IBM Cloud Manager with OpenStack

Administrator Guide, version 4.1

IBM Cloud Manager with OpenStack

Administrator Guide, version 4.1

Note

Before using this information and the product it supports, read the information in "Notices" on page 281.

Sixth Edition (June 2014)

© Copyright IBM Corporation 2012, 2014.
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Chapter 1. Overview of IBM Cloud Manager with OpenStack 1

IBM Cloud Manager with OpenStack technical overview 1
Self-service management portal (optional) 2
Overview of access-based roles 3
Terminology 3
Worksheets: Common scenarios 4
Worksheet: Getting started with a minimal topology 5
Worksheet: Controller +*n* compute or distributed database topology 5
What's new in IBM Cloud Manager with OpenStack 5
License information 6
Accessibility 6

Chapter 2. Planning for IBM Cloud Manager with OpenStack 7

IBM Cloud Manager with OpenStack prerequisites . 7
IBM Cloud Manager with OpenStack hardware prerequisites 7
IBM Cloud Manager with OpenStack operating system and software prerequisites 8
IBM Cloud Manager with OpenStack virtualization environment prerequisites 10
Microsoft Hyper-V prerequisites 10
KVM prerequisites 11
PowerKVM prerequisites 11
IBM Power Virtualization Center prerequisites . . 11
VMware prerequisites 12
z/VM prerequisites 13
Supported IBM Cloud Manager with OpenStack and IaaS gateway support matrix 14
Network considerations 16
Scalability and performance considerations 17
Server and concurrent user maximums 17
Neutron plug-in deployment considerations . . 18

Chapter 3. Installing and uninstalling IBM Cloud Manager with OpenStack . . 19

Installing IBM Cloud Manager with OpenStack on Linux 19
Installing IBM Cloud Manager with OpenStack on Linux by using console installation 19
Installing IBM Cloud Manager with OpenStack on Linux by using graphical installation 21
Installing IBM Cloud Manager with OpenStack on Linux by using silent installation 23
Installing IBM Cloud Manager with OpenStack self-service portal on Windows 25
Graphical installation (default) 25
Console installation 26
Silent installation 26
Installing and uninstalling IBM Cloud Manager with OpenStack agents 27
Installing and uninstalling the IBM Cloud Manager with OpenStack Hyper-V Agent . . . 27
Configuring operating system yum repositories on the deployment server 36
Applying fixes and updates 38
Applying fixes and updates for IBM Cloud Manager with OpenStack 38
Applying fixes and updates for the self-service portal on Windows 39
Applying fixes and updates for DB2 41
Uninstalling IBM Cloud Manager with OpenStack 42
Uninstalling IBM Cloud Manager with OpenStack from deployment server 42
Uninstalling the self-service portal on Linux . . 42
Uninstalling IBM Cloud Manager with OpenStack self-service portal from Windows . . 43

Chapter 4. Migrating IBM Cloud Manager with OpenStack 45

Migrating from IBM SmartCloud Entry 3.2 web interface to 4.1 self-service portal 45
Migrating your configuration 46
Migrating your data 47
Migrating configuration manually 48
Migrating between two 4.1 self-service portals . . . 48
Migrating self-service portal to another system 48
Migrating a Derby database to DB2 database . . 49
Migrating cinder volume groups from the appliance to the OpenStack controller 49
Identifying cinder-volumes LVM volume groups 50
Identifying corresponding devices to move . . . 51
Preparing appliance to detach devices 52
Importing volume groups on the new controller 54
Verifying cinder volume group migration . . . 54
Migrating OpenStack controller from 3.2 to 4.1 . . 55
Migrating OpenStack controller prerequisites . . 56
Deploying a new 4.1 server for migration . . . 57
Migration process 58
Migrating the configuration manually 61
Upgrading the IBM Cloud Manager with OpenStack Hyper-V Agent 62
Graphical Upgrade 63
Silent upgrade 63
Migrating the KVM Agent and CLI client 65
Preparing to migrate the KVM Agent and CLI client 65
Upgrading the IBM Cloud Manager with OpenStack KVM Agent and CLI Client 66
Completing the migration process for KVM agent or CLI client 71

© Copyright IBM Corp. 2012, 2014

Chapter 5. Deploying IBM Cloud Manager with OpenStack 73

Topology overview 73
Minimal deployment 74
Controller +*n* compute deployment 74
Distributed database deployment 76
Deploying the cloud environment 76
Deploying prerequisites 76
Deploying a minimal topology 78
Deploying a controller +*n* compute or distributed database topology 80
Deployment customization options 100
Verifying Open vSwitch configuration 110
Adding a compute node to a deployed topology 111
Changing passwords and secrets 112
Redeploying a node cleanup 113
Updating a deployed topology 113

Chapter 6. Configuring web interface properties 117

Configuring IBM Cloud Manager with OpenStack properties 117
Changing IBM Cloud Manager with OpenStack properties 117
Creating initial networks 117
Configuring SSH and ICMP access to deployed virtual machines 118
Configuring migration for a KVM node with NFS shared storage 119
Installing a network time service 122
Configuring multiple block storage backends 122
Configuring IBM Cloud Manager with OpenStack self-service portal properties 124
Configuring secure shell (SSH) communication 125
Configuring user registry authentication . . . 127
Configuring database. 132
Configuring email notifications 133
Configuring common cloud properties 134
Configuring global image deployment 135
Configuring multiple instances for a single deployment 155
Configuring logging 156
Configuring a network 156
Configuring billing 158
Configuring metering 162
Configuring capacity and overcommit rates . . 163
Configuring web user interface 164
Configuring IBM Cloud Manager with OpenStack dashboard properties 166
Rebranding the dashboard 167

