can delete a VM while it is running. The process stops the running VM and then deletes it.



The screenshot shows the IBM PowerVC interface for managing virtual machines. The left sidebar contains navigation options like Dashboard, Logs, Virtual machines, VM list, Compute templates, Collocation rules, Images, Networks, Hosts, Storages, Optimization, and Global settings. The main area displays a table of virtual machines with columns for Name, Host, IP, State / Health, Resources, and Running task. The 'MyApp01' VM is selected, and the 'Delete' button is highlighted with a red box. Other actions like Resize, Migrate, Clone VM, Capture, Snapshot, Unmanage, and Manage existing are also visible. The table shows four VMs: MyApp01, MyApp02, nimlpar, and powervc02, all with 'Active' status and 'OK' health.

Name	Host	IP	State / Health	Resources	Running task
<input checked="" type="checkbox"/> MyApp01	Redbooks04	9.47.76.114	Active OK	2 vCPU / 8 GB / 1 PU	
<input type="checkbox"/> MyApp02	Redbooks04	9.47.76.116	Active OK	2 vCPU / 2 GB / 0.2 PU	
<input type="checkbox"/> nimlpar	Redbooks01	9.47.76.107, 10.10.0.10	Active OK	2 vCPU / 2 GB / 0.2 PU	
<input type="checkbox"/> powervc02	Redbooks01	9.47.76.117, 10.10.0.11	Active OK	2 vCPU / 24 GB / 2 PU	

Figure 4-143 Delete a VM

2. Click **Delete** and the confirmation windows opens as shown in Figure 4-144.

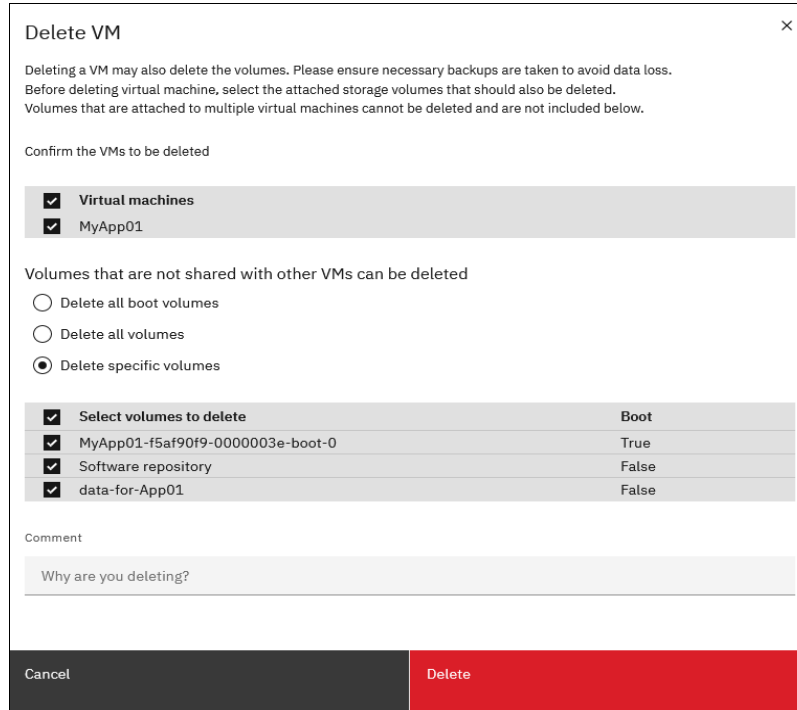


Figure 4-144 Delete VM - confirmation

In the confirmation window you can select to delete all volumes that belong to the VM(s), none of the volumes or specific. By selecting **Delete specific volumes** as shown in Figure 4-144 the volumes will show up and you can select the volumes to delete.

Note: Delete specific volumes is not available when you selected multiple VMs for deletion.

PowerVC will only delete volumes that are not additionally attached to another VM (a shared volume). In the example shown in Figure 4-144 you see that the volume with the name Software repository is marked for deletion. As this is a shared volume that is also attached to another VM, it will not be deleted although it is selected to be deleted. The volume remains in the list of manageable data volumes and it also remains attached to the second VM.

VMs that are not marked for deletion will also stay in the list of volumes and can be used for other purposes.

If the attachment method for the volumes was NPIV, the zoning to the storage and also the host definition in the storage will be deleted.

4.16 Images

Images are a central concept in PowerVC. An image consists of one or more volumes with at least one bootable volume that has an installed operating system. Images are specially prepared and can be used many times to be deployed on different hosts as new virtual machines. An image is a blueprint of a virtual machine that can be duplicated.

Supported operating systems for an image are: AIX, IBM i, Red Hat Enterprise Linux, SUSE Linux Enterprise Server, Ubuntu and also Red Hat CoreOS to be used in an OpenShift environment.

The following sections describe how to get or create an image and also how to use it:

- ▶ 4.16.1, “Preparation of a new image” on page 210
- ▶ 4.16.2, “Capturing a virtual machine image” on page 224
- ▶ 4.16.3, “Import an image from OVA or raw image” on page 226
- ▶ 4.16.4, “Define an image with existing volumes” on page 231
- ▶ 4.16.5, “Export images into OVA” on page 232
- ▶ 4.16.6, “Import and export with backup service” on page 235
- ▶ 4.16.7, “Deploying a new virtual machine” on page 239
- ▶ 4.16.8, “Deleting images” on page 246

4.16.1 Preparation of a new image

There are many ways to get an image as basis to deploy new virtual machines. One way is to create an image from scratch:

1. Create an LPAR using the HMC or NovaLink with all required virtual adapters.
2. Create SAN zones (in case of NPIV attached storage).
3. Create a volume or multiple volumes in the storage.
4. Attach the volumes in the storage to the storage host.
5. Create the necessary mappings in VIOS.
6. Install the operating system.
7. Do some customizing, for example install agents or tools that are required in your environment.
8. Install cloud-init as the necessary deployment tool.

Note: Do not use any physical devices such as local disk, SR-IOV adapters (without vNIC) or PCIe adapters in the VM, because PowerVC only accepts fully virtualized VMs for management.

An image may also consist of several volumes. You can for example also add data disks or a mirrored operating system disk if you want.

When you finished your manual virtual machine for the image, then onboard it into PowerVC as discussed in 4.15.1, “Virtual machine onboarding” on page 158.

If you have PowerVC already running with hosts, storage and network configured you can also let PowerVC do steps 1 to 5 by using an empty image as discussed in the next section.

Define an empty image

You can define an empty image with the following steps:

1. Create a data volume used for the operating system in **Data volumes** under **Storages** as shown in Figure 4-145. Select the size according to you desired operating disk size.

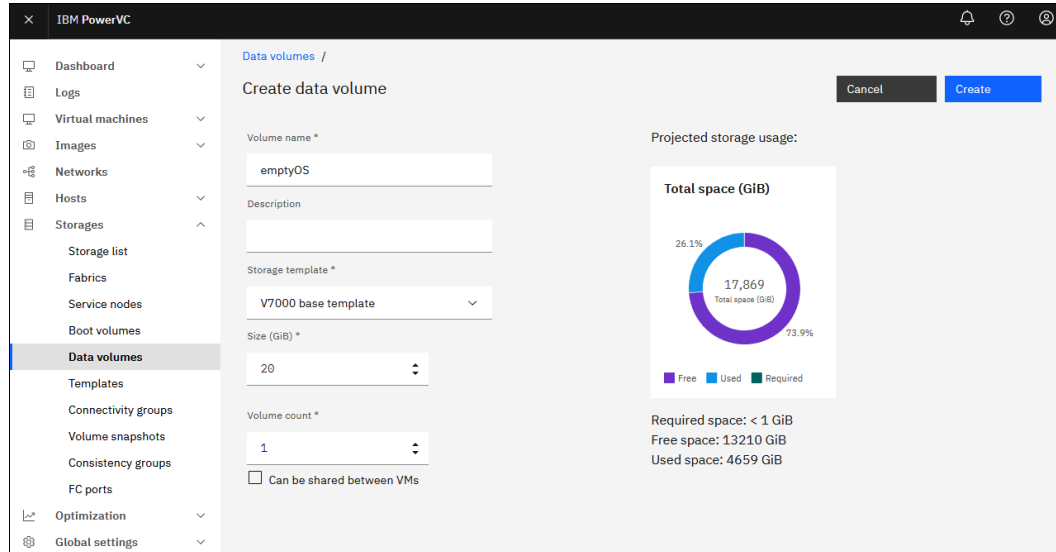


Figure 4-145 Volume for an empty image

2. Define the empty image.

In **Image list** under **Images** click **Create**. Now enter in the details tab the name of the image and select the operating system that you are planning to install. Also select the correct Endianness as shown in Figure 4-146.

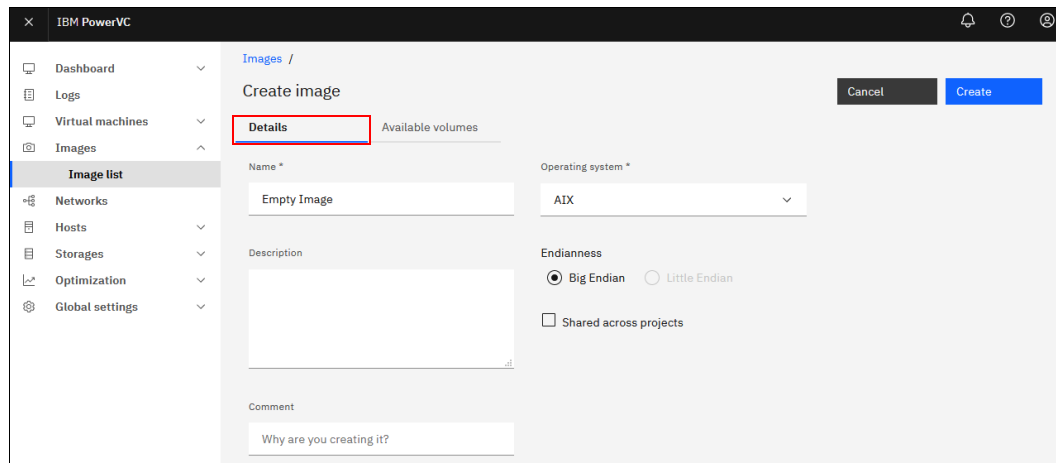


Figure 4-146 Define an empty image - details

On the Available volumes tab select the volume that you created in step 1. Figure 4-147 shows an example:

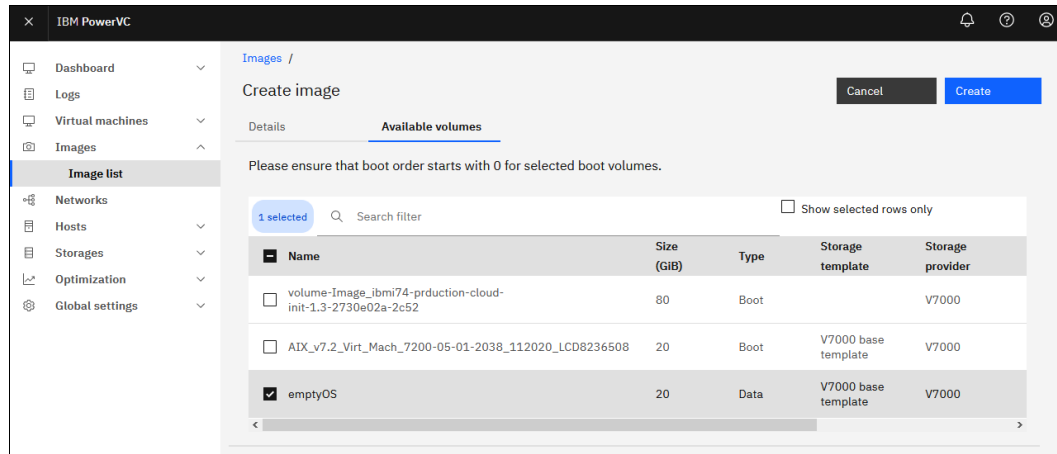


Figure 4-147 Define an empty image - volumes

Do not forget to scroll to the right and enter 0 in the Boot order column (not visible in Figure 4-147). After that click **Create**.

3. In the next step deploy the empty image.

In the Image list under Images select the empty image you just created and click **Deploy**. In the tabs enter all necessary information to build a VM. You can also add additional new volumes, if you want more than one volume in you image. An example is shown in Figure 4-148.

Details on how to deploy a VM using an image you can find in 4.16.7, “Deploying a new virtual machine” on page 239.

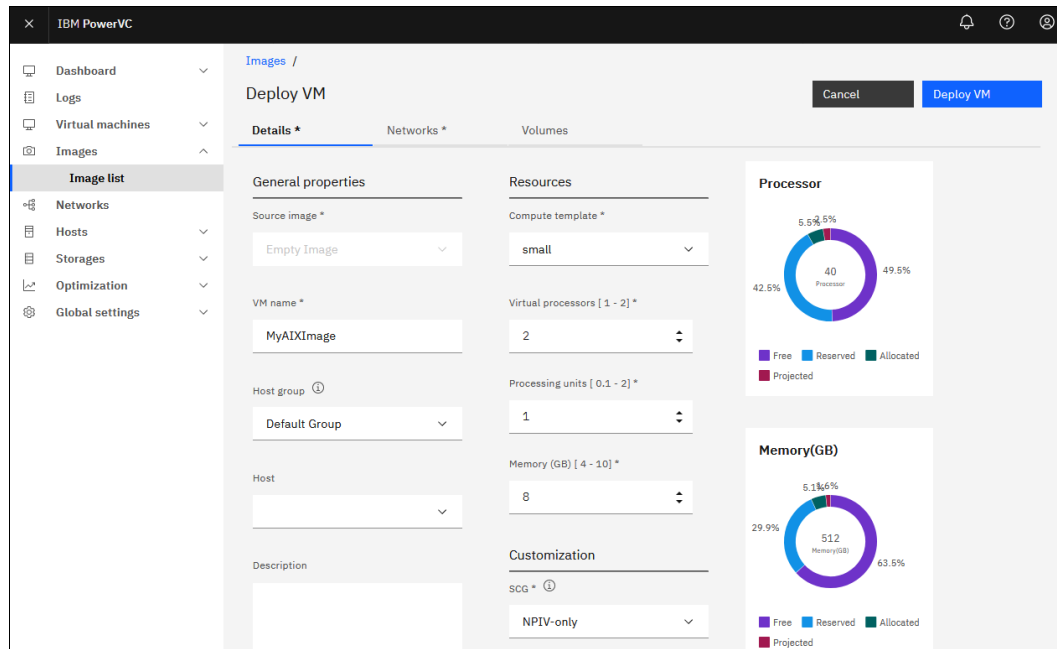


Figure 4-148 Deployment of the empty image

- After a while the new VM shows up, with its assigned IP address and the volume connected as defined during the deployment. The Health state of the VM will stay Warning, as no operating system is installed. Figure 4-149 shows the new created VM.

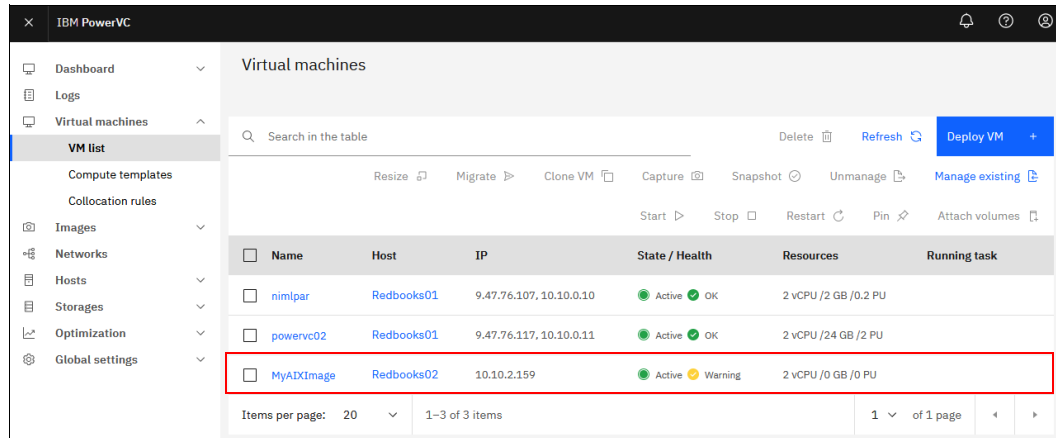


Figure 4-149 VM with empty image is ready

- Now use the assigned IP address to do a network install of the operating system or attach a virtual DVD with the operating ISO to the VM and install the operating system.
- After the installation of the operating system and the configuration of the network, the Health state should turn to OK for AIX and IBM i. Linux operating systems also need the IBM Service and Productivity tools installed to have RMC available as requirement for a Health state of OK. This is discussed below.
- You can optionally add your tools and do your customizing of the VM that should be included into the image.

Preparing a virtual machine with cloud-init

The `cloud-init` script enables VM activation and initialization, and it is widely used for OpenStack. Before you capture a VM, install the `cloud-init` initialization package. Cloud-init takes care in the first start of a new deployed VM that for example IP addresses, hostname or ssh keys are changed. This package is available in the `/opt/ibm/powervc/images/cloud-init` path in the IBM PowerVC host and is also part of all supported Linux distros.

Important: If you are installing the `cloud-init` package on AIX or Linux to capture a VM on which the Activation Engine (AE) is already installed, you must first uninstall the AE. The Activation Engine was the predecessor tool before cloud-init. It still works, but is no longer supported.

To check whether the AE Red Hat Package Managers (RPMs) are installed, run the following command on the VM (IBM i does not require you to remove AE, just install the cloud-init PTF):

```
# rpm -qa | grep activation
```

Cloud-init preparation for AIX

In AIX the preferred method of installing cloud-init is the usage of yum.

Information about how to install yum on AIX could be found here:

<https://public.dhe.ibm.com/aix/freeSoftware/aixtoolbox/eziinstall/ppc/README-yum>

After yum is ready, install cloud-init with:

```
# yum install cloud-init
```

The main configuration file for cloud-init is the cloud.cfg, that you can find in AIX under /opt/freeware/etc/cloud/cloud.cfg.

Note: For Linux the cloud.cfg is located in /etc/cloud/cloud.cfg.

In this file you find several modules to activate:

Table 4-6 Modules and descriptions

Module	Description
restore_volume_group	This module restores non-rootVG volume groups when you deploy a new VM.
set_multipath_hcheck_interval	Use this module to set the hcheck interval for multipath. If you deploy a multidisk VM and this module is enabled, you can deploy specifying a cloud-config data entry that is named multipath_hcheck_interval and give it an integer value that corresponds to seconds. Post-deployment, each of the VM's disks must have their hcheck_interval property set to the value that was passed through the cloud-config data. Run the <code>lsattr -E1 hdisk# -a hcheck_interval</code> command for verification. If you do not specify the value within the cloud-config data, the module sets each disk's value to 60 seconds.
set_hostname_from_dns	Use this module to set your VM's host name by using the host name values from your Domain Name Server (DNS). To enable this module, add this line to the cloud_init_modules section: - set_hostname_from_dns Then, remove these lines: - set_hostname - update_hostname
set_hostname_from_interface	Use this module to choose the network interface and IP address to be used for the reverse lookup. The valid values are interface names, such as eth0 and en1. On Linux, the default value is eth0. On AIX, the default value is en0.
set_dns_shortcode	This module specifies whether to use the short name to set the host name. Valid values are True to use the short name or False to use the fully qualified domain name. The default value is False.

More information about the modules could be found in IBM Knowledge Center:

https://www.ibm.com/support/knowledgecenter/SSVSPA_2.0.0/com.ibm.powervc.cloud.help.doc/powervc_custom_modules_aix_cloud.html

Important: For an AIX image with multiple Volume Groups (VGs), run the /opt/freeware/lib/cloud-init/create_pvid_to_vg_mappings.sh command to save the information that is used to restore custom volume groups on all VMs that are deployed from the image that is captured. Saving this information is useful if you have a multidisk VM that has a dataVG volume group defined. The module restore_volume_group restores the VGs after the deployment.

You can also integrate custom scripts into cloud-init. For that use the following directories:

- ▶ /opt/freeware/var/lib/cloud/scripts/per-boot
Scripts in this directory will be executed in every boot.
- ▶ /opt/freeware/var/lib/cloud/scripts/per-instance
Scripts in this directory will be executed once after the deployment.
- ▶ /opt/freeware/var/lib/cloud/scripts/per-once
Scripts in this directory will be executed once after the deployment.

Note: The directories will not be created during the installation of cloud-init. After a reboot the directories will be created. Alternatively you can create the directories with the following commands:

```
# mkdir -p /opt/freeware/var/lib/cloud/scripts/per-boot
# mkdir /opt/freeware/var/lib/cloud/scripts/per-instance
# mkdir /opt/freeware/var/lib/cloud/scripts/per-once
```

Tip: An AIX image that will be deployed will always have the same PVID(s) for the rootvg.

To get a new PVID, the following needs to be done before you capture the VM:

```
# chdev -l sys0 -a ghostdev=2
```

An image with the ghostdev attribute set to 2 will assign a new PVID in the boot process and afterwards set the ghostdev parameter back to 0 to avoid assigning a new PVID on every boot.

Cloud-init preparation for Red Hat Enterprise Linux

Before you install cloud-init in Red Hat Enterprise Linux, it is recommended to install the IBM Service and Productivity tools. This is the basis for RMC and with that to get a Health state of OK in PowerVC:

1. Get the latest version of the ibm-power-repo package from the following website:
<https://www.ibm.com/support/pages/service-and-productivity-tools>
2. Install the RPM as in Example 4-3.

Example 4-3 Installation of IBM Service and Productivity tools for Red Hat Enterprise Linux

Install RPM:

```
# rpm -Uvh /tmp/ibm-power-repo-latest.noarch.rpm
warning: /tmp/ibm-power-repo-latest.noarch.rpm: Header V4 DSA/SHA1 Signature, ...
Preparing... #####
[100%]
1:ibm-power-repo-3.0.0-19 #####
[100%]
```

Configure the repo:

```
# /opt/ibm/lop/configure
```

Install the tools:

```
# yum install ibm-power-managed-rhel7
```

or

```
# yum install ibm-power-managed-rhel8
```

Now cloud-init can be installed and configured:

1. For Red Hat Enterprise Linux 8 all requisites for cloud-init are part of the Base and AppStream repositories. Install them with:

```
# yum install e2fsprogs iproute net-tools procps python3-configobj
python3-jinja2 python3-jsonpatch python3-jsonschema python3-libselinux
python3-oauthlib python3-policycoreutils python3-prettytable python3-requests
python3-six python3-yaml
```

2. Install the cloud-init version that is shipped with PowerVC. It could be found under `/opt/ibm/powervc/images/cloud-init/rhel` on the PowerVC server:

```
# rpm -Uvh cloud-init-19.1-11.ibm.el8.noarch.rpm
```

3. For every `/etc/sysconfig/network-scripts/ifcfg-*` file remove any MAC or UUID addresses.

4. For Red Hat Enterprise Linux 7 and Red Hat Enterprise Linux 8.0, Red Hat Enterprise Linux 8.1 and Red Hat Enterprise Linux 8.2 no further changes are necessary in the `/etc/cloud/cloud.cfg`

5. For Red Hat Enterprise Linux 8.3 additional configuration is needed:

- a. Add the following data into the `/etc/cloud/cloud.cfg` file as shown in Example 4-4:

Example 4-4 cloud.cfg additions for Red Hat Enterprise Linux 8.3

```
ssh_pwauth: true
manage_resolv_conf: false

datasource_list: [ ConfigDrive, NoCloud, None ]
datasource:
  ConfigDrive:
    dsmode: local
```

Add `resolv_conf` in cloud-init module section after `update_etc_hosts`.

Note: Ensure that the `cloud.cfg` file has the `reset_rmc` module listed. If you do not see the module listed, then make sure you update the OS to the latest version of cloud-init.

In case the `reset_rmc` module is missing even after OS update, make sure you add the module in `cloud_final_modules` section of the `/etc/cloud/cloud.cfg` as follows:

```
- reset_rmc
```

- b. Create a new file `/etc/cloud/ds-identify.cfg` and add following data.

```
policy: search,found=all,maybe=all,notfound=disabled
```

- c. Update the NetworkManager package to the latest provided by Red Hat Enterprise Linux 8.3.

```
# yum upgrade NetworkManager
```

- d. Make sure firewall is configured correctly to open up RMC port 657 for inbound and outbound traffic.

6. PowerVC is responsible for only bringing up the virtual machines. You can leverage cloud-init for specific customization. An example of generating new machine ID is given here. Update the file at /etc/cloud/cloud.cfg. Add the code shown in Example 4-5:

Example 4-5 Generating a new machine ID in cloud-init for Red Hat Enterprise Linux

```
runcmd:  
- rm -f /etc/machine-id  
- systemd-machine-id-setup
```

The information above and additional information about the cloud-init installation for Red Hat Enterprise Linux 7 could be found in IBM Knowledge Center:

https://www.ibm.com/support/knowledgecenter/en/SSXK2N_2.0.0/com.ibm.powervc.standa.rd.help.doc/powervc_install_cloudinit_rhel_hmc.html

Tip: The Red Hat Enterprise Linux update using yum may install a newer version of cloud-init that is not supported and may have issues in a PowerVC environment. Therefore you can exclude cloud-init from the update with:

```
# yum update --exclude=cloud-init
```

Cloud-init preparation for SUSE Linux Enterprise Server

Before you install cloud-init in SUSE Linux Enterprise Server, it is recommended to install the IBM Service and Productivity tools. This is the basis for RMC and with that to get a Health state of OK in PowerVC:

1. Get the latest version of the ibm-power-repo package from the following website:

<https://www.ibm.com/support/pages/service-and-productivity-tools>

2. Install the RPM as in Example 4-6:

Example 4-6 Installation of IBM Service and Productivity tools for SUSE Linux Enterprise Server

Install RPM:

```
# zypper install ibm-power-repo-3.0.0-19.noarch.rpm  
Loading repository data...  
Reading installed packages...  
Resolving package dependencies...
```

The following NEW package is going to be installed:

ibm-power-repo

The following package is not supported by its vendor:

ibm-power-repo

1 new package to install.

Configure the repo:

```
# /opt/ibm/lop/configure
```

Install the tools:

```
# yum install ibm-power-managed-sles12
```

or

```
# yum install ibm-power-managed-sles15
```

Now cloud-init can be installed and configured:

1. For SUSE Linux Enterprise Server12 or SUSE Linux Enterprise Server15 all requisites for cloud-init are part of the Base and Module Public Cloud Base and Update repositories. For SUSE Linux Enterprise Server 15 install them with:

```
# zypper install xfsprogs python3-oauthlib python3-Jinja2 python3-PyYAML
python3-distro python3-jsonpatch python3-jsonschema python3-PrettyTable
python3-requests python3-configobj
```

2. Install the cloud-init version that is shipped with PowerVC that could be found under /opt/ibm/powervc/images/cloud-init/sles on the PowerVC server. For SUSE Linux Enterprise Server15 install:

```
# rpm -Uvh cloud-init-19.1-11.sles15.ibm.noarch.rpm
```

3. PowerVC is responsible for only bringing up the virtual machines. You can leverage cloud-init for specific customization. An example of generating new machine ID is given here. Update the file at /etc/cloud/cloud.cfg. Add the code shown in Example 4-7:

Example 4-7 Generating a new machine ID in cloud-init for SUSE Linux Enterprise Server

```
runcmd:
- rm -f /etc/machine-id
- rm -f /var/lib/dbus/machine-id
- /usr/bin/dbus-uuidgen --ensure
- systemd-machine-id-setup
```

Tip: The SUSE Linux Enterprise Server update using zypper may install a newer version of cloud-init that is not supported and may have issues in a PowerVC environment. Therefore you can exclude cloud-init from the update with:

```
# zypper addlock cloud-init
```

The information above and additional information about the cloud-init installation for SUSE Linux Enterprise Server12 could be found in IBM Knowledge Center:
https://www.ibm.com/support/knowledgecenter/en/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_install_cloudfinit_sles_hmc.html

Cloud-init preparation for Ubuntu

PowerVC only supports Ubuntu 16.04 as guest operating system. Before you install cloud-init in Ubuntu, it is recommended to install the IBM Service and Productivity tools. This is the basis for RMC and with that to get a Health state of OK in PowerVC:

For information how to install the IBM Service and Productivity tools see the following website: <https://www.ibm.com/support/pages/service-and-productivity-tools>

One way of installing the tools is shown in Example 4-8:

Example 4-8 Installation of IBM Service and Productivity tools for Ubuntu

```
# add-apt-repository ppa:ibmpackages/rsct

# apt-get update
```

```
# apt-get install src rsct.core.utils rsct.core rsct.basic
devices.chrp.base.ServiceRM sqlite sg3-utils iprutils lsvpd servicelog ppc64-diag
dynamicrm
```

Now cloud-init can be installed and configured:

1. For Ubuntu 16.04 all requisites for cloud-init are part of the Base and supplemental repositories.

2. Uninstall any older versions of cloud-init:

```
# apt-get remove cloud-init
# apt-get purge cloud-init
```

3. Install the cloud-init version that is shipped with PowerVC that could be found under /opt/ibm/powervc/images/cloud-init/rhel on the PowerVC server:

```
# apt-get install downloaded_path/cloud-init_19.1-3.ibm.ubuntu.deb
```

Important: Make sure you provide the complete Debian package path during installation.

4. The virtual Ethernet adapters are detected as “ibmveth<n>”, but cloud-init expects “eth<n>”. Therefore the naming needs to be changed (according to IBM Knowledge Center):

a. In /lib/udev/rules.d/73-special-net-names.rules, comment out the line containing DRIVERS=="ibmveth".

b. # update-initramfs -u

c. Modify /etc/network/interfaces to change references to ibmvethN to ethN. Remove hwaddr if it is present.

5. Modify the cloud.cfg file available at /etc/cloud/cloud.cfg with the following values:

a. Modify these values, if they exist. Add any of these values that are missing:

```
ssh_password: true
```

```
disable_root: false
```

b. Add these values after preserve_hostname and remove disable_ec2_metadata module:

```
datasource_list: [ ConfigDrive, None ]
```

```
datasource:
```

```
ConfigDrive:
```

```
    dsmode: local
```

6. Enable and start the cloud-init related services by issuing the following commands:

```
# systemctl enable cloud-init-local.service
```

```
# systemctl enable cloud-init.service
```

```
# systemctl enable cloud-config.service
```

```
# systemctl enable cloud-final.service
```

```
# systemctl start cloud-init.service
```

```
# systemctl start cloud-init-local.service
```

```
# systemctl start cloud-config.service
```

```
# systemctl start cloud-final.service
```

Tip: The Ubuntu update using apt-get can install a newer version of cloud-init that is not supported and may have issues in a PowerVC environment. Therefore you can exclude cloud-init from the update with:

```
# apt-mark hold cloud-init
```

The information above and additional information can be found in IBM Knowledge Center: https://www.ibm.com/support/knowledgecenter/en/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_install_cloudinit_ubuntu_hmc.html

Custom scripts with cloud-init for all supported Linux distros

Also in Linux you can also integrate custom scripts into cloud-init. For that use the following directories:

- ▶ `/var/lib/cloud/scripts/per-boot`
Scripts in this directory will be executed in every boot.
- ▶ `/var/lib/cloud/scripts/per-instance`
Scripts in this directory will be executed once after the deployment.
- ▶ `/var/lib/cloud/scripts/per-once`
Scripts in this directory will be executed once after the deployment.

Note: The directories will not be created during the installation of cloud-init. After a reboot the directories will be created. Alternatively you can create the directories with the following commands:

```
# mkdir -p /var/lib/cloud/scripts/per-boot
# mkdir /var/lib/cloud/scripts/per-instance
# mkdir /var/lib/cloud/scripts/per-once
```

Cloud-init preparation IBM i

To install cloud-init on IBM i do the following steps:

1. First install the following requisites that all can be found in the B-GROUP1 ISOs:
 - 5770DG1 with *ALL option
 - 5770SS1 with Option 30 and 33
 - 5770SC1 *BASE and option 1
2. Make sure that the SSHD is running, or start it with:
====> STRTCPSVR *SSHD
3. Install the Open source environment using the IBM i Access Client Solutions:
 - a. Enter the Open Source Package Management as shown in Figure 4-150 on page 221:

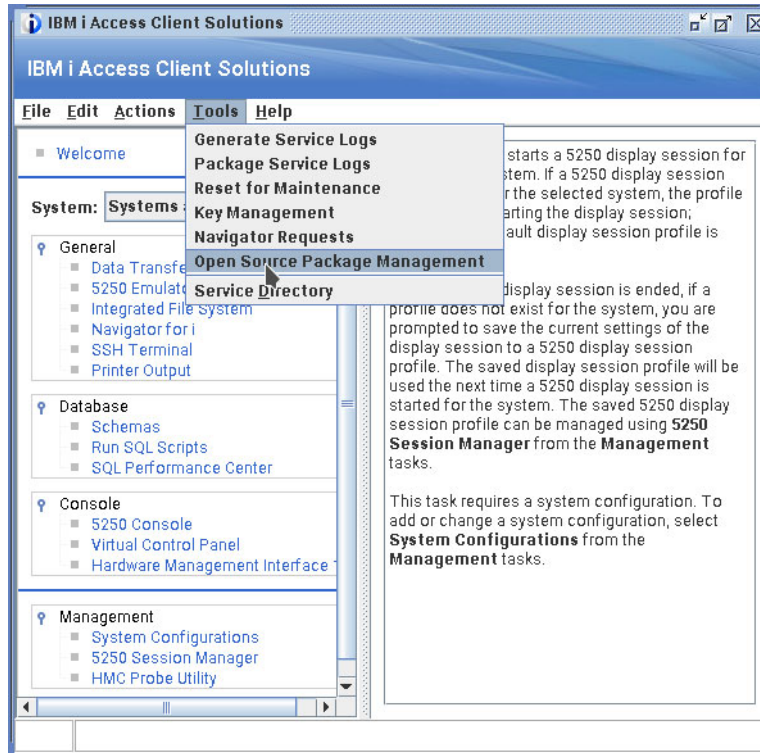


Figure 4-150 Access the IBM i Open Source Package Management

- b. Now select the system and enter the user and password to connect as shown in Figure 4-151.

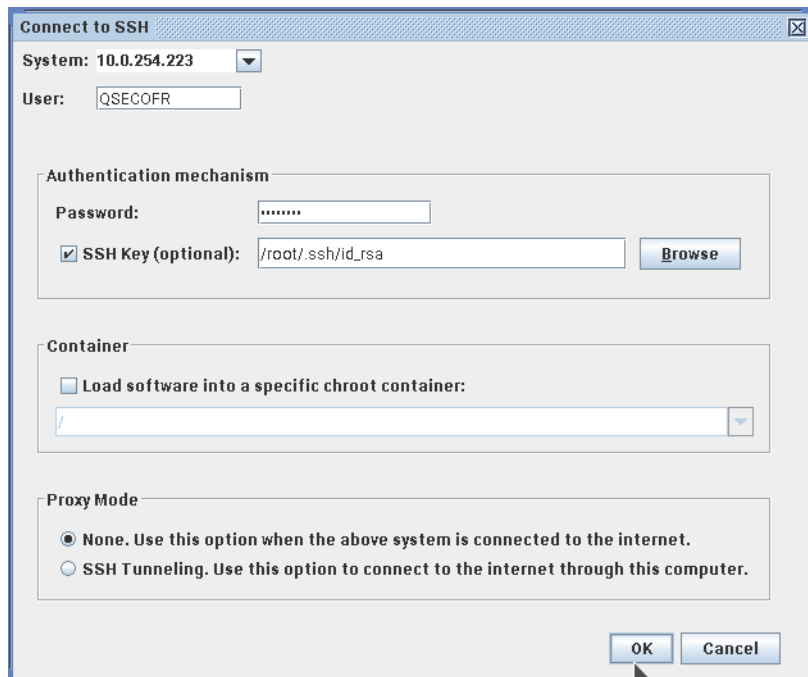


Figure 4-151 Connection information for the IBM i system

- c. After verifying the RSA key fingerprint, you will be asked if you want to install the open source environment. Select **Yes**. Now the open source environment will be installed as shown in Figure 4-152.




Figure 4-152 Installing the IBM i open source environment

4. To install cloud-init open the **Available packages** tab. There select the following packages as shown in Figure 4-153.

- cloud-init
- python2-ibm_db
- python2-six

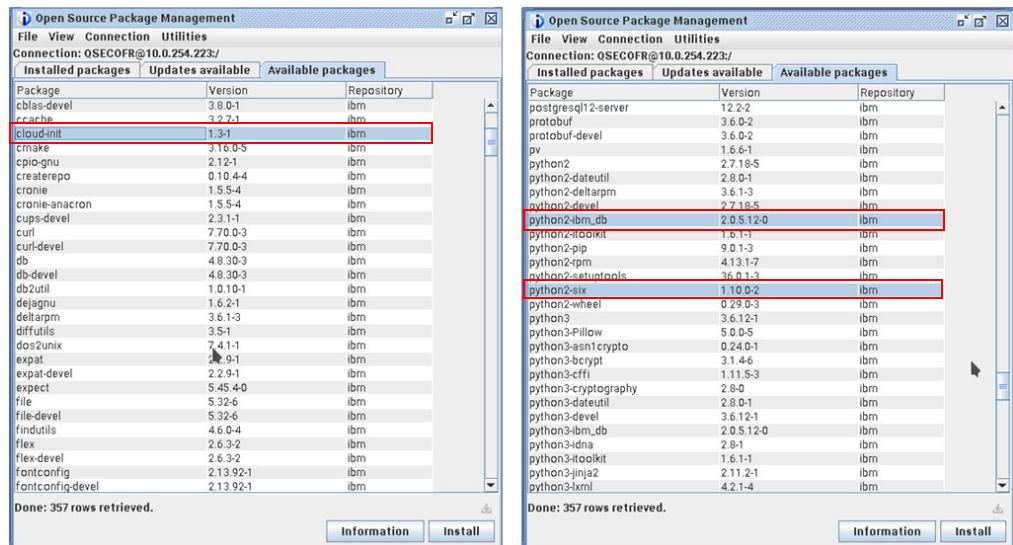


Figure 4-153 IBM i packages to install for cloud-init

After clicking **Install**, the installation starts and you must confirm the selection as shown in Figure 4-154 on page 223.

```

Package Installation
-> Package python2-ibm_db.ppc64 0:2.0.5.12-0 will be installed
-> Package python2-six.noarch 0:1.10.0-2 will be installed
-> Running transaction check
-> Package python3.ppc64 0:3.6.12-1 will be installed
-> Processing Dependency: update-alternatives for package: python3-3.6.12-1.ppc64
-> Processing Dependency: lib:/QOpenSys/pkgs/lib/libreadline.so.8(shr_64.o)(ppc64) for package: python3-3.6.12-1.ppc64
-> Processing Dependency: update-alternatives for package: python3-3.6.12-1.ppc64
-> Package python3-ibm_db.ppc64 0:2.0.5.12-0 will be installed
-> Processing Dependency: python3-six for package: python3-ibm_db-2.0.5.12-0.ppc64
-> Running transaction check
-> Package libreadline8.ppc64 0:8.0-1 will be installed
-> Package python3-six.noarch 0:1.10.0-1 will be installed
-> Package update-alternatives.ppc64 0:1.19.7-1 will be installed
-> Finished Dependency Resolution

Dependencies Resolved

=====
Package      Arch  Version      Repository Size
=====
Installing:
cloud-init   ppc64  1.3-1        ibm        947 k
python2-ibm_db ppc64  2.0.5.12-0  ibm        179 k
python2-six  noarch 1.10.0-2     ibm         16 k
Installing for dependencies:
libreadline8 ppc64  8.0-1       ibm        517 k
python3       ppc64  3.6.12-1    ibm       37 M
python3-ibm_db ppc64  2.0.5.12-0  ibm        189 k
python3-six  noarch 1.10.0-1     ibm         16 k
update-alternatives ppc64  1.19.7-1    ibm         84 k

Transaction Summary
=====
Install 8 Packages

Total download size: 38 M
Installed size: 146 M
Is this ok [y/N]: y

```

Figure 4-154 Confirmation for IBM i package install

Enter **y** and click **Enter**.

Figure 4-155 shows the successful installation.

```

Package Installation
Is this ok [y/N]: y
Downloading Packages:
(1/8): cloud-init-1.3-1.ibm7.2.ppc64.rpm | 947 kB 00:03
(2/8): libreadline8-8.0-1.ibm7.2.ppc64.rpm | 517 kB 00:01
(3/8): python2-ibm_db-2.0.5.12-0.ibm7.2.ppc64.rpm | 179 kB 00:00
(4/8): python2-six-1.10.0-2.ibm7.2.noarch.rpm | 16 kB 00:00
(5/8): python3-3.6.12-1.ibm7.2.ppc64.rpm | 37 MB 01:45
(6/8): python3-ibm_db-2.0.5.12-0.ibm7.2.ppc64.rpm | 189 kB 00:01
(7/8): python3-six-1.10.0-1.ibm7.2.noarch.rpm | 16 kB 00:00
(8/8): update-alternatives-1.19.7-1.ibm7.2.ppc64.rpm | 84 kB 00:00

Total 338 kB/s | 38 MB 01:56
Running Transaction Check
Running Transaction Test
Transaction Test Succeeded
Running Transaction
Installing : libreadline8-8.0-1.ppc64 1/8
Installing : python2-six-1.10.0-2.noarch 2/8
Installing : update-alternatives-1.19.7-1.ppc64 3/8
Installing : python3-3.6.12-1.ppc64 4/8
update-alternatives: using /QOpenSys/pkgs/bin/python3.6 to provide /QOpenSys/pkgs/bin/python3 (python3) in auto mode
update-alternatives: using /QOpenSys/pkgs/bin/python3.6 to provide /QOpenSys/pkgs/bin/python (python) in auto mode
Installing : python3-six-1.10.0-1.noarch 5/8
Installing : python3-ibm_db-2.0.5.12-0.ppc64 6/8
Installing : cloud-init-1.3-1.ppc64 7/8
Installing : python2-ibm_db-2.0.5.12-0.ppc64 8/8

Installed:
cloud-init.ppc64 0:1.3-1 python2-ibm_db.ppc64 0:2.0.5.12-0
python2-six.noarch 0:1.10.0-2

Dependency Installed:
libreadline8.ppc64 0:8.0-1 python3.ppc64 0:3.6.12-1
python3-ibm_db.ppc64 0:2.0.5.12-0 python3-six.noarch 0:1.10.0-1
update-alternatives.ppc64 0:1.19.7-1

Complete!

```

Figure 4-155 IBM i - cloud-init installed

5. Now activate cloud-init with the following command:

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

6. After the activation, shutdown the system. With that the VM is ready to be captured as an image.

Information on how to install cloud-init can be also found in IBM Knowledge Center:

https://www.ibm.com/support/knowledgecenter/en/SSXK2N_2.0.0/com.ibm.powervc.standa rd.help.doc/powervc_install_cloudinit_ibmi_hmc.html

4.16.2 Capturing a virtual machine image

To capture a virtual machines means to create an image of virtual machine. After you prepared an image VM for capture, log on to the IBM PowerVC GUI.

1. Go to the VM list under Virtual Machines. Select the VM that you want to capture and click **Capture** on the top menu as shown in Figure 4-156.

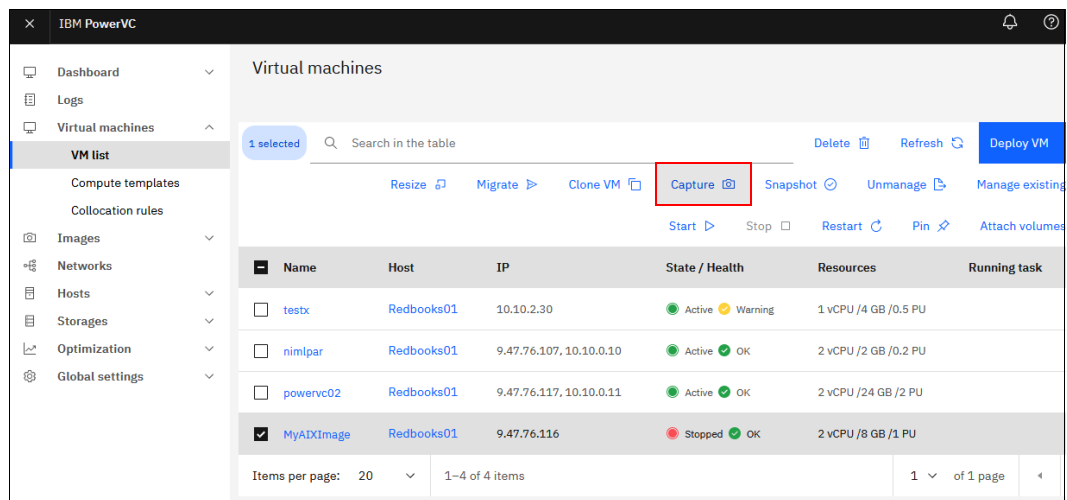


Figure 4-156 Capture an image

Note: It is supported to capture a VM while it is running. In this case make sure to have consistent data. To be sure that the image is consistent do the capture from an VM that is stopped.

On the next screen you see a warning message saying that the VM requires preparation. This is simply a reminder if you have installed cloud-init. If you did, select I have prepared the VM.

2. Now enter the name for the image. The VM shown in Figure 4-157 on page 225 has two volumes, one bootable with the operating system and a volume called monitoring (as an example). You can deselect the additional volumes if you want, but in this example it should be part of the image.

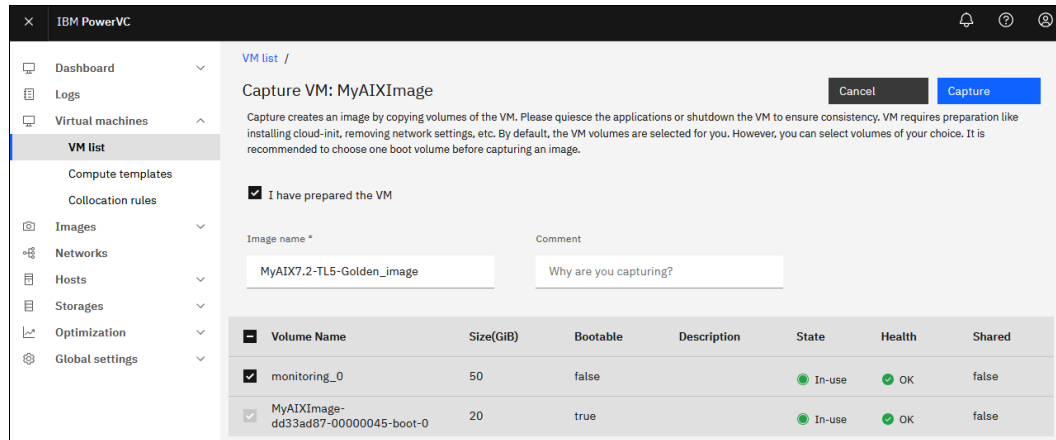


Figure 4-157 Details of the image to capture

3. Now click **Capture**.

After a while the images shows up in the Image list under Images as shown in Figure 4-158.

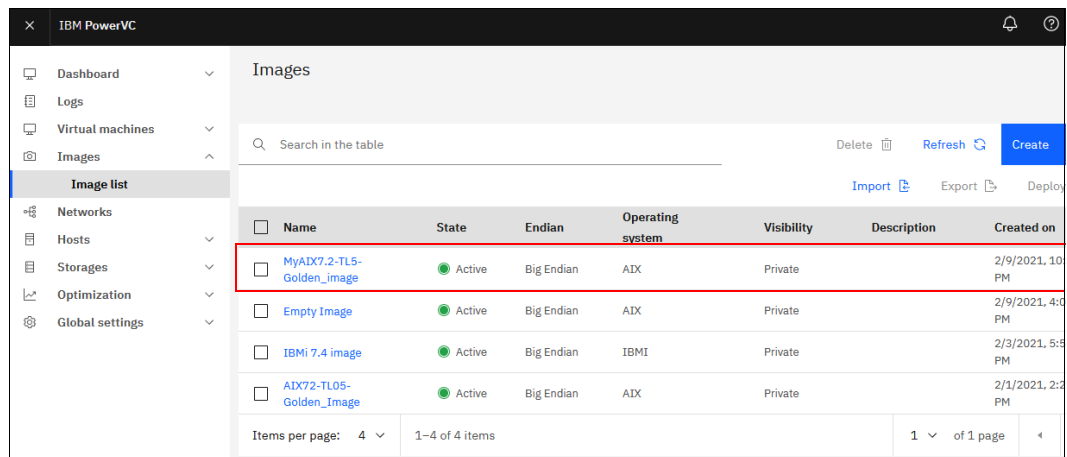


Figure 4-158 Captured image is available

4. When you click on the image name you can see the image details as shown in Figure 4-159 on page 226, the volumes related to the image and also virtual machines that have been deployed from this image.

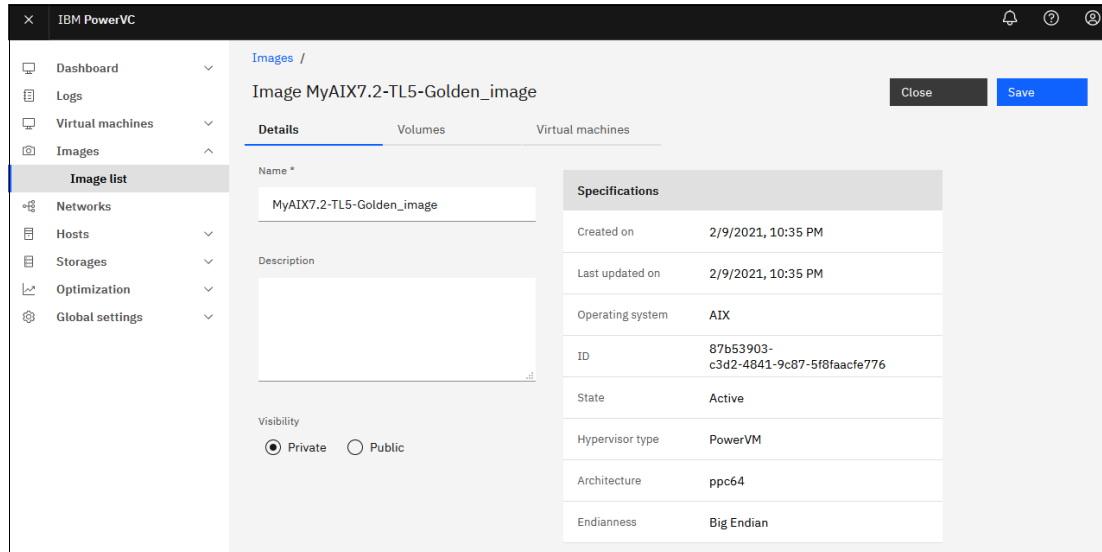


Figure 4-159 Image details

In the image details you can change the Visibility of the image from Private to Public. A public image is visible and can be used in all PowerVC projects, while a private image is only visible in the current project. For more information on project see section 2.8.2, “Projects and role management planning” on page 62.

Here you can also enter a description of the image. On the right you find Specifications of the image that are explained in Table 4-7:

Table 4-7 Description of the Specifications

Field	Description
Created on	Date and time of the image capture
Last updated	Date and time of the last modification of the image
Operating system	The operating system on the image
ID	Unique identifier number for the resource This ID is useful, when you work on the CLI or the REST-API.
State	State of the image
Hypervisor type	Hypervisor for the image
Architecture	Architecture of the image
Endianness	Big Endian or Little Endian

4.16.3 Import an image from OVA or raw image

PowerVC can import existing images if they are in a raw, Open Virtualization Archive (OVA) or OVA.gz format. OVA would be the desired format to exchange images with the IBM Public Cloud or other cloud providers. You can import OVAs from an Power Virtual Server (PowerVS) environment in the IBM Public cloud. When importing an OVA, the archive file is automatically extracted to a temporary staging directory.

IBM provides prepared raw images for AIX 7.1 and AIX 7.2. These images contain cloud-init and can be used for deployment without any further changes. A client with a software maintenance agreement (SWMA) may download these images from the ESS website <https://www.ibm.com/servers/eserver/ess/OpenServlet.wss>.

Some things to consider, before doing an import:

- ▶ As the image will be imported into volumes of the backend storage, PowerVC needs a Fibre Channel adapter inside the PowerVC virtual machine. This adapter can also be a virtualized adapter using NPIV.
- ▶ Images tend to be very big, therefore consider the available disk space in PowerVC and the directories where the image handling will take place.

Use the **powervc-image config** command to change the location of the staging directory, if you plan to import big images as shown in the following Example 4-9.

Example 4-9 Change the PowerVC staging directory for image imports

```
# powervc-image config
List powervc-image configuration properties:
+-----+
| Configuration Property | Value |
+-----+
| default_ova_export_dir | /var/opt/ibm/powervc/ova |
| enforce_unique_image_names | True |
| enforce_unique_volume_names | False |
| polling_timeout | 86400 |
| powervc_image_log | /opt/ibm/powervc/log/powervc-image.log |
| temp_staging_parent_dir | /var/opt/ibm/powervc/imgstaging |
+-----+

# powervc-image config --set temp_staging_parent_dir --value
/powervc-backups/staging
The configuration value has been set.

# powervc-image config
List powervc-image configuration properties:
+-----+
| Configuration Property | Value |
+-----+
| default_ova_export_dir | /var/opt/ibm/powervc/ova |
| enforce_unique_image_names | True |
| enforce_unique_volume_names | False |
| polling_timeout | 86400 |
| powervc_image_log | /opt/ibm/powervc/log/powervc-image.log |
| temp_staging_parent_dir | /powervc-backups/staging |
+-----+
```

Note: Alternatively you can also soft link the default directory to another directory.

Before you can start with the import, you need the ID of the storage template that should be used. The storage template defines the backend storage and the storage attributes to be used. To find out the ID you can enter the desired storage template by clicking on the name of the template in Templates under Storage, as shown in Figure 4-160.

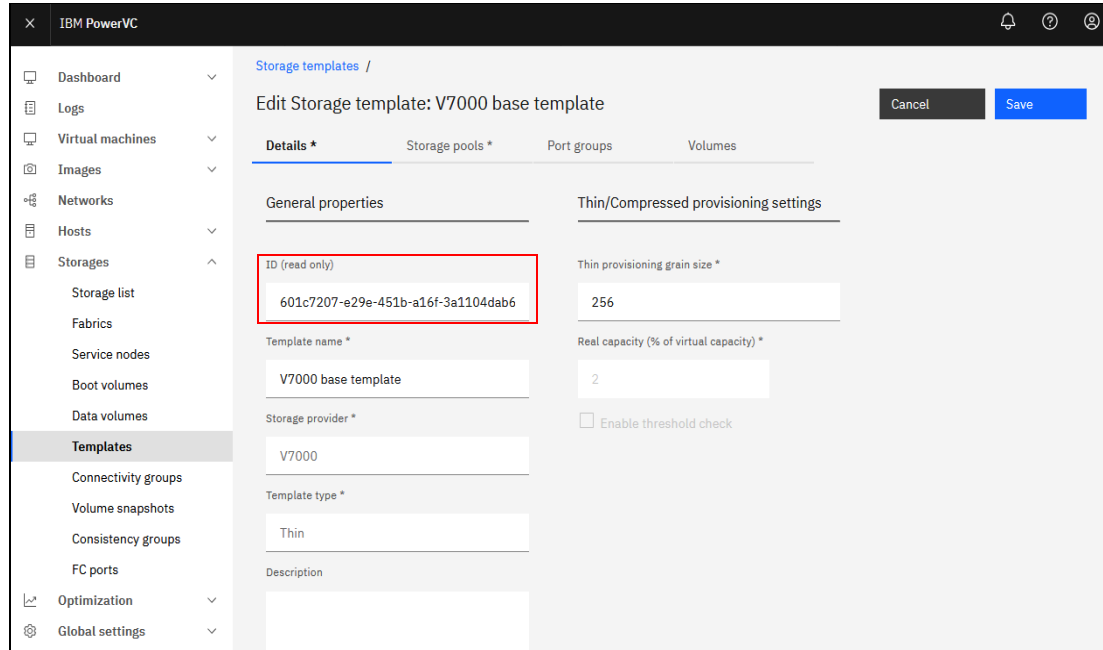


Figure 4-160 Find the ID of a storage template

As you are working on the CLI for importing the image, it may be convenient to use the CLI to get the ID as shown in Example 4-10.

Example 4-10 Find the ID of a storage template

```
# cinder type-list
```

ID	Name	Description	Is_Public
14b6ee5e-04b2-4647-bcee-28ea37d63322	V7000 base template 2	-	True
601c7207-e29e-451b-a16f-3a1104dab602	V7000 base template	-	True

Hint: To work with the CLI, the PowerVC environment needs to be defined. For more information see section “Useful OpenStack commands” on page 318.

Note: The import of images is not supported into an Shared Storage Pool (SSP).

For more information on storage templates see section 4.10, “Storage template setup” on page 140.

1. To import an OVA image, do the following as shown in Example 4-11.

Example 4-11 Import an AIX virtual machine image from ESS

```
# powervc-image import -n AIX72-TL05-VM_image -p
/home/stephen/AIX72-TL05-VM_image.ova.gz -t
601c7207-e29e-451b-a16f-3a1104dab602
```

```

Created temporary staging directory /powervc-backups/staging/tmpxqk9f71o
Compressed archive detected.
Extracting package contents to temporary directory.
Parsing the OVF descriptor.
Register temporary file-copy volume driver.
Registered temporary driver PVC-Lite-File_tmpxqk9f71o servicing location
/powervc-backups/staging/tmpxqk9f71o
Begin processing image volume
'AIX_v7.2_Virt_Mach_7200-05-01-2038_112020_LCD8236508'.
Create temporary copy-source volume for staging file
'AIX_v7.2_Virt_Mach_7200-05-01-2038_112020_LCD8236508'....Done creating.
Migrate volume data for 'AIX_v7.2_Virt_Mach_7200-05-01-2038_112020_LCD8236508'
from 'File-system' to the target storage template 'V7000 base template'.
The size to copy is 20.00 GiB.
Creating and attaching target volume.....

Copying AIX_v7.2_Virt_Mach_7200-05-01-2038_112020_LCD8236508 [100%] Rate:
159.29 MiB-per-S, ETA: 0:00:00 [H:MM:SS]
Detaching volume and finalizing metadata...
Copy complete after waiting 0:02:08 [H:MM:SS]

GiBs remaining to copy for image: 0
Creating image resource 'AIX72-TL05-VM_image'.
/usr/bin/powervc-volume-image-import --name "AIX72-TL05-VM_image" --os-distro
aix --endianness big --boot-set b3c66e6b-84e3-467f-8c81-b1fb329772f4
Created new image with ID 955c10d5-09a3-4eb3-bb97-26368a200a76.
De-register the temporary file driver 'PVC-Lite-File_tmpxqk9f71o'.
Cleaning up the temporary staging directory...
Time spent: 0:08:32 [H:MM:SS]
Successfully created new image 'AIX72-TL05-VM_image' with ID
'955c10d5-09a3-4eb3-bb97-26368a200a76'.

```

2. To import an raw image, you need to specify some more information, as just the raw image has no metadata about the operating system and the architecture. In this case the import command must include os-type and architecture like this:

```

# powervc-image import -n AIX72-TL05-VM_image -p
/home/stephen/AIX_v7.2_Virt_Mach_7200-05-01-2038_112020_LCD8236508 -m
os-type=aix architecture=ppc64 -t 601c7207-e29e-451b-a16f-3a1104dab602

```

When the import has completed you find the image in the GUI as shown in Figure 4-161.

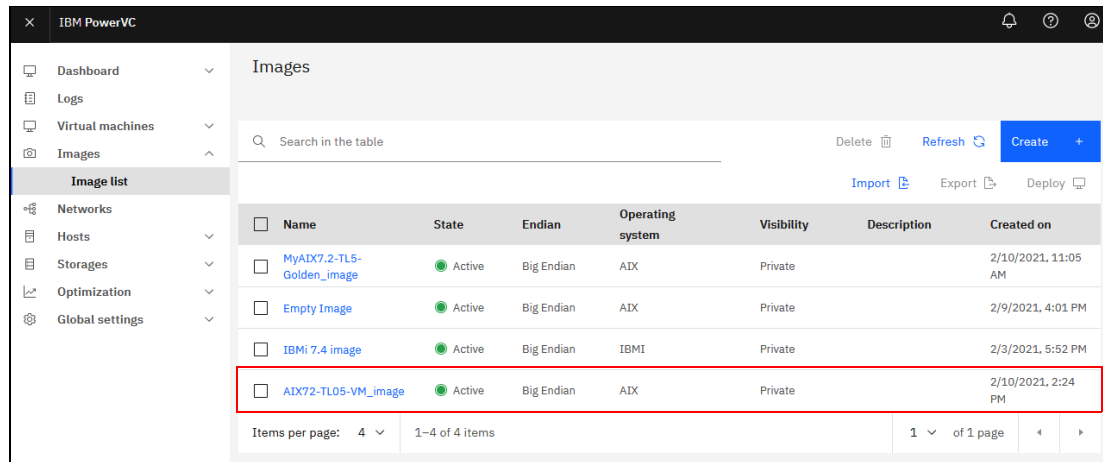


Figure 4-161 Imported image in the GUI of PowerVC

Alternative way of importing an image into PowerVC

If your PowerVC system has no Fibre Channel adapters, or you want to import an image into a Shared Storage Pool (SSP), you can use the alternative way as follows:

1. If you have an OVA, extract the raw image using tar as in the following Example 4-12 for the AIX virtual machine image downloaded from ESS.

Example 4-12 Extract the raw image from an OVA

```
# tar -zxvf AIX72-TL05-VM_image.ova.gz
AIX72-TL05-VM_image.ovf
AIX72-TL05-VM_image.meta
AIX_v7.2_Virt_Mach_7200-05-01-2038_112020_LCD8236508

# ls -al
total 20971528
drwxr-xr-x 2 root root 129 Feb 10 07:30 .
drwx----- 4 stephen stephen 197 Feb 10 07:29 ..
-rw-r----- 1 root root 117 Feb 10 05:52 AIX72-TL05-VM_image.meta
-rw-r----- 1 root root 1727 Feb 10 05:52 AIX72-TL05-VM_image.ovf
-rw-r----- 1 cinder cinder 21474836480 Feb 10 05:52
AIX_v7.2_Virt_Mach_7200-05-01-2038_112020_LCD8
```

AIX_v7.2_Virt_Mach_7200-05-01-2038_112020_LCD8236508 is the raw image file

2. Create a volume that is big enough for the image. In this case $21.474.836.480 / 1.024^3 = 20$ GB.
3. Attach the volume to any running system with the **dd** command available.
4. Copy the raw image onto the volume. In this case on a Linux system to a disk named `/dev/sdb` like this:

```
# dd if=AIX_v7.2_Virt_Mach_7200-05-01-2038_112020_LCD8 of=/dev/sdb bs=4M conv=fsync
```
5. Detach the volume from the VM
6. Manage the volume in PowerVC, if you did not create it with PowerVC (in step 2).

- Define the volume as an image (discussed in 4.16.4, “Define an image with existing volumes” on page 231).

4.16.4 Define an image with existing volumes

If you imported an image using the alternative way as described in “Alternative way of importing an image into PowerVC” on page 230 or if you have the image volumes from an previous installation of PowerVC you can use these volumes to define an image. Follow these steps:

- As it is not possible to import (manage existing) boot volumes, make the volume(s) available in PowerVC as data volume(s). If the desired volume does not exist in the list of data volumes, import it by using the **Manage existing** button in the Data volumes list under Storages.
- Now select the volumes that should be part of the image and click **Create** in Data volumes under Storages as shown in Figure 4-162.

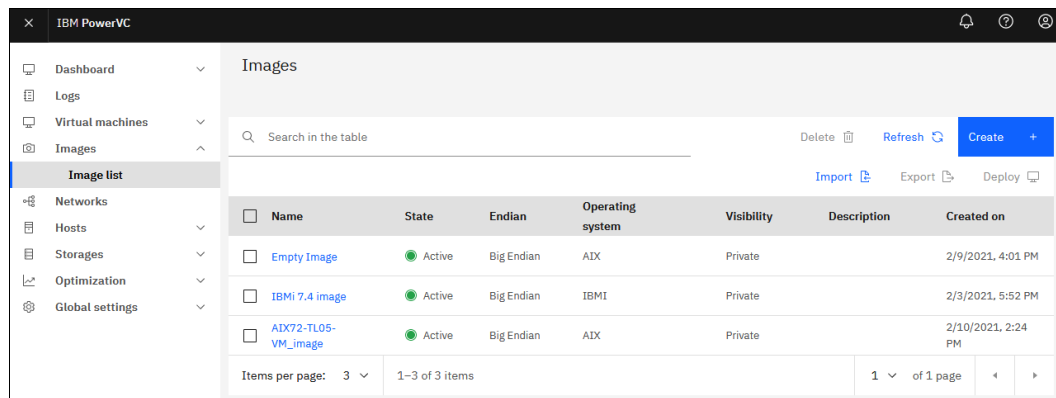


Figure 4-162 Define an image based on existing volumes

- On the next screen on the details tab enter an image name, select the operating system and the correct Endianness. You may also want to make the image public for all projects by selecting Shared across projects as shown in Figure 4-163.

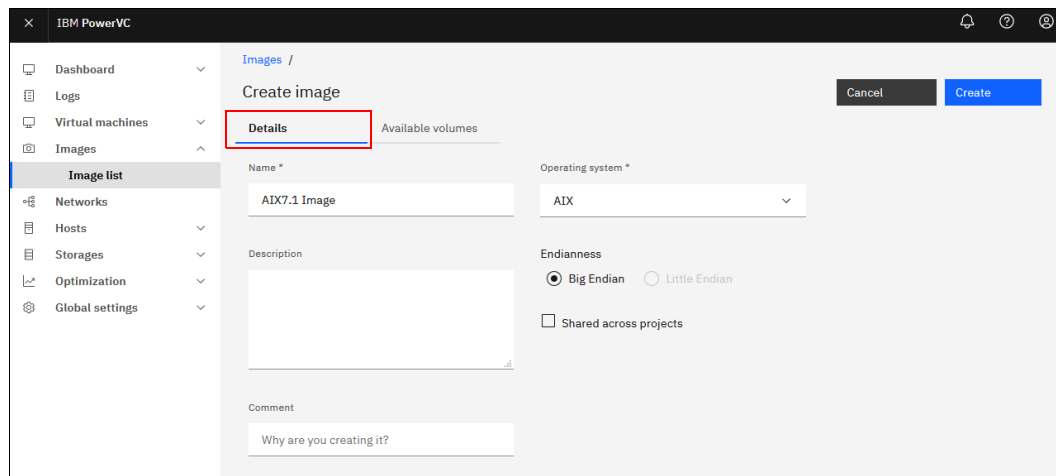


Figure 4-163 Details for the image to create

- Switch to the Available volumes tab and select all volumes that should be part of the image. This example shows the import of an previously created AIX 7.1 image. Don't forget to scroll to the right and define the boot order. Assign 0 to the first boot disk and ascending numbers if you have more than one boot disk as shown in Figure 4-164.

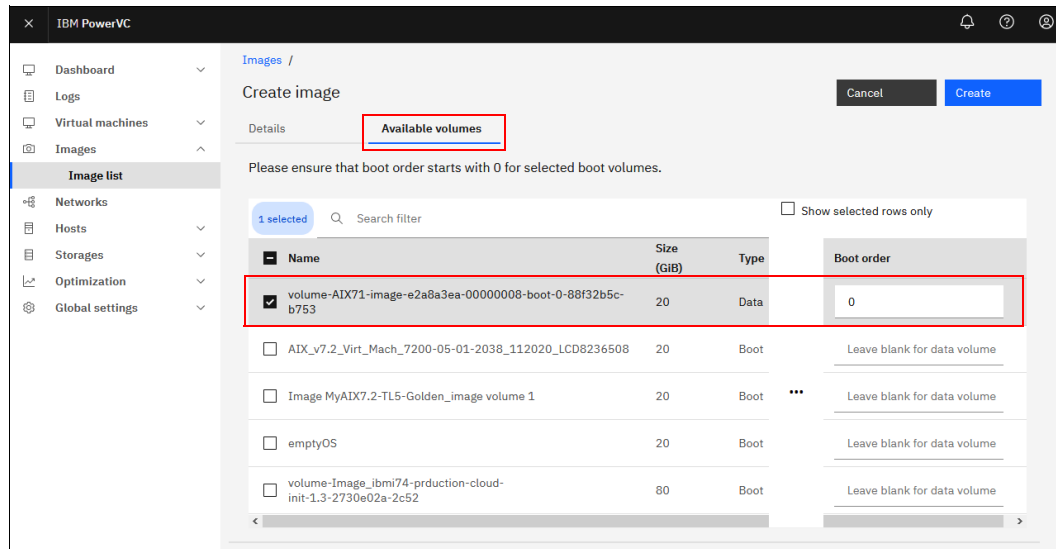


Figure 4-164 Volumes for the image to be defined

- Now the image is ready for use as shown in Figure 4-165

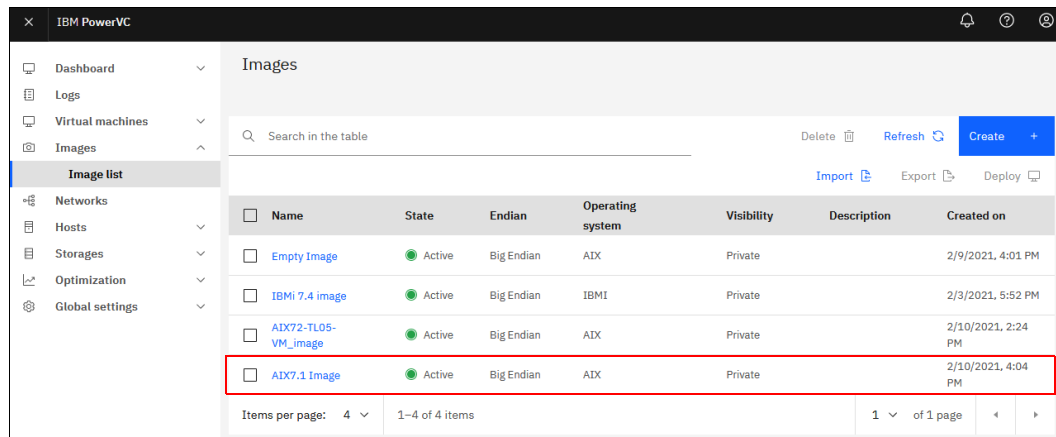


Figure 4-165 Created image ready for use

Note: It is also supported to create an image without any volumes. Use such an image for deployments of a new VM with an existing boot volume that is managed by PowerVC. The existing boot volume can be attached in the deployment details under existing volumes.

4.16.5 Export images into OVA

The export of an image into an Open Virtualization Archive (OVA) can be done using the command line. An OVA may be used to deploy a VM in the IBM Public Cloud (PowerVS) or at other cloud providers.

Images tend to be very big, therefore consider the available disk space in PowerVC before you start the export. Therefore use the `powervc-image config` command to find the default export location and change it if needed as shown in Example 4-13.

Example 4-13 Change the OVA export directory

```
# powervc-image config
List powervc-image configuration properties:
+-----+
| Configuration Property | Value |
+-----+
| default_ova_export_dir | /var/opt/ibm/powervc/ova |
| enforce_unique_image_names | True |
| enforce_unique_volume_names | False |
| polling_timeout | 86400 |
| powervc_image_log | /opt/ibm/powervc/log/powervc-image.log |
| temp_staging_parent_dir | /powervc-backups/staging |
+-----+

# powervc-image config --set default_ova_export_dir --value /powervc-backups/ova
The configuration value has been set.

# powervc-image config
List powervc-image configuration properties:
+-----+
| Configuration Property | Value |
+-----+
| default_ova_export_dir | /powervc-backups/ova |
| enforce_unique_image_names | True |
| enforce_unique_volume_names | False |
| polling_timeout | 86400 |
| powervc_image_log | /opt/ibm/powervc/log/powervc-image.log |
| temp_staging_parent_dir | /powervc-backups/staging |
+-----+
```

Note: Alternatively you can also soft link the default directory to another directory.

Although you can specify a different directory for the OVA in the `export` command, the staging area is still needed. Therefore make sure that the staging directory has at least 105% of the image size.

To export an image, the name of the image is sufficient, you do not need the ID. As you are working on the CLI for exporting the image, it may be convenient to use the CLI to get a list of the images as shown in Example 4-14.

Example 4-14 Find the ID of the image to be exported

```
# powervc-image list
+-----+-----+-----+-----+-----+-----+-----+-----+
| Name | ID | Status | Volumes | Size | ... | Exportable |
+-----+-----+-----+-----+-----+-----+-----+-----+
| AIX72-TL05-VM_image | 2dfb40e9-40d7-4099-9ffd-299390eab6df | active | 1 | 20 | ... | True |
| Empty Image | 895193f6-cd97-413a-9511-16e0368bdb23 | active | 1 | 20 | ... | True |
| IBMi 7.4 image | e42ec1a5-9015-4476-871c-04ae48628222 | active | 1 | 80 | ... | True |
| MyAIX7.2-TL5-Golden_image | 4f2e6d68-9efe-4389-b90f-3c60a869dbf3 | active | 1 | 20 | ... | True |
+-----+-----+-----+-----+-----+-----+-----+-----+
4 managed images returned.
```

Hint: To work with the CLI, the PowerVC environment needs to be defined. For more information see section “Useful OpenStack commands” on page 318.

With `powervc-image list` you see also in the last column if the image is exportable. If the image is stored on an Shared Storage Pool (SSP), the value will show False.

Now export the an image as shown in Example 4-15.

Example 4-15 Export an image into OVA

```
# powervc-image export -i "AIX72-TL05-VM_image" -c
Created temporary staging directory /powervc-backups/staging/tmp2n_qj9e
Found image with ID '2dfb40e9-40d7-4099-9ffd-299390eab6df' and name
'AIX72-TL05-VM_image'.
The export directory and the staging directory are on the same file system. Double
image size space is required.
Using the image name 'AIX72-TL05-VM_image' for the default OVA name.
Register temporary file-copy volume driver.
Registered temporary driver PVC-Lite-File_tmp2n_qj9e servicing location
/powervc-backups/staging/tmp2n_qj9e
Cloning 'AIX_v7.2_Virt_Mach_7200-05-01-2038_112020_LCD8236508' into temporary
volume 'AIX_v7.2_Virt_Mach_7200-05-01-2038_112020_LCD8236508_tmp2n_qj9e'.
The size to clone is 20 GiBs...Done cloning.
Migrate volume data for
'AIX_v7.2_Virt_Mach_7200-05-01-2038_112020_LCD8236508_tmp2n_qj9e' from 'V7000' to
the target storage template 'PVC-Lite-File_tmp2n_qj9e base template'.
Attaching volume.....

Copying AIX_v7.2_Virt_Mach_7200-05-01-2038_112020_LCD8236508 [100%] Rate: 117.57
MiB-per-S, ETA: 0:00:00 [H:MM:SS]
Detaching volume and finalizing metadata...
Copy complete after waiting 0:02:54 [H:MM:SS]

GiBs remaining to copy for image: 0
Creating image package with 1 volumes.
Creating OVF: /powervc-backups/staging/tmp2n_qj9e/AIX72-TL05-VM_image.ovf
Creation of OVF completed.
Adding OVF to OVA /powervc-backups/ova/AIX72-TL05-VM_image.ova
Adding volume 'AIX_v7.2_Virt_Mach_7200-05-01-2038_112020_LCD8236508' to OVA.
Cleaning up Lite-Volume export resources...
De-register the temporary file driver 'PVC-Lite-File_tmp2n_qj9e'.
Cleaning up the temporary staging directory...
Compressing OVA.
Exported OVA /powervc-backups/ova/AIX72-TL05-VM_image.ova.gz size: 6.33 GiB
Time spent: 0:24:12 [H:MM:SS]
Successfully finished creating image package
/powervc-backups/ova/AIX72-TL05-VM_image.ova.gz
```

Hints:

With the `-p` flag you can specify a different directory where the OVA should be written.

With the `-c` flag the OVA will be compressed using GZIP.

The export of an image into an OVA does not delete the image in PowerVC. The image is still usable for deployments.