Chapter 7. Managing IBM Cloud Manager with OpenStack as an Administrator 169

Configuring the license key. 169
Managing IBM Cloud Manager with OpenStack services 169
IBM Cloud Manager with OpenStack services 169
Checking status of OpenStack services 171
Restarting IBM Cloud Manager with OpenStack services 172
Managing the IaaS gateway 173
Managing with the IBM Cloud Manager with OpenStack self-service portal (Administrator access). 173
Starting and stopping IBM Cloud Manager with OpenStack self-service portal 174
Configuring the default administrator user account 175
Configuring LDAP authentication using the web interface 175
Configuring the license key for the self-service portal running on Windows 177
Managing images 178
Managing projects. 187
Managing approval policies 190
Managing requests 193
Managing expiration policies 194
Managing flavors (OpenStack only) 195
Managing clouds 197
Managing network configurations 201
Managing instances 207
Managing storage volumes 211
Managing users 211
Managing accounts 213
Clearing or archiving events 214
Viewing capacity statistics 215
Managing with OpenStack dashboard 216
Backing up and restoring IBM Cloud Manager with OpenStack 216
Backing up and restoring the deployment server 216
Backing up and restoring the OpenStack databases. 218
Backing up and restoring the self-service portal 219

Chapter 8. Security 221

Port usage 221
User names and passwords. 223

Chapter 9. Reference 225

Commands 225
Data bags 229
Mapping attributes to services. 230
Qpid messaging service attributes 230
OpenStack identity service attributes 231
OpenStack image service attributes 231
OpenStack block storage service attributes. . . 231
OpenStack orchestration service attributes . . . 232
OpenStack telemetry service attributes 233
OpenStack compute service attributes 233
OpenStack network service attributes 234
IBM OpenStack IaaS gateway attributes . . . 236
IBM OpenStack PowerVC driver service attributes 236
IBM OpenStack z/VM driver service attributes 236
Roles 237
Best practices when using the self-service portal 238
IBM Cloud Manager with OpenStack FAQ . . 238
Using the **screen** command 240

Using the **nohup** command 240

Chapter 10. Troubleshooting and support for IBM Cloud Manager with OpenStack 243

Techniques for troubleshooting problems 243
Logging tasks 245
Change logging levels from the OSGi command prompt 245
Retrieve log and system files 246
Troubleshooting using the OSGi console . . . 247
Logging chef-client output 247
Searching knowledge bases. 248
Getting fixes from Fix Central 249
Contacting IBM Support. 250
Exchanging information with IBM 251
Sending information to IBM Support 251
Receiving information from IBM Support . . . 251
Subscribing to Support updates 252
Troubleshooting errors when deploying topologies 253
Data bag item not found 253
Node verification failed 254
Deploying a client node fails 254
Cookbook not found 255
Network interface not found 255
Operating system repository not found. . . . 255
Package failed to install or upgrade 256
PowerVC driver deploys successfully; later is dead 256
Compute node is deployed or updated successfully; later Open vSwitch agent is down . 257
Fingerprint error during deploy 257
Error restarting neutron when deploying a topology 258
Cinder error occurs when you deploy a controller node 259
Known problems and solutions for an Administrator 259
Known issues 259
Limitations 270
Known problems and solutions for a User. . . . 274
Cached display. 275
Saving image exception 275
Error message language 275
Internet Explorer display 275
Login fails 277

Accessibility. 279

Notices 281

Trademarks 282
Privacy policy considerations 283

Chapter 1. Overview of IBM Cloud Manager with OpenStack

IBM® Cloud Manager with OpenStack is an easy to deploy, simple to use cloud management software offering that is based on OpenStack with IBM enhancements that feature a self-service portal for workload provisioning, virtual image management, and monitoring. It is an innovative, cost-effective approach that also includes automation, metering, and security for your virtualized environment.

IBM Cloud Manager with OpenStack technical overview

The IBM Cloud Manager with OpenStack solution, provides the architecture, software, and mechanisms that deliver OpenStack clouds to your environment quickly.

By installing the IBM Cloud Manager with OpenStack product, you receive the tools to easily deploy an OpenStack cloud, made up of an OpenStack controller node and one or more compute hosts. There are various topologies that are provided to help you deploy the nodes, which create your cloud. After you deploy a cloud, the IBM Cloud Manager with OpenStack solution includes the following management tools.

- IBM Cloud Manager with OpenStack **self-service portal**: Optional and intended for management by cloud users.
- OpenStack **dashboard**: Intended for use by cloud administrators.
- **APIs and knife commands**: Intended for cloud administrator use.

Before you begin, you must understand the concepts and terminology of OpenStack and Chef server technology. There is significant documentation that is provided in the community. Review some of the key concepts like *cookbooks*, *recipes*, *environments*, and *roles*. For more information, see the following terminology resources:

- All about Chef ...
  - Overviews
  - Key Concepts
- Get started with OpenStack

IBM Cloud Manager with OpenStack is designed to get you started quickly, yet provide you flexibility for production level installations. It uses Chef server technology to provide a robust installation and configuration method by using *cookbooks*, *recipes*, *environments*, and *roles*. Chef provides a client that communicates with the deployment server to install software on remote nodes as well.

This image shows the communication between the IBM Cloud Manager with OpenStack components in your cloud.