4.16.6 Import and export with backup service

If you want to move an image from a PowerVC managed storage provider to another storage provider managed by the same PowerVC, export it first to the backup service. After that import it back into the destination storage provider.

Another usecase for the backup services would be to archive images no longer needed into an inexpensive storage.

Export into the backup service

Also for the backup service the images tend to consume a lot of storage. Therefore first make sure you use a directory that has enough free space. You can change the directory as shown in Example 4-16.

Example 4-16 Storage path

```
# powervc-config storage backup-service show
+-----+
| Enabled backup service type: posix |
+-----+
| Property | Value |
+-----+
| "backup_compression_algorithm" | "none" |
| "backup_driver" | "cinder.backup.drivers.posix.PosixBackupDriver" |
| "backup_file_size" | "268435456" |
| "backup_posix_path" | "/var/opt/ibm/powervc/image-backups" |
| "backup_sha_block_size_bytes" | "<default>" |
+-----+

# powervc-config storage backup-service posix --posix-path /powervc-backups/image-backups
--restart
```

Take Note: This command will restart the cinder service. Therefore execute that command when no storage related tasks are running.

The Posix backup service configuration was updated.
Restart of cinder services requested.

```
+-----+
| Property | Value |
+-----+
| "backup_compression_algorithm" | "none" |
| "backup_driver" | "cinder.backup.drivers.posix.PosixBackupDriver" |
| "backup_file_size" | "268435456" |
| "backup_posix_path" | "/powervc-backups/image-backups" |
| "backup_sha_block_size_bytes" | "32768" |
+-----+
```

Note: Make sure that the directory used for the image backups is owned by cinder.cinder.

1. To export an image into the backup service, first select the image to be exported in the Image list under Images and click **Export** as shown in Figure 4-166.

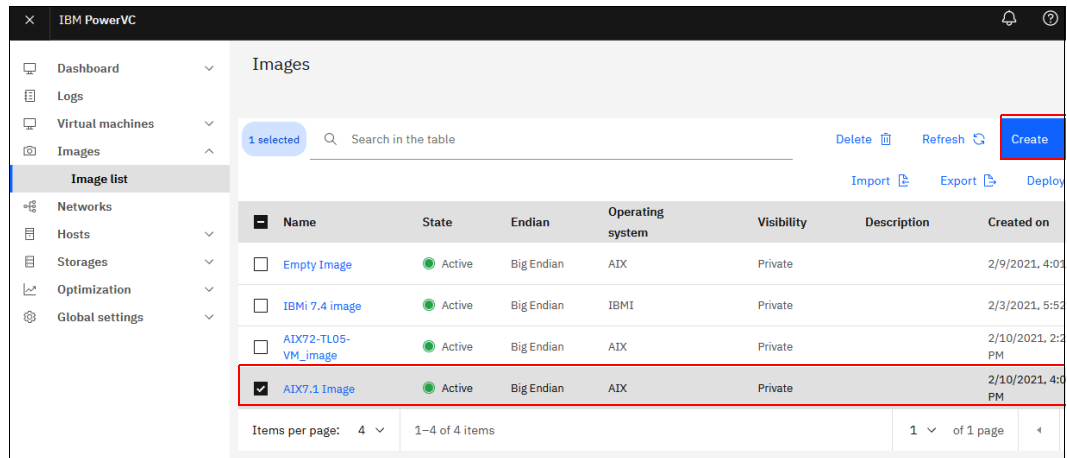


Figure 4-166 Select image to be exported

2. In the next screen (Figure 4-167), you may want to change the name to a name as it will show up in the backup service or add some additional description.

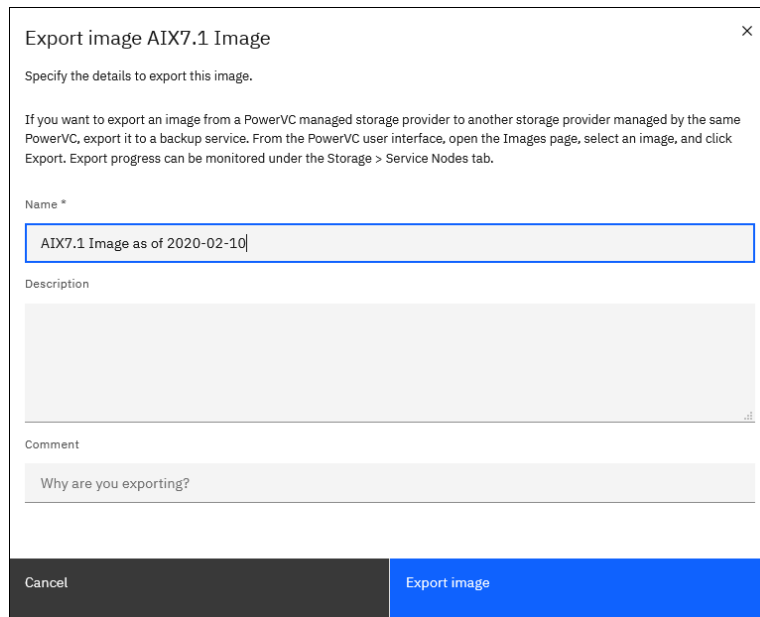


Figure 4-167 Name for image to be exported

The backups service is represented in PowerVC as a service node that can be found under Storages as shown in Figure 4-168 on page 237.

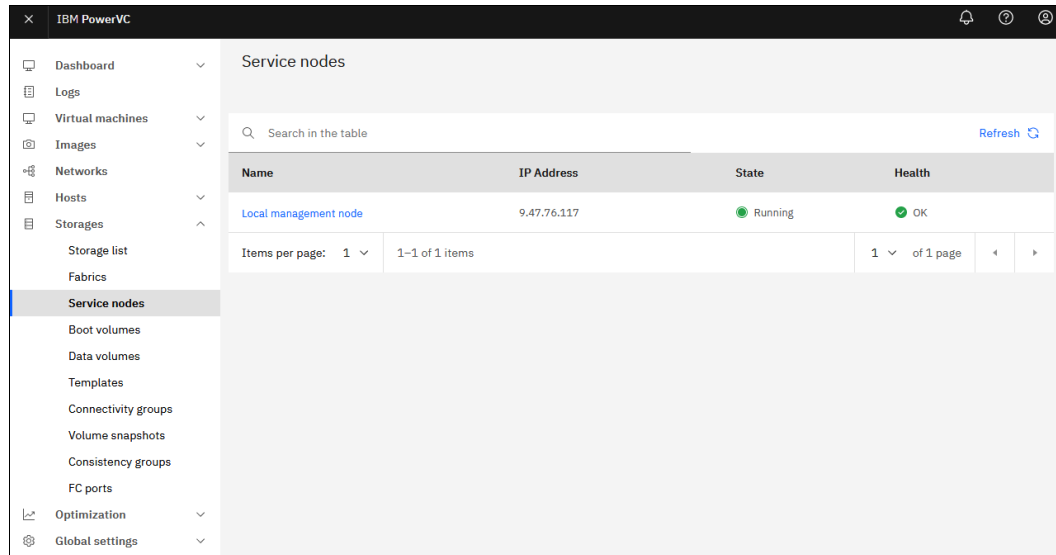


Figure 4-168 Service nodes

3. Then click on Local management node as shown in Figure 4-169.

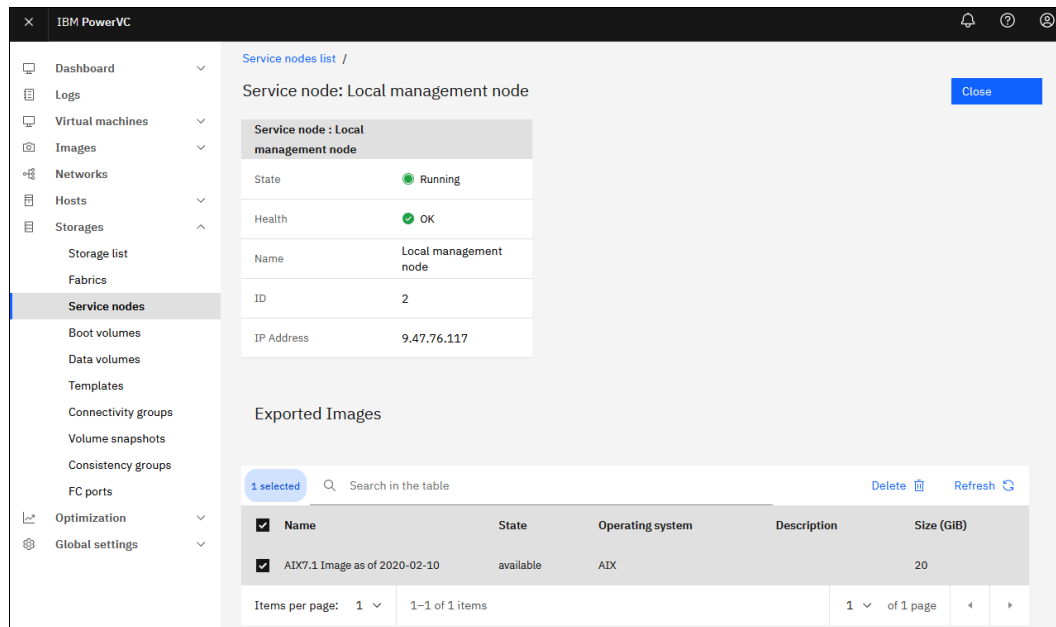


Figure 4-169 Exported images in the service node

4. You can see the images that have been exported. If you want to delete an image from the backup services, this is the place to go. Select it and click **Delete**.

Notes:

The images stored in the backup services are not compressed and take about the same amount of space as the original.

An exported image will not be deleted in the image list. If you do not need the image active and want to delete it, follow the steps in 4.16.8, “Deleting images” on page 246.

Import an image from the backup service

The following steps are used to import an image from the backup service.

1. To import an image from the backup service enter the Image list under Images and click **Import** as shown in Figure 4-170.

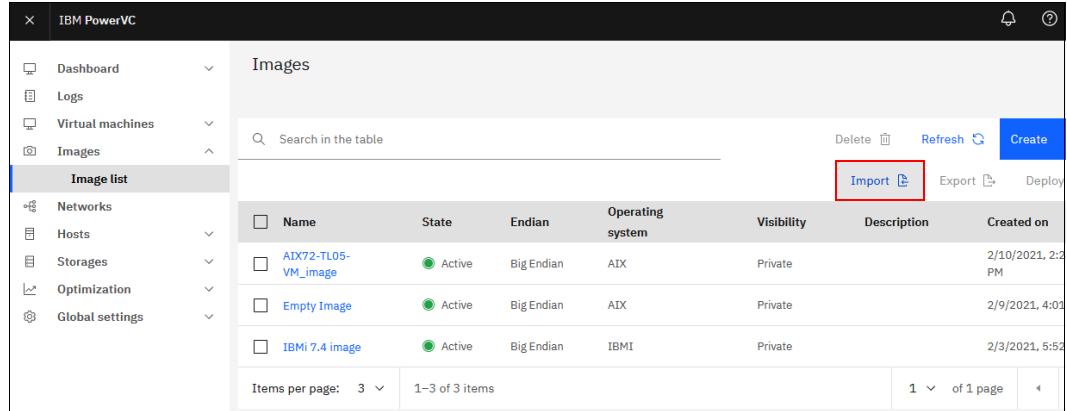


Figure 4-170 Import an image from backup services

2. In the next window (Figure 4-171) you can select the image you want to restore and click **Import**.

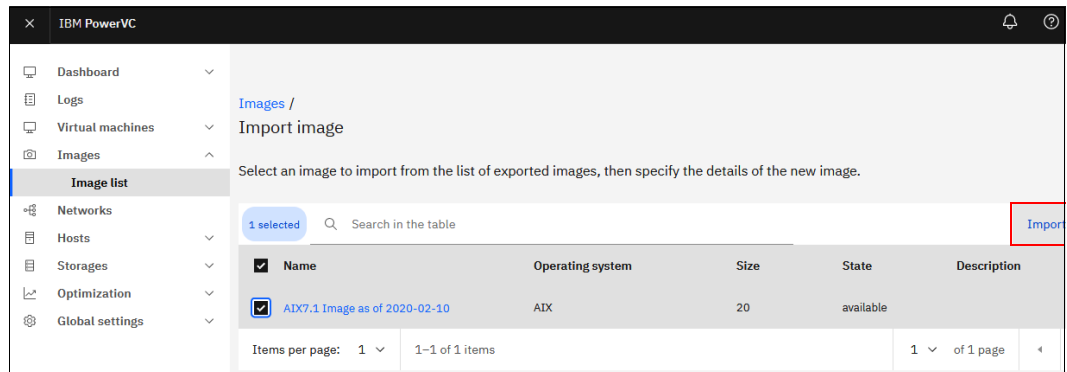


Figure 4-171 Select image to be imported from backup services

3. After that you can change the name of the image, as it should be made available in the image list and select the Storage template where the image should be imported, as shown in Figure 4-172 on page 239.

Figure 4-172 Details of the image import from backup services

4. After you clicked **Import image** the import starts. Figure 4-173 shows the image appear in the state of importing. After some time this status changes to Active, meaning the image restore has finished and it is active to be used.

Name	State	Endian	Operating system	Visibility	Description	Created on
AIX72-TL05-VM_image	Active	Big Endian	AIX	Private		2/10/2021, 2:24 PM
Empty Image	Active	Big Endian	AIX	Private		2/9/2021, 4:01 PM
IBMi 7.4 image	Active	Big Endian	IBMI	Private		2/3/2021, 5:52 PM
AIX7.1 Image as of 2020-02-10	Importing	Big Endian	AIX	Private	Imported image	2/10/2021, 6:03 PM

Figure 4-173 Restore of image is running

Note: After the restore, the image stays in backup services and will not be deleted. If you want to delete the image enter the service node as shown in Figure 4-169 to delete the image.

4.16.7 Deploying a new virtual machine

In the previous section we showed how to create or import an image. The reason for having images in PowerVC is to create new virtual machines with the images created. A resource usage definition plus an image is the basis for a new VM as shown in this section.

The deployment of a new VM starts in the Image list under Images.

1. There select the image with the operating system that you want to use and click **Deploy** as shown in Figure 4-174 on page 240.

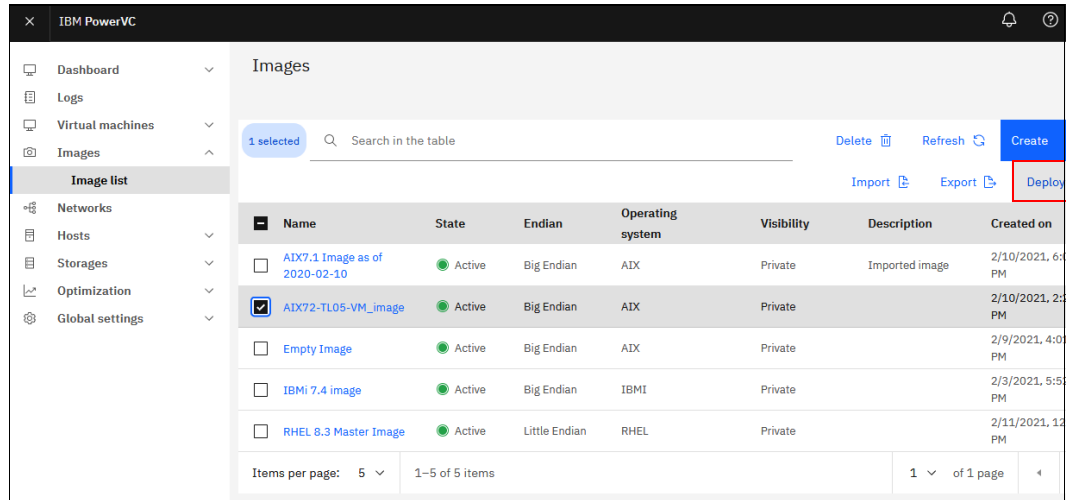


Figure 4-174 Deploy the selected image

- Now define the resources that you want to use and how the new VM should be connected to the storage as shown in Figure 4-175.

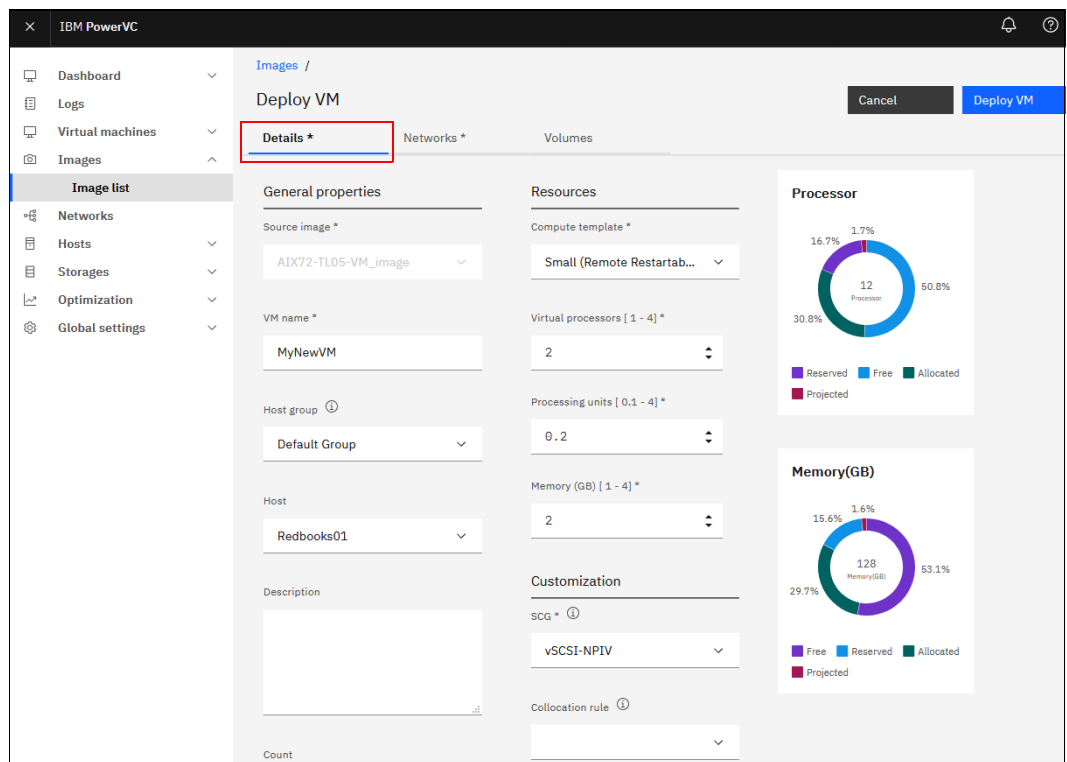


Figure 4-175 Image deployment details

The following attributes of the new VM can be set:

- General properties

- VM name

The VM name will be also used for the hostname in the operating system, if cloud-init does not use the hostname from DNS. This can be set in a cloud-init module. You find more information on the module in Table 4-6 on page 214.

- Host group

Select the host group so that IBM PowerVC can choose the host based on the configured policy. For details about the automatic placement of partitions, see section 4.5.5, “Host groups” on page 126.
- Host

You can also select a specific host that is part of the selected host group. The shown host all belong to the host group selected. If you change the host group, the list of available hosts will also change.

An example for selecting a specific host would be, if you want to deploy an HA cluster then select different hosts for the cluster partner.
- Description

You may enter a description for the new VM.
- Count

If you need multiple VMs using that image, you can here specify the number of VMs you want to create. In this case you can not specify an specific IP address for the VMs. The IP addresses will be picked from the pool or defined by DHCP, depending on the setup of the used network(s).

Multiple VMs created at once will get a suffix -<number>.
- ▶ Activation input

In this field you can enter or paste a script that will be executed from cloud-init during the deployment. You can also put scripts already in the image. For information on cloud-image scripts refer to 4.16.1, “Preparation of a new image” on page 210.
- ▶ Resources
 - Compute template

Select the compute template that you want to use to deploy the new VM. This defines the resources below the Compute template. For detailed information about planning for CPU and memory resources by using templates, see section 2.4.4, “Information that is required for compute template planning” on page 34.
 - Virtual processors

The value of the virtual processors will be set according to the desired value in the compute template. In brackets you find the minimum and maximum values. You can overwrite the default with a number in the range of minimum and maximum.
 - Processing units

The value of processing units (entitled capacity) will be set according to the desired value in the compute template. In brackets you find the minimum and maximum values. You can overwrite the default with a value in the range of minimum and maximum.
 - Memory (GB)

The amount of memory in GB will be set according to the desired value in the compute template. In brackets you find the minimum and maximum values. You can overwrite the default with a value in the range of minimum and maximum.

► Customization

– SCG (Storage connectivity group)

Select one storage connectivity group for the new VM to access its storage. IBM PowerVC can use a storage connectivity group to determine the use of vSCSI or NPIV to access SAN storage and the desired redundancy. For details about the selection of the storage path and FC ports to use, see section 4.9, “Storage connectivity group setup” on page 138.

– Collocation rule

If you have defined collocation rules, for example an anti-affinity rule that defines that members of this rule are not allowed to run on the same system, you can select the rule already at deployment time. Make sure that a host you have selected fits to that rule.

– SSH Key

When deploying a VM from IBM PowerVC, you can inject an SSH public key into the deployed VM. Users can then use the corresponding private key to connect to the deployed VM. This removes the need to have a known password embedded in the image.

To add a key, click the user symbol on the right upper corner, then click **User settings**. After that the SSH keys settings open, as shown in Figure 4-176. There you can add or delete SSH keys to be used during the deployment.

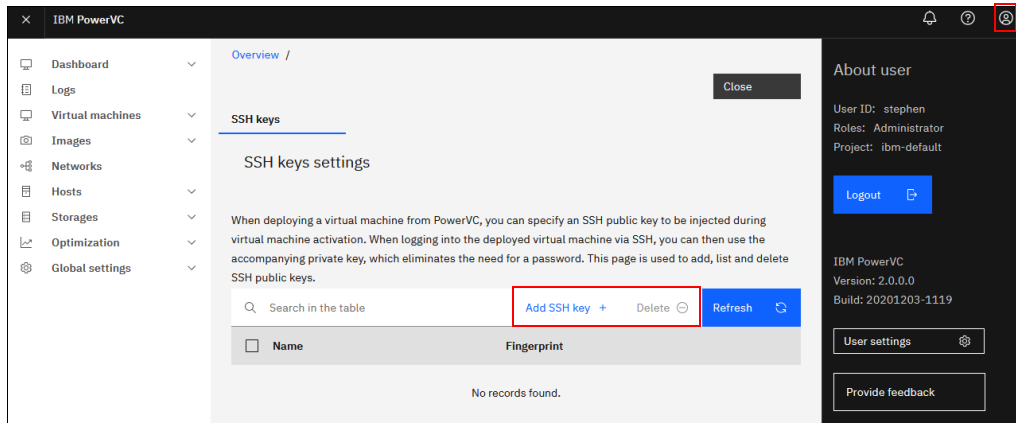


Figure 4-176 SSH keys for injection

► Persistent memory

Persistent memory is a new concept available with POWER9 servers. Persistent memory will be available to the LPAR as volumes that are persistent when the VM is shutdown. Persistent memory volumes can be defined as part of a compute template, or individually in the deployment of a new VM.

To Add a persistent memory volume, click **Add a persistent memory** to add a new line. Then enter the size and check Affinity if you want the physical used memory near the physical used processors. To delete a persistent memory volume click the trash symbol.

The processor and memory graphs on the right show the available resources of the selected host group or if you selected a specific host the resources of that host. From the projected values, you can see how the new VM(s) fit into the host group or the specific host.

3. Now click on the Networks tab to specify the network details as shown in Figure 4-177 on page 243.

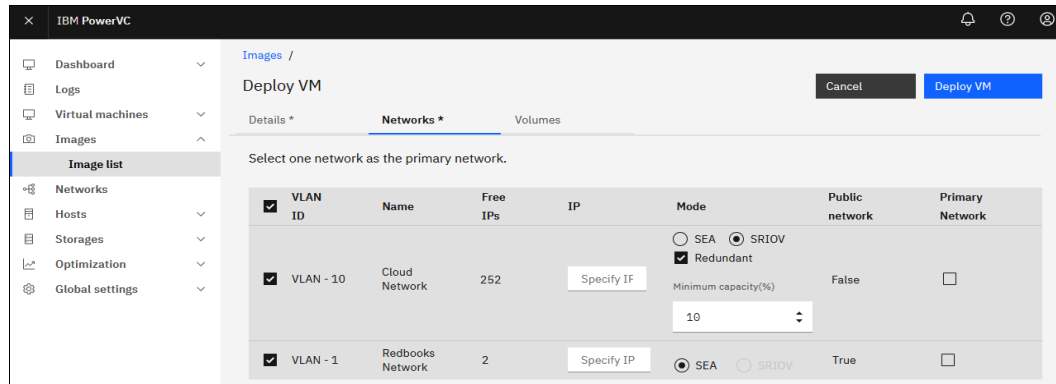


Figure 4-177 VM deployment - Networks tab

In the networks tab to can select all networks that you want to use and specify the following details:

► IP address

Depending on the network configuration you can:

- Enter a specific IP address. This is not available if you want to deploy multiple VMs in one deployment operation.
- Let PowerVC pick an IP address from the pool of available IP addresses. You can see in the column Free IP addresses, how many IP addresses are available.
- Let a DHCP server (if specified in the network definition) select an IP address.

Note: IBM PowerVC verifies according to its database that the IP address that you provide is not already used for another VM, even if the IP address is used in a VM that is powered off.

► Mode

If the VLAN is provided by SEA or SR-IOV you can select the connection type you want.

Note: SR-IOV in PowerVC always means SR-IOV with vNIC, as a VM in PowerVC must be fully virtualized. Do deploy an SR-IOV adapter without vNIC is not possible.

► Maximum capacity (%) - only visible for SR-IOV

Select the virtual NIC capacity to specify quality of service (QoS). This is the minimum but guaranteed capacity for this network adapter in the VM. If the requested capacity is not available, the deployment is not allowed.

Important: The maximum capacity value is calculated based on the amount of capacity currently available for a host/network combination.

► Redundant - only visible for SR-IOV

Specify **Redundant** if your network is set up for redundancy. You can check only the **Redundant** check box if two physical ports are available for a host/network combination. Redundancy cannot be used with one physical port because two ports are required to provide alternative paths for the Ethernet traffic. If this is selected, the vNIC adapter that is created for the VM includes multiple logical ports.

► Public network

This attribute can not be changed in a deployment. It shows you, if the network is available as public network in all defined projects in PowerVC.

► Primary Network

You must specify one network as primary. PowerVC will take the DNS and gateway information from the primary network.

4. After you selected the network you can click the Volumes tab to get information on the volumes used in the image as shown in Figure 4-178.

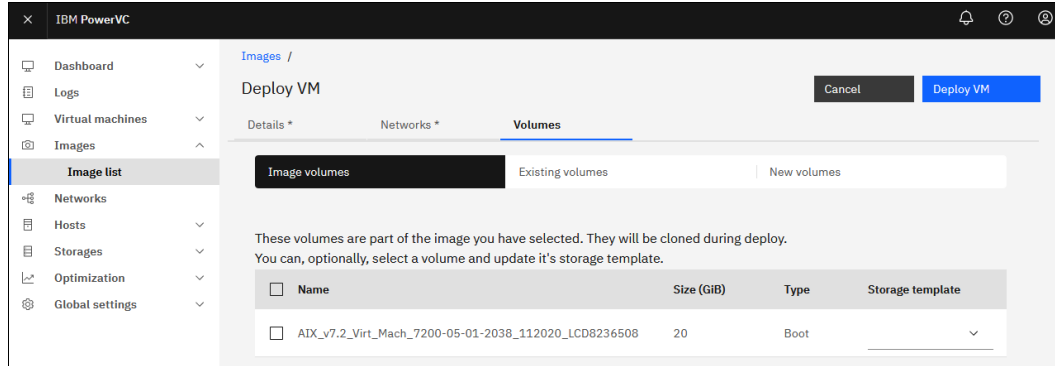


Figure 4-178 Volumes in image to be deployed

Under the Image volumes tab you find the volumes that belong to the image. You do not need to select them, as these volumes will be cloned for the new VM anyway. But you can select a Storage template for each image volume. If you do not specify a storage template, then PowerVC select the one, where the image volume is stored.

With the storage template you can specify which storage, storage pool should be used, or you may have a storage template that also defines a storage mirroring. You can for example put the operating system on spinning disks and the data on very fast flash volumes. More information on storage templates could be found in 4.10, “Storage template setup” on page 140.

5. Under the Existing volumes tab you find existing volumes that could be attached to the new VM. These volumes are either unused or defined as shared and with that may be attached to multiple VMs. An example could be found in Figure 4-179.

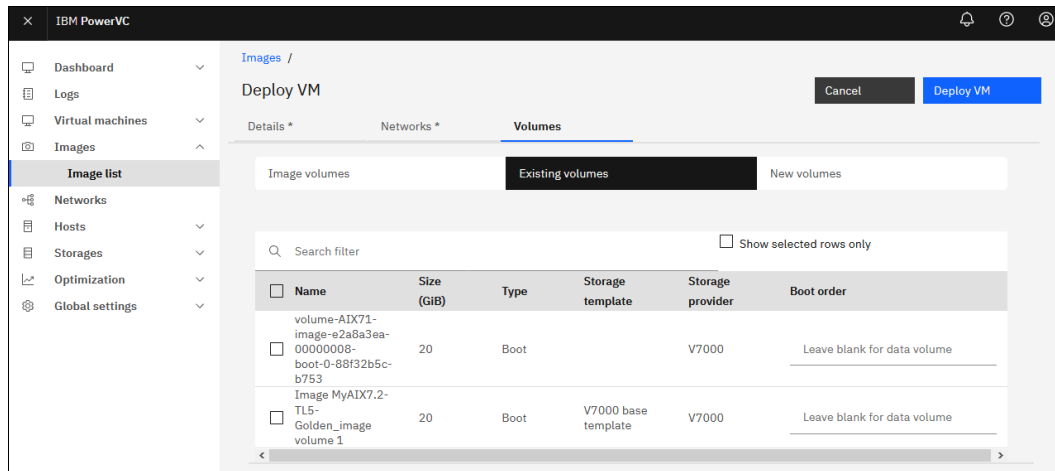


Figure 4-179 Existing volumes to be attached during image deployment

- If you want a volume to be attached, just select it in front of the volume.

If you deploy an image without volumes under the Image volumes tab as shown in Figure 4-178 on page 244 you must select at least one boot volume from the list of existing volumes. In this case specify the boot order for the volume with 0. If you have multiple boot volumes continue with 1, 2 and so on.

- Under the last volumes tab you can define additional new volumes to be attached. In the example in Figure 4-180 you see two additional volume named data2-<1 or 2> with 20 GB each created using the storage template V7000 base template. The new volumes are not shared and therefore can only be used by the new VM.

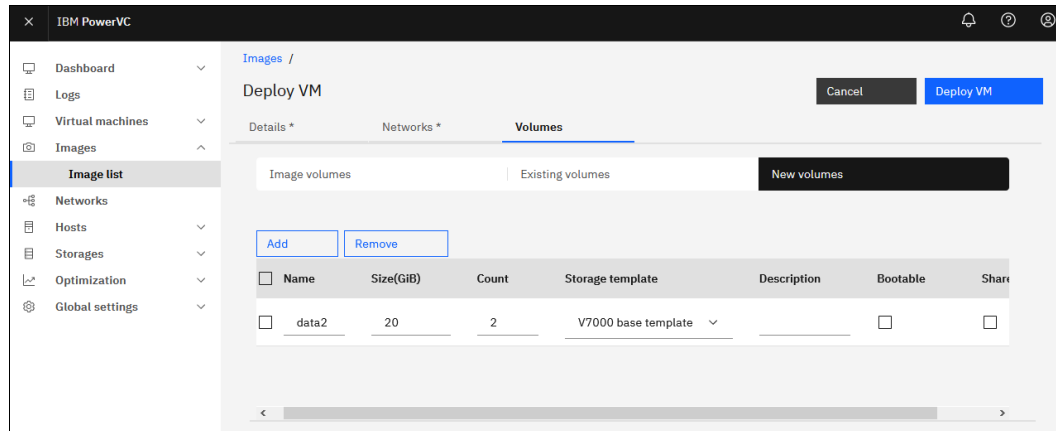


Figure 4-180 New volumes for a new VM to be deployed

- As the last step click **Deploy VM**. To see the progress of the VM deployment, you can read the messages showing up in the notifications (click the bell symbol in the upper right corner), or have a look into the VM list under Virtual machines as shown in Figure 4-181.

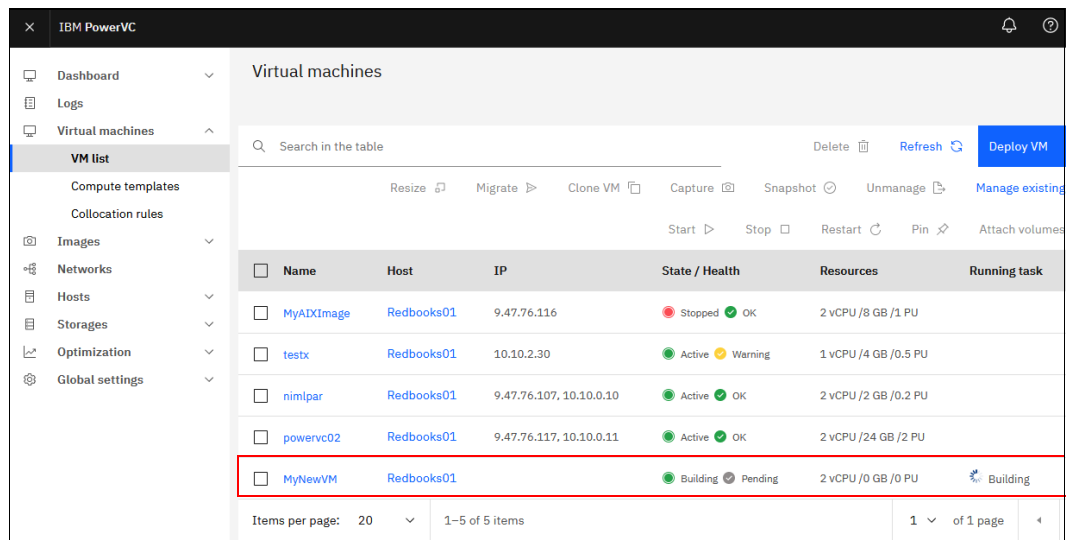


Figure 4-181 Running deployment

You can see in the Running task column the status of Building. During that state PowerVC creates the volumes, zones and the LPAR and connect everything together. After a short time, Building disappears and the VM changes its state to Active / Warning. Warning is normal, as no RMC connect is available yet. Now the VM also shows the assigned IP addresses.

IBM PowerVC Version 2.0.0 has the following limits regarding deployments:

- ▶ IBM PowerVC supports a maximum of 50 concurrent deployments. As a preferred practice, do not exceed eight concurrent deployments for each host.
- ▶ Running more than 10 concurrent deployment operations might require additional memory and processor capacity on the IBM PowerVC management host.
- ▶ If you use only SAN storage and you plan to batch-deploy over 100 VMs that are based on one image, you must make multiple copies of that image and deploy the VMs in batches of 10.

The following settings might increase the throughput and decrease the duration of deployments:

- ▶ Use the striping policy instead of the packing policy.
- ▶ Limit the number of concurrent deployments to match the number of hosts.

4.16.8 Deleting images

To delete an image that is not in use, enter the Image list under Images, select the image to remove and click **Delete** as shown in Figure 4-182.

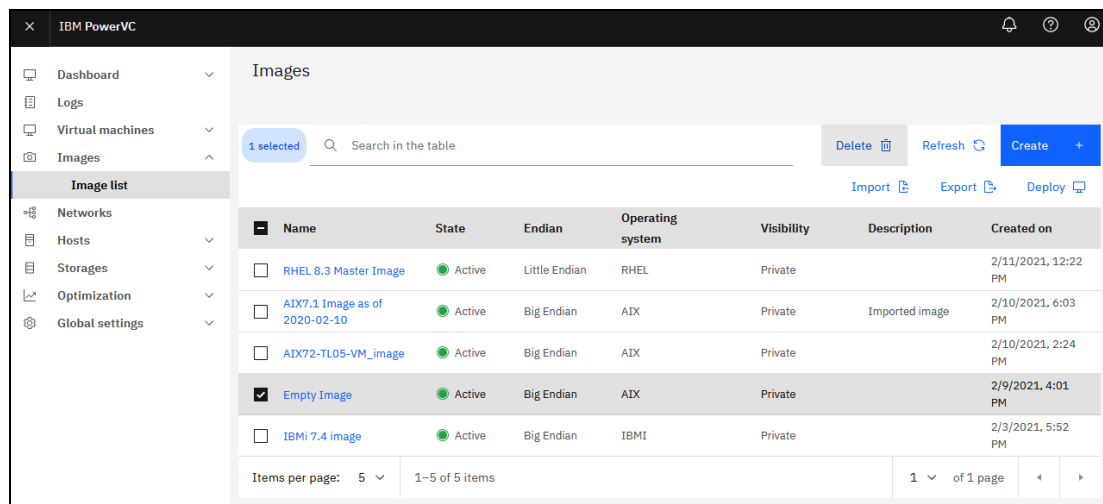


Figure 4-182 Delete an image

The IBM PowerVC management server displays a confirmation window to delete this image, as shown in Figure 4-183.

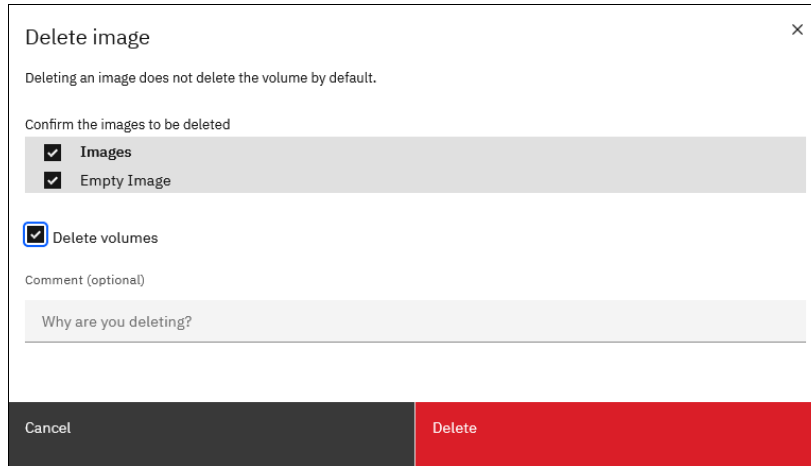


Figure 4-183 Delete an image - confirmation

You can select in the confirmation window, if PowerVC shall also delete the volumes in that image. If you don't check Delete volumes, the volumes will stay in PowerVC and can be used for other purposes.

4.17 Users, groups, and roles setup

During the PowerVC installation, the root user is assigned the administrator role on the ibm-default project. The root user can log in to PowerVC and start performing operations such as assigning roles to other users or groups.

By default, PowerVC uses the local operating system to manage users and groups. To avoid exposing all of the system's users and groups, such as those for system accounts, PowerVC supports both user and group filtering. Only users and groups which match the corresponding filter will be exposed in PowerVC.

Operating system-based user management requires command-line experience, but it is easy to maintain. No dependency exists on other servers or services. To see user accounts in the IBM PowerVC, click the **Global settings** tab, then the **Users and groups** section.

4.17.1 Adding user accounts

You can add additional users with different roles in IBM PowerVC. To add a user account to the operating systems on the IBM PowerVC management host, run the `useradd` command as root from the Linux command-line interface (CLI).

There is a group named `powervc-filter` that is created and visible to IBM PowerVC, so any user that had a PowerVC role assignment is added to that group. The default filters are configured so that these users and groups are visible to PowerVC:

- ▶ The `powervc-filter` group and all of its members
- ▶ Any group that already had a PowerVC role assignment and all of their members

Tip: If you want to create a filter so that PowerVC can see all of the users that are members of the group `mygroup`, and can also see `mygroup` group itself, you can use `powervc-config` command as shown in Example 4-17 below:

Example 4-17 Add a new group filter

```
# powervc-config identity repository -t os --user-filter\  
"(|(memberOf=powervc-filter)(memberOf=mygroup))" --group-filter\  
"(|(name=powervc-filter)(name=mygroup))"
```

Configuring PowerVC for local OS.

```
Stopping httpd services...  
Starting httpd services...
```

```
Stopping ui-server services...  
Starting ui-server services...  
No new role assignments were created. Pre-existing role assignments were  
preserved.
```

```
# powervc-config identity repository  
Type: os  
User filter: (|(memberOf=powervc-filter)(memberOf=mygroup))  
Group filter: (|(name=powervc-filter)(name=mygroup))
```

```
# openstack group list  
985ad84d4d7a3232985 powervc-filter  
7f2540fe9ad3f9d4c60 mygroup
```

As an example, if you want to create a user ID for a system administrator who is new to IBM PowerVC. You want to allow this administrator to view the IBM PowerVC environment only, not to act on any of the managed objects. Therefore, you want to give this administrator only a viewer privilege.

1. By using the following command, you can create the user `monitor` (just for viewing the configuration as viewer) with `/home/monitor` home directory, the `monitor` group as the main group, and a comment with additional information, such as IBM PowerVC Monitoring.

```
# useradd -d /home/monitor -G powervc-filter -m -c "IBM PowerVC Monitor"  
monitor
```

2. Add a proper password for the newly created user:

```
# passwd monitor
```

3. Add the viewer role to the user:

```
# openstack role add --project ibm-default --user monitor viewer
```

If you experienced any issues running the previous OpenStack command shown, like the following two errors:

- Missing value `auth-url` required for auth plugin password
- Missing parameter(s), etc.

Then you will have to do two steps:

- Edit your `/opt/ibm/powervc/powervcrc` environment file with correct information as shown below in Example 4-18.

Example 4-18 Edit /opt/ibm/powervc/powervcrc environment file

```
# vi /opt/ibm/powervc/powervcrc
export OS_IDENTITY_API_VERSION=3
export OS_AUTH_URL=https://9.47.76.117:5000/v3/
export OS_CACERT=/etc/pki/tls/certs/powervc.crt
export OS_REGION_NAME=RegionOne
export OS_PROJECT_DOMAIN_NAME=Default
export OS_PROJECT_NAME=ibm-default
export OS_TENANT_NAME=$OS_PROJECT_NAME
export OS_USER_DOMAIN_NAME=Default
#export OS_USERNAME=root
#export OS_PASSWORD=Pa$$w0rd
export OS_COMPUTE_API_VERSION=2.46
export OS_NETWORK_API_VERSION=2.0
export OS_IMAGE_API_VERSION=2
export OS_VOLUME_API_VERSION=3
```

- Rerun the `/opt/ibm/powervc/powervcrc` environment file using **source** command so the OpenStack can have the authentication information updated:

```
# source /opt/ibm/powervc/powervcrc
```

- Then run the **openstack** command again.

4. Make sure the user has been added. See Example 4-19

Example 4-19 Check the user creation

```
# openstack user list | grep monitor
# openstack user show monitor
```

See the below link for more information about IBM PowerVC ready roles and their meaning:

https://www.ibm.com/support/knowledgecenter/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_supported_roles_hmc.html

4.17.2 Updating the user accounts

You can modify a user account in the operating systems on the IBM PowerVC by running **usermod** command:

```
# usermod [options] login_name
```

If you are using the user role assignment approach, then you must run it through an OpenStack command:

```
# openstack role add --user <user> --project <project> <role_name>
```

If you decided to delete a role assigned for some reason, you can follow this command:

```
# openstack role delete --user <user> --project <project> <role_name>
```

4.17.3 Disabling the root user account from IBM PowerVC

During the PowerVC installation, the root user is assigned the **administrator** role on the **ibm-default** project. The root user can log in to PowerVC and start performing operations such as assigning roles other users or groups.

As part of best practices, It is always recommended not to use the root user account on IBM PowerVC. It is a security preferred practice to remove it from the admin group.

Run the following **openstack** command to remove the **admin** role from root user:

```
# openstack role delete --user root --project ibm-default admin
```

4.17.4 Working with roles

Roles are used to specify what actions a user can perform since they are assigned to a user or group. A user or group can have more than one role, in which case they are able to perform any action that at least one of their roles allows. At least one user must have the admin role, since otherwise any action that requires the admin role such as creating role assignments, would never be possible.

When PowerVC is installed, root is initially assigned the admin role. It is recommended that you assign the admin role to another user or group and then remove the admin role assignment from root as mentioned earlier in 4.17.3, “Disabling the root user account from IBM PowerVC” on page 250.

If you want to list the roles, use **openstack** command as shown below in Example 4-20.

Example 4-20 List the available OpenStack roles

```
# openstack role list
```

ID	Name
0175da7e73f24232b5cb5ab483bd9291	deployer
0371f5d7cbd047b5ac2a4da3f6af0f97	service
7120b52f9d054d0f91491320107c18a7	admin_asst
982422f932674490a816c075fc8a2ea9	vm_user
a0bcacd28e6f4e9a97bc5a5ff8d8e797	admin
a1fc10f835af423883089771e391865a	storage_manager
bcfa0a2d1e5e414081a2d2b97d8001e4	image_manager
c3b2d9e44aaa47d9b68b5f427fd7d2e7	vm_manager
f454f3075ae54dd7bda731c7e24dac3b	viewer

You can also list the OpenStack role arguments by running **openstack role -h** command as shown in Example 4-21 below:

Example 4-21 List the OpenStack role arguments

```
# openstack role -h
Command "role" matches:
  role add
  role assignment list
  role create
  role delete
  role list
  role remove
  role set
  role show
```

4.17.5 Setting project policies

A project, sometimes referred to as a tenant, is a unit of ownership. Most resources, such as virtual machines, volumes, and images, belong to a specific project. Only users with a role assignment for a given project can work with the resources belonging to that project. The `ibm-default` project is created during installation, but PowerVC supports the creation of additional projects for resource segregation. To work with projects, an admin can log in to the `ibm-default` project and click **Projects** from the Configuration page.

IBM PowerVC for Private Cloud administrators can set several project-specific policies. These policies apply only to users with `self_service` authority. You can set properties such as whether users require administrator approval to perform certain tasks.

To set a project's policies, go to **Global settings**, then select **Enterprise Policies**.

These policies that can be set:

- ▶ **Deploys before approval is required**

The number of virtual machines each self service user can own without requiring approval for deploys. If a user tries to deploy an additional virtual machine, that deploy would require approval. If this policy is disabled, virtual machine deploys always require approval. When editing this policy, you can specify the number of deploys or deselect Specify value if approvals are never required.

- ▶ **Days before virtual machine expiration**

The number of days before virtual machines created by self service users in the project expire. When a virtual machine is deployed in this project, this value is used to set its expiration date. The default is 30 days. A self service user can request an extension to the expiration date after deploy. When requesting an extension, users cannot select a date further out than the value set in this policy. For example, if this value is set to 90, users cannot select a date more than 90 days away from the current expiration date. This policy cannot be disabled.

- ▶ **Days before expired virtual machine deletion**

Specifies the number of days before an expired virtual machine is automatically deleted. This value applies to all virtual machines created by self service users in this project. If this policy is disabled, expired virtual machines are never deleted.

- ▶ **Expiration extensions before approval is required**

The combined number of virtual machine expiration date extensions that are automatically approved. After this limit is reached, further expiration date extension requests require approval. If this policy is disabled, virtual machine expiration extension requests always require approval.

- ▶ **Days before automatic approval of pending expiration extension requests**

The number of days that virtual machine expiration extension requests can be in pending state. If the request is not approved after the specified number of days, it is automatically approved. If this policy is disabled, expiration extension requests are never automatically approved.

- ▶ **Clone virtual machine approval limit**

This policy controls processing of clone virtual machine requests by a self service user. If this policy is disabled, virtual machine clones require approval. When editing this policy, the only possible value is: `Approval not required` which is auto-selected. By default, clone virtual machines requests need approval.

- ▶ Virtual machine snapshots approval limit

This policy controls processing of virtual machine snapshot requests by a self service user. If this policy is disabled, virtual machine snapshots require approval. When editing this policy, the only possible value is: `Approval not required` which is auto-selected.

If you want to list the policy types through PowerVC cloud edition, use the `powervc-cloud-config` and `powervc-cloud-config --help` commands.

You can run `powervc-cloud-config` command to create a policy if it does not exist or update the existing policy. If you do not specify a value for `-p` option, the policy is created for the `ibm-default` project. The following example specifies only 10 VM captures or snapshots that can be performed without administrator approval. After this limit is reached, further capture requests require approval:

```
# powervc-cloud-config set-policy snapshot_approval_limit 10
```

4.17.6 Lightweight Directory Access Protocol

IBM PowerVC uses the local operating system to manage users and groups by default, however you can optionally configure it to work with an existing Lightweight Directory Access Protocol (LDAP) server. To be able to do that, you have to initially ensure that a supported LDAP server is installed and running in your environment if you plan to configure PowerVC with an LDAP server.

LDAP can always handle as much information as you need, but it is commonly used to associate names with phone numbers and addresses. LDAP is a client/server solution. The client requests information and the server answers the request. LDAP can be used as an authentication server.

Use the `powervc-config identity repository` command to configure LDAP or to switch back to using the local operating system. For instructions to use this command, run `powervc-config identity repository --help`. When you run this command, you must provide information about your LDAP configuration. Be sure to specify a user or a group for the initial administrator role assignment so that at least one user has administrator permissions. Administrators can then create additional role assignments as required.

See the following links for details:

<https://www.youtube.com/watch?v=pXQyhbIZZ4M>

<https://community.ibm.com/community/user/power/blogs/divya-k-konoor1/2018/07/11/ibm-powervc-ldap-configuration-faqs?CommunityKey=6f9fc4cd-7614-41b1-98d8-2941c3497e37&tab=recentcommunityblogsdashboard>

https://www.ibm.com/support/knowledgecenter/en/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_ldap_hmc.html

4.17.7 User and Group filtering

By default, PowerVC uses the local operating system to manage users and groups. To avoid exposing all of the system's users and groups, such as those for system accounts, PowerVC supports both user and group filtering. Only users and groups which match the corresponding filter will be exposed in PowerVC. When you install or upgrade PowerVC, default user and group filters are created. On a new installation, a new group named **powervc-filter** is created, the default user (root) is added to that group, and the filters are configured so that only that the **powervc-filter** group and its members are visible to PowerVC.

Default values on a freshly installed PowerVC appear as below. The **Type** indicates that the identity backend has been by default configured to use the local OS registry. The below Example 4-22 shows filters which indicate that only OS users who are members of OS group named **powervc-filter** and OS group with name **powervc-filter** are be visible to PowerVC. With the below configuration, for a new OS user to be visible, add the user to the **powervc-filter** OS group.

Example 4-22 Check the identity repository filter

```
# powervc-config identity repository
Type: os
User filter: (memberOf=powervc-filter)
Group filter: (name=powervc-filter)
```

Note: If you backed up your PowerVC, your backup should include information about the configured user and group filters. However, PowerVC does not create operating system users or groups, or adjust group memberships at the time of restore. If operating system users and groups are configured differently when the backup is restored, this might lead to issues.

Before performing any PowerVC restore operation, ensure that the system where restore is performed has the same set of PowerVC OS users or groups with the right group membership as on that of the PowerVC system from where backup has been taken, and make sure always that keystone user must be able to run the **su** command and login as PowerVC end users.

To create or update a filter, run the **powervc-config identity repository** command with the appropriate attributes. Every time you run this command, it replaces the existing filter. Therefore, it is recommended that you view the current filter before updating the filter by running **powervc-config identity repository** command.

Example 4-23 below shows the full command to create a filter so that PowerVC can see all of the users that are members of the groups mygroup1, mygroup2, or mygroup3, and can also see any groups that are named mygroup1, mygroup2, or mygroup3:

Example 4-23 Create a new user group filter

```
# powervc-config identity repository -t os --user-filter
"(| (memberOf=mygroup1) (memberOf=mygroup1) (memberOf=mygroup1))" --group-filter
"(| (name=mygroup1) (name=mygroup1) (name=mygroup1))"
```

You could change filters to make a new user and group (that the user is a member to) visible to PowerVC (Example 4-24 on page 254):

Example 4-24 Change user group filters

```
# useradd testuser
# passwd testuser
# groupadd testgroup
# usermod -a -G testgroup testuser
# powervc-config identity repository -t os --user-filter
"(|(memberOf=powervc-filter)(memberOf=testgroup))" --group-filter
"(|(name=powervc-filter)(name=testgroup))"
```

Note: If the user-filter is not modified to update the group, the users within the new group will not be visible.

If you decided to remove user and group filters, you will have to run **powervc-config identity repository** command:

```
# powervc-config identity repository -t os --user-filter None --group-filter
None
```

Although Example 4-24 shows how to remove the filter anyway, however it is not a best practice where the outcome it makes all OS users and groups becoming visible to PowerVC without any filtering.



IBM PowerVC for Private Cloud

This chapter describes exclusive cloud features of IBM PowerVC for Private cloud which allows the self-service users to easily deploy and manage one or several virtual machines (VMs).

The user of the cloud services is not autonomous when deploying VMs, as policies set by the project administrator guides the process. This is key in ensuring self-service users maintain their quota and deploy VMs on the right hosts. Project policies specify what users can do and whether administrator approval is required for each action. In general, a self-service user can perform the following actions:

- ▶ Manage VMs that are owned by that user
- ▶ Deploy VMs by using a deploy template
- ▶ Review and withdraw action requests
- ▶ View their own metering data

This chapter will focus on the self-service user experience, covering tasks such as:

- ▶ Logging in to PowerVC for Private Cloud
- ▶ Exploring the PowerVC for Private Cloud user interface
- ▶ Deploying VMs from a template
- ▶ Reviewing and withdrawing action requests

This chapter will also cover administrative supportive tasks for the self-service user such as:

- ▶ Creating an account for the self service user
- ▶ Creating deploy templates
- ▶ Configuration of cloud policies
- ▶ Approving and denying deployment requests
- ▶ Extending expiration dates of VMs and changing owners of VMs

5.1 Creating a self-service user

The method of creating a user account for a self-service user depends on settings made by the project administrator. IBM PowerVC can be configured to query an LDAP server for authentication, or the operating system user accounts authentication. To add a self-service user account to the operating systems on the IBM PowerVC management host, login to the IBM PowerVC management host and run the following commands:

```
# useradd [options] login name
# openstack role add --user <user> --project <project> <role>
```

Example 5-1 shows usage of these commands to add the user account `clouduser`.

Example 5-1 Adding a self-service user account called clouduser

```
# useradd -d /home/clouduser -g powervc-filter -m -c "Self Service User" clouduser
# openstack role add --user clouduser --project ibm-default self_service
```

5.2 Creating a deploy template

A Deploy template is a set of properties that define the deployment of an image in IBM PowerVC for Private Cloud. Properties such as the CPU, memory, network, storage connectivity group and others are defined in a deployment template to ensure the deployment process is fast and easy for the self-service user. A Deploy template is a feature exclusive to IBM PowerVC for Private Cloud.

1. In order to create a deploy template, images must exist in the image library. A project administrator can create a deploy template by navigating to the **Image list**, selecting an image and clicking **Create deploy template** as shown in Figure 5-1.

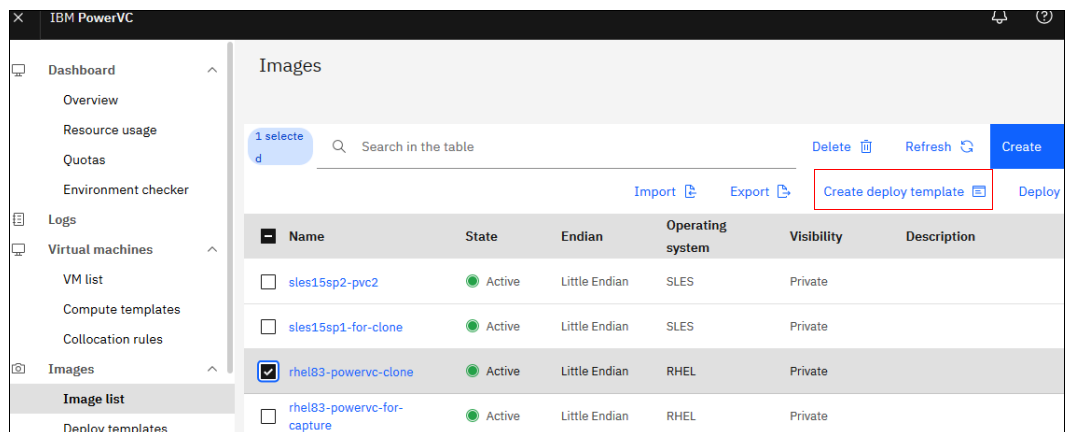


Figure 5-1 Creating a deploy template from the Red Hat Enterprise Linux 8.3 image

2. The deploy template has three tabs: details, image volumes, and network. Figure 5-2 on page 257, Figure 5-3 on page 257 and Figure 5-4 on page 258 show the deploy template creation process with filled fields providing example of values that can be used by an administrator.

Figure 5-2 The details section when creating a deploy template

- The next tab is the **Image volumes**. This section is not editable during the creation of a deploy template. This tab displays the boot and data volumes contained in the selected image.

Name	Size(GB)	Status	Storage template	Storage provider	Bootable
rhel83-powervc-for-capture	190	Available	storwize-v7000 base template	9_114_181_189	Yes

Figure 5-3 image volume tab

- The next tab when creating a deploy template is the **Network** tab. In this tab the administrator selects which network segments the IP addresses will be selected from. It is worth noting that manual IP addresses cannot be assigned when creating deploy templates. In Figure 5-4 on page 258, the Redbooks Network is selected.

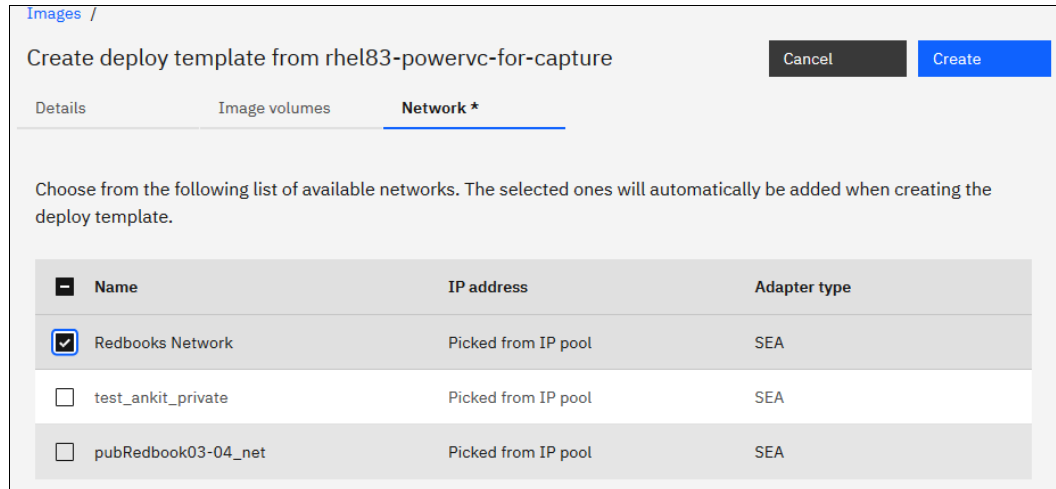


Figure 5-4 Deploy template network tab

5. Deploy templates can be listed, copied and edited from the deploy template window as shown in Figure 5-5.

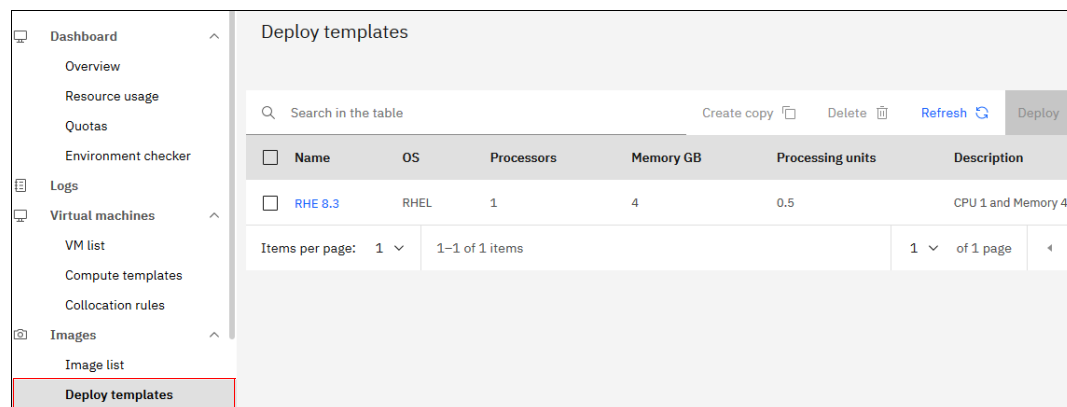


Figure 5-5 List of deploy templates

5.3 Setting cloud policies

The project administrator can set policies to define the self-service user's interaction with IBM PowerVC for Private Cloud. These policies set by the administrator govern all self-service users. The policies are as follows:

1. Deploys before approval is required: The number of virtual machines that can be deployed by a self-service user before the project administrator's approval is required.
2. Days before virtual machine expiration: The number of days before a VM deployed by a self-service user expires. This value can be edited after the deployment of the VM, it can be reduced or increased.
3. Days before expired virtual machine deletion: The number of days before expired VMs are automatically deleted. This policy can be disabled, or a particular number of days can be set. When the expiration date of a VM arrives, the VM is shut down and cannot be restarted by the self-service user.

4. Expiration extensions before approval is required: The number of times a self-service user can extend the expiration of a VM before the project administrator’s approval is required.
5. Days before automatic approval of pending expiration extension requests: The number of days before an automatic approval of pending expiration extension requests by the self-service user.
6. Virtual machine clone approval limit: This policy controls whether clone VM action needs approval or not. By default, Approval is always required. During edit, it can be changed to Approval not required value which is auto filled.
7. Virtual machine snapshots approval limit: This policy controls whether snapshot VM action needs approval or not. By default, Approval is always required. During edit, it can be changed to Approval not required value which is auto filled.

The cloud policy can be set by navigating to the “Global Policies” drop down and selecting Enterprise policies as shown in Figure 5-6.

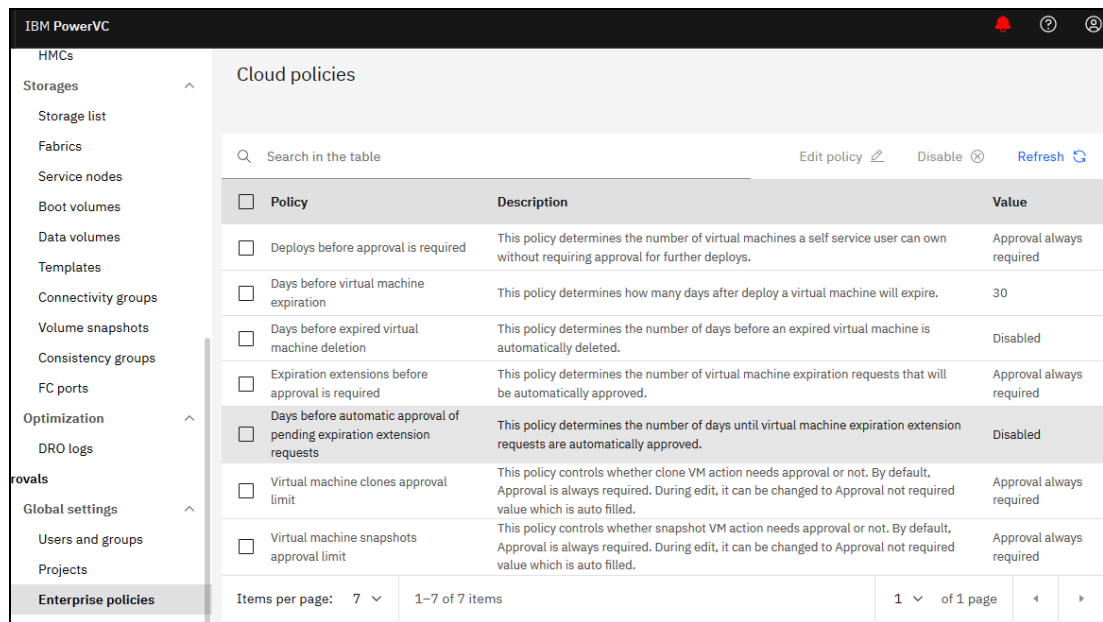


Figure 5-6 Setting cloud policies

In a recap of the actions performed by the project administrator in this and preceding sections, a self-service user account was created, a deploy template was created and cloud policies were set to govern self-service users. The next section will take a look at the self-service user accessing IBM PowerVC for Private Cloud.

5.4 Logging in as a self-service user

After the project administrator has created the self-service user account and it has been shared with the user, the user can log into IBM PowerVC for Private Cloud.

1. Open a web browser on your workstation and point it to the IBM PowerVC address:
`https://<ip address or hostname>/`
2. Log in to IBM PowerVC with your given user ID and password, as shown in Figure 5-7.

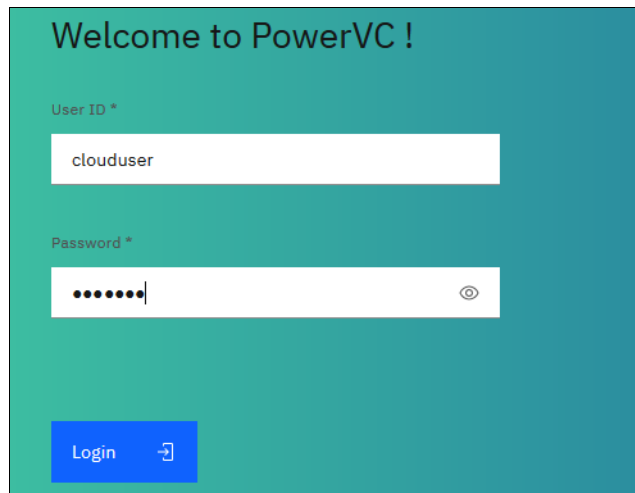


Figure 5-7 IBM PowerVC for Private Cloud login window

5.5 Self-service portal graphical user interface

The self service portal view varies from the administrator's GUI. Figure 5-8 shows the self-service portal with limited capabilities relative to the administrative window. If at this stage VMs have been deployed by the self-service user, they will be seen in the overview window.

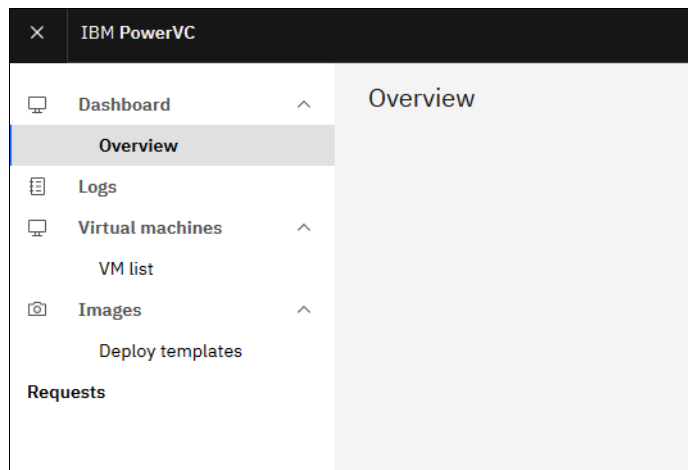


Figure 5-8 Self-service portal

Functions that can be performed by the self-service user include the following:

- ▶ Viewing of logs and status of requests
- ▶ Listing deployed VMs
- ▶ Deploying VMs through deploy template

5.6 Deploying virtual machines from a template

In order for the self-service user to create a VM, the user must select a deploy template, as shown in Figure 5-9.

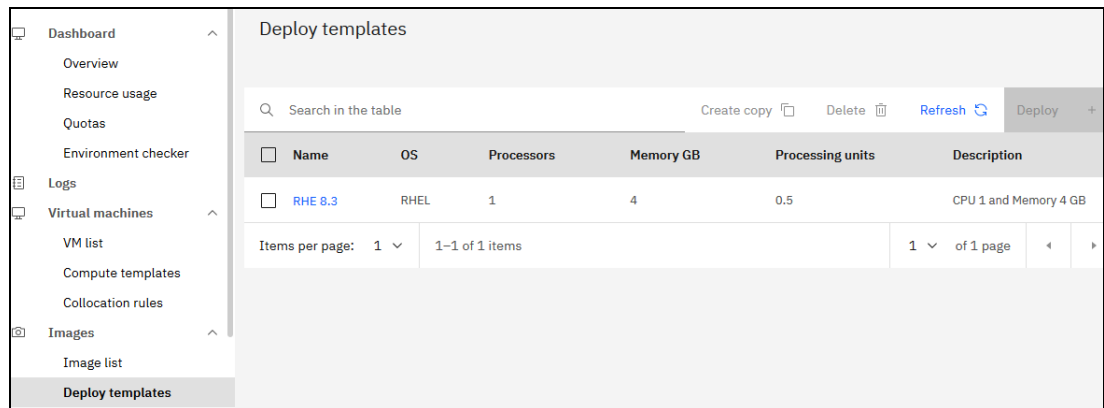


Figure 5-9 Viewing deploy templates from the self-service portal

Prior to deploying VMs, a self-service user can specify an SSH public key to be injected during virtual machine activation. When logging into the deployed virtual machine via SSH, you can then use the accompanying private key, which eliminates the need for a password. Public keys can be added by navigating to about user, then user settings.

When deploying a VM, the user navigates to the deploy templates and selects a template. The user must specify a name for the VM and how many instances to create. Additionally, the user can add a key pair, a description, and a message for the administrator. The message to the administrator can be to request for approval or customized resources for the VM(s). Figure 5-10 shows the deployment window.

The screenshot shows a deployment window titled "Deploy RHE 8.3" with a "Cancel" button and a "Deploy" button. The form includes the following fields:

- Virtual machine name ***: Redbook_test
- Processors**: 1
- Key pair**: None
- Processing units**: 0.5
- Instances**: 1
- Memory**: 4096
- Description**: VM deployed as an example for the PowerVC 2.0 Redbook
- Operating system**: RHEL
- Message to administrator**: Kindly approve my VM request

Figure 5-10 Deployment of a VM by a self-service user

After the deployment of the VM, if the cloud policy is set such that every deployed VM has to be approved by the administrator, the deployed VM will be seen in a pending state when the user navigates to the request window in the self-service user portal as shown in Figure 5-11.

The screenshot shows a "Requests" window with a sidebar menu containing "Dashboard", "Overview", "Logs", "Virtual machines", "Images", and "Requests". The main content area displays a table of requests:

Request type	Status	Virtual machine	Expiration	Request date
Deploy	Pending	Redbook_test		2/17/2021, 9:16:22 AM

Below the table, there is a pagination control showing "Items per page: 1" and "1-1 of 1 items".

Figure 5-11 Viewing deployed requests

The request is seen from the administrator approval window as shown in Figure 5-12. The project administrator can choose to approve or deny the request.

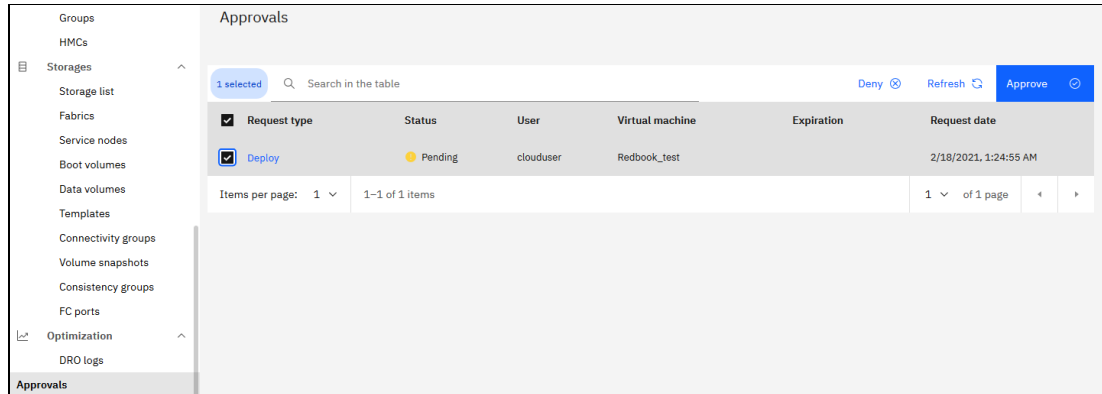


Figure 5-12 Administrator approval page

After the project administrator approves the deployment of the VM, the VM appears in the VM list and overview pages of the self-service user. On the VM list page, the self service user can perform the following actions:

1. Start, restart, stop, and delete the VM
2. Clone the VM
3. Take a snapshot of the VM
4. Edit the expiration date of the VM

Figure 5-13 shows the listed VM and the tasks which can be performed by the self-service user.

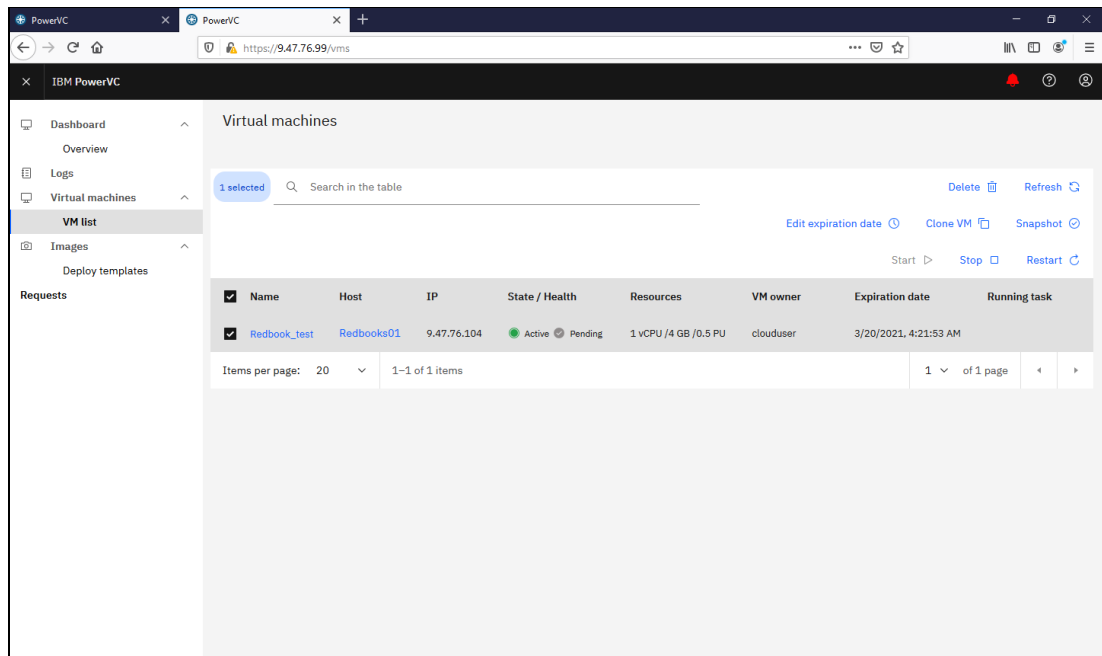


Figure 5-13 Self-service user VM list page

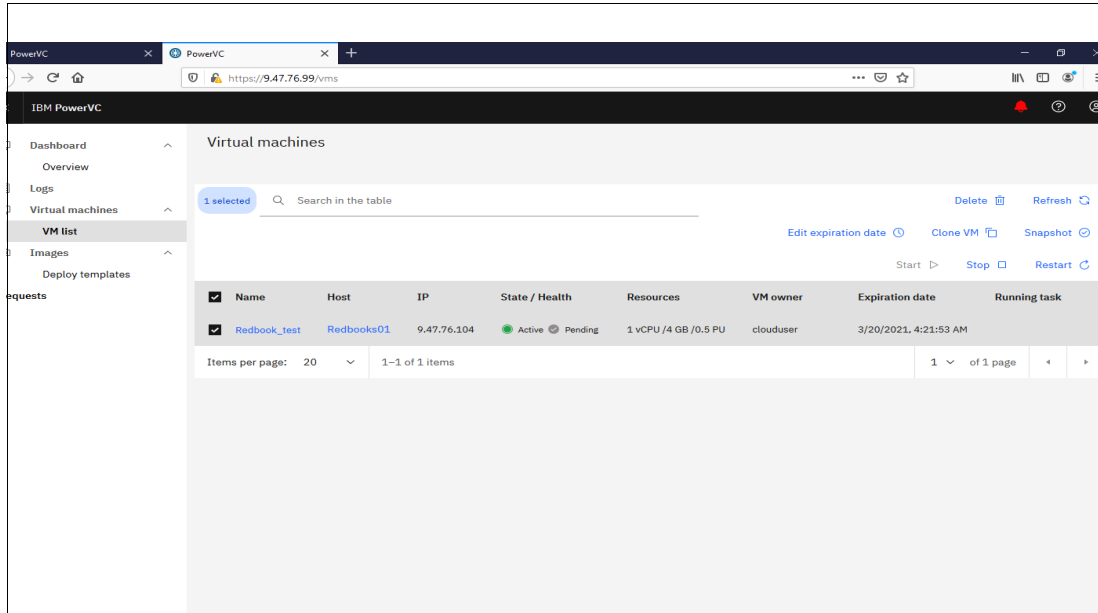


Figure 5-14 Self-service user VM list page

5.6.1 Ownership and expiration dates of VMs

VMs that have been previously deployed by the administrator or other self-service users, can be reallocated to another self-service user. Figure 5-15 shows the changing of ownership from the administrator to the *clouduser* of a VM called *nimlpar*. Click on the dropdown box indicated with the red square and select the new owner of the VM.