© Copyright IBM Corp. 2012, 2014

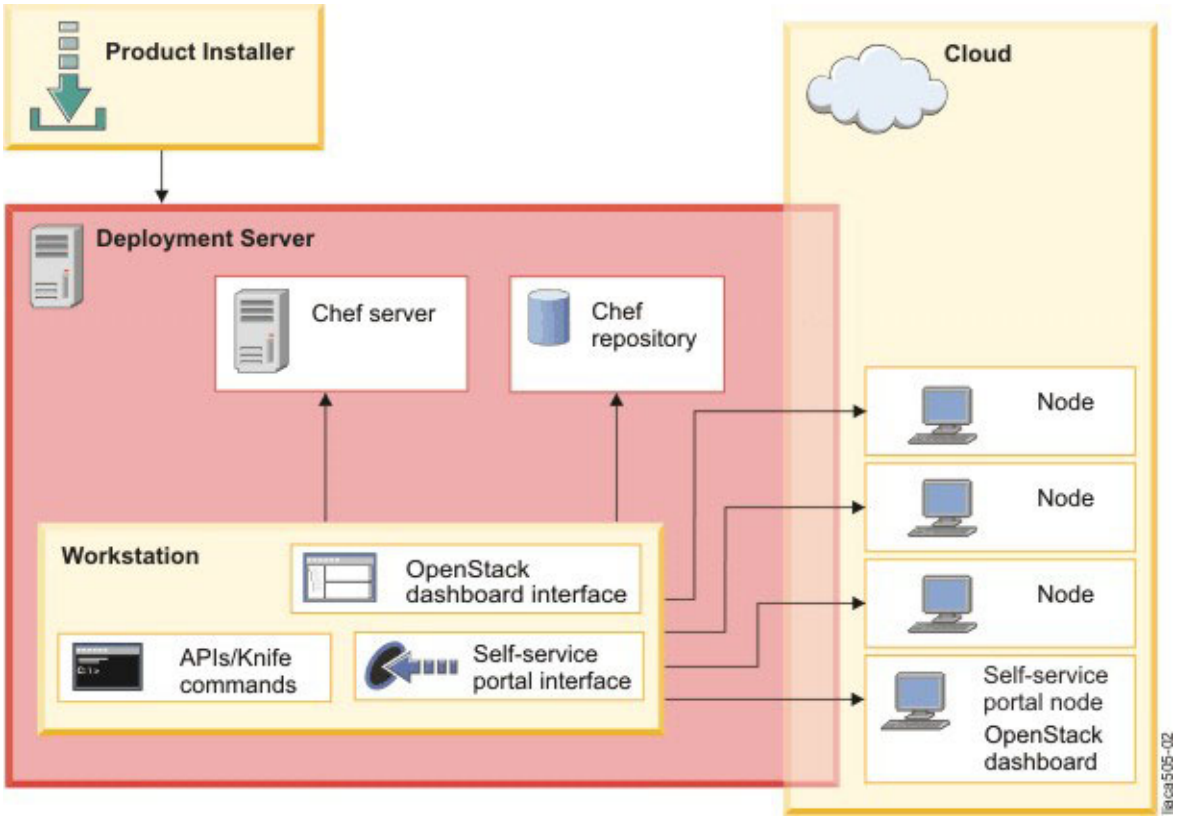

The deployment server is where you install the IBM Cloud Manager with OpenStack solution. This server becomes your Chef server and it stores cookbooks and templates (for example) that are applied to nodes. The nodes use the recipes, templates, and file distributions to do as much of the configuration work as possible on the nodes themselves (and not on the Chef server).

For more detail about specific deployment options, see the "Topology overview" on page 73.

Related reference:

OpenStack (http://www.openstack.org/)

Chef Overview (http://docs.opscode.com/chef_overview.html)

Self-service management portal (optional)

Easy to deploy and easy to use, IBM Cloud Manager with OpenStack features a self-service portal for performing cloud operations.

With the IBM Cloud Manager with OpenStack self-service portal, you can perform the following public and private cloud operations:

- Provisioning and de-provisioning virtual servers on OpenStack (KVM, PowerKVM, Hyper-V, PowerVC, z/VM®), and VMware vSphere using vCenter virtualization environments
- Providing access to multiple clouds from a single portal
- Drafting and cloning instances
- Capturing instances
- Starting and stopping servers as part of an instance
- Resizing existing virtual machines
- Creating projects to give team-specific access to instances

- Providing network configurations, which set unique network properties to different instances
- Creating expiration polices to reduce abandoned virtual machines
- Providing request and approval workflow support
- Monitoring resource allocations and billing for services

IBM Cloud Manager with OpenStack self-service portal uses the following FIPS 140-2 approved cryptographic providers:

- IBMJCEFIPS (certificate 376)
- IBMJSSEFIPS (certificate 409)
- IBM Crypto for C (certificate 384)

The certificates are listed on the NIST website at http://csrc.nist.gov/cryptval/140-1/1401val2004.htm.

Overview of access-based roles

The IBM Cloud Manager with OpenStack self-service portal supports role-based access control.

Role-based access control is enforced by assigning one or more roles to the user, and allowing each role to perform a given set of operations.

The self-service portal provides the following predefined roles:

User The main responsibility of an end user is to request the provisioning of deployed virtual machines.

- Deploy an image.
- Manage instances.
- View requests.
- View projects.
- View activity reports.

Administrator

Users with this role are responsible for administering all resources in the cloud. The typical tasks of administrators are:

- Configuring the cloud.
- Creating a project, which manages cloud access.
- Manage images.
- Manage instances.
- Manage requests.
- View activity reports.

Content that is specific to one role or the other is designated with the following flags:

- User
- Administrator

Terminology

IBM Cloud Manager with OpenStack supports many different types of virtualization infrastructure environments. These environments sometimes use different terminology for the same concepts.

The terminology differences are described in the following table.

Note: IBM Cloud Manager with OpenStack is aligning more closely with OpenStack terminology. For example, workload and appliance are now referred to as an instance and image. OpenStack is an open source cloud-computing platform for private and public clouds. For information about OpenStack, see http://www.openstack.org/

Table 1. A terminology comparison between the virtualization infrastructure type and the IBM Cloud Manager with OpenStack equivalent term

Virtualization infrastructure type	Term	Definition	IBM Cloud Manager with OpenStack equivalent
VMware	Template	A blueprint of a virtual machine containing the metadata and one or more disk images that can be used to create new virtual machines.	Image
VMware	Virtual machine	A runnable instance of a virtual computer, similar to a physical computer that runs an operating system and applications.	Instance
OpenStack	Flavor	A flavor is a defined size for a provisioned virtual machine. Each flavor has a unique combination of resource configurations and sizes.	Flavor

In addition to terminology differences between environments, there are key concepts that you must understand, if you use the self-service portal.