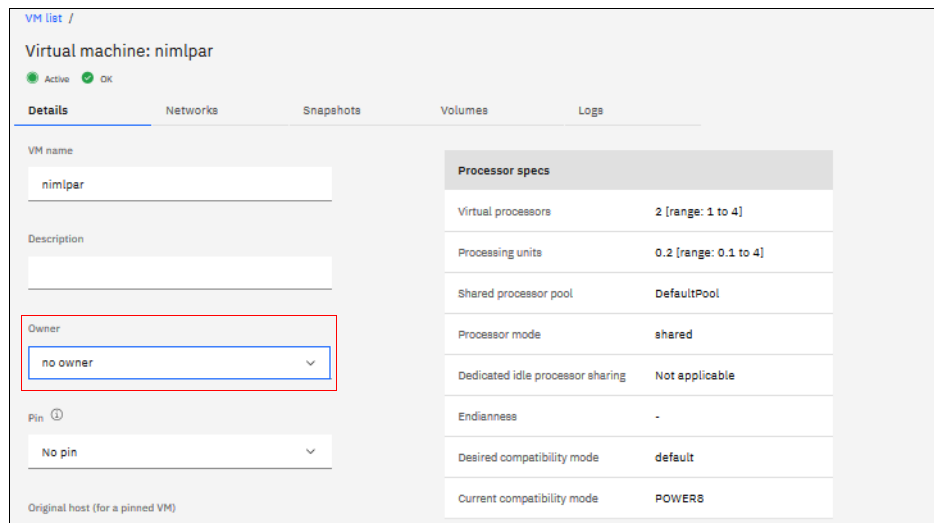


Figure 5-15 Changing the owner of a VM

After the change is made by the administrator, the VM is instantly seen on the VM list and overview pages of the self-service user. This is shown in Figure 5-16.




Figure 5-16 Edit the expiration date

The self-service user may require more time to work on a VM, therefore the user can request the expiration date of the VM be extended or decreased. This will need to be approved by the administrator or not, depending on the cloud policy set for self service users. In order to edit the expiration date of the VM, do the following:

1. Navigate to the VM list page
2. Select the VM
3. Click on **Edit expiration date** (as shown in Figure 5-16)
4. Select the new date and click **Edit expiration date** (as shown in Figure 5-17)

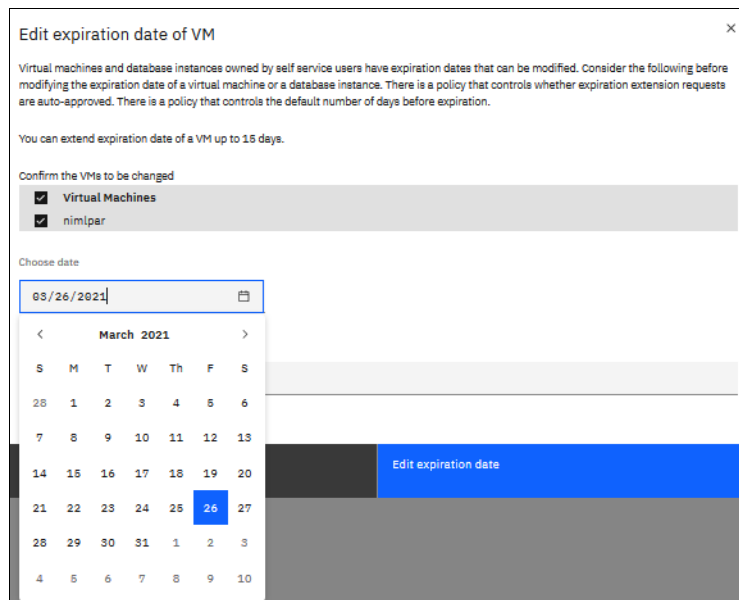


Figure 5-17 Selecting a new expiration date

The request is sent to the administrator where it can be approved or denied as shown in Figure 5-18.




Figure 5-18 Administrator approval page

5. After the administrator approves the request, the VM's expiration date is extended as shown in Figure 5-19.



The screenshot shows the 'Virtual machines' management interface. At the top, there is a search bar and a 'Refresh' button. Below the search bar, there are several action buttons: 'Delete', 'Edit expiration date', 'Clone VM', 'Snapshot', 'Start', 'Stop', and 'Restart'. The main part of the interface is a table with the following columns: Name, Host, IP, State / Health, Resources, VM owner, Expiration date, and Running task. The table contains one entry for a VM named 'nimipar' on host 'Redbook01'. The VM is in an 'Active' state with a health status of 'OK'. Its resources are '2 vCPU / 2 GB / 0.2 PU' and it is owned by 'clouduser'. The expiration date is '3/31/2021, 12:00:00 AM'. At the bottom of the table, there is a pagination control showing 'Items per page: 20' and '1-1 of 1 items'.

Name	Host	IP	State / Health	Resources	VM owner	Expiration date	Running task
<input type="checkbox"/> nimipar	Redbook01	-	Active ● OK	2 vCPU / 2 GB / 0.2 PU	clouduser	3/31/2021, 12:00:00 AM	

Figure 5-19 VM with a changed expiration date

This chapter has taken a look at the IBM PowerVC for Private Cloud exclusive cloud capabilities feature. It provides users with a public cloud experience but on premise. Clients can take advantage of the cloud capability features to improve their infrastructure delivery.



IBM PowerVC integration with Terraform and Ansible

This chapter describes options that are available for automating your deployments and management including.

- ▶ Terraform
- ▶ Ansible

6.1 Terraform

Terraform is an infrastructure as code tool for building, changing, and managing infrastructure in a safe, repeatable way.

Infrastructure as Code is the process of managing infrastructure in a file or files rather than manually configuring resources in a user interface. A resource in this instance is any piece of infrastructure in a given environment, such as a virtual machine, network interface, for example.

It is described using a high-level configuration syntax called the HashiCorp Configuration Language (HCL) for human-readable, automated deployments. This allows a blueprint of your datacenter to be versioned and treated as you would do with any other code.

Terraform has a planning step where it generates an execution plan. The execution plan shows what Terraform will do when you call apply. This lets you avoid any surprises when Terraform manipulates infrastructure.

6.1.1 Install and configure Terraform

We installed the Terraform CLI on the IBM PowerVC Server by following this instructions: <https://github.com/ocp-power-automation/terraform-providers-power/blob/main/README.md>

1. Download the latest Terraform binary for IBM Power

```
# curl -L
https://oplab9.parqtec.unicamp.br/pub/ppc64el/terraform/latest/terraform-0.14.5
-o /usr/bin/terraform && chmod +x /usr/bin/terraform
```

2. Verify that the Terraform version is v0.14.5

```
# terraform --version
Terraform v0.14.5
```

3. Download the Terraform provider plugins that will be used in the scripts to deploy a VM as shown in Example 6-1.

Example 6-1 Download the Terraform provider

```
# mkdir /home/terraform
# cd /home/terraform
# curl -L
https://github.com/yussufsh/providers_ppc64le/releases/download/v0.7/archive.zip
p -o archive.zip
# unzip archive.zip
```

4. Validate that the plugins are present in the current directory, as in Example 6-2.

Example 6-2 Terraform plugins

```
# ls -lrt registry.terraform.io/
total 0
drwxr-xr-x. 3 root root 22 Dec 18 17:26 community-terraform-providers
drwxr-xr-x. 3 root root 22 Dec 18 17:26 terraform-providers
drwxr-xr-x. 4 root root 32 Dec 18 17:27 hashicorp
drwxr-xr-x. 3 root root 23 Dec 18 17:28 terraform-provider-openstack
drwxr-xr-x. 3 root root 17 Dec 18 17:29 IBM-Cloud
```

- Copy the OpenStack plugin in the correct directory.


```
# cp -r ./registry.terraform.io/terraform-provider-openstack/openstack/
registry.terraform.io/hashicorp/openstack
```
- Create a simple Terraform script to validate that the plugins are correctly installed as shown in Example 6-3.

Example 6-3 Simple Terraform script to test

```
# vi main.tf

provider "openstack" {
  user_name   = "admin1"
  password    = "abcd1234"
  tenant_name = "powervc"
  domain_name = "Default"
  auth_url    = "https://10.3.54.17:5000/v3/"
  insecure    = true
}

resource "openstack_compute_instance_v2" "my-VM" {
  name       = "terra1"
  image_id   = "1111"
  flavor_id  = "1111"

  network {
    uuid = "1111"
    name = "VLAN"
  }
}
```

- Run the **terraform init** command to use the local plugins as in the following Example 6-4.

Example 6-4 Terraform init

```
# terraform init --plugin-dir ./
```

```
Initializing the backend...
```

```
Initializing provider plugins...
```

- Reusing previous version of hashicorp/openstack from the dependency lock file
- Installing hashicorp/openstack v1.32.0...
- Installed hashicorp/openstack v1.32.0 (unauthenticated)

```
Terraform has been successfully initialized!
```

You may now begin working with Terraform. Try running "terraform plan" to see any changes that are required for your infrastructure. All Terraform commands should now work.

If you ever set or change modules or backend configuration for Terraform, rerun this command to reinitialize your working directory. If you forget, other commands will detect it and remind you to do so if necessary.

```
auth_url    = "https://10.3.54.17:5000/v3/"
```

Now you can run the Terraform commands to create your VM on PowerVC.

6.1.2 Create a template for an AIX deployment

The fundamental component for the deployment of an image within Terraform is a template. A *template* defines a set of resources that should be provisioned during deployment. We can use Terraform templates to define how we want a new logical partition / virtual machine (LPAR / VM) deployed through IBM PowerVC by providing information, such as:

- ▶ A compute template to use
- ▶ A host group to which to deploy
- ▶ Network adapters to attach
- ▶ Extra disks to allocate

You can use the OpenStack CLI command to retrieve this information on the PowerVC server. The OpenStack CLI needs several pieces of information about the PowerVC deployment in order to work. You can provide this information by using CLI arguments, but it is easier to set the environment variables before using the CLI. PowerVC includes the `/opt/ibm/powervc/powervcrc` file to simplify this task. Copy the file to the `/home/terraform` directory as shown in Example 6-5 and edit it as necessary. In particular, set these values:

- ▶ `OS_USERNAME` - Your user name
- ▶ `OS_PASSWORD` - Your password
- ▶ `OS_PROJECT_NAME` - Your project
- ▶ `OS_AUTH_URL` - PowerVC URL

Example 6-5 Copy powervcrc

```
# cp /opt/ibm/powervc/powervcrc /home/terraform/powervcrc
# cd /home/terraform
# vi powervcrc

export OS_IDENTITY_API_VERSION=3
export OS_AUTH_URL=https://10.3.54.17:5000/v3/
export OS_CACERT=/etc/pki/tls/certs/powervc.crt
export OS_REGION_NAME=RegionOne
export OS_PROJECT_DOMAIN_NAME=Default
export OS_PROJECT_NAME=powervc
export OS_TENANT_NAME=$OS_PROJECT_NAME
export OS_USER_DOMAIN_NAME=Default
export OS_USERNAME=admin1
export OS_PASSWORD=abcd1234
export OS_COMPUTE_API_VERSION=2.46
export OS_NETWORK_API_VERSION=2.0
export OS_IMAGE_API_VERSION=2
export OS_VOLUME_API_VERSION=3
```

Before using the OpenStack CLI run `source powervcrc` to set the environment variables.

```
# source powervcrc
```

Retrieve information from PowerVC as shown in Example 6-6 on page 271.

Example 6-6 Retrieve information from PowerVC

The Compute Template:

```
# openstack
(openstack) flavor list -c ID -c Name
+-----+
| ID                                     | Name |
+-----+-----+
| 03fe52be-9f40-4c9d-aa9e-3574587115a5 | vscsi |
| 12697f5b-624c-4668-b42d-b32f5f447009 | xlarge |
| 1fa27097-f521-498e-9d8c-2ce686cf5660 | medium |
| 2072be21-8e1c-4fe2-bd35-f93c7bd39d87 | small |
| 561e4132-99c7-42cb-8f0f-d52ce09e84bd | for-deploy |
| 78a5f874-487a-4488-b0ae-7a204581c0f8 | xxlarge |
| 7917f79e-8d1f-4f28-94bf-487ce22b9468 | large |
| 90b9fa41-fe80-4351-a226-2501d95f1492 | tiny |
| beca1dfa-c7df-49e2-ad3d-fe84bd59e6dd | was |
+-----+-----+
```

We will use for-deploy compute template / flavor.

The Network:

```
(openstack) network list -c ID -c Name
+-----+-----+
| ID                                     | Name |
+-----+-----+
| 4c6fe8ae-e845-4c40-9be2-c1c4adb14313 | VLAN354 |
| efb5bc05-0575-4ec2-9da3-8eca4e666dd9 | VL354 |
+-----+-----+
```

We will use VL354 Network.

The Image:

```
(openstack) image list -c ID -c Name
+-----+-----+
| ID                                     | Name |
+-----+-----+
| 1ad0f2bb-266b-4efc-af03-f32264a12226 | AIX71 |
| 86c82681-9dff-419f-a138-a6af2704f94c | AIX71TL4 |
| f234cb56-3776-4eb9-8100-4bf736d13d1 | RHEL83 |
+-----+-----+
```

We will use AIX71 Image.

A basic Terraform template to deploy a VM uses the `OpenStack_compute_instance_v2` resource. A sample Terraform “all in one” template is shown in Example 6-7, which creates a single VM with a single user-defined network.

Example 6-7 Terraform template to create an IBM PowerVC partition

```
provider "openstack" {
  user_name     = "admin1"
  password     = "abcd1234"
  tenant_name  = "powervc"
  domain_name  = "Default"
  auth_url     = "https://10.3.54.17:5000/v3/"
  insecure     = true
}
```

```

}

resource "openstack_compute_instance_v2" "my-VM" {
  name      = "terra1"
  image_name = "AIX71"
  flavor_name = "for-deploy"

  network {
    name = "VL354"
  }
}

```

Terraform has a planning step where it generates an execution plan. The execution plan shows what Terraform will do when you call apply. This lets you avoid any surprises when Terraform manipulates the infrastructure. The **terraform plan** command is shown in Example 6-8.

Example 6-8 Terraform execution plan

```
# terraform plan
```

An execution plan has been generated and is shown below.
Resource actions are indicated with the following symbols:
+ create

Terraform will perform the following actions:

```

# openstack_compute_instance_v2.my-VM will be created
+ resource "openstack_compute_instance_v2" "my-VM" {
  + access_ip_v4      = (known after apply)
  + access_ip_v6      = (known after apply)
  + all_metadata      = (known after apply)
  + all_tags          = (known after apply)
  + availability_zone  = (known after apply)
  + flavor_id         = (known after apply)
  + flavor_name       = "for-deploy"
  + force_delete      = false
  + id                = (known after apply)
  + image_id          = (known after apply)
  + image_name        = "AIX71"
  + name              = "terra1"
  + power_state       = "active"
  + region            = (known after apply)
  + security_groups   = (known after apply)
  + stop_before_destroy = false

  + network {
    + access_network = false
    + fixed_ip_v4    = (known after apply)
    + fixed_ip_v6    = (known after apply)
    + floating_ip    = (known after apply)
    + mac            = (known after apply)
    + name           = "VL354"
    + port           = (known after apply)
    + uuid           = (known after apply)
  }
}

```

Plan: 1 to add, 0 to change, 0 to destroy.

Note: You didn't specify an "-out" parameter to save this plan, so Terraform can't guarantee that exactly these actions will be performed if "terraform apply" is subsequently run.

The **terraform apply** command is used to apply the changes required to reach the desired state of the configuration, or the pre-determined set of actions generated by a terraform plan execution plan. How to apply the template is shown in Example 6-9.

Example 6-9 Apply the Terraform template

```
# terraform apply
```

An execution plan has been generated and is shown below.
Resource actions are indicated with the following symbols:
+ create

Terraform will perform the following actions:

```
# openstack_compute_instance_v2.my-VM will be created
+ resource "openstack_compute_instance_v2" "my-VM" {
  + access_ip_v4      = (known after apply)
  + access_ip_v6      = (known after apply)
  + all_metadata      = (known after apply)
  + all_tags          = (known after apply)
  + availability_zone = (known after apply)
  + flavor_id         = (known after apply)
  + flavor_name       = "for-deploy"
  + force_delete      = false
  + id                = (known after apply)
  + image_id          = (known after apply)
  + image_name        = "AIX71"
  + name              = "terra1"
  + power_state       = "active"
  + region            = (known after apply)
  + security_groups   = (known after apply)
  + stop_before_destroy = false

  + network {
    + access_network = false
    + fixed_ip_v4    = (known after apply)
    + fixed_ip_v6    = (known after apply)
    + floating_ip    = (known after apply)
    + mac            = (known after apply)
    + name           = "VL354"
    + port           = (known after apply)
    + uuid           = (known after apply)
  }
}
```

Plan: 1 to add, 0 to change, 0 to destroy.

Do you want to perform these actions?
Terraform will perform the actions described above.
Only 'yes' will be accepted to approve.

Enter a value: yes

```
openstack_compute_instance_v2.my-VM: Creating...
openstack_compute_instance_v2.my-VM: Still creating... [10s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [20s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [30s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [40s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [50s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [1m0s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [1m10s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [1m20s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [1m30s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [1m40s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [1m50s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [2m0s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [2m10s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [2m20s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [2m30s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [2m40s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [2m50s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [3m0s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [3m10s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [3m20s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [3m30s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [3m40s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [3m50s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [4m0s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [4m10s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [4m20s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [4m30s elapsed]
openstack_compute_instance_v2.my-VM: Creation complete after 4m33s
[id=15a93a3f-8ca7-45ed-a406-5686b0e708af]
```

Apply complete! Resources: 1 added, 0 changed, 0 destroyed.

The **terraform state list** command is used to list resources within a Terraform state:

```
# terraform state list
openstack_compute_instance_v2.my-VM
```

The **terraform show** command is used to provide human-readable output from a state. This can be used to inspect the current state of a specific resource as shown in Example 6-10.

Example 6-10 Inspect current state of a resource

```
# terraform state show openstack_compute_instance_v2.my-VM
# openstack_compute_instance_v2.my-VM:
resource "openstack_compute_instance_v2" "my-VM" {
  access_ip_v4      = "10.3.54.126"
  all_metadata     = {
    "enforce_affinity_check" = "false"
    "move_pin_vm"           = "false"
  }
}
```

```

        "original_host"          = "900942A_7800980"
    }
    all_tags                    = []
    availability_zone           = "Default Group"
    flavor_id                   = "561e4132-99c7-42cb-8f0f-d52ce09e84bd"
    flavor_name                  = "for-deploy"
    force_delete                 = false
    id                           = "15a93a3f-8ca7-45ed-a406-5686b0e708af"
    image_id                     = "1ad0f2bb-266b-4efc-af03-f32264a12226"
    image_name                   = "AIX71"
    name                         = "terra1"
    power_state                  = "active"
    region                       = "RegionOne"
    security_groups              = []
    stop_before_destroy         = false

    network {
        access_network = false
        fixed_ip_v4    = "10.3.54.126"
        mac            = "fa:ac:23:99:c4:20"
        name           = "VL354"
        uuid           = "efb5bc05-0575-4ec2-9da3-8eca4e666dd9"
    }
}

```

Check that the VM has been correctly deployed in PowerVC with the **openstack server list** command as shown in Example 6-11.

Example 6-11 List VMs in PowerVC

```

#(openstack) server list --name terra1 -c Name -c Status -c Networks -c Image -c
Flavor
+-----+-----+-----+-----+-----+
| Name   | Status | Networks           | Image | Flavor   |
+-----+-----+-----+-----+-----+
| terra1 | ACTIVE | VL354=10.3.54.127 | AIX71 | for-deploy |
+-----+-----+-----+-----+-----+

```

You can use the **terraform destroy** command to destroy the VM that has been deployed in PowerVC as shown in Example 6-12.

Example 6-12 Destroy VM with Terraform

```

# terraform destroy

```

An execution plan has been generated and is shown below.
Resource actions are indicated with the following symbols:
- destroy

Terraform will perform the following actions:

```

# openstack_compute_instance_v2.my-VM will be destroyed
- resource "openstack_compute_instance_v2" "my-VM" {
  - access_ip_v4      = "10.3.54.127" -> null
  - all_metadata     = {

```

```

    - "enforce_affinity_check" = "false"
    - "move_pin_vm"           = "false"
    - "original_host"         = "900942A_7800980"
  } -> null
- all_tags                   = [] -> null
- availability_zone          = "Default Group" -> null
- flavor_id                  = "561e4132-99c7-42cb-8f0f-d52ce09e84bd" -> null
- flavor_name                 = "for-deploy" -> null
- force_delete               = false -> null
- id                         = "57829ea1-1fe5-44a5-a5c4-b942564bd5bd" -> null
- image_id                   = "1ad0f2bb-266b-4efc-af03-f32264a12226" -> null
- image_name                  = "AIX71" -> null
- name                       = "terra1" -> null
- power_state                 = "active" -> null
- security_groups             = [] -> null
- stop_before_destroy        = false -> null

- network {
  - access_network = false -> null
  - fixed_ip_v4    = "10.3.54.127" -> null
  - mac            = "fa:a2:01:3b:b3:20" -> null
  - name           = "VL354" -> null
  - uuid           = "efb5bc05-0575-4ec2-9da3-8eca4e666dd9" -> null
}
}

```

Plan: 0 to add, 0 to change, 1 to destroy.

Do you really want to destroy all resources?

Terraform will destroy all your managed infrastructure, as shown above.
There is no undo. Only 'yes' will be accepted to confirm.

Enter a value: yes

```

openstack_compute_instance_v2.my-VM: Destroying...
[id=57829ea1-1fe5-44a5-a5c4-b942564bd5bd]
openstack_compute_instance_v2.my-VM: Still destroying...
[id=57829ea1-1fe5-44a5-a5c4-b942564bd5bd, 10s elapsed]
openstack_compute_instance_v2.my-VM: Still destroying...
[id=57829ea1-1fe5-44a5-a5c4-b942564bd5bd, 20s elapsed]
openstack_compute_instance_v2.my-VM: Still destroying...
[id=57829ea1-1fe5-44a5-a5c4-b942564bd5bd, 30s elapsed]
openstack_compute_instance_v2.my-VM: Still destroying...
[id=57829ea1-1fe5-44a5-a5c4-b942564bd5bd, 40s elapsed]
openstack_compute_instance_v2.my-VM: Still destroying...
[id=57829ea1-1fe5-44a5-a5c4-b942564bd5bd, 50s elapsed]
openstack_compute_instance_v2.my-VM: Still destroying...
[id=57829ea1-1fe5-44a5-a5c4-b942564bd5bd, 1m0s elapsed]
openstack_compute_instance_v2.my-VM: Still destroying...
[id=57829ea1-1fe5-44a5-a5c4-b942564bd5bd, 1m10s elapsed]
openstack_compute_instance_v2.my-VM: Destruction complete after 1m15s

```

Destroy complete! Resources: 1 destroyed.

6.1.3 Specify the Storage Connectivity Group

You can specify which Storage Connectivity Group (SCG) you will use to deploy your VM by using Flavors extra-specs **powervm:storage_connectivity_group** as describe in the documentation:

https://www.ibm.com/support/knowledgecenter/en/SSVSPA_2.0.0/com.ibm.powervc.cloud.help.doc/powervc_pg_flavorsextraspecs_cloud.html

1. Retrieve the ID ebf26f02-1cb0-4828-afd8-f2de1f380ef6 that corresponds to Any host, all VIOS SCG from PowerVC GUI as illustrated in Figure 6-1.

Storage connectivity group: Any host, all VIOS				
Details *	VIOS members *	Fabrics	IPG	Member health
Name * Any host, all VIOS				
Boot volume connectivity type * NPIV				
Data volume connectivity type NPIV				
Match FC Port tag No restriction				
You may create a new port tag here .				
<input type="checkbox"/> Include untagged ports				
Details				
ID	ebf26f02-1cb0-4828-afd8-f2de1f380ef6			
System generated	Yes			
Boot health	OK			
Data health	OK			
Fabric access requirements met	Yes			
Data VIOS redundancy requirements met	Yes			
Boot VIOS redundancy requirements met	Yes			

Figure 6-1 Storage connectivity group

2. Modify the Terraform template to create a new flavor102 flavor that specify the SCG and deploy a VM using it as shown in Example 6-13.

Example 6-13 Add Storage Connectivity Group in Flavor

```
# vi main.tf

provider "openstack" {
  user_name   = "admin1"
  password   = "abcd1234"
  tenant_name = "powervc"
  domain_name = "Default"
  auth_url    = "https://10.3.54.17:5000/v3/"
  insecure    = true
}

resource "openstack_compute_flavor_v2" "flavor102" {
  name      = "flavor102"
  ram       = 2048
  vcpus     = 1
  disk      = 1
}
```

```

        extra_specs = {"powervm:storage_connectivity_group":
"ebf26f02-1cb0-4828-afd8-f2de1f380ef6" }
    }

resource "openstack_compute_instance_v2" "my-VM" {
    name      = "terra2"
    image_id  = "86c82681-9dff-419f-a138-a6af2704f94c"
    flavor_id = openstack_compute_flavor_v2.flavor102.id
    network {
        uuid = "4c6fe8ae-e845-4c40-9be2-c1c4adb14313"
        name = "VL354"
    }
}

```

3. Create an execution plan as shown in Example 6-14.

Example 6-14 SCG Terraform execution plan

```
# terraform plan
```

An execution plan has been generated and is shown below.

Resource actions are indicated with the following symbols:

```
+ create
```

Terraform will perform the following actions:

```

# openstack_compute_flavor_v2.flavor102 will be created
+ resource "openstack_compute_flavor_v2" "flavor102" {
  + disk          = 1
  + extra_specs  = {
    + "powervm:storage_connectivity_group" =
"ebf26f02-1cb0-4828-afd8-f2de1f380ef6"
  }
  + id            = (known after apply)
  + name         = "flavor102"
  + ram          = 2048
  + region       = (known after apply)
  + rx_tx_factor = 1
  + vcpus        = 1
}

# openstack_compute_instance_v2.my-VM will be created
+ resource "openstack_compute_instance_v2" "my-VM" {
  + access_ip_v4      = (known after apply)
  + access_ip_v6      = (known after apply)
  + all_metadata      = (known after apply)
  + all_tags          = (known after apply)
  + availability_zone = (known after apply)
  + flavor_id         = (known after apply)
  + flavor_name       = (known after apply)
  + force_delete      = false
  + id                = (known after apply)
  + image_id          = "86c82681-9dff-419f-a138-a6af2704f94c"
  + image_name        = (known after apply)
  + name              = "terra2"
  + power_state       = "active"
  + region            = (known after apply)
}

```

```

+ security_groups      = (known after apply)
+ stop_before_destroy = false

+ network {
  + access_network = false
  + fixed_ip_v4    = (known after apply)
  + fixed_ip_v6    = (known after apply)
  + floating_ip    = (known after apply)
  + mac            = (known after apply)
  + name           = "VL354"
  + port           = (known after apply)
  + uuid           = "4c6fe8ae-e845-4c40-9be2-c1c4adb14313"
}
}

```

Plan: 2 to add, 0 to change, 0 to destroy.

Note: You didn't specify an "-out" parameter to save this plan, so Terraform can't guarantee that exactly these actions will be performed if "terraform apply" is subsequently run.

4. Run **terraform apply** to make the changes in the infrastructure as defined in the plan, the **-auto-approve** option tells terraform not to require interactive approval of the plan before applying it, as shown in Example 6-15.

Example 6-15 Apply with auto approve option the Terraform plan

```

# terraform apply --auto-approve

openstack_compute_flavor_v2.flavor102: Creating...
openstack_compute_flavor_v2.flavor102: Still creating... [10s elapsed]
openstack_compute_flavor_v2.flavor102: Still creating... [20s elapsed]
openstack_compute_flavor_v2.flavor102: Still creating... [30s elapsed]
openstack_compute_flavor_v2.flavor102: Creation complete after 34s
[id=628a9bf4-215b-4fb3-bc60-caf01ba027ee]
openstack_compute_instance_v2.my-VM: Creating...
openstack_compute_instance_v2.my-VM: Still creating... [10s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [20s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [30s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [40s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [50s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [1m0s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [1m10s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [1m20s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [1m30s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [1m40s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [1m50s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [2m0s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [2m10s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [2m20s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [2m30s elapsed]
openstack_compute_instance_v2.my-VM: Still creating... [2m40s elapsed]
openstack_compute_instance_v2.my-VM: Creation complete after 2m43s
[id=3b4d523d-9c79-4c1c-b33e-fdb7e647c1f1]

```

Apply complete! Resources: 2 added, 0 changed, 0 destroyed.

The **terraform state list** command is used to list resources within a Terraform state:

```
# terraform state list
openstack_compute_flavor_v2.flavor102
openstack_compute_instance_v2.my-VM
```

Inspect the current state of the flavor resource as shown in Example 6-16.

Example 6-16 Inspect current state of the flavor

```
# terraform state show openstack_compute_flavor_v2.flavor102
# openstack_compute_flavor_v2.flavor102:
resource "openstack_compute_flavor_v2" "flavor102" {
  disk          = 1
  ephemeral     = 0
  extra_specs   = {
    "powervm:storage_connectivity_group" =
"ebf26f02-1cb0-4828-afd8-f2de1f380ef6"
  }
  id            = "628a9bf4-215b-4fb3-bc60-caf01ba027ee"
  is_public     = false
  name          = "flavor102"
  ram           = 2048
  region        = "RegionOne"
  rx_tx_factor  = 1
  swap          = 0
  vcpus         = 1
}
```

6.1.4 Specify the Storage Template for the boot volume

You can specify which Storage Template ID to use for virtual machine boot volume by using the Flavors extra-specs **powervm:boot_volume_type** as describe in the documentation https://www.ibm.com/support/knowledgecenter/en/SSVSPA_2.0.0/com.ibm.powervc.cloud.help.doc/powervc_pg_flavorsextraspecs_cloud.html

1. Retrieve the ID **d69b95d3-9eeb-4e9f-ba68-72662de2e5ca** that correspond to the SVC base template Storage Template as shown in Example 6-17.

Example 6-17 Retrieve Storage Template ID

```
# openstack volume type list
```

ID	Name	Is Public
0c9b2b17-5094-450c-bdd8-3d6ddd8f854ec	Pool_F840	True
5a838f6f-db1a-422e-902f-1530102af628	VMRCLS base template	True
d69b95d3-9eeb-4e9f-ba68-72662de2e5ca	SVC base template	True

2. Modify the Terraform template to create a new flavor102 flavor that specifies the Storage Template and deploy a VM using it as shown in Example 6-18 on page 281.

Example 6-18 Add Storage Template in Flavor

```
# vi main.tf

resource "openstack_compute_flavor_v2" "flavor102" {
  name      = "flavor102"
  ram       = 2048
  vcpus     = 1
  disk      = 1
  extra_specs = {"powervm:storage_connectivity_group":
"ebf26f02-1cb0-4828-afd8-f2de1f380ef6",
                "powervm:proc_units" : 0.1,
                "powervm:boot_volume_type":
"d69b95d3-9eeb-4e9f-ba68-72662de2e5ca"
}
}
```

3. Create an execution plan as shown in Example 6-19.

Example 6-19 Storage Template Terraform execution plan

```
# terraform plan
An execution plan has been generated and is shown below.
Resource actions are indicated with the following symbols:
+ create
terraform plan
```

An execution plan has been generated and is shown below.
Resource actions are indicated with the following symbols:
+ create

Terraform will perform the following actions:

```
# openstack_compute_flavor_v2.flavor102 will be created
+ resource "openstack_compute_flavor_v2" "flavor102" {
+ disk      = 1
+ extra_specs = {
+   "powervm:boot_volume_type" =
"d69b95d3-9eeb-4e9f-ba68-72662de2e5ca"
+   "powervm:proc_units" = "0.1"
+   "powervm:storage_connectivity_group" =
"ebf26f02-1cb0-4828-afd8-f2de1f380ef6"
}
+ id          = (known after apply)
+ name        = "flavor102"
+ ram         = 2048
+ region      = (known after apply)
+ rx_tx_factor = 1
+ vcpus       = 1
}
...
Output truncated
...
Plan: 2 to add, 0 to change, 0 to destroy.
```

Note: You didn't specify an "-out" parameter to save this plan, so Terraform can't guarantee that exactly these actions will be performed if "terraform apply" is subsequently run.

4. Run terraform apply to make the changes in the infrastructure as defined in the plan, as shown in Example 6-20.

Example 6-20 Create a VM using a specific Storage Template

```
# terraform apply --auto-approve

openstack_compute_flavor_v2.flavor102: Creating...
openstack_compute_flavor_v2.flavor102: Creation complete after 2s
[id=5d923886-44a7-4f60-9283-dc1654c0d689]
openstack_compute_instance_v2.my-VM: Creating...
openstack_compute_instance_v2.my-VM: Still creating... [10s elapsed]
openstack_compute_flavor_v2.flavor102: Creating...
...
Output truncated
...
openstack_compute_instance_v2.my-VM: Still creating... [2m0s elapsed]
openstack_compute_instance_v2.my-VM: Creation complete after 2m3s
[id=80cded7d-4739-4b3d-8518-06c2bcd7ee47]

Apply complete! Resources: 2 added, 0 changed, 0 destroyed.
```

5. List the resources and show their current state as shown in Example 6-21.

Example 6-21 List resources in terraform state

```
# terraform state list
openstack_compute_flavor_v2.flavor102
openstack_compute_instance_v2.my-VM

# terraform state show openstack_compute_flavor_v2.flavor102
# openstack_compute_flavor_v2.flavor102:
resource "openstack_compute_flavor_v2" "flavor102" {
  disk          = 1
  ephemeral    = 0
  extra_specs  = {
    "powervm:boot_volume_type" =
"d69b95d3-9eeb-4e9f-ba68-72662de2e5ca"
    "powervm:proc_units"      = "0.1"
    "powervm:storage_connectivity_group" =
"ebf26f02-1cb0-4828-afd8-f2de1f380ef6"
  }
}
...
Output truncated
...
vcpus          = 1
}
```

6.2 Ansible automation

Ansible is an open source IT configuration management, deployment, and orchestration tool. It delivers productivity gains via automation, to enable faster time-to-market for IT projects with predictable and consistent deployments. Ansible supports the orchestration of multi-tier workloads, including infrastructure, networking, operating systems, applications, and services across hundreds of nodes. Ansible's automation language Yet Another Markup Language (YAML) is written in easy-to-understand Playbooks that allow your entire organization to benefit from automation.

Ansible is agent-less and performs actions on a set of servers from an Ansible control node where the Ansible engine is installed. It communicates with your servers securely by using the OpenSSH protocol.

As PowerVC is built on top of OpenStack technology, you can use the available OpenStack modules included with Ansible to deploy VMs as illustrated in Figure 6-2.

Openstack modules included with Ansible:

- ▶ `os_image_info` to retrieve the image information
- ▶ `os_flavor_info` to retrieve flavor information
- ▶ `os_networks_info` to retrieve network information
- ▶ `os_server` to provision a new virtual machine (VM)
- ▶ `os_volume` to provision a new volume
- ▶ `os_server_volume` to attach a volume to a VM

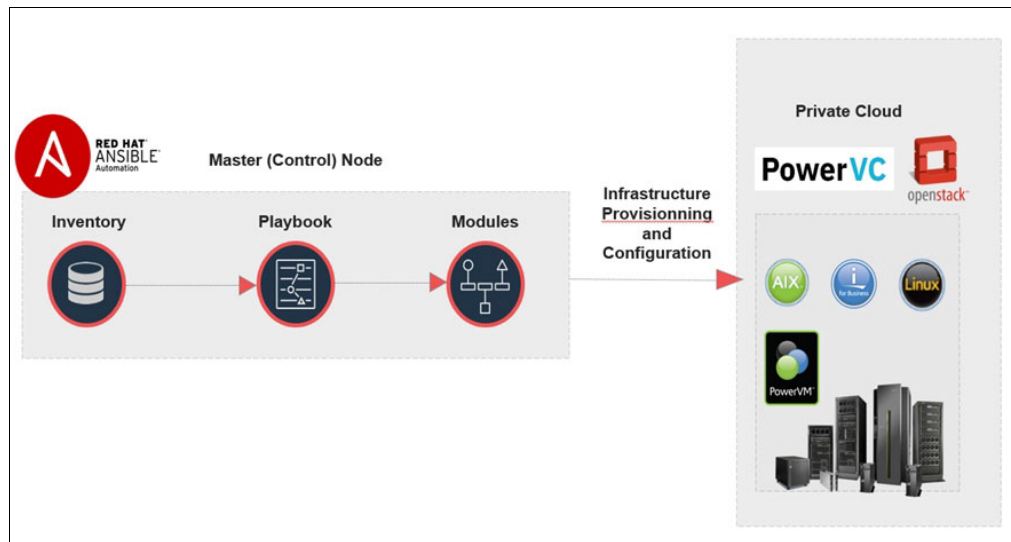


Figure 6-2 Ansible and PowerVC

6.2.1 Install and configure Ansible

We install the Ansible engine on the IBM PowerVC Server by following this instructions:

1. Enable the Ansible Engine repository

```
# subscription-manager repos --enable=ansible-2.9-for-rhel-8-ppc64le-rpms  
Repository 'ansible-2.9-for-rhel-8-ppc64le-rpms' is enabled for this system.
```

2. Install Ansible 2.9 as shown in Example 6-22.

Example 6-22 Installation of Ansible

```
# yum install ansible
Red Hat Enterprise Linux 8 for Power, little endian - BaseOS (RPMs)          14 kB/s | 4.1 kB    00:00
Red Hat Ansible Engine 2.9 for RHEL 8 ppc64le (RPMs)                    325 kB/s | 1.2 MB   00:03
Red Hat Enterprise Linux 8 for Power, little endian - AppStream (RPMs)   17 kB/s | 4.5 kB    00:00
Dependencies resolved.
=====
Package           Architecture  Version                Repository              Size
=====
Installing:
ansible           noarch       2.9.17-1.el8ae        ansible-2.9-for-rhel-8-ppc64le-rpms 17 M

Transaction Summary
=====
Install 1 Package

Total download size: 17 M
Installed size: 96 M
Is this ok [y/N]: y
Downloading Packages:
ansible-2.9.17-1.el8ae.noarch.rpm                                         4.9 MB/s | 17 MB    00:03
-----
Total                                                                       4.9 MB/s | 17 MB    00:03
Running transaction check
Transaction check succeeded.
Running transaction test
Transaction test succeeded.
Running transaction
Preparing :                                                                1/1
Installing : ansible-2.9.17-1.el8ae.noarch                               1/1
Running scriptlet: ansible-2.9.17-1.el8ae.noarch                        1/1
Verifying  : ansible-2.9.17-1.el8ae.noarch                               1/1
Installed products updated.

Installed:
  ansible-2.9.17-1.el8ae.noarch

Complete!
```

Verify that the Ansible Engine version is 2.9.17:

```
# ansible --version
ansible 2.9.17
  config file = /etc/ansible/ansible.cfg
  configured module search path = ['/root/.ansible/plugins/modules', '/usr/share/ansible/plugins/modules']
  ansible python module location = /usr/lib/python3.6/site-packages/ansible
  executable location = /usr/bin/ansible
  python version = 3.6.8 (default, Dec 5 2019, 16:11:43) [GCC 8.3.1 20191121 (Red Hat 8.3.1-5)]
```

6.2.2 Create a playbook to retrieve PowerVC informations

The fundamental component for the deployment of an image within Ansible is a playbook. Ansible playbooks are lists of tasks that are automatically execute against hosts. Each module within an Ansible playbook performs a specific task. Each module contains metadata that determines when and where a task is executed. We can use Ansible Playbook to define how we want new logical partition / Virtual machine (LPAR / VM) deployed through IBM PowerVC by providing information, such as:

- ▶ A compute template to use
- ▶ A host group to which to deploy
- ▶ Network adapters to attach
- ▶ Extra disks to allocate

You can use a playbook to retrieve this information on the PowerVC Server, The Ansible OpenStack module needs several pieces of information about the PowerVC deployment in order to work. You can set environment variables before using the Ansible playbook. PowerVC includes the `/opt/ibm/powervc/powervcrc` file to simplify this task. Copy the file to `/home/ansible` directory and edit it as necessary as shown in Example 6-23. In particular, set these values:

- ▶ `OS_USERNAME` - Your user name
- ▶ `OS_PASSWORD` - Your password
- ▶ `OS_PROJECT_NAME` - Your project
- ▶ `OS_AUTH_URL` - PowerVC URL

Example 6-23 Configure the OpenStack environment for Ansible

```
# mkdir /home/ansible
# cd /home/ansible
# cp /opt/ibm/powervc/powervcrc /home/ansible/powervcrc
# vi powervcrc

export OS_IDENTITY_API_VERSION=3
export OS_AUTH_URL=https://10.3.54.17:5000/v3/
export OS_CACERT=/etc/pki/tls/certs/powervc.crt
export OS_REGION_NAME=RegionOne
export OS_PROJECT_DOMAIN_NAME=Default
export OS_PROJECT_NAME=powervc
export OS_TENANT_NAME=$OS_PROJECT_NAME
export OS_USER_DOMAIN_NAME=Default
export OS_USERNAME=admin1
export OS_PASSWORD=abcd1234
export OS_COMPUTE_API_VERSION=2.46
export OS_NETWORK_API_VERSION=2.0
export OS_IMAGE_API_VERSION=2
export OS_VOLUME_API_VERSION=3
```

Before using Ansible playbooks run **source powervcrc** to set the environment variables:

```
# source powervcrc
```

A sample playbook `all in one` is shown in Example 6-24, which retrieves information about the AIX71 image, the for-deploy flavor and the VL354 network.

Example 6-24 Playbook to retrieve PowerVC information

```
# vi list_powervc.yml
- name: List available PowerVC Networks
  hosts: localhost
  tasks:
    - name: Retrieve VL354 network
      os_networks_info:
        register: result

    - name: Print VL354 Network
      debug:
        msg: "{{ result | json_query('openstack_networks[?name==`VL354`].
          {name: name, id: id}') }}"
```

```

- name: Retrieve for-deploy flavor
  os_flavor_info:
  register: result

- name: Print for-deploy flavor
  debug:
    msg: "{{ result | json_query('openstack_flavors[?name==`for-deploy`].
      {name: name, id: id, ram: ram, vcpus: vcpus}') }}"

- name: Retrieve AIX71 image
  os_image_info:
  properties:
    os_distro: aix
  register: result

- name: Print AIX71 image
  debug:
    msg: "{{ result | json_query('openstack_image[?name==`AIX71`].
      {name: name, id: id, os_distro: os_distro,
      project: location.project.name}') }}"

```

After the YAML file is created, run the Ansible playbook, as shown in Example 6-25.

Example 6-25 Execution output of playbook

```

# ansible-playbook list_powervc.yml
[WARNING]: provided hosts list is empty, only localhost is available. Note that
the implicit localhost does not
match 'all'

PLAY [List available PowerVC Networks]
*****

TASK [Gathering Facts]
*****
**
ok: [localhost]

TASK [Retrieve VL354 network]
*****
**
ok: [localhost]

TASK [Print VL354 Network]
*****
**
ok: [localhost] => {
  "msg": [
    {
      "id": "efb5bc05-0575-4ec2-9da3-8eca4e666dd9",
      "name": "VL354"
    }
  ]
}

TASK [Retrieve for-deploy flavor]

```

```

ok: [localhost]

TASK [Print for-deploy flavor]
*****
ok: [localhost] => {
  "msg": [
    {
      "id": "561e4132-99c7-42cb-8f0f-d52ce09e84bd",
      "name": "for-deploy",
      "ram": 8192,
      "vcpus": 1
    }
  ]
}

TASK [Retrieve AIX71 image]
*****
ok: [localhost]

TASK [Print AIX71 image]
*****
ok: [localhost] => {
  "msg": [
    {
      "id": "1ad0f2bb-266b-4efc-af03-f32264a12226",
      "name": "AIX71",
      "os_distro": "aix",
      "project": "powervc"
    }
  ]
}

PLAY RECAP
*****
localhost                : ok=7    changed=0    unreachable=0    failed=0
skipped=0    rescued=0    ignored=0

```

6.2.3 Create a playbook for an AIX deployment

Use the `os_server` module to deploy an `ansible1` VM using the AIX71 image, the `for-deploy` flavor and the VL354 Network as shown in Example 6-26.

Example 6-26 VM Deploy playbook

```

# vi deployvm.yml
- name: Create a PowerVC Virtual Machine
  hosts: localhost
  tasks:
    - name: Create a new VM instance
      os_server:
        state: present
        name: ansible1
        image: AIX71
        flavor: for-deploy
        nics:
          - net-name: VLAN354

```

Run the playbook to deploy the ansible1 VM as shown in Example 6-27.

Example 6-27 Deploy VM with Ansible playbook

```
# ansible-playbook deployvm.yml

PLAY [Create a PowerVC Virtual Machine]
*****

TASK [Gathering Facts]
*****
ok: [localhost]

changed: [localhost]

PLAY RECAP
*****
localhost      : ok=2    changed=1    unreachable=0    failed=0
skipped=0     rescued=0    ignored=0
```

Check that the VM has been correctly deployed on PowerVC with the **openstack server list** command as shown in Example 6-28.

Example 6-28 List VM on PowerVC

```
#(openstack) server list --name ansible1 -c Name -c Status -c Networks -c Image -c
Flavor
+-----+-----+-----+-----+-----+
| Name   | Status | Networks           | Image | Flavor   |
+-----+-----+-----+-----+-----+
| ansible1 | ACTIVE | VL354=10.3.54.134 | AIX71 | for-deploy |
+-----+-----+-----+-----+-----+
```

6.2.4 Create an Ansible ready VM with an extra disk and ssh keys

Use the `os_server` module to deploy an `ansible2` VM using the `AIX71` image, the `for-deploy` flavor, the `VL354` Network, use the `os_volume` and `os_server_volume` modules to attach a new disk. This VM is ready to be used as an Ansible target as shown in Example 6-29.

Example 6-29 Playbook to create a VM with extra disk and ssh keys

```
# vi vmdiskssh.yml
- name: Create a PowerVC Virtual Machine
  hosts: localhost
  tasks:

  - name: Create a PowerVC SSH Key Pair
    os_keypair:
      state: present
      name: powervc-ssh-key
      public_key_file: "{{ ansible_env.HOME }}/.ssh/id_rsa.pub"

  - name: Create a new VM instance
    os_server:
      timeout: 300
```

```

    state: present
    name: ansible2
    image: AIX71
    flavor: for-deploy
    key_name: powervc-ssh-key
    nics:
      - net-name: VL354
  register: vm

- name: create a new 40G volume
  os_volume:
    state: present
    size: 40
    display_name: ansible2disk
  register: disk

- name: attach volume to VM
  os_server_volume:
    state: present
    server: ansible2
    volume: ansible2disk

- name: Print VM's public IP address
  debug:
    var: "{{ vm.openstack.public_v4 }}"

- name: Waits for SSH port 22 to open
  wait_for:
    delay: 10
    host: "{{ vm.openstack.public_v4 }}"
    port: 22
    timeout: 300

- name: Add VM host key to known hosts
  shell: "ssh-keyscan -H {{ vm.openstack.public_v4 }} >>
/root/.ssh/known_hosts"

```

Run the playbook to deploy the ansible2 VM as shown in Example 6-30.

Example 6-30 Deploy VM with extra disk and ssh keys

```

# ansible-playbook vmdiskssh.yml
[WARNING]: provided hosts list is empty, only localhost is available. Note that
the implicit localhost does not match
'all'

PLAY [Create a PowerVC Virtual Machine]
*****

TASK [Gathering Facts]
*****
ok: [localhost]

TASK [Create a PowerVC SSH Key Pair]
*****
ok: [localhost]

```

```

TASK [Create a new VM instance]
*****
changed: [localhost]

TASK [create a new 40G volume]
*****
changed: [localhost]

TASK [attach volume to VM]
*****
changed: [localhost]

TASK [Print VM's public IP address]
*****
ok: [localhost] => {
    "10.3.54.123"
}

TASK [Waits for SSH port 22 to open]
*****
ok: [localhost]

TASK [Add VM host key to known hosts]
*****
changed: [localhost]

PLAY RECAP
*****
localhost      : ok=8   changed=4   unreachable=0   failed=0
skipped=0     rescued=0   ignored=0

PLAY RECAP
*****
localhost      : ok=2   changed=1   unreachable=0   failed=0
skipped=0     rescued=0   ignored=0

```

Check that the VM has been correctly deployed on PowerVC with the **openstack server list** command as shown in Example 6-31.

Example 6-31 List VM on PowerVC

```

# openstack server list --name ansible2 -c Name -c Status -c Networks -c Image -c
Flavor
+-----+-----+-----+-----+-----+
| Name   | Status | Networks          | Image | Flavor   |
+-----+-----+-----+-----+-----+
| ansible2 | ACTIVE | VL354=10.3.54.123 | AIX71 | for-deploy |
+-----+-----+-----+-----+-----+

```

Check that you can connect to the VM from the PowerVC server with **ssh** and that there is a 40 GB extra disk available as shown in Example 6-32 on page 291.

Example 6-32 Connect to the deployed VM

```
# ssh 10.3.54.123 cfmgr
# ssh 10.3.54.123 lspv
hdisk0          00cff3b40a6acc17          rootvg          active
hdisk1          none                               None
# ssh 10.3.54.123 bootinfo -s hdisk1
40960
```



IBM PowerVC in the Hybrid Cloud

This chapter provides information about the IBM PowerVC requirements for Red Hat OpenShift on IBM Power Systems. Other topics covered include how to use IBM PowerVC storage with IBM Cloud Private and how to deploy a custom image on IBM Public Cloud within a Power Systems Virtual Server. It also provides an overview on IBM Cloud Pak® for Multicloud Management.

This chapter includes the following topics:

- ▶ 7.1, “Hybrid cloud overview” on page 294
- ▶ 7.2, “Red Hat OpenShift Container Platform” on page 294
- ▶ 7.3, “IBM PowerVC and IBM Power Systems Virtual Server” on page 308
- ▶ 7.4, “IBM Cloud Pak for Multicloud Management Version 2.2” on page 312

7.1 Hybrid cloud overview

The technology world changes day by day and requires companies and IT departments to adopt new approaches and different mindsets. The separation of the different cloud types and deployment models tends to turn into an obstacle instead of being a feature itself. Hence the need of different combinations and agile approaches. The integration of the traditional IT environments is now enhanced with the opportunities provided by the cloud computing models such as infrastructure as a service (IaaS), platform as a service (PaaS), and software as a service (SaaS), and the mix of cloud types – private, public, and hybrid.

The power of hybrid cloud lays on the philosophy of taking the best of both worlds – the security of the private cloud and the scalability of the public cloud. Technical leaders seek to use the public cloud services for higher amounts of data, and private cloud services for sensitive loads and higher security. Computing resources are no more a challenge as the hybrid cloud allows access to any public cloud services, provided and paid only when needed. To deliver value with the hybrid cloud, IBM uses the approach of a common container platform called *Red Hat OpenShift* which takes care of consistency, ease of management and ensures high security standards. IBM PowerVC, together with Red Hat OpenShift and IBM Cloud Paks is a combination to be designed once and then can be managed from anywhere in the world.

7.2 Red Hat OpenShift Container Platform

The Red Hat OpenShift Container Platform is one of the most reliable enterprise-grade hybrid cloud container platform. It is designed and optimized to easily deploy web applications and cloud services anywhere and from any cloud. Categorized as a cloud development Platform as a Service (PaaS), it is based on industry standards such as Kubernetes, Open Container Initiative (OCI), and Container Runtime Interface (CRI). With the help of IBM PowerVC you can rapidly deploy OpenShift clusters on IBM Power Systems enterprise servers to help modernize your existing workloads.

The following sections provide IBM PowerVC requirements for Red Hat OpenShift (version 4.3 and above) on IBM Power Systems and describe how to create Red Hat Enterprise Linux 8.2 (or later) and Red Hat CoreOS (RHCOS) images in PowerVC.

Red Hat OpenShift for IBM Power Systems is an enterprise-grade platform that provides a secure, private platform-as-a-service cloud on IBM Power Systems servers and was officially announced by IBM on April 28, 2020. For additional information, see the [announcement letter](#) for Red Hat OpenShift V4.3 on Power Systems.

For more information on Red Hat OpenShift and IBM Cloud Paks on IBM Power Systems, see

Red Hat OpenShift V4.3 on IBM Power Systems Reference Guide, REDP-5599 and *Red Hat OpenShift and IBM Cloud Paks on IBM Power Systems: Volume 1*, SG24-8459.

PowerVC prerequisites

As a requirement for Red Hat OpenShift, you will need to create Red Hat Enterprise Linux 8.2 (or later) and Red Hat CoreOS (RHCOS) images in your IBM PowerVC environment.

For creating a Red Hat Enterprise Linux 8.2 (or later) image in IBM PowerVC Version 2.0.0, see section 4.16, “Images” on page 210.

Follow these steps to create a Red Hat CoreOS (RHCOS) image via IBM PowerVC Version 2.0.0:

1. Download the latest RHCOS image with suffix "**-openstack.ppc64le.qcow2.gz**" (at the time of writing, the latest RHCOS version is 4.7.0) from the following [link](#) on a system with public internet access. You will need to transfer this image to an existing Red Hat Enterprise Linux 8.x VM, or a VM that you will be guided to create in the next step.
2. Log into PowerVC menu. You will then have the option to create and attach additional empty volumes with minimum size of 120 GB to an existing Red Hat Enterprise Linux 8.x VM. You may alternatively choose to deploy a new Red Hat Enterprise Linux 8.x VM having an additional empty volume with minimum size of 120 GB, as shown in Figure 7-1 and Figure 7-2. Make a note of the new volume name.

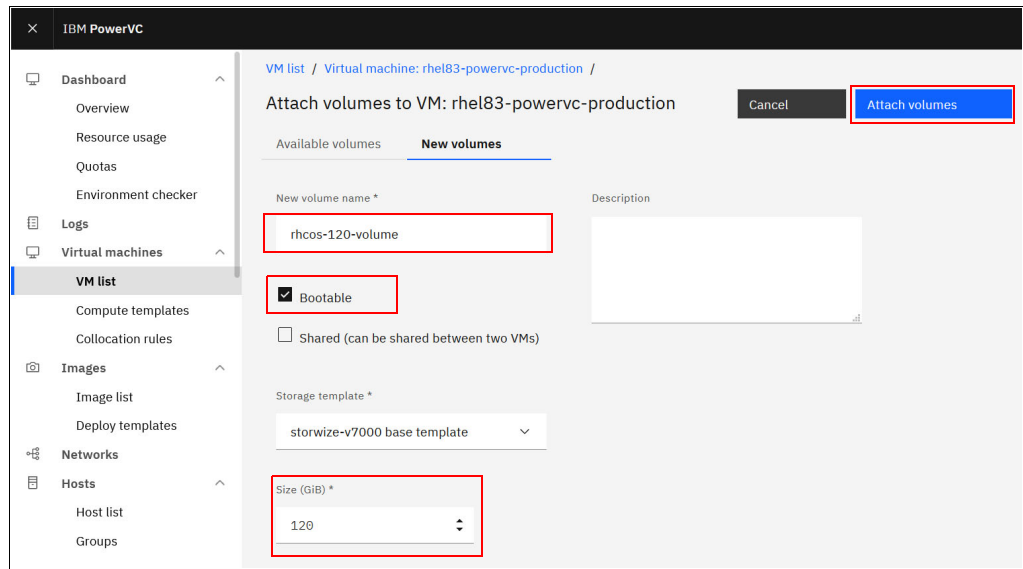


Figure 7-1 Create and attach an additional empty volume with 120 GB

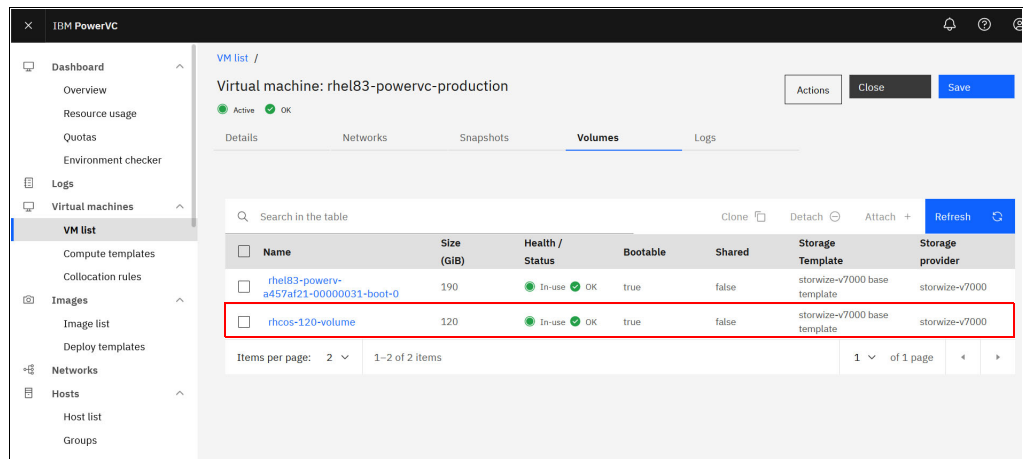


Figure 7-2 Volume successfully attached

3. Log into the VM and follow the steps below:

- a. Install `wget`, `qemu-img`, `parted` and `gzip` packages
- ii. Transfer the downloaded RHCOS image to this VM
- iii. Extract the image by using the command as shown in Example 7-1:

Example 7-1 Extract CoreOS image

```
# gunzip rhcos-openstack.ppc64le.qcow2.gz

# ls -l
total 2430656
-rw-r--r-- 1 root root 2488991744 Feb 18 08:56 rhcos-openstack.ppc64le.qcow2
```

- b. Convert the CoreOS `qcow2` image to raw image as shown in Example 7-2:

Example 7-2 Convert CoreOS qcow2 to raw image

```
# qemu-img convert -f qcow2 -o raw rhcos-openstack.ppc64le.qcow2
rhcos-openstack.ppc64le.raw
```

- c. Identify the disk device representing the additional empty volume that is attached to the VM. Run the commands as shown in Example 7-3:

Example 7-3 Rescanning the bus and listing the disk devices

```
# sudo rescan-scsi-bus.sh -a -m -r

# lsblk
NAME                MAJ:MIN RM  SIZE RO TYPE  MOUNTPOINT
loop0                 7:0    0   7.6G  1 loop  /mnt
loop1                 7:1    0    2G   0 loop  /srv/node/partition1
sda                   8:0    0  190G   0 disk
  ··mpatha            253:0    0  190G   0 mpath
    ··mpatha1         253:1    0    4M   0 part
    ··mpatha2         253:2    0    1G   0 part  /boot
    ··mpatha3         253:3    0  189G   0 part
      ··rhel-root     253:4    0    70G   0 lvm   /
      ··rhel-swap     253:5    0    13G   0 lvm   [SWAP]
      ··rhel-home     253:6    0  106G   0 lvm   /home
sdb                   8:16    0  190G   0 disk
  ··mpatha            253:0    0  190G   0 mpath
    ··mpatha1         253:1    0    4M   0 part
    ··mpatha2         253:2    0    1G   0 part  /boot
    ··mpatha3         253:3    0  189G   0 part
      ··rhel-root     253:4    0    70G   0 lvm   /
      ··rhel-swap     253:5    0    13G   0 lvm   [SWAP]
      ··rhel-home     253:6    0  106G   0 lvm   /home
sdc                  8:32    0  120G  0 disk
  ··mpathd            253:7    0  120G   0 mpath
sdd                   8:48    0  120G   0 disk
  ··mpathd            253:7    0  120G   0 mpath
sr0                   11:0    1   7.6G   0 rom
```

- d. Verify the newly created raw disk image, and dump it to the newly added disk. In our case we use multipath mappings, the newly discovered devices are `sdc` or `sdd`, as shown in Example 7-4 on page 297:

IMPORTANT: Pay special attention when using the `dd` command and its syntax. Failure to do so might permanently destroy the disk data.

Example 7-4 Verifying the raw disk and dd the image

```
# file -k rhcos-openstack.ppc64le.raw
rhcos-latest.raw: DOS/MBR boot sector; partition 1 : ID=0xee, start-CHS
(0x0,0,2), end-CHS (0x3ff,255,63), startsector 1, 33554431 sectors, extended
partition table (last) DOS/MBR boot sector GPT partition table, version 1.0,
GUID: 00000000-0000-4000-a000-000000000001, disk size: 33554432 sectors of
512 bytes (Lepton 3.x), scale 0-0, spot sensor temperature 0.000000, unit
celsius, color scheme 0, calibration: offset 0.000000, slope 0.000000
(Lepton 2.x), scale 0-0, spot sensor temperature 0.000000, unit celsius,
color scheme 0, calibration: offset 0.000000, slope 0.000000\012- data

# dd if=rhcos-openstack.ppc64le.raw of=/dev/sdc bs=4M
4096+0 records in
4096+0 records out
17179869184 bytes (17 GB, 16 GiB) copied, 26.1357 s, 657 MB/s

# lsblk
NAME                MAJ:MIN RM  SIZE RO TYPE  MOUNTPOINT
loop0                 7:0    0   7.6G  1 loop  /mnt
loop1                 7:1    0    2G   0 loop  /srv/node/partition1
sda                   8:0    0  190G   0 disk
  ··mpatha            253:0   0  190G   0 mpath
    ··mpatha1         253:1   0    4M   0 part
    ··mpatha2         253:2   0    1G   0 part  /boot
    ··mpatha3         253:3   0  189G   0 part
      ··rhel-root     253:4   0    70G   0 lvm   /
      ··rhel-swap     253:5   0    13G   0 lvm   [SWAP]
      ··rhel-home     253:6   0  106G   0 lvm   /home
sdb                   8:16   0  190G   0 disk
  ··mpatha            253:0   0  190G   0 mpath
    ··mpatha1         253:1   0    4M   0 part
    ··mpatha2         253:2   0    1G   0 part  /boot
    ··mpatha3         253:3   0  189G   0 part
      ··rhel-root     253:4   0    70G   0 lvm   /
      ··rhel-swap     253:5   0    13G   0 lvm   [SWAP]
      ··rhel-home     253:6   0  106G   0 lvm   /home
sdc                   8:32   0  120G   0 disk
  ··sdc1              8:33   0    4M   0 part
  ··sdc3              8:35   0  384M   0 part
  ··sdc4              8:36   0   3.1G   0 part
  ··mpathd            253:7   0  120G   0 mpath
sdd                   8:48   0  120G   0 disk
  ··mpathd            253:7   0  120G   0 mpath
sr0                   11:0    1   7.6G   0 rom
```

- e. Detach the volume from the VM by selecting it, then click **Detach** as shown in Figure 7-3 on page 298:

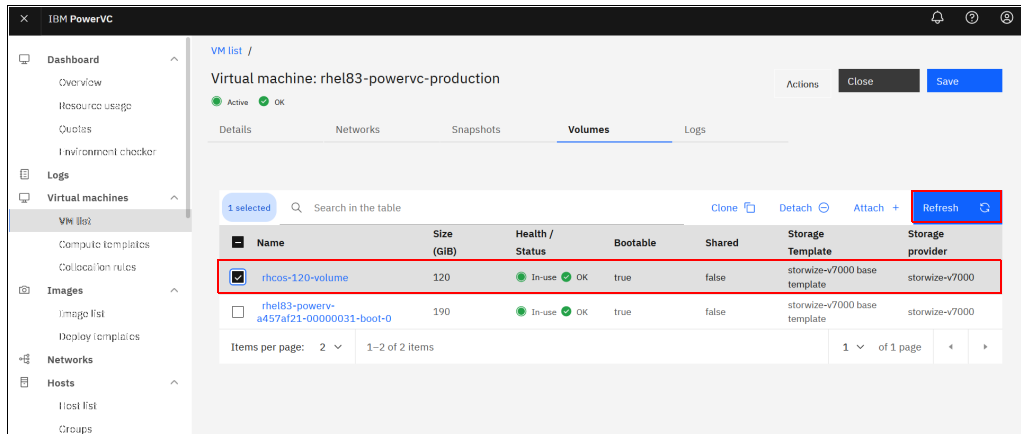


Figure 7-3 Detaching a volume from PowerVC 2.0.0 UI.

- f. Run the commands shown in Example 7-5 to rescan the bus and list the disk devices.

Example 7-5 Rescanning the bus and listing the disk devices

```
# sudo rescan-scsi-bus.sh -a -m -r
Syncing file systems
Scanning SCSI subsystem for new devices and remove devices that have
disappeared
.....
Attempting to update multipath devices...
Removing multipath mappings for removed devices if all paths are now
failed...
Flushing multipath device mpathd... Done (0 retries)
0 new or changed device(s) found.
0 remapped or resized device(s) found.
2 device(s) removed.
    [1:0:0:1]
    [1:0:1:1]

# lsblk
NAME                MAJ:MIN RM  SIZE RO TYPE  MOUNTPOINT
loop0                 7:0    0  7.6G  1 loop  /mnt
loop1                 7:1    0    2G   0 loop  /srv/node/partition1
sda                   8:0    0  190G  0 disk
├─mpatha              253:0    0  190G  0 mpath
│ ├─mpatha1           253:1    0    4M   0 part
│ ├─mpatha2           253:2    0    1G   0 part  /boot
│ └─mpatha3           253:3    0  189G  0 part
│   ├─rhel-root       253:4    0    70G  0 lvm   /
│   ├─rhel-swap       253:5    0    13G  0 lvm   [SWAP]
│   └─rhel-home       253:6    0  106G  0 lvm   /home
sdb                   8:16    0  190G  0 disk
├─mpatha              253:0    0  190G  0 mpath
│ ├─mpatha1           253:1    0    4M   0 part
│ ├─mpatha2           253:2    0    1G   0 part  /boot
│ └─mpatha3           253:3    0  189G  0 part
│   ├─rhel-root       253:4    0    70G  0 lvm   /
│   ├─rhel-swap       253:5    0    13G  0 lvm   [SWAP]
│   └─rhel-home       253:6    0  106G  0 lvm   /home
```

- g. Go to Images and click **Create** to create a new image from the PowerVC 2.0.0 GUI, as shown in Figure 7-4.

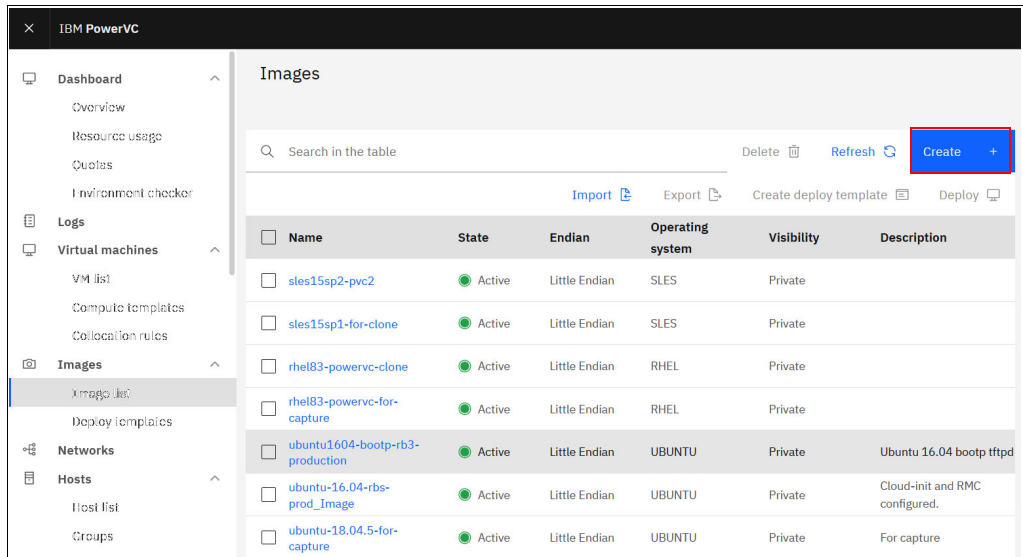


Figure 7-4 Creating a new Image for CoreOS

- h. Specify the image name and choose CoreOS from **Operating system** dropdown menu. In PowerVC Version 2.0.0, Little Endian is selected by default. See Figure 7-5.

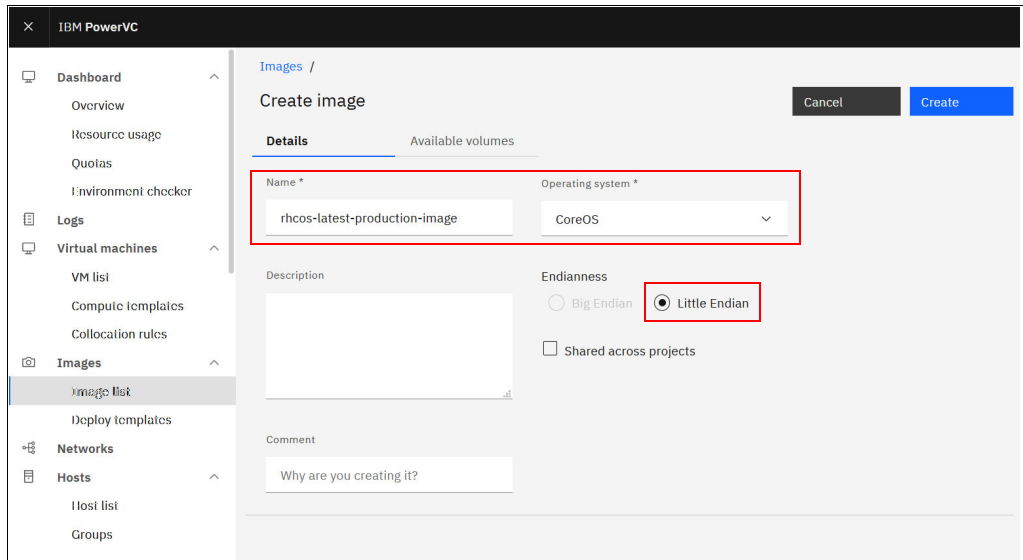


Figure 7-5 CoreOS image name and operating system

- i. Select **Available volumes** to add the volume where the RHCOS raw image was defined. Then search for the specific volume name and select it. Set the boot order to start with a 0 for the selected boot volumes, as shown in Figure 7-6 on page 300.

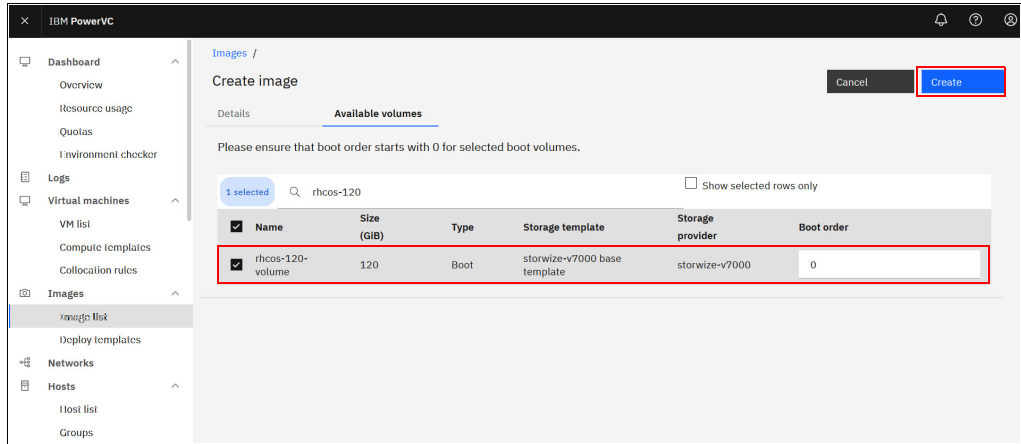


Figure 7-6 Attaching the volume where the CoreOS was defined

- j. Create the image by clicking on the **Create** button. See Figure 7-6.

Note: The Reliable Scalable Cluster Technology (RSCT) for Resource Monitoring Control (RMC) connections is supported on Red Hat Enterprise Linux CoreOS (RHCOS) as a DaemonSet in OpenShift cluster starting with OCP version 4.5 and above. The primary purpose of RMC in RHCOS is to ensure successful LPM and Active Health Status in dashboard. To dynamically (DLPAR) resize the processor, memory and I/O, we suggest rebooting the nodes to avoid issues with wrong resource accounting in OCP cluster (statement valid as of the date when this document was created). For additional details refer to the following [playbook](#).

Compute templates

You will need to create compute templates for bastion, bootstrap, control plane and compute nodes. For more information on how to create compute templates see Figure 7-7 on page 301, Figure 7-8 on page 301 and 2.4.4, “Information that is required for compute template planning” on page 34.

The following is the minimum recommended LPAR configuration (see Table 7-1) that you can use when creating the compute templates for different types of nodes (bastion, bootstrap, control plane, compute nodes):

Table 7-1 OpenShift cluster minimum resource requirements

Machine	Operating System	vCPU ^a	Virtual RAM	Storage
Bastion	Red Hat Enterprise Linux 8.x	2	32 GB	200 GB
Bootstrap	RHCOS	2	32 GB	120 GB
Control plane	RHCOS	2	32 GB	120 GB
Compute	RHCOS	2	32 GB	120 GB

- a. PowerVM LPARs by default use SMT=8. With 2 vCPUs, the number of logical CPUs seen by the operating system will be 16 (2 vCPUs x 8 SMT). This configuration is suitable for the majority of scenarios.