Projects
Projects, within the self-service portal, provide a management realm to group images and instances that only the members of that project can see and manage.

Requests
Requests are any actions that require administrator approval before they can complete. The self-service portal sends an approval request when a user attempts an operation that an administrator has set up to require approvals.

Accounts
Enabling the billing operation in the self-service portal activates the account feature. An account includes a balance, an owner, an account balance threshold, account members, and invoices. The account members are charged for the instances that they deploy.

Note: Only administrators can create accounts, but a user can be assigned as an account owner.

Basic and advanced deployments
Users deploy an image by using the basic deployment form. Project owners or administrators can use the basic or the advanced deployment forms. They can also configure which deployment settings are shown on the basic deployment form.

Worksheets: Common scenarios

Use these worksheets to review common actions that a IBM Cloud Manager with OpenStack user or administrator might perform.

Worksheet: Getting started with a minimal topology

Use this worksheet to review common tasks required to get started with IBM Cloud Manager with OpenStack. This is a common deployment if you want to evaluate and learn more about the product.

Table 2. Getting started

	Tasks	Description
___ 1.	Chapter 2, "Planning for IBM Cloud Manager with OpenStack," on page 7 __ 1. Review hardware and software support. __ 2. Review scalability and performance considerations.	
___ 2.	"Installing IBM Cloud Manager with OpenStack on Linux" on page 19	
___ 3.	"Configuring operating system yum repositories on the deployment server" on page 36	
___ 4.	Chapter 5, "Deploying IBM Cloud Manager with OpenStack," on page 73 __ 1. Select "Minimal deployment" on page 74 __ 2. Deploy the minimal topology to create your cloud environment.	
___ 5.	"Managing with OpenStack dashboard" on page 216	

Worksheet: Controller +*n* compute or distributed database topology

Use this worksheet to review common tasks required to get started with IBM Cloud Manager with OpenStack. This is a common deployment if you want to start using the product in production environments and perform some customization.

Table 3. Getting started

	Tasks	Description
___ 1.	Chapter 2, "Planning for IBM Cloud Manager with OpenStack," on page 7 __ 1. Review hardware and software support. __ 2. Review scalability and performance considerations.	
___ 2.	"Installing IBM Cloud Manager with OpenStack on Linux" on page 19	
___ 3.	"Configuring operating system yum repositories on the deployment server" on page 36	
___ 4.	"Installing a network time service" on page 122	
___ 5.	Chapter 5, "Deploying IBM Cloud Manager with OpenStack," on page 73 __ 1. Review "Controller +*n* compute deployment" on page 74. __ 2. Review "Distributed database deployment" on page 76. __ 3. Complete "Deploying a controller +*n* compute or distributed database topology" on page 80 to create your cloud environment.	
___ 6.	"Managing with OpenStack dashboard" on page 216	

What's new in IBM Cloud Manager with OpenStack

IBM Cloud Manager with OpenStack is a cloud solution with new features and support in this release.

The new features include the following:

- Additional hypervisor support
  - Added support for z/VM and PowerKVM virtualization through OpenStack technologies.
- Chef deployment server
  - The IBM Cloud Manager with OpenStack Chef server provides greater flexibility and control over how you deploy OpenStack in your cloud.
- You can now manage storage volumes for individual virtual machines with IBM Cloud Manager with OpenStack self-service portal. You can create, delete, edit, and attach or detach storage volumes for a corresponding instance.
- An administrator can now lock a user account to prevent access to IBM Cloud Manager with OpenStack self-service portal.

License information

Learn about license information for the IBM Cloud Manager with OpenStack product.

The IBM Cloud Manager with OpenStack product includes a default trial license, with a 90-day trial period. You can use this license to investigate IBM Cloud Manager with OpenStack. However, to have full use of IBM Cloud Manager with OpenStack, you need to update the trial license to a permanent license.

Related tasks:

"Configuring the license key" on page 169
The IBM Cloud Manager with OpenStack product includes a default trial license, with a 90-day trial period. You can use this license to investigate IBM Cloud Manager with OpenStack.

Accessibility

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

Accessibility features

IBM® strives to provide products with usable access for everyone, regardless of age or ability. This product uses standard browser navigation keys.

The following list includes the major accessibility features in IBM Cloud Manager with OpenStack:

- Keyboard-only operation
- Interfaces that are commonly used by screen readers
- Keys that are discernible by touch but do not activate just by touching them
- Industry-standard devices for ports and connectors
- The attachment of alternative input and output devices

The IBM Cloud Manager with OpenStack Information Center, and its related publications, are accessibility-enabled.

Keyboard navigation

The product use standard browser navigation keys.

IBM and accessibility

See the IBM Human Ability and Accessibility Center for more information about the commitment that IBM has to accessibility.

Chapter 2. Planning for IBM Cloud Manager with OpenStack

Review the following information for planning to install and using IBM Cloud Manager with OpenStack. The information includes hardware and software requirements for the various IBM Cloud Manager with OpenStack components.

IBM Cloud Manager with OpenStack prerequisites

Review the following information for planning to install and using IBM Cloud Manager with OpenStack.

The information includes hardware and software requirements for the various IBM Cloud Manager with OpenStack components.

IBM Cloud Manager with OpenStack hardware prerequisites

This section describes hardware prerequisites for various platforms and components that IBM Cloud Manager with OpenStack supports.

The information provided is a general guideline and actual requirements can vary from installation to installation. Specific sizings should be done to meet your installation requirements.

The IBM Cloud Manager with OpenStack server is supported on the following platforms:

- Intel x86-64 (Linux)
- POWER6®, POWER7®, or POWER8™ (PowerLinux™)

This table describes both the minimum hardware requirements and recommended minimum production hardware requirements for the IBM Cloud Manager with OpenStack deployment server, controller, database, and compute node components.

The minimum requirements listed indicate the absolute minimum hardware levels needed when running with only 5-10 concurrent users.

The recommended minimum production requirements are recommendations to support a small cloud. As with any software solution, hardware needs to be properly sized for a specific customer scenario.