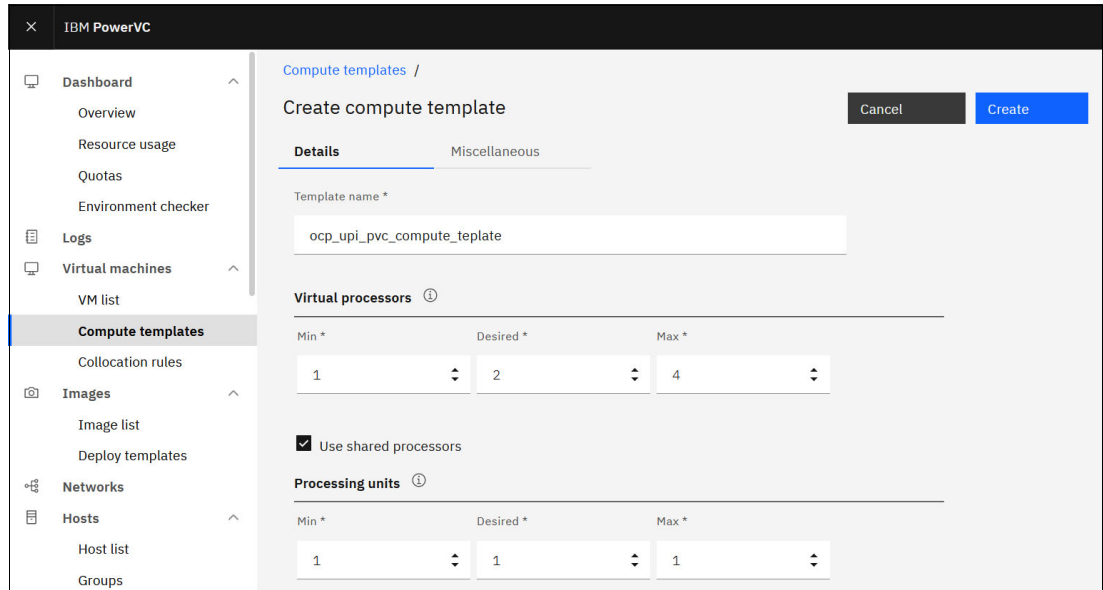


Figure 7-7 Creating a compute template from PowerVC 2.0.0 UI

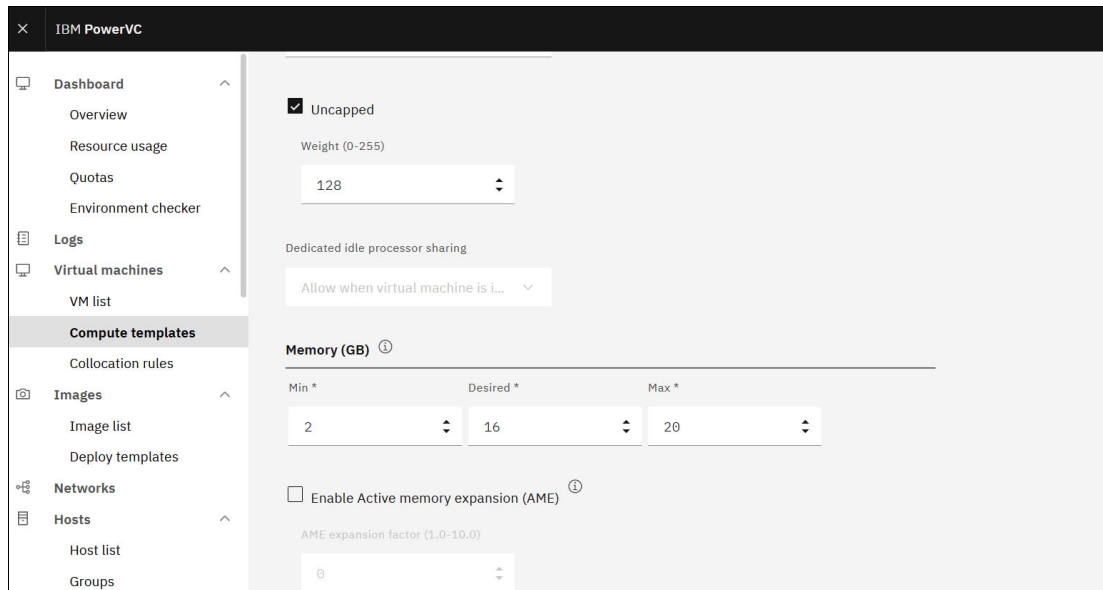


Figure 7-8 Configuring memory resources on a compute template

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 by utilizing the following operations: creating volumes, deleting volumes, and attaching or detaching volumes.

PowerVC CSI driver has two main components:

- ▶ Node plug-in
- ▶ Controller plug-in

The **Node plug-in** is deployed as a **DaemonSet** and runs on the node where the volumes are provisioned and consumed (on the worker node where application pods run). **Controller plug-in** provides the functionality of volume management outside of worker node provision. This plug-in implements the CSI controller service and sidecar containers like external-provisioner and external-attacher. External provisioner triggers create or delete volume operations based on Persistent Volume Claims (PVC). External **attacher** watches volume attachment objects and triggers volume attach (publish) or detach (unpublish) operations.

Requirements

Your system must meet the following requirements to use the PowerVC CSI pluggable driver:

- ▶ OpenShift environment on top of PowerVC (version 1.4.4.1 or later) based IBM Power infrastructure with storage registered.
- ▶ Red Hat OpenShift Container Platform version 4.3 or later.

Installing and configuring PowerVC CSI driver on Red Hat OpenShift

Follow these steps to configure PowerVC CSI driver on Red Hat OpenShift Container Platform:

1. Download the latest **yaml** files from [GitHub](#) location:
ibm-powervc-csi-driver-template.yaml, scc.yaml, and secret.yaml
2. Login to Red Hat OpenShift Container Platform.
3. Use the following command to list all nodes and ensure that they are working.

Example 7-6 List all nodes

```
#oc get nodes
NAME          STATUS    ROLES    AGE   VERSION
master-0     Ready    master   3h51m v1.19.0+2f3101c
master-1     Ready    master   3h53m v1.19.0+2f3101c
master-2     Ready    master   3h53m v1.19.0+2f3101c
worker-0     Ready    worker   3h40m v1.19.0+2f3101c
worker-1     Ready    worker   3h41m v1.19.0+2f3101c
```

4. Create a new project and switch to it.

Example 7-7 Creating a new project by using OCP command-line tool

```
# oc create namespace redbooks
namespace/redbooks created

# oc project redbooks
Now using project "redbooks" on server
"https://api.ocp-10.5.33.138.nip.io:6443".
```

5. Run these commands to allow cluster role for the project administrator.

Example 7-8 This command provides access to system service account

```
# oc adm policy add-cluster-role-to-user cluster-admin
system:serviceaccount:openshift-infra:template-instance-controller
clusterrole.rbac.authorization.k8s.io/cluster-admin added:
"system:serviceaccount:openshift-infra:template-instance-controller"
```

Example 7-9 This command provides access to system service account default

```
# oc adm policy add-cluster-role-to-user cluster-admin
system:serviceaccount:redbooks:default
clusterrole.rbac.authorization.k8s.io/cluster-admin added:
"system:serviceaccount:redbooks:default"
```

6. Download and apply Security Context Constraints (SCC) template.

Example 7-10 Downloading and applying scc.yaml file

```
# curl
https://raw.githubusercontent.com/IBM/power-openstack-k8s-volume-driver/master/
template/scc.yaml > scc.yaml
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload   Total   Spent    Left   Speed
100  998    100  998    0     0   7984      0  --:--:--  --:--:--  --:--:--  7984

# oc apply -f <directory path>/scc.yaml
securitycontextconstraints.security.openshift.io/csiaccess created
```

7. Next download **ibm-powervc-csi-driver-template.yaml**, make the following changes to the **yaml** template file and apply them:
 - a. For **OPENSTACK_CERT_DATA**, add attribute **value: ""** and copy / paste the contents of **powervc.crt** file between the quotation marks.
 - b. For **DRIVER_VOLUME_TYPE**, add attribute **value: ""** and enter the name or UUID of the storage template between the quotation marks.

Example 7-11 Downloading, editing and applying ibm-powervc-csi-driver-template.yaml file

```
# curl
https://raw.githubusercontent.com/IBM/power-openstack-k8s-volume-driver/master/
template/ibm-powervc-csi-driver-template.yaml >
ibm-powervc-csi-driver-template.yaml
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload   Total   Spent    Left   Speed
100 24420    100 24420    0     0  162k      0  --:--:--  --:--:--  --:--:--  162k

# vi ibm-powervc-csi-driver-template.yaml

oc apply -f ibm-powervc-csi-driver-template.yaml
template.template.openshift.io/ibm-powervc-csi created
```

8. Download **secret.yaml** and replace **OS_USERNAME** and **OS_PASSWORD** parameters of secret file with your PowerVC username and password encrypted in **base64** format.

Example 7-12 Downloading secret.yaml and replace parameters of secret file

```
# curl
https://raw.githubusercontent.com/IBM/power-openstack-k8s-volume-driver/master/
template/secret.yaml > secret.yaml
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
```

```

                                Dload Upload Total Spent Left Speed
100 130 100 130 0 0 460 0 ---:---:-- --:---:-- --:---:-- 460

# base64 <<< <enter-your-powervc-username>

# base64 <<< <enter-your-powervc-password>

# vi secret.yaml

```

-
9. Use the following command to apply this secret file into OpenShift.

Example 7-13 Applying secret.yaml file

```

# oc apply -f secret.yaml
secret/ocp-csi-secret created

```

-
10. Execute the following command to validate all parameters and default values.

Example 7-14 Validating all parameters and default values

```

# oc process --parameters ibm-powervc-csi

```

-
11. Process the template file.

12. Run the following command to create <a>.yaml consisting of all artifacts.

Note: Use PowerVC IP address and the **secret** that you created in steps 8 and 9, in our case we used “**ocp-csi-secret**”.

Example 7-15 Command below will create “a.yaml” that has all the artifacts in it

```

# oc process ibm-powervc-csi -p OPENSTACK_IP_OR_HOSTNAME=< PowerVC IP address >
-p OPENSTACK_CRED_SECRET_NAME=ocp-csi-secret > a.yaml

```

-
13. Install the OpenShift artifacts, by running the following command described below. See Example 7-16 and Figure 7-9 on page 305.

Example 7-16 Installing the OpenShift artifacts

```

# oc apply -f a.yaml
configmap/ibm-powervc-config created
storageclass.storage.k8s.io/ibm-powervc-csi-volume-default created
serviceaccount/ibm-powervc-csi-attacher created
serviceaccount/ibm-powervc-csi-resizer created
clusterrole.rbac.authorization.k8s.io/ibm-powervc-csi-attacher created
clusterrole.rbac.authorization.k8s.io/ibm-powervc-csi-resizer created
clusterrolebinding.rbac.authorization.k8s.io/ibm-powervc-csi-attacher-role
created
clusterrolebinding.rbac.authorization.k8s.io/ibm-powervc-csi-resizer-role
created
service/ibm-powervc-csi-attacher-plugin created
statefulset.apps/ibm-powervc-csi-attacher-plugin created

```

```

statefulset.apps/ibm-powervc-csi-resizer-plugin created
serviceaccount/ibm-powervc-csi-provisioner created
clusterrole.rbac.authorization.k8s.io/ibm-powervc-csi-provisioner created
clusterrolebinding.rbac.authorization.k8s.io/ibm-powervc-csi-provisioner-role
created
service/ibm-powervc-csi-provisioner-plugin created
statefulset.apps/ibm-powervc-csi-provisioner-plugin created
serviceaccount/ibm-powervc-csi-node created
clusterrole.rbac.authorization.k8s.io/ibm-powervc-csi-node created
clusterrolebinding.rbac.authorization.k8s.io/ibm-powervc-csi-node-role created
daemonset.apps/ibm-powervc-csi-plugin created
W0315 15:08:54.961193 29693 warnings.go:67] storage.k8s.io/v1beta1 CSIDriver
is deprecated in v1.19+, unavailable in v1.22+; use storage.k8s.io/v1 CSIDriver
W0315 15:08:54.975690 29693 warnings.go:67] storage.k8s.io/v1beta1 CSIDriver
is deprecated in v1.19+, unavailable in v1.22+; use storage.k8s.io/v1 CSIDriver
csidriver.storage.k8s.io/ibm-powervc-csi created

```

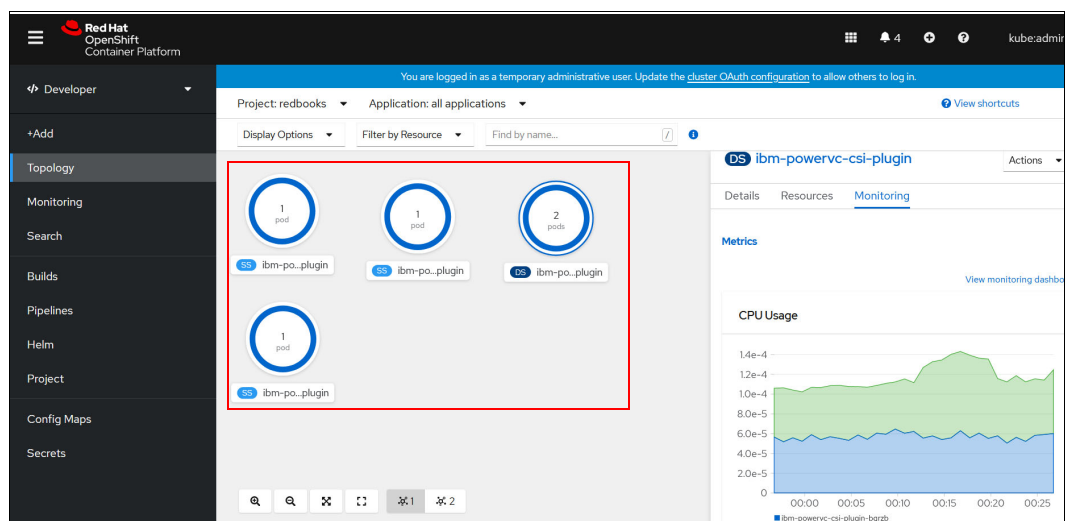


Figure 7-9 Configured PowerVC CSI driver on Red Hat OpenShift Container Platform.

14. Run the following command and wait until all CSI nodes are listed.

Example 7-17 List all CSI nodes (in our case worker-0 and worker-1)

```

# oc get csinodes
NAME          DRIVERS  AGE
master-0     0        4h45m
master-1     0        4h47m
master-2     0        4h47m
worker-0     1        4h35m
worker-1     1        4h35m

```

15. Run the following commands to ensure that no output errors are returned.

Example 7-18 Listing all pod and verifying that there are no errors in the output

```

# oc get pods
NAME                                     READY  STATUS   RESTARTS  AGE

```

```

ibm-powervc-csi-attacher-plugin-0      2/2      Running  0          72s
ibm-powervc-csi-plugin-bqrzb          3/3      Running  0          72s
ibm-powervc-csi-plugin-gknnn         3/3      Running  0          71s
ibm-powervc-csi-provisioner-plugin-0  2/2      Running  0          72s
ibm-powervc-csi-resizer-plugin-0     2/2      Running  0          72s

```

```
# oc logs <plugin> -c <container>
```

```
# oc logs ibm-powervc-csi-attacher-plugin-0 -c csi-attacher
```

```
# oc logs ibm-powervc-csi-plugin-bqrzb -c ibm-powervc-csi
```

```
# oc logs ibm-powervc-csi-plugin-gknnn -c ibm-powervc-csi
```

```
# oc logs ibm-powervc-csi-provisioner-plugin-0 -c csi-provisioner
```

```
oc logs ibm-powervc-csi-resizer-plugin-0 -c csi-resizer
```

16. Refer to the example scripts for PowerVC-CSI drivers available at [GitHub](#)

Note: RSCT package support is available for Red Hat CoreOS as a DaemonSet in OpenShift cluster starting with version 4.5 and above (statement valid as of the date when this document was created). The following command can be used in PowerVC to skip the RMC check during volume attach or detach operations:
powervc-config compute disable-rmc-check --enable --restart

17. To test the volume creation, download `static-pvc.yaml` file available at [GitHub](#)

Example 7-19 Downloading static-pvc.yaml

```

# curl
https://raw.githubusercontent.com/IBM/power-openstack-k8s-volume-driver/master/
csi_examples/static-pvc.yaml > static-pvc.yaml
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload   Total   Spent    Left   Speed
100  208  100  208    0     0    753      0  --:--:-- --:--:-- --:--:--   753

```

18. Run the following command to create a volume.

Example 7-20 Applying static-pvc.yaml file

```
# oc apply -f static-pvc.yaml
persistentvolumeclaim/my-claim created
```

19. Run the following commands to list all Persistent Volumes and all Persistent Volume Claims. See Figure 7-10 on page 307, Figure 7-11 on page 307 and Figure 7-12 on page 308.

Example 7-21 Listing all Persistent Volumes and Persistent Volume Claims

```
# oc get pv
```

```
# oc get pvc
```

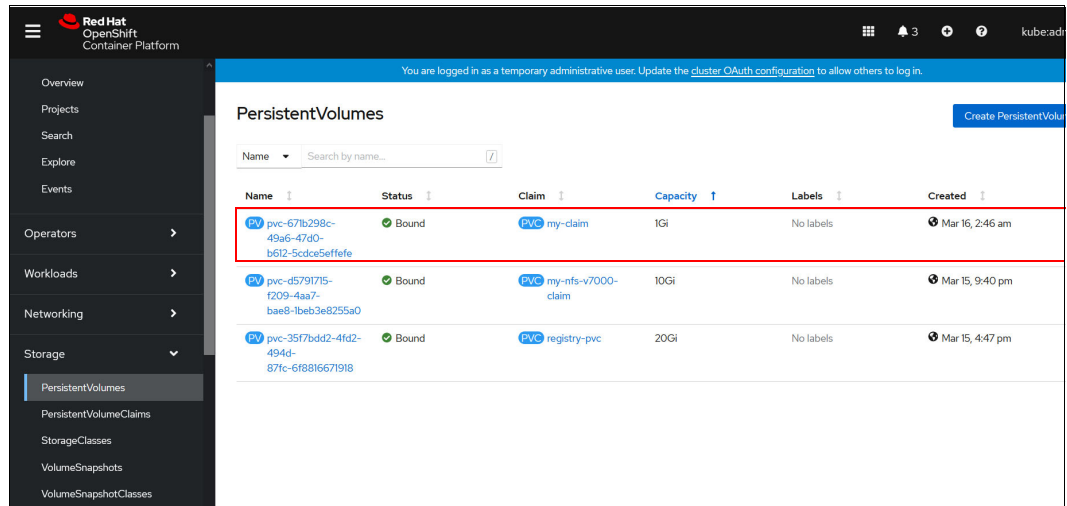


Figure 7-10 Listing all Persistent Volumes from OCP UI

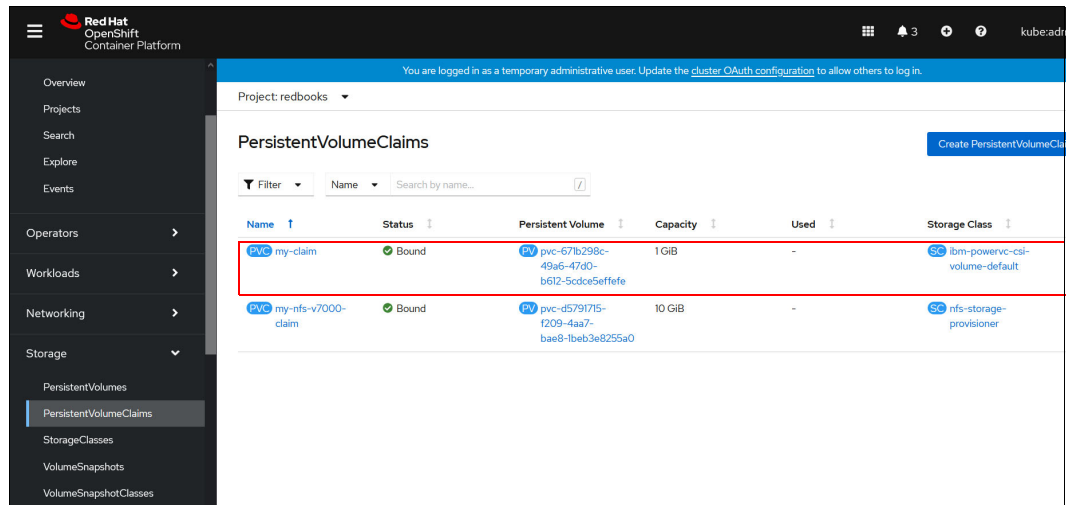


Figure 7-11 Listing all Persistent Volume Claims from OCP UI

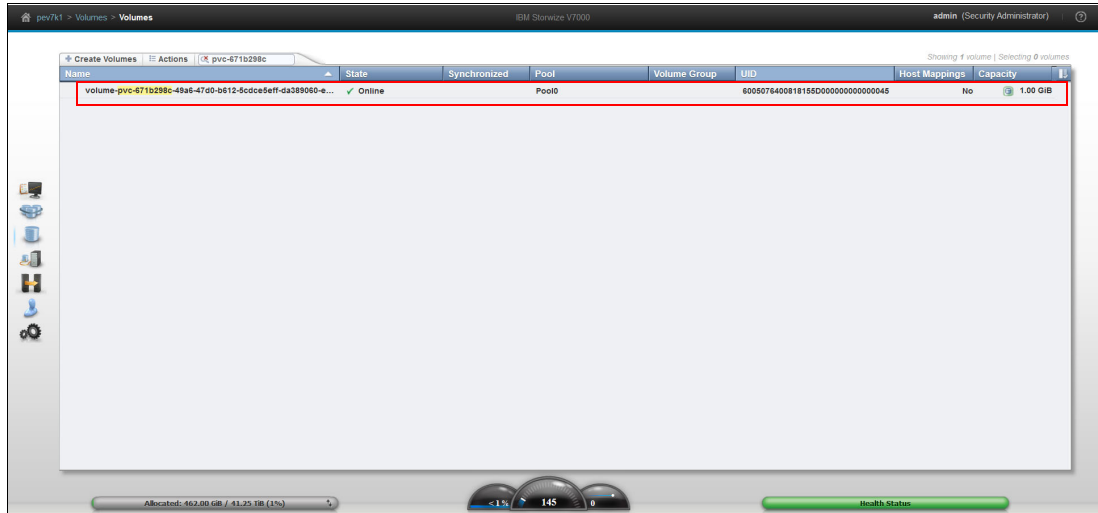


Figure 7-12 Listing the newly created Volume on Storage device

20. To delete the Persistent Volume Claim use the following command.

Example 7-22 Deleting the Persistent Volume Claim

```
oc delete pvc my-claim
persistentvolumeclaim "my-claim" deleted
```

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

https://www.ibm.com/support/knowledgecenter/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_csi_storage_install.html

For complete instructions on planning, configuring and deployment of Red Hat OpenShift and IBM Cloud Paks on IBM Power Systems, see:

- ▶ *Red Hat OpenShift V4.3 on IBM Power Systems Reference Guide*, REDP-5599
- ▶ *Red Hat OpenShift and IBM Cloud Paks on IBM Power Systems: Volume 1*, SG24-8459

7.3 IBM PowerVC and IBM Power Systems Virtual Server

IBM PowerVC for Private Cloud edition provides the feel and features of the public cloud. However, clients who have typically relied upon on-premises-only infrastructure can now quickly and economically extend their Power IT resources onto IBM Power Systems Virtual Server (IBM PowerVS). IBM PowerVS is a Power Systems enterprise Infrastructure as a service (IaaS) offering. IBM PowerVS is colocated and connected with IBM Cloud, integrating IBM AIX, IBM i and Linux capabilities. You can use IBM PowerVS to deploy VMs in a matter of minutes. Clients can benefit from the easy to scale up and out, and the self-service provisioning services that IBM PowerVS provides.

The combination of IBM PowerVC and IBM PowerVS provides clients with a true hybrid cloud solution for enterprise workloads. IBM Power Systems Virtual Server supports bringing your own custom AIX, IBM i or Linux image that you have tested and deployed. This means workloads being managed by IBM PowerVC can easily be captured, exported to and deployed on IBM Power Systems Virtual Server. In the following section we are going to be migrating an IBM AIX server to IBM Power Systems Virtual Server.

7.3.1 Migrating VMs to IBM PowerVS

There are several ways workloads can be migrated from on premise to IBM PowerVS. Some of the available options include IBM Cloud Object Storage (COS), IBM Mass Data Migration (MDM), IBM Aspera® Technologies. Application-specific replication, OS-based tools, third-party tools etc. can be also be used to migrate on-premise data to IBM PowerVS and back. This book shows how to achieve the same with IBM PowerVC.

IBM PowerVC images and COS can be used together to migrate workloads to IBM PowerVS and back using the export and import of images feature of IBM PowerVC discussed in 4.16.5, “Export images into OVA” on page 232 and 4.16.3, “Import an image from OVA or raw image” on page 226.

You can capture OVA images to easily migrate your data. The PowerVS offering allows you to provision a new virtual server based on an OVA image. To accomplish this, regardless of the operating system (OS), you must complete the following steps:

- ▶ Create an OVA image on your local system
- ▶ Copy the OVA image to your Cloud Object Storage account
- ▶ Deploy the OVA image and provision a new Power Systems Virtual Server

Before embarking on the migration certain steps have to be completed:

- ▶ IBM PowerVC host server must be zoned to all storage providers where volumes of the images are located
- ▶ The image of the VM to be migrated has to be captured
- ▶ You must have an IBM PowerVS account
- ▶ You should have a COS configured, with credentials and a bucket created

Note: Creation of an IBM PowerVS account and configuration of a COS are out of the scope of this book. For more information visit

<https://cloud.ibm.com/docs/power-iaas?topic=power-iaas-getting-started>

The steps involved in migrating a workload to IBM PowerVS from IBM PowerVC using COS are as follows:

1. Export captured imaged to an OVA file using the IBM PowerVC host command line interface
2. Install awscli
3. configure awscli
4. Access and transfer ova file to COS
5. Import OVA file in IBM PowerVS
6. Deploy OVA file as a VM in IBM PowerVS

upload: ./AIX72-TL05-VM_image.ova to s3://808ibmredbook/AIX72-TL05-VM_image.ova

The image file has been transferred to the COS. It can now be imported in IBM PowerVS and deployed as a virtual machine. The image is imported into IBM PowerVS by navigating to **Resource list-> services-> boot images**. Then click on **Import image** and fill in the self-explanatory details. Figure 7-13 shows the import image process.

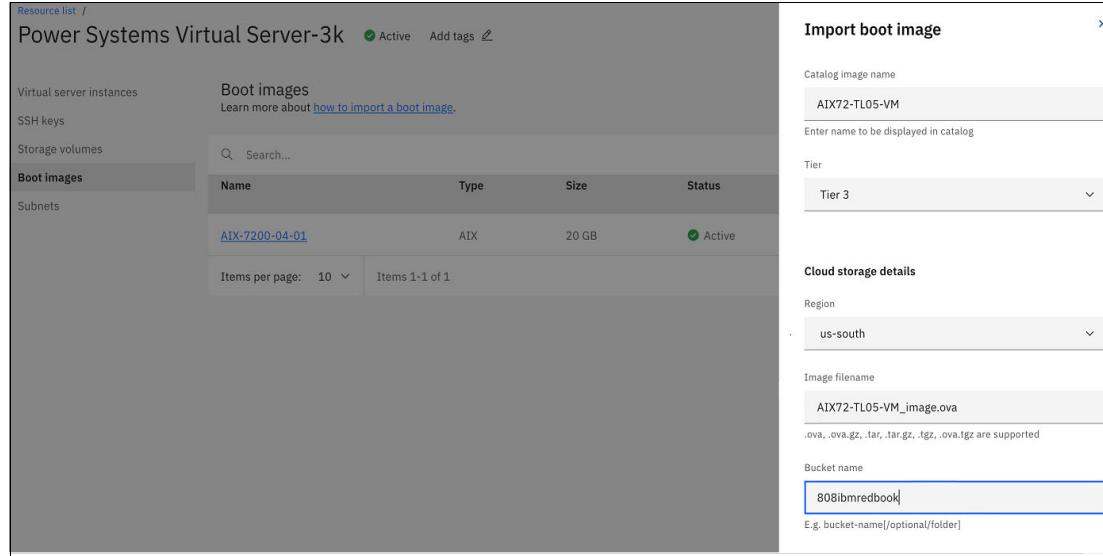


Figure 7-13 Importing an OVA file from COS to IBM PowerVS

4. The image has been successfully imported into IBM PowerVS and can be deployed. This is shown in Figure 7-14.

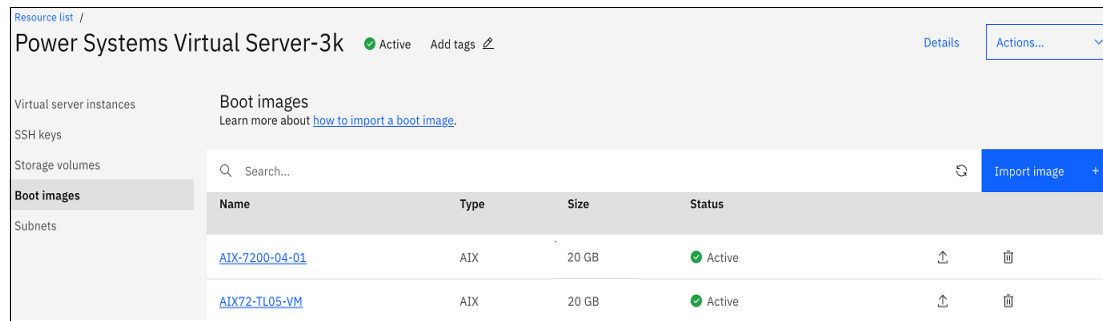


Figure 7-14 List of boot images in IBM PowerVS

In this section we have covered how workloads running on IBM PowerVC can easily be prepared and migrated to IBM's IaaS offering IBM PowerVS.

7.4 IBM Cloud Pak for Multicloud Management Version 2.2

IBM Cloud Pak for Multicloud Management provides consistent visibility, governance, and automation across a multitude of platforms. Whether you're looking at managing VMs, containers, public clouds, or an on-premises infrastructure, IBM Cloud Pak for Multicloud Management has the tools to ensure that your enterprise runs smoothly. IBM Cloud Pak for Multicloud Management includes Infrastructure management and Service management components.

Infrastructure management delivers the insight, control, and automation enterprises need to address the challenges of managing virtual environments, which are far more complex than physical ones. The Infrastructure management module was previously called IBM Red Hat CloudForms.

The Managed services module was previously called Terraform & Service Automation or IBM Cloud Automation Manager.

7.4.1 Configure PowerVC as an Openstack Cloud provider

In Infrastructure Management, a cloud provider is a cloud computing environment that you can add to Infrastructure Management to manage and interact with the resources in that environment. As PowerVC is based on Openstack you can add it to Infrastructure Management.

1. To add PowerVC as a Cloud Provider, navigate to menu: **Compute > Clouds > Providers**
2. Click **Configuration**, then click **Add a New Cloud Provider** as illustrated in Figure 7-15

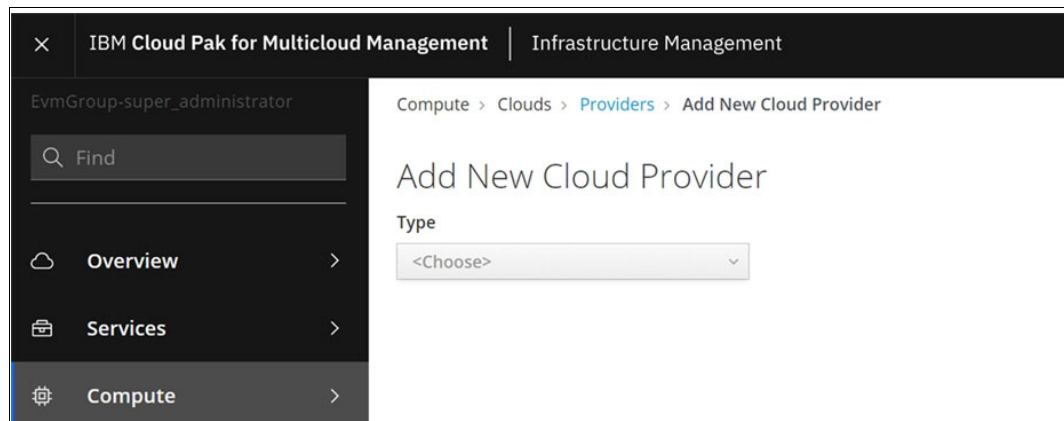


Figure 7-15 Add a New Cloud Provider

3. From the Type list, select **OpenStack**, enter a Name for the provider, select **keystone v3** for the API Version, enter **default** for the Domain ID as illustrated in Figure 7-16 on page 314.

Add New Cloud Provider

Type

* Name

* Zone

Provider Region

Openstack Infra Provider

* API Version

* Domain ID

Tenant Mapping Enabled
 OFF

Figure 7-16 PowerVC cloud provider 1/3

4. Select **SSL without validation** for the Security Protocol, enter IP address or hostname of PowerVC server, enter 5000 for the API Port, enter Username and Password to connect to PowerVC server, click **Validate** then **Add** as illustrated in Figure 7-17 on page 315.

Endpoints

* Security Protocol

SSL without validation

* Hostname (or IPv4 or IPv6 address)

10.3.54.17

* API Port

5000

* Username

admin1

* Password

.....

Validation successful

Figure 7-17 PowerVC cloud provider 2/3

The PowerVC server is now in the Cloud provider list as illustrated in Figure 7-18.

Cloud Providers

20 per page 1-1 of 1 1 of 1

<input type="checkbox"/>	Name	Type	EVM Zone	Instances	VCpus	Memory	Images	Tenant	Region
<input type="checkbox"/>	PowerVCdemo	OpenStack	default	4	4	36 GB	6	My Company	Region 0

Figure 7-18 PowerVC cloud provider 3/3

7.4.2 Manage VMs with Infrastructure management

To manage VMs with Infrastructure Management, navigate to menu: **Compute > Clouds > Instances** as illustrated in Figure 7-19.

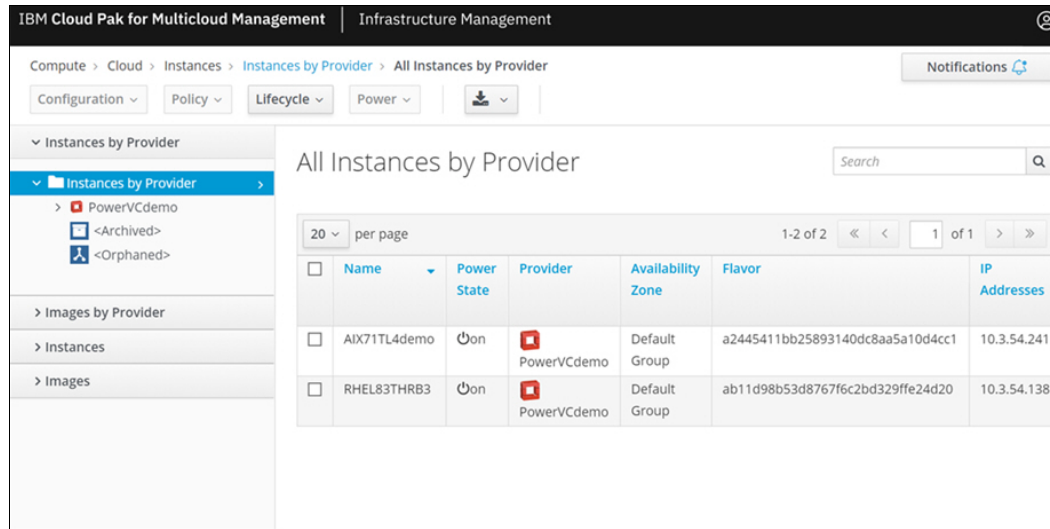


Figure 7-19 Manage VM

You can do VM Stop, Start, Delete, Reboot. Select a VM and Click **Stop** as illustrated in Figure 7-20.

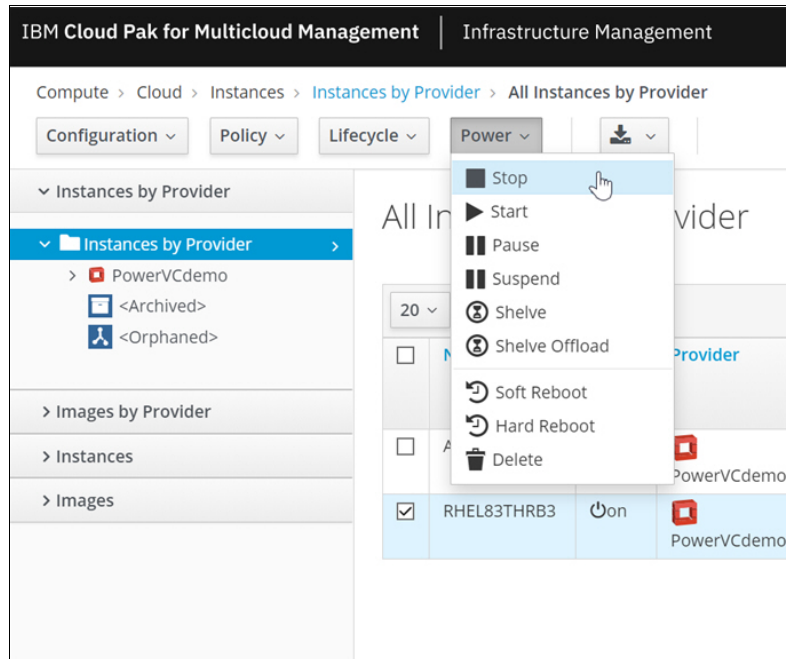


Figure 7-20 Stop VM



A

Useful commands and scripts

This appendix provides some useful commands and scripts, and covers the following topics:

- ▶ “Useful OpenStack commands” on page 318
- ▶ “Useful powervc-config commands” on page 320
- ▶ “Useful scripts” on page 324
- ▶ “REST APIs example” on page 325

Useful OpenStack commands

This section includes several OpenStack commands that you might find useful.

Prepare environment to use OpenStack command

You can use the **openstack** CLI command to retrieve information on the PowerVC server. It needs several pieces of information about the PowerVC in order to work. PowerVC includes the `/opt/ibm/powervc/powervcrc` file to simplify this task.

Before using the OpenStack CLI, edit the `powervcrc` file and run **source powervcrc** to set the environment variables as shown in Example A-1.

Example A-1 Source the OpenStack environment

```
# cat powervcrc

export OS_IDENTITY_API_VERSION=3
export OS_AUTH_URL=https://10.3.54.17:5000/v3/
export OS_CACERT=/etc/pki/tls/certs/powervc.crt
export OS_REGION_NAME=RegionOne
export OS_PROJECT_DOMAIN_NAME=Default
export OS_PROJECT_NAME=powervc
export OS_TENANT_NAME=$OS_PROJECT_NAME
export OS_USER_DOMAIN_NAME=Default
export OS_USERNAME=admin1
export OS_PASSWORD=abcd1234
export OS_COMPUTE_API_VERSION=2.46
export OS_NETWORK_API_VERSION=2.0
export OS_IMAGE_API_VERSION=2
export OS_VOLUME_API_VERSION=3

# source /opt/ibm/powervc/powervcrc
```

Adding a user to access the GUI instead of root

You can run the commands that are shown in Example A-2 to access the PowerVC GUI. Change the password to something much more secure than what is shown in the example.