Table 4. Minimum hardware requirements

Component	Minimum hardware requirements	Recommended minimum hardware production requirements
IBM Cloud Manager with OpenStack deployment server [1]	• 4 CPUs • Free disk space: 4 GB for `/opt/ibm/cmwo` • Free disk space: 4.5 GB of temporary space while installing • Chef server: 5.0 GB of free disk space in `/opt`; 5.0 GB of free disk space in `/var`. • 4 GB physical memory	• 8 CPUs • 25 GB free disk space • 8 GB physical memory

© Copyright IBM Corp. 2012, 2014

Table 4. Minimum hardware requirements (continued)

Component	Minimum hardware requirements	Recommended minimum hardware production requirements
OpenStack controller	• 4 CPUs • 16 GB physical memory • 6 GB free disk space for installed code and empty databases	• 4 CPUs • 16 GB physical memory • 6 GB free disk space for installed code and empty databases (4 GB for DB2® and databases)
Standalone DB2	4 GB for DB2 and empty databases	4 GB for DB2 and empty databases
Compute nodes	The following compute nodes are installed from IBM Cloud Manager with OpenStack. For specific requirements, see information about the applicable virtualization environment. • "KVM prerequisites" on page 11 • "PowerKVM prerequisites" on page 11 • "Microsoft Hyper-V prerequisites" on page 10 • "z/VM prerequisites" on page 13 Other supported virtualization environments: • "IBM Power Virtualization Center prerequisites" on page 11 • "VMware prerequisites" on page 12	
[1]Requirements are for the IBM Cloud Manager with OpenStack deployment server only (for example, Chef server). If other servers are installed on the same system, the requirements would need to be higher to account for the additional needs of the other components installed and running there.		

IBM Cloud Manager with OpenStack operating system and software prerequisites

This section lists the software prerequisites for IBM Cloud Manager with OpenStack.

The software versions listed are current at the time of publication. See the IBM Cloud Manager with OpenStack wiki for any updates.

Supported IBM Cloud Manager with OpenStack operating systems

This section lists the operating systems and versions that IBM Cloud Manager with OpenStack supports.

The following table lists the operating systems that are supported by the IBM Cloud Manager with OpenStack server, controller, and stand-alone DB2 nodes.

Table 5. Supported operating systems

Operating system	Versions	Notes
Red Hat Enterprise Linux	Version 6.4 and 6.5 (64-bit)	With latest fix pack

Note: If you use the GNOME Desktop Environment (gnote) or KDE Desktop Environment, some package conflicts might exist with IBM Cloud Manager with OpenStack. For more information about resolving the package conflicts, see "Error occurs when installing qpid" on page 266.

Supported IBM Cloud Manager with OpenStack databases

This sections lists the database and versions that IBM Cloud Manager with OpenStack supports.

The following table lists the databases that are supported by the IBM Cloud Manager with OpenStack server.

Table 6. Supported databases

Database	Versions	Notes
DB2	10.5	With service pack 3
MySQL	5.1.x	For Red Hat Enterprise Linux versions 6.4 and 6.5 on x86 systems

Supported IBM Cloud Manager with OpenStack web browsers

This section lists the web browsers and versions that the IBM Cloud Manager with OpenStack self-service portal supports.

The following table lists the web browsers that are supported for clients that access the IBM Cloud Manager with OpenStack servers.

Note: Web browsers or versions that are not listed here might also work.

Table 7. Web browser compatibility

Browser	Versions	Notes
Internet Explorer	9.0 or 10.0	With latest fix pack Minimum resolution of 1024x768 (or greater) Internet Explorer 9 or 10 compatibility view is not supported
Firefox ESR	18, 24	With latest fix pack Minimum resolution of 1024x768 (or greater)
Chrome	24, 30	With latest fix pack
Safari	6	With latest fix pack

Supported IBM Cloud Manager with OpenStack user registries

Review the user registries that are supported by the IBM Cloud Manager with OpenStack server.

The following table lists the user registries that are supported by the IBM Cloud Manager with OpenStack server.

Table 8. Supported user registries

User registry	Versions	Notes
IBM Cloud Manager with OpenStack	Local identity storage and authentication	• The IBM Cloud Manager with OpenStack database is used to store identity artifacts including credentials. • Intended for small-scale usage, such as proof-of-concept scenarios, demonstrations, or environments with up to 30 users and projects
Lightweight Directory Access Protocol (LDAP) Version 3	IBM Tivoli® Directory Server Version 6.1 OpenLDAP Version 2.4.x	• Intended for production environments to provide the highest level of security. • Scales to hundreds or thousands of users and projects. • Supports TLS (transaction level security).

IBM Cloud Manager with OpenStack virtualization environment prerequisites

IBM Cloud Manager with OpenStack depends on one of several providers for platform management and virtualization services. These providers have unique software prerequisites that depend on the type of cloud provider that you use.

Microsoft Hyper-V prerequisites

View the prerequisites for using IBM Cloud Manager with OpenStack with Microsoft Hyper-V.

IBM Cloud Manager with OpenStack is compatible with the following versions of these products:

Table 9. Supported Microsoft Hyper-V products

Microsoft Hyper-V products	Versions
Microsoft Hyper-V Server 2012 R2 [1]	
Windows Server 2012 R2 with Hyper-V role enabled	• Standard Edition • Datacenter edition

1. Microsoft Hyper-V Server 2012 R2 does not provide the APIs that are needed for IBM Cloud Manager with OpenStack to create the ISO image file that provides customization data to virtual machines. To useMicrosoft Hyper-V Server 2012 R2 with IBM Cloud Manager with OpenStack, you must install a compatible ISO generation utility such as genisoimage from Cygwin. For more information, see "Enabling Microsoft Hyper-V Server 2012 R2 systems for ISO generation."

Note: All operating systems must have the latest fix pack applied.

The management of Microsoft Hyper-V hypervisor is outside the scope of this document. Refer to the product documentation for troubleshooting instructions.

Enabling Microsoft Hyper-V Server 2012 R2 systems for ISO generation

If you are using Microsoft Hyper-V Server 2012 R2 with IBM Cloud Manager with OpenStack, you must install a compatible ISO generation utility such as genisoimage from Cygwin. After you install a

compatible ISO generation utility such as genisoimage from Cygwin to use on Microsoft Hyper-V Server 2012 R2 systems, you must update the `nova.conf` file on each system where you installed the IBM Cloud Manager with OpenStack Hyper-V Agent.

Procedure

1. Use a text editor to edit the `nova.conf` file that is located at `C:\Program Files (x86)\IBM\Cloud Manager with OpenStack\Hyper-V Agent\etc\nova`.
2. Find the line `mkisofs_cmd=C:\Program Files (x86)\IBM\Cloud Manager with OpenStack\Hyper-V Agent\bin\ibmgenisoimage.exe` and change the path and file name of the **mkisofs_cmd** property to the ISO generation utility that you installed. For example: mkisofs_cmd=C:\cygwin\bin\genisoimage.exe
3. Restart the IBM Cloud Manager with OpenStack Hyper-V Compute Agent Service by running the following commands:
   - net stop "IBM Cloud Manager with OpenStack Hyper-V Compute Agent Service"
   - net start "IBM Cloud Manager with OpenStack Hyper-V Compute Agent Service"

KVM prerequisites

View the prerequisites for using IBM Cloud Manager with OpenStack with KVM.

The KVM host must satisfy the following requirements:

- Red Hat Enterprise Linux 6.4 or 6.5
- A libvirt version 0.9.11 or later
- Open vSwitch version 2.0
- Python version 2.6.x.

Certify that the KVM host device drivers work with Open vSwitch and update the host device drivers to the latest level. IBM testing with Red Hat Enterprise Linux 6.4 or 6.5 on an IBM PureFlex™ system required that the following device drivers be updated:

kmod-be2iscsi-4.6.267.4-1.x86_64
kmod-tg3-3.129d-1.x86_64
kmod-elx-lpfc-8.3.7.29-1.x86_64
kmod-be2net-4.6.267.4-1.x86_64
kmod-brocade-bfa-3.2.1.1-0.x86_64
kmod-qlgc-qla2xxx-8.04.00.12.06.0_k3-1.x86_64

Updated IBM device drivers are available on the following IBM Support website:

http://ibm.com/support

PowerKVM prerequisites

View the prerequisites for using IBM Cloud Manager with OpenStack with PowerKVM.

IBM Cloud Manager with OpenStack supports IBM PowerKVM compute nodes. PowerKVM compute nodes must be hosts; nested virtual machines are not supported.

The PowerKVM host must satisfy the following requirements:

- Operating system: IBM_PowerKVM release 2.1.0
- Hardware: Power8 Server with the PowerKVM hypervisor configured.

IBM Power Virtualization Center prerequisites

View the prerequisites for using IBM Cloud Manager with OpenStack with IBM Power® Virtualization Center.

IBM Cloud Manager with OpenStack is compatible with IBM Power Virtualization Center Standard version 1.2.0.1, 1.2.0.2, and 1.2.x, a comprehensive virtualization management tool for the PowerVM® platform. To take advantage of the latest features of IBM Power Virtualization Center Standard, version 1.2.1 is required.

IBM Cloud Manager with OpenStack along with PowerVC provides the following capabilities:

- Support for IBM Power Systems™ hosts that are managed by Hardware Management Console.
- Support for storage area networks.
- Support for multiple Virtual I/O Server virtual machines on each host.
- Support for multiple storage connectivity groups, which enable you to deploy images so that they have access to storage that is dedicated to a particular purpose. For more information on PowerVC and storage connectivity groups, refer to the IBM PowerVC documentation.

Limitations for PowerVC support from IBM Cloud Manager with OpenStack:

- PowerVC Express Edition is not supported.
- Image capture is not supported for any existing virtual machines that are added to PowerVC by using the Manage Existing function.
- Restarting and config drive are not supported for PowerVC 1.2.0.1 or 1.2.0.2.

For information about IBM Power Virtualization Center Standard, see the IBM Power Virtualization Center information center at the following website.

http://www.ibm.com/support/knowledgecenter/SSXK2N/welcome

VMware prerequisites

View the prerequisites for using IBM Cloud Manager with OpenStack with VMware.

IBM Cloud Manager with OpenStack is compatible with the following versions of VMware products:

Table 10. Supported VMware products

VMware products	**Versions**
VMware vCenter Server 4	Standard edition (version 4.1 update 1) Essentials (version 4.1 update 1)
VMware vCenter Server 5	Standard edition Essentials edition Editions that are listed support the following versions: • 5.1.0 • 5.1 update 1 • 5.5
VMware vSphere 4	Standard edition (version 4.1 update 1) Advanced edition (version 4.1 update 1) Enterprise edition (version 4.1 update 1) Essentials Plus (version 4.1 update 1)

Table 10. Supported VMware products (continued)

VMware products	Versions
VMware vSphere 5	Standard edition Essentials Plus edition Enterprise edition Enterprise Plus edition Editions that are listed support the following versions: • 5.0.0 • 5.1 • 5.1 update 1 • 5.5

IBM Cloud Manager with OpenStack is compatible with existing installations of VMware vSphere managed by VMware vCenter. Ensure that the VMware vCenter Server product is installed, operational, and managing a VMware vSphere environment.

The configuration of the VMware products is outside the scope of this document. Refer to the product documentation for configuration and troubleshooting instructions.