Example A-2 Commands to add user to access ibm-default project

```
# adduser admin1 -p abcd1234
# usermod -a -G powervc-filter admin1
# openstack role add --project ibm-default --user admin1 admin
```

Adding and listing a project

To create a project and list it, run the commands that are shown in Example A-3.

Example A-3 Create and list a project

```
# openstack project create --description "Redbook project" redbook
+-----+-----+
| Field      | Value                                |
+-----+-----+
| description | Redbook project                      |
| domain_id  | default                              |
| enabled    | True                                  |
| id         | 18a0914e913e4c0e9a8c79483c7cf07a    |
| is_domain  | False                                 |
```

```

| name      | redbook |
| options   | {}      |
| parent_id | default |
| tags      | []      |
+-----+

```

```
# openstack project list
```

```

+-----+
| ID                | Name      |
+-----+
| 135537788461455183a90a4886142a65 | powervm   |
| 1f0ca545a81a4fffa8af6cf5fd2bea91 | ibm-default |
| 18a0914e913e4c0e9a8c79483c7cf07a | redbook   |
| 4e5ed9491daf44319eab573ebd0825e6 | six       |
| aa0ed38f28154f0aaac3c0e60decfb65 | service   |
+-----+

```

Create users and roles for access to a project

You can create users and add roles to them by running the commands that are shown in Example A-4.

Example A-4 Create users and roles for access to a project

```

# adduser admin1 -p abcd1234
# adduser user1 -p abcd1234
# usermod -a -G powervc-filter admin1
# usermod -a -G powervc-filter user1
# openstack
(openstack) role add --project powervc --user admin1 admin
(openstack) role add --project powervc --user user1 self_service
(openstack) user list

```

```

+-----+
| ID                | Name      |
+-----+
| 846f722a69c01e833645918d2806ed2909291b87c115bd4b60152daf3f24a3d0 | user1     |
| a1637d18a94ccdcade35e93cf47ecf559a4f7b6d11db51948e88384189174016 | admin1    |
+-----+

```

Changing a role

To change a role for an existing user, run the commands that are shown in Example A-5.

Example A-5 Change a role

```

# openstack role remove --project six --user halfdozen admin
# openstack role add    --project six --user halfdozen deployer

```

Useful powervc-config commands

This section is a quick reference to several powervc-config subcommands to configure the PowerVC management server. The commands display the current settings with example output from IBM PowerVC. Most of the settings are the default values except the network ones. Adding **-h** after the command shows you additional help information.

General subcommands

This section lists general **powervc-config** subcommands.

Set the network domain

The command that is shown in Example A-6 can be used to set the network domain as the root user.

Example A-6 Set the network domain

```
# /usr/bin/powervc-config compute server-domain --set ats.uk.ibm.com
```

The default domain name was updated for the specified hosts. Verify that no tasks are in-progress, then run the '/opt/ibm/powervc/bin/powervc-services nova restart' command to restart the nova services and complete the operation.

Listing an interface configuration

To list an interface configuration, run the command that is shown in Example A-7.

Example A-7 List an interface configuration

```
# powervc-config general ifconfig
powervc IP version : 4
powervc management address: vm17.ats.uk.ibm.com
powervc hostname: vm17
    ifconfig          Change the hostname or IP address of the IBM PowerVC
controller
```

Listing an IP address

To list an IP address, run the command that is shown in Example A-8.

Example A-8 List an IP address

```
# powervc-config general ipaddress
Current value: 9.137.62.17
    ipaddress          The IP address of the IBM PowerVC controller Warning: This
option will be deprecated in future releases.
```

Listing a host name

To list a host name, run the command that is shown in Example A-9.

Example A-9 List a host name

```
# powervc-config general hostname
Current value: vm17
    hostname          The hostname of the IBM PowerVC controller Warning: This
option will be deprecated in future releases.
```

Identity subcommands

The following commands are identity subcommands for the **powervc-config** command.

Listing a token expiration

To list a token expiration, run the command that is shown in Example A-10.

Example A-10 List a token expiration

```
# powervc-config identity token_expiration
Current value: 6:00:00
Default value: 6:00:00
    token_expiration    Configure the expiration interval for the identity token
```

The output is in Hours:Minutes:Seconds format.

Listing a repository

To list a repository, run the command that is shown in Example A-11.

Example A-11 List a repository

```
# powervc-config identity repository
IBM PowerVC is configured for operating system authentication.
    repository          Configure the repository for authentication credentials,
for example an LDAP server.

# powervc-config image image_size_cap
Current value: 1099511627776 B
Default value: 1099511627776 B
    image_size_cap      Configure the maximum image size, which can be uploaded
through IBM PowerVC controller
```

Image subcommand

Example A-12 shows the **image** subcommand.

Example A-12 Image subcommand

```
# powervc-config image user_storage_quota
Current value: 0 B
Default value: 0 B
    user_storage_quota  Configure the maximum amount of per-user image storage.
```

Storage subcommand

Example A-13 shows the **storage** subcommand.

Example A-13 Storage subcommand

```
# powervc-config storage portgroup
portgroup          Configure the EMC VMAX PortGroup list
```

Compute subcommand

Example A-14 shows the **compute** subcommand.

Example A-14 Compute subcommand

```
# powervc-config compute server-domain
[global]                (none)
828641A_214423W         localdomain
828641A_212B8BW         localdomain
824721L_213C93A         pok.stglabs.ibm.com
824721L_2142B2A         pok.stglabs.ibm.com
828642A_21C1B6V         pok.stglabs.ibm.com
```

Metering subcommands

The following commands are useful subcommands to display the metering in your IBM PowerVC environment.

Working with meter data

You can show the current time to live value for the meter data by running the command that is shown in Example A-15.

Example A-15 Current time to live value command

```
# powervc-config metering meter_ttl
Current value: 14 days, 0:00:00
Default value: 14 days, 0:00:00
```

Note: By default, the meter data is stored for 14 days.

To change it (for example, to 30 days - $30 \times 24 = 720$), run the following command:

```
# powervc-config metering meter_ttl --set 720 --unit hr
Setting metering_time_to_live to 720 hr
```

To get metering data by using “the curl method”, complete the following steps:

1. Source the environment (for a specific project; the example here is `ibm-default`) by running the following command:

```
# source /opt/ibm/powervc/powervcrc
```

2. Get the token by running the following command:

```
# openstack token issue
```

Because the Token ID is a long value and might be difficult to extract because of the ASCII table, you can parse the command output for easy copy by running the following command:

```
# openstack token issue | head -5 | grep id | awk '{print $4}'
```

3. Use the obtained Token ID and get all data for a given project for all users in that project (in JSON Format) by running the following command:

```
# curl -1 -k -X GET \
"https://localhost:5000/powervc/openstack/metering/v2/samples?q.field=project_id&q.value=<Project ID>" -H "X-Auth-Token:<Token ID>" | python-m json.tool
```

You also can get CPU EC values for a given project for all users in that project (in JSON Format) by running the following command:

```
# curl -1 -k -X GET \  
"https://localhost:5000/powervc/openstack/metering/v2/meters/total_vccpu?q.field  
=project_id&q.value=<Project ID>" -H "X-Auth-Token:<Token ID>" | python-m  
json.tool
```

Running data queries

You can run data queries by using the following command examples.

To run a data query to obtain the total number of volumes, run the command that is shown in Example A-16.

Example A-16 Data query to obtain the total number of volumes

```
# curl -1 -k -X GET \  
"https://localhost:5000/powervc/openstack/metering/v2/meters/total_volumes?q.field  
=project_id&q.value=<Project ID>" -H "X-Auth-Token:<Token ID>" | python -m  
json.tool | grep counter_name | wc -l
```

Note: By default, a query for data gives back a maximum of 100 entries.

To run a data query to obtain the total number of volumes and show up to 1000 entries if available, run the command that is shown in Example A-17.

Example A-17 Data query to obtain the total number of volumes and show up to 1000 entries

```
# curl -1 -k -X GET \  
"https://localhost:5000/powervc/openstack/metering/v2/meters/total_volumes?limit=1  
000q.field=project_id&q.value=<Project ID>" -H "X-Auth-Token:<Token ID>" |  
python-mjson.tool | grep counter_name | wc -l
```

Web subcommands

The following subcommands control inactivity time outs and password prompts.

inactivity_timeout

Example A-18 shows the **inactivity_timeout** subcommand.

Example A-18 The inactivity_timeout command

```
# powervc-config web inactivity_timeout  
Current value: 2:00:00  
Default value: 2:00:00  
    inactivity_timeout Time that UI users will be allowed to idle before being  
prompted and logged out. Value 0 or less will disable the timer.
```

token_expiration_warning_time

Example A-19 shows the **token_expiration_warning_time** subcommand.

Example A-19 The token_expiration_warning_time subcommand

```
# powervc-config web token_expiration_warning_time  
Current value: 0:15:00  
Default value: 0:15:00
```

token_expiration_warning_time Time before token expiration to prompt user for password to obtain new token. Value 0 or less will disable the timer.

Useful scripts

The scripts that are provided in this section format output as report quality output.

Listing the defined projects and user assignments

To list the projects and users that are defined, run the `openstack role assignment list` command. A script is helpful to align the output into something more readable.

Note: The following highlight some important command usage guidelines.

- ▶ The `-f` option stops TTY style output.
- ▶ The `-c column-name` option selects what is output.
- ▶ The `--name` option removes the hexadecimal D and provides the names (note the double "-").
- ▶ Removing the `@Service` lines removes internal users.

By project, role, and user name

The script that is shown in Example A-20 produces list quality output from the `openstack` command.

Example A-20 List by project, role, and user name

```
# openstack role assignment list -c Project -c Role -c User --names -f value \  
| grep -v @Service | sed 's/@Default//g' | \  
awk '{ printf "%-15s %-15s %-15s\n", $3, $1, $2 }' | sort  
eight          admin          eightadmin  
eight          admin          nig  
eight          admin          powervc  
eight          deployer      eightdep  
eight          self_service eightself  
eight          self_service scotts  
ibm-default    admin          nig  
ibm-default    admin          powervc  
ibm-default    admin          root  
ibm-default    self_service  scotts  
seven          admin          nig  
seven          admin          powervc  
seven          admin          sevenadmin  
seven          deployer      sevendep  
seven          self_service scotts  
seven          self_service powrpro  
seven          self_service sevenself  
six            admin          nig  
six            admin          powervc  
six            admin          sixadmin  
six            deployer      halfdozen  
six            self_service scotts  
six            self_service sixtine
```

six self_service sixtus

By user name, project, and role

The script that is shown in Example A-21 produces list quality output from the **openstack** command.

Example A-21 List by user name, project, and role

```
# openstack role assignment list -c Project -c Role -c User --names -f value |
grep -v @Service | sed 's/@Default//g' | awk '{ printf "%-15s %-15s %-15s\n", $2,
$3, $1 }' | sort
eightadmin        eight            admin
eightdep          eight            deployer
eightself        eight            self_service
halfdozen        six              deployer
scotts            eight            self_service
scotts            ibm-default     self_service
scotts            seven            self_service
scotts            six              self_service
nig                eight            admin
nig                ibm-default     admin
nig                seven            admin
nig                six              admin
powervc           eight            admin
powervc           ibm-default     admin
powervc           seven            admin
powervc           six              admin
root              ibm-default     admin
powrpro           seven            self_service
sevenadmin        seven            admin
sevendep          seven            deployer
sevself           seven            self_service
sixadmin          six              admin
sixtine           six              self_service
sixtus            six              self_service
```

REST APIs example

There is no CLI for working with IBM PowerVC on a script basis, but you can use several types of interfaces to build solutions on top of IBM PowerVC:

- ▶ **Supported OpenStack APIs:** These APIs are a subset of the APIs that are provided by OpenStack and can be used with IBM PowerVC without any modifications.
- ▶ **Extended OpenStack APIs:** These APIs are a subset of the APIs that are provided by OpenStack, but their functions are extended by IBM PowerVC.
- ▶ **IBM PowerVC APIs:** These APIs do not exist in OpenStack and are exclusive to IBM PowerVC.

Note: For more information about IBM PowerVC APIs, see IBM Knowledge Center:

https://www.ibm.com/support/knowledgecenter/en/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_pg_kickoff_hmc.html

Here is an example of removing a volume by using the supported OpenStack APIs.

Complete the following steps:

1. After IBM PowerVC services are installed and running on a managing system, the first step is to request an authorization token by using the **openstack token issue** command, as shown in Example A-22. This token is required to use the required APIs.

Example A-22 Token request from the command line

```
# openstack token issue
+-----+
| Field | Value |
+-----+
| expires | 2021-01-21 13:11:33+00:00 |
| id | gAAAAABYmFWIDo33jqA1S7xs24ajCkTq1Gmhcss1NHwVI daR2c0gTFK04yAGY0A2YanERAw7vmoSdYUOAmzYn-5ns_k37EL4FjW0j49aQFqd44XcFGxHIKm6kP16_wyyYwqrD390XJB11XBT_hh0z-AqYBpQS9PZrYt-gv39KgK0q4pRFoUtXd0ffwNMIFin36hY5MFBtLDXDRRZWKb481M0dGAMGHMgRdUGsfbKcwCHOavf_Z54 |
| project_id | c0c4c1559ef24cc2b6541e80c43858d4 |
| user_id | 0688b01e6439ca32d698d20789d52169126fb41fb1a4ddafacebb97d854e836c9 |
+-----+
```

Note: Access tokens are valid for 6 hours. Request a new token at least 1 hour before it expires to ensure that the token does not expire in the middle of an operation.

2. To use the OpenStack APIs, you need the tenant ID. For Version 3.0 and later APIs, the terms *tenant* and *project* are interchangeable. To get the ID, create a python script or use a REST plug-in in your browser, such as RESTED 1.0.0. You can accomplish this task by using either the CLI or the GUI.

To use the CLI, run the commands that are shown in Example A-23.

Example A-23 Tenant-id request from command line

```
# cat tenants
#!/usr/bin/python

import httplib
import json
import os
import sys

def main():
    token = raw_input("Please enter PowerVC token : ")
    print "PowerVC token used = "+token

    conn = httplib.HTTPSConnection('localhost')
    headers = {"X-Auth-Token":token, "Content-type":"application/json"}
    body = ""

    conn.request("GET", "/powervc/openstack/identity/v3/projects", body, headers)
    response = conn.getresponse()
    raw_response = response.read()
    conn.close()
    json_data = json.loads(raw_response)
    print json.dumps(json_data, indent=4, sort_keys=True)

if __name__ == "__main__":
    main()

# chmod u+x tenants
```

```

# ./tenants
Please enter PowerVC token :
gAAAAABYMy1-006kjXd9oQaZV4W0uZL6uyeobItJ5HR_SjcGVc0pzZGdp_loP2hTJIEz1KgS8qwRK4f
mJy3WoDUzVx8u_fDLArW1U86rkUYFshHregGnE_Dskny0SO_ymLATZwcuGeEQ01cBpijeRbZIyP2A0u
qjRxOwnVtKEMmeVd48J-Zd7PsziqT8W6EdEDClrQG_hmIiEKcMwqpf0CrsObrtoV_YIno95df5psd0f
IkKJCXnk20
PowerVC token used =
gAAAAABYMy1-006kjXd9oQaZV4W0uZL6uyeobItJ5HR_SjcGVc0pzZGdp_loP2hTJIEz1KgS8qwRK4f
mJy3WoDUzVx8u_fDLArW1U86rkUYFshHregGnE_Dskny0SO_ymLATZwcuGeEQ01cBpijeRbZIyP2A0u
qjRxOwnVtKEMmeVd48J-Zd7PsziqT8W6EdEDClrQG_hmIiEKcMwqpf0CrsObrtoV_YIno95df5psd0f
IkKJCXnk20
{
  "links": {
    "next": null,
    "previous": null,
    "self": "https://9.47.76.108/powervc/openstack/identity/v3/projects"
  },
  "projects": [
    {
      "description": "IBM Service Tenant for service users and groups",
      "domain_id": "7ca0bd908588480486e9399a238e3f4b",
      "enabled": true,
      "id": "17cd686591c948d3a970d8ac204eaf6e",
      "is_domain": false,
      "links": {
        "self":
          "https://9.47.76.108/powervc/openstack/identity/v3/projects/17cd686591c948d3a97
          0d8ac204eaf6e"
      },
      "name": "service",
      "parent_id": "7ca0bd908588480486e9399a238e3f4b"
    },
    {
      "description": "IBM Tenant for storing VM NVRAM data for remote restart",
      "domain_id": "7ca0bd908588480486e9399a238e3f4b",
      "enabled": true,
      "id": "473511b7bb5e4733b650e00454e59f47",
      "is_domain": false,
      "links": {
        "self":
          "https://9.47.76.108/powervc/openstack/identity/v3/projects/473511b7bb5e4733b65
          0e00454e59f47"
      },
      "name": "powervm",
      "parent_id": "7ca0bd908588480486e9399a238e3f4b"
    },
    {
      "description": "IBM Default Tenant",
      "domain_id": "default",
      "enabled": true,
      "id": "c0c4c1559ef24cc2b6541e80c43858d4",
      "is_domain": false,
      "links": {

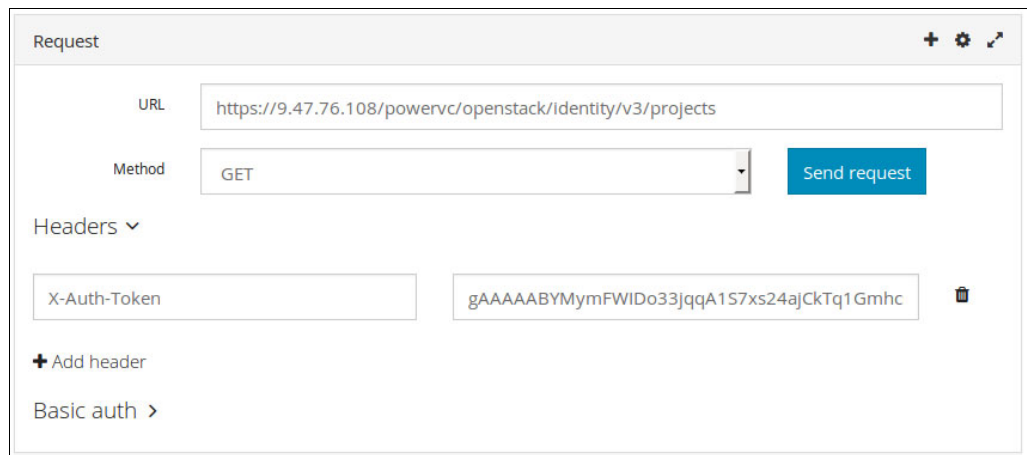
```

```
"self":  
"https://9.47.76.108/powervc/openstack/identity/v3/projects/c0c4c1559ef24cc2b65  
41e80c43858d4"  
},  
"name": "ibm-default",  
"parent_id": "default"  
}  
]  
}
```

To use the GUI, complete the following steps:

- a. Open your browser plug-in and provide the following information, as shown in Figure A-1.

URL	https://powervc.ip/powervc/openstack/identity/v3/projects
Method	GET
Header	X-Auth-Token



The screenshot shows a REST client interface with the following fields:

- URL:** https://9.47.76.108/powervc/openstack/identity/v3/projects
- Method:** GET (selected from a dropdown menu)
- Send request:** A blue button to execute the request.
- Headers:** A section with a dropdown arrow, containing one header:
 - X-Auth-Token:** gAAAAABYmFWIDo33JqqA1S7xs24ajCkTq1Gmhc
- + Add header:** A button to add a new header.
- Basic auth >:** A link to expand basic authentication options.

Figure A-1 REST request to get the tenant IDs

- b. After submitting the request, you will get an answer with the details as shown in Figure A-2 on page 329.

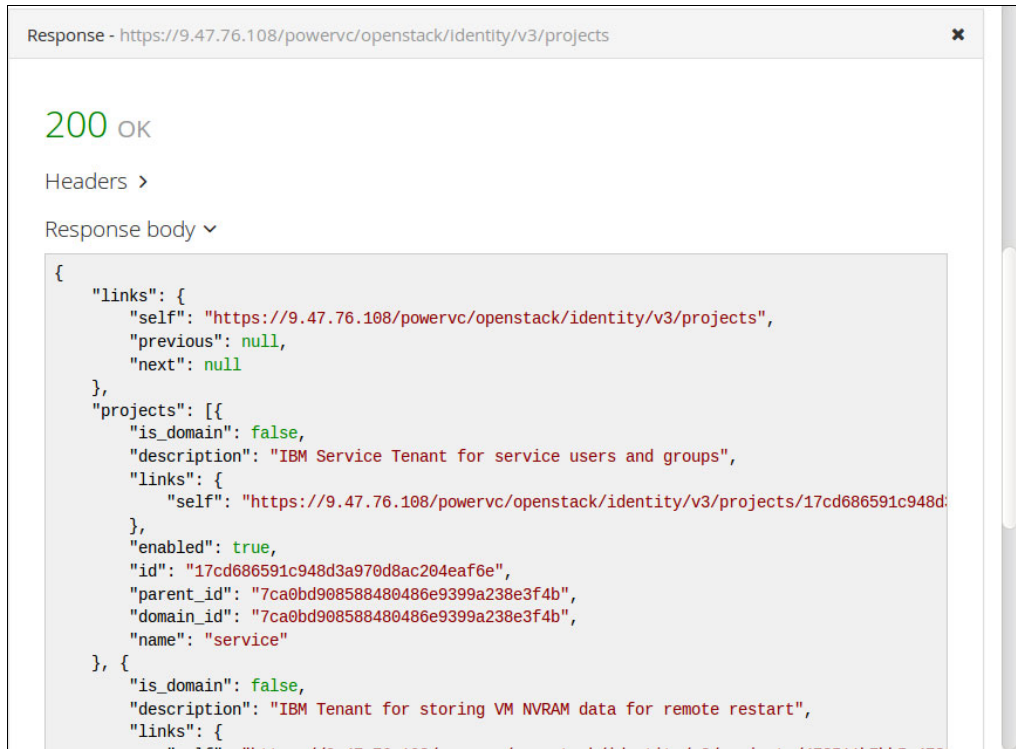


Figure A-2 API response example

3. To remove a specific volume, use the supported OpenStack Block storage (Cinder) API. The details are shown in Table A-1.

Table A-1 OpenStack API to remove a volume

Method	Path	Description
DELETE	<code>/v2/{tenant_id}/volumes/{volume_id}</code>	Deletes a specified volume.

You can get the volume_id by calling another Cinder API, or by looking at the volume attributes on IBM PowerVC GUI as shown in Figure A-3.

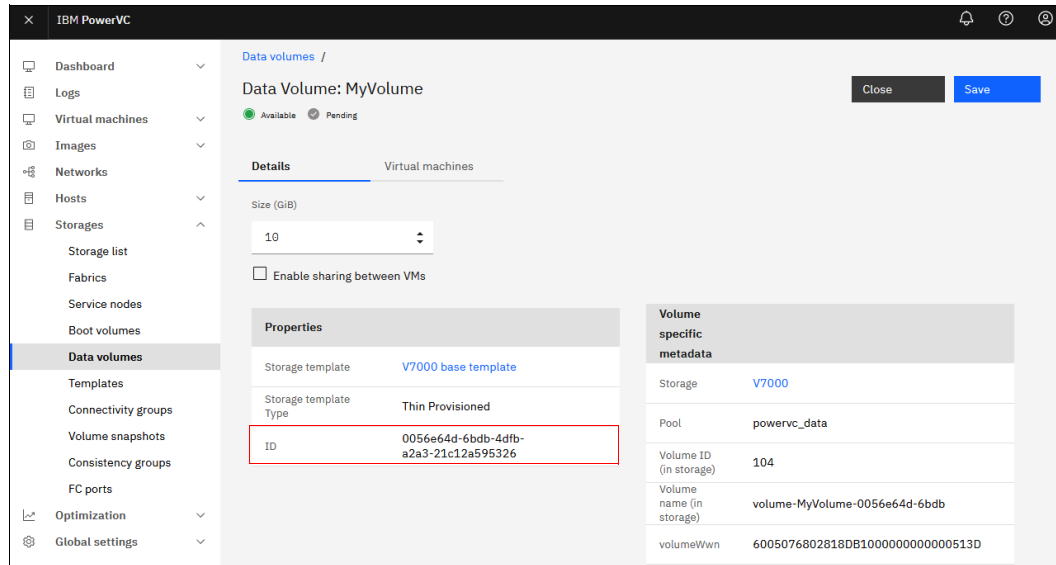


Figure A-3 Get volume ID attribute

After you get the volume ID, you can call the delete volume API by using a script from the command line or by using any REST plug-in on your browser. You can check the task status on the **Logs** tab in IBM PowerVC, as shown in Figure A-4.




Figure A-4 Volume deletion status message

Related publications

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

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- ▶ *IBM Power Systems HMC Implementation and Usage Guide*, SG24-8334
- ▶ *IBM AIX Enhancements and Modernization*, SG24-8453
- ▶ *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940
- ▶ *IBM PowerVM Virtualization Managing and Monitoring*, SG24-7590
- ▶ *Implementing the IBM System Storage SAN Volume Controller with IBM Spectrum Virtualize V8.2.1*, SG24-7933
- ▶ *Building a SAN-less Private Cloud with IBM PowerVM and IBM PowerVC*, REDP-5455
- ▶ *Red Hat OpenShift V4.3 on IBM Power Systems Reference Guide*, REDP-5599
- ▶ Red Hat OpenShift and Cloud Paks on IBM Power Systems: Volume 1, SG24-8459.

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

These websites are also relevant as further information sources:

- ▶ IBM Knowledge Center:
<https://www.ibm.com/support/knowledgecenter/>
- ▶ PowerVC Community Blog
<https://community.ibm.com/community/user/power/communities/community-home/recent-community-blogs?communitykey=6f9fc4cd-7614-41b1-98d8-2941c3497e37&tab=recentcommunityblogsdashboard>
- ▶ Managing PowerVC as an administrator
https://www.ibm.com/support/knowledgecenter/en/SSVSPA_2.0.0/com.ibm.powervc.cloud.help.doc/powervc_cloud_managing_admin.html
- ▶ Red Hat
<https://www.redhat.com/en>
<https://access.redhat.com/>
- ▶ SUSE
<https://www.suse.com/>

- ▶ SUSE Linux Enterprise Server 15 SP1
<https://documentation.suse.com/sles/15-SP1/>
- ▶ SUSE Linux Enterprise Server 15 SP1 Deployment Guide
<https://documentation.suse.com/sles/15-SP1/html/SLES-all/book-sle-deployment.html>
- ▶ SUSE Linux Enterprise Server 15 SP2
<https://documentation.suse.com/sles/15-SP2/>
- ▶ OpenStack
<https://openinfra.dev/>
https://wiki.openstack.org/wiki/Main_Page
- ▶ PowerVC NovaLink Overview
<https://community.ibm.com/community/user/power/blogs/archive-user/2015/11/01/powervc-novalink-overview>
- ▶ IBM Software lifecycle
<https://www.ibm.com/support/lifecycle/>
- ▶ IBM PowerVC Edition v2.0.0
https://www.ibm.com/support/knowledgecenter/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/kc_welcome-standard-supermap.html
- ▶ IBM PowerVC for Private Cloud v2.0.0
https://www.ibm.com/support/knowledgecenter/SSVSPA_2.0.0/com.ibm.powervc.cloud.help.doc/kc_welcome-cloud-supermap.html
- ▶ Recommended Fixes - HMC Code Upgrades
<https://www.ibm.com/support/pages/recommended-fixes-hmc-code-upgrades>
- ▶ IBM Fix Level Recommendation Tool
<https://www14.software.ibm.com/webapp/set2/flrt/home>
- ▶ IBM Fix Central
<https://www.ibm.com/support/fixcentral/>
- ▶ Planning for storage
https://www.ibm.com/support/knowledgecenter/en/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_planning_storage_hmc.html
- ▶ PowerVC Standard Edition 2.0.0 Remote restart
https://www.ibm.com/support/knowledgecenter/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_remote_restart_hmc.html
- ▶ Active Memory Expansion (AME)
https://www.ibm.com/support/knowledgecenter/en/ssw_aix_72/performance/intro_ame_process.html
- ▶ Processor Compatibility modes on IBM POWER9 based systems
<https://www.ibm.com/support/pages/processor-compatibility-modes-ibm-power9-based-systems>
- ▶ SAP HANA and PowerVM Virtual Persistent Memory
<https://www.ibm.com/support/pages/system/files/inline-files/vPMEM-SAPHANA-White-paper-V1.2.pdf>

- ▶ Planning for storage providers
https://www.ibm.com/support/knowledgecenter/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_storage_providers.html
- ▶ Creating Simple SSP among two VIO servers
<https://www.ibm.com/support/pages/creating-simple-ssp-among-two-vio-servers>
- ▶ Shared storage pools
https://www.ibm.com/support/knowledgecenter/9119-MME/p8eew/p8eew_shared_storage.htm
- ▶ PowerVM vNIC and vNIC Failover FAQs
<https://community.ibm.com/community/user/power/viewdocument/vnic-and-vnic-failover-faqs?CommunityKey=71e6bb8a-5b34-44da-be8b-277834a183b0&tab=librarydocuments>
- ▶ vNIC Functionality Guide
<https://www.ibm.com/support/pages/node/6357843>
- ▶ Ports used on the management server
https://www.ibm.com/support/knowledgecenter/bs/SSVSPA_2.0.0/com.ibm.powervc.cloud.help.doc/powervc_planning_security_firewall_cloud.html#powervc_planning_security_firewall_hmc__d264e87
- ▶ Ports used by IBM PowerVC on the management server
https://www.ibm.com/support/knowledgecenter/bs/SSVSPA_2.0.0/com.ibm.powervc.cloud.help.doc/powervc_planning_security_firewall_cloud.html#powervc_planning_security_firewall_hmc__d264e567
- ▶ Ports used by PowerVM NovaLink managed host
https://www.ibm.com/support/knowledgecenter/bs/SSVSPA_2.0.0/com.ibm.powervc.cloud.help.doc/powervc_planning_security_firewall_cloud.html#powervc_planning_security_firewall_hmc__d264e734
- ▶ Replacing the certificate for PowerVC access
https://www.ibm.com/support/knowledgecenter/SSVSPA_2.0.0/com.ibm.powervc.cloud.help.doc/powervc_certificate_cloud.html
- ▶ IBM support
<https://www.ibm.com/support/home/>
- ▶ IBM Support Insider
<https://www.ibm.com/support/pages/node/796356>
- ▶ IBM Support Guide
<https://www.ibm.com/support/pages/node/733923>
- ▶ Offering Information
https://www-01.ibm.com/common/ssi/SearchResult.wss?request_locale=en&MPPEFFTR=CONTENTS&MPPEFSCH=5765-VC2&MPPEFSRT=2
- ▶ IBM Entitled Systems Support
<https://www.ibm.com/servers/eserver/ess/ProtectedServlet.wss>
- ▶ IBM Enterprise Support and Preferred Care
<https://www.ibm.com/support/pages/node/738889>

- ▶ Installing Red Hat Enterprise Linux on the PowerVC management server or host
https://www.ibm.com/support/knowledgecenter/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_installing_rhel_mgmt_vm_hmc.html
- ▶ Installing SUSE Linux on the management server or host
https://www.ibm.com/support/knowledgecenter/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_installing_sles_mgmt_vm.html
- ▶ PowerVM NovaLink command-line interface
https://www.ibm.com/support/knowledgecenter/9009-42A/p9eig/p9eig_cli.htm
- ▶ Linux information for IBM systems:
<https://www.ibm.com/support/knowledgecenter/en/linuxonibm/liaaa/ic-homepage.html>
- ▶ Managing OpenBMC-based systems
https://www.ibm.com/support/knowledgecenter/POWER9/p9eih/p9eih_openbmc_kickoff.htm?pos=3
- ▶ Upgrading PowerVM NovaLink software to version 2.0.0.0
https://www.ibm.com/support/knowledgecenter/POWER9/p9eig/p9eig_upgrade_novalink.htm
- ▶ IBM cloud repo
<http://public.dhe.ibm.com/systems/virtualization/powervc/>
- ▶ IBM My Notifications
<https://www.ibm.com/support/pages/stay-date-my-notifications>
- ▶ IBM Installation Toolkit for Linux on Power
<https://www14.software.ibm.com/support/customercare/sas/f/lopdiags/installtools/home.html>
- ▶ Set up an RMC connection
https://www.ibm.com/support/knowledgecenter/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_configuring_vlans_hmc.html
- ▶ Installing and configuring cloud-init on Red Hat Enterprise Linux
https://www.ibm.com/support/knowledgecenter/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_install_cloudinit_rhel_hmc.html
- ▶ Installing and configuring cloud-init on SUSE Linux Enterprise Server
https://www.ibm.com/support/knowledgecenter/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_install_cloudinit_sles_hmc.html
- ▶ Installing and configuring cloud-init on Ubuntu
https://www.ibm.com/support/knowledgecenter/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_install_cloudinit_ubuntu_hmc.html
- ▶ Installing and configuring cloud-init on AIX
https://www.ibm.com/support/knowledgecenter/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_install_cloudinit_aix_hmc.html
- ▶ IBM i Open Source Technologies
<https://bitbucket.org/ibmi/opensource/src/master/docs/yum/>

- ▶ IBM i Technology Updates
<https://www.ibm.com/support/pages/node/1119129>
- ▶ Installing and configuring cloud-init on IBM i
https://www.ibm.com/support/knowledgecenter/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_install_cloudinit_ibmi_hmc.html
- ▶ Carbon Design System Search
<https://www.carbondesignsystem.com/>
- ▶ PowerVC Version 2.0.0 UI known issues and limitations
https://www.ibm.com/support/knowledgecenter/SSVSPA_2.0.0/com.ibm.powervc.cloud.help.doc/powervc_2.0.0_known_issues_limitations_cloud.html
- ▶ Hardware and software requirements
https://www.ibm.com/support/knowledgecenter/SSVSPA_2.0.0/com.ibm.powervc.cloud.help.doc/powervc_hwandsw_reqs_cloud.html
- ▶ Virtual Machine Capture requirements
https://www.ibm.com/support/knowledgecenter/en/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_capture_reqs_hmc.html
- ▶ HMC V8 R8.7.0 SP1 Enhancements
<https://community.ibm.com/community/user/power/blogs/hariganesh-muralidharan1/2020/06/08/hmc-v8-r870-sp1-enhancements>
- ▶ AIX Toolbox ReadMe
<https://public.dhe.ibm.com/aix/freeSoftware/aixtoolbox/eziinstall/ppc/README-yum>
- ▶ Custom PowerVC cloud-init modules
https://www.ibm.com/support/knowledgecenter/SSVSPA_2.0.0/com.ibm.powervc.cloud.help.doc/powervc_custom_modules_aix_cloud.html
- ▶ IBM Power Systems service and productivity tools
<https://www.ibm.com/support/pages/service-and-productivity-tools>
- ▶ Installing and configuring cloud-init on Red Hat Enterprise Linux
https://www.ibm.com/support/knowledgecenter/en/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_install_cloudinit_rhel_hmc.html
- ▶ Installing and configuring cloud-init on SUSE Linux Enterprise Server
https://www.ibm.com/support/knowledgecenter/en/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_install_cloudinit_sles_hmc.html
- ▶ Installing and configuring cloud-init on Ubuntu
https://www.ibm.com/support/knowledgecenter/en/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_install_cloudinit_ubuntu_hmc.html
- ▶ Managing roles
https://www.ibm.com/support/knowledgecenter/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_supported_roles_hmc.html
- ▶ PowerVC LDAP VideoBlog
<https://www.youtube.com/watch?v=pXQyhbIZZ4M>

- ▶ IBM PowerVC LDAP Configuration FAQs
<https://community.ibm.com/community/user/power/blogs/divya-k-konoor1/2018/07/11/ibm-powervc-ldap-configuration-faqs?CommunityKey=6f9fc4cd-7614-41b1-98d8-2941c3497e37&tab=recentcommunityblogsdashboard>
- ▶ Configuring LDAP
https://www.ibm.com/support/knowledgecenter/en/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_ldap_hmc.html
- ▶ Terraform
<https://www.terraform.io/>
- ▶ Terraform automation projects on IBM Power systems.
<https://github.com/ocp-power-automation/terraform-providers-power/blob/main/README.md>
- ▶ Announcement Letters “Red Hat OpenShift Container Platform 4.3 and Red Hat OpenStack Platform 16 are available for IBM Power Systems servers”
<https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=AN&subtype=CA&htmlfid=897/ENUS220-193&appname=USN>
- ▶ Download Red Hat Enterprise Linux CoreOS (RHCOS) image
<https://mirror.openshift.com/pub/openshift-v4/ppc64le/dependencies/rhcos/pre-release/latest/>
- ▶ IBM Power Virtualization Center APIs
https://www.ibm.com/support/knowledgecenter/en/SSXK2N_2.0.0/com.ibm.powervc.standard.help.doc/powervc_pg_kickoff_hmc.html
- ▶ Installing a OCP cluster on IBM Power
https://docs.openshift.com/container-platform/4.3/installing/installing_ibm_power/installing-ibm-power.html
- ▶ Terraform based automation code to help the deployment of OpenShift Container Platform (OCP) 4.x on PowerVM systems managed by PowerVC
<https://github.com/ocp-power-automation/ocp4-upi-powervm>
- ▶ OpenShift 4.3 now available on IBM Power Systems
<https://www.openshift.com/blog/openshift-4.3-now-available-on-ibm-power-systems-1>
- ▶ Install OpenShift on Power with user-provisioned infrastructure
<https://cloud.redhat.com/openshift/install/power/user-provisioned>
- ▶ Persistent storage using the Container Storage Interface (CSI)
https://docs.openshift.com/container-platform/4.3/storage/persistent_storage/persistent-storage-csi.html
- ▶ IBM IT Infrastructure
<https://www.ibm.com/blogs/systems/ibm-power9-scale-out-servers-deliver-more-memory-better-price-performance-vs-intel-x86/>
- ▶ IBM Power Systems hybrid cloud
<https://www.ibm.com/support/pages/power-hybrid-cloud>

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Redbooks

IBM PowerVC Version 2.0 Introduction and Configuration

SG24-8477-00

ISBN 0738459739



(0.5" spine)

0.475" x 0.873"

250 <-> 459 pages



SG24-8477-00

ISBN 0738459739

Printed in U.S.A.

Get connected