Links

- VMware vCenter Server website at http://www.vmware.com/products/vcenter-server/overview.html.
- VMware vSphere website at http://www.vmware.com/products/vsphere/overview.html.
- VMware Documentation website at http://www.vmware.com/support/pubs/

IBM Cloud Manager with OpenStack supports only Windows and Linux guest operating systems, which are supported by vCenter and ESXi and allow guest customizations. For more information, see the following resources:

- VMware Compatibility Guide for Guest Operating Systems
- Guest Operating Customization Guide

Customization of certain Microsoft Windows operating systems requires Microsoft Sysprep Tools. See the information for your version of vCenter for detailed instructions about obtaining the Sysprep Tools and where to store the tools on the vCenter Servers file systems.

Configuration considerations for VMware

- Use DRS-enabled clusters for advanced placement

  Allow vCenter to place the user workload on the best host machine by using a DRS-enabled cluster within vCenter and setting the appliance target to use the cluster or a resource pool that is defined in the cluster. This allows vCenter to manage the available host resources. Otherwise, the appliance target is an individual host machine or a resource pool on a host machine.
- Place vCenter server and IBM Cloud Manager with OpenStack self-service portal on the same network.

  For best performance, it is recommended the vCenter server and the IBM Cloud Manager with OpenStack self-service portal are on the same network.

z/VM prerequisites

View the prerequisites for using IBM Cloud Manager with OpenStack with z/VM.

IBM Cloud Manager with OpenStack is compatible with z/VM version 6.3, which is a comprehensive virtualization management tool for the System z® platform.

IBM Cloud Manager with OpenStack with z/VM provides the following capabilities:

- Support for IBM System z hosts that are managed by Extreme Cloud Administration Toolkit.
- Support for storage area networks.
- Support for multiple virtual machines on each host.

For more information about enabling z/VM for OpenStack, see Enabling z/VM for OpenStack (Support for OpenStack Icehouse Release).

Ensure that z/VM 6.3 and Extreme Cloud Administration Toolkit (xCAT) are configured correctly and working well.

The configuration of z/VM and xCAT is outside the scope of this document. Refer to the following product documentation for configuration and troubleshooting information:

- z/VM 6.3 web site:

  http://www.vm.ibm.com/zvm630/
- z/VM 6.3 product information in IBM Knowledge Center:

  http://www.ibm.com/support/knowledgecenter/SSB27U_6.3.0/com.ibm.zvm.v630/zvminfoc03.htm?lang=en
- z/VM Systems Management Application Programming Guide, version 6.3http://www.vm.ibm.com/library/hcsl8c20.pdf

Supported IBM Cloud Manager with OpenStack and IaaS gateway support matrix

The following information describes the OpenStack configurations that are supported through the IBM Cloud Manager with OpenStack.

Table 11. IBM Cloud Manager with OpenStack and IaaS gateway support matrix

Controller node	Compute (Nova)	Scheduler	Filter-scheduler IBM Platform Resource Scheduler
		Nova -network	Not supported, use Neutron
		quota_drive	nova.quota.DbQuotaDriver
			nova.quota.NoopQuotaDriver (default)
	Network (Neutron)	IP version	IPv4
			IPv4/IPv6 dual (not supported by the PowerVC driver)
		Network type	Flat
			GRE (one compute hypervisor must be KVM)
			VLAN
		Virtual switch	Open vSwitch 2.0
		Plug-in	Ml2Plugin (default) OVSNeutronPluginV2
		IP address assignment for deployed instances	DHCP - provided by Neutron DHCP agent / dnsmasq service. Not supported on PowerVC Static - provided to instance by config_drive. Metadata service not supported. L3 agent or floating - Not enabled or configured in the example environments
		Security groups	Default neutron security groups do not allow ICMP or SSH into deployed instance.
	Storage (Cinder)	Driver	cinder.volume.drivers.ibm.storwize_svc.StorwizeSVCDriver to SAN Volume Controller 6.4.1/6.4.2 for iSCSI and Fibre Channel
			cinder.volume.drivers.LVMVolumeDriver on Red Hat Enterprise Linux version 6.4 or 6.5 for iSCSI
			cinder.volume.drivers.ibm.ibmnas.IBMNAS_NFSDriver for SONAS and IBM Storwize® V7000 Unified
			cinder.volume.drivers.ibm.gpfs.GPFSDriver
			powervc.volume.driver.powervc.PowerVCDriver
	Authentication (Keystone)	Identity backend	keystone.identity.backends.sql.Identity
			keystone.identity.backends.ldap.Identity
		Token format	UUID
			PKI
	Image (Glance)	Image activation	config drive+cloud init (guest operating system: Linux or Windows)
			config drive+VSAE (guest operating system: Linux or Windows)
			config drive+sysprep (guest operating system: Windows)
			ovf+VSAE (for PowerVC)
	Others	Operating System	Red Hat Enterprise Linux version 6.4 Red Hat Enterprise Linux version 6.5
		Database	DB2 version 10.5 SP3
		Queue	Qpid 0.26 Transient in memory

Table 11. IBM Cloud Manager with OpenStack and IaaS gateway support matrix (continued)

Compute node	Hypervisor	Type	KVM in Red Hat Enterprise Linux version 6.4 KVM in Red Hat Enterprise Linux version 6.5
			Hyper-V in Microsoft Hyper-V Server 2012 R2, or Microsoft Windows Server 2012 Standard, 2012 data center, 2012 R2 Standard, 2012 R2 data center
			PowerKVM 2.1
			PowerVC in Red Hat Enterprise Linux version 6.4 PowerVC in Red Hat Enterprise Linux version 6.5
		Driver	KVM: Nova.virt.libvirt.LibvirtDriver
			Hyper-V: Nova.virt.hyperv.HyperVDriver
			PowerVC: powervc.nova.driver.virt.powervc.driver.PowerVCDriver
			z/VM: nova.virt.zvm.ZVMDriver
	Network (Neutron)	Agent	KVM: OVSNeutronAgent
			Hyper-V: HyperVNeutronAgent
			PowerKVM: OVSNeutronAgent
			PowerVC: PowerVCNeutronAgent (Running on controller node)

Network considerations

Consider the following information before you deploy your cloud environment.

There are a number of factors to consider when you decide which network type is applicable. If you require a web interface for managing network aspects of your environment, you must consider the following web interface support:

Table 12. Web interface management capabilities

Web interface	**Network types**
Self-service portal	VLAN, Flat **Note:** DHCP cannot be configured or managed through the self-service portal.
Dashboard (based on Horizon)	VLAN, Flat, GRE, and DHCP

The network type that you can define is also limited by your network configuration and the type of hypervisor on the nodes in your environment.

Table 13. Supported network configuration

Hypervisor type	**Number of network interface cards per node**	**Network type**
Hyper-V	Two[1]	VLAN[2] or Flat
KVM	One	• Local • GRE[3]
	Two[1]	• Local • GRE[3] • Flat • VLAN[2]

Table 13. Supported network configuration (continued)

Hypervisor type	Number of network interface cards per node	Network type
PowerKVM	Two[1]	• VLAN[2] • Flat
PowerVC	Two[1]	VLAN **Note:** PowerVC can be configured without a virtual machine data network (only a single network card). DHCP is not supported.
z/VM	Two[1]	VLAN[2] or Flat

- [1] One network (default to eth0) must be dedicated as the management network. The second network (default to eth1) must be defined as the virtual machine data network. When you modify your environment file, you only need to update the *eth0* and *eth1* settings (in all occurrences) if they do not match your environment.
- [2] DHCP can be configured with GRE and flat networks. If using a VLAN network with DHCP, ensure that the single controller is running on dedicated physical hardware, instead of a virtual machine.
- [3] If one or more of the compute hypervisors is not KVM, then GRE cannot be used.

Note:

- The virtual machine data network must be on a dedicated interface. If there is an IP address on this data network interface, it is no longer ping-able after the node is deployed. Communication to the node must be done through the management network or another interface on the node.
- The `local` network type can be configured; however, the network traffic is limited to the current node. The minimum topology uses the `local` network option, by default.

Scalability and performance considerations

IBM Cloud Manager with OpenStack offers considerations with regard to scalability and performance within the cloud environment.

Server and concurrent user maximums

IBM Cloud Manager with OpenStack allows management of a configurable number of OpenStack compute nodes from one OpenStack controller node and self-service portal interface.

Depending on your cloud target, IBM Cloud Manager with OpenStack supports a different number of users and server instances in the environment. The number of concurrent servers that are supported per OpenStack compute node is highly variable based on the hardware configuration of the host and the resource consumption of the workload being deployed.

The following data represents the maximum scale with which this framework was tested.

Table 14. Server and concurrent user maximums for IBM Cloud Manager with OpenStack

Concurrent users	Concurrent instances	Compute nodes
50	1250	11
Environment specifications where testing was completed: • Controller +*n* compute topology with distributed database • Red Hat Enterprise Linux 6.5 guest image using 2048 MB of memory and 20 GB of disk space • Red Hat Enterprise Linux 6.5 x86 KVM hosts		

The VMware cloud manager allows management of one vCenter without the use of OpenStack, from the self-service portal interface. The following data represents the maximum scale with which this framework was tested.

Table 15. Server and concurrent user maximums for VMware cloud manager

Concurrent users	Concurrent instances
50	3000
For more information about VMware configuration maximums, see the following information: • VMware configuration maximums (v4) PDF at http://www.vmware.com/pdf/vsphere4/r40/vsp_40_config_max.pdf • VMware configuration maximums (v5) PDF at http://www.vmware.com/pdf/vsphere5/r50/vsphere-50-configuration-maximums.pdf • VMware configuration maximums (v5.5) PDF at http://www.vmware.com/pdf/vsphere5/r55/vsphere-55-configuration-maximums.pdf	

Neutron plug-in deployment considerations

The default chef deployment environment of IBM Cloud Manager with OpenStack creates an implementation of the Neutron service using ML2 as the default plug-in. This implementation is used to manage, for example Neutron Open vSwitch agents on PowerKVM or KVM compute nodes.

With IBM Cloud Manager with OpenStack 4.1, testing showed that using the Open vSwitch plug-in allowed for higher instance densities on the compute nodes. You might consider implementations using the Neutron Open vSwitch plug-in if you require maximum instance density on the compute nodes.

Review the Chef `environment.json` file that is documented in the "Deploying with KVM or QEMU compute nodes" on page 82 topic and note the **`openstack.network.core_plugin`** section.

```
"If you change the openstack.network.core_plugin property
from "neutron.plugins.ml2.plugin.Ml2Plugin" to
"neutron.plugins.openvswitch.ovs_neutron_plugin.OVSNeutronPluginV2",
you must remove the openstack.network.service_plugins property or
set the value of the property to [ ]. In the example environment,
the openstack.network.service_plugins property defaults to
"neutron.services.l3_router.l3_router_plugin.L3RouterPlugin".
If the core_plugin is set to OVS, the ML2 specific properties
in the environment are ignored."
```

The Chef Neutron cookbook of the IBM Cloud Manager with OpenStack 4.1 deployment server does not handle the change of the symbolic link of the Neutron `/etc/neutron/plugin.ini` correctly.

Make sure that it links to the appropriate version that you intend to use.

- For the Open vSwitch plug-in, you can accomplish that on the Neutron server:

  ```
  ln -sf /etc/neutron/plugins/openvswitch/ovs_neutron_plugin.ini /etc/neutron/plugin.ini
  ```

- For the ML2 plug-in:

  ```
  ln –sf /etc/neutron/plugins/ml2/ml2_conf.ini /etc/neutron/plugin.ini
  ```

- In case you set the symbolic link after the chef-client run, restart the Neutron service using the following command:

  ```
  service neutron-server restart
  ```

Consider the comments on message queue tuning in the "Customizing message queue attributes" on page 106 topic.

Chapter 3. Installing and uninstalling IBM Cloud Manager with OpenStack

You can install the IBM Cloud Manager with OpenStack application on 64-bit Linux by using one of three installation types: console, graphical, or silent.

The default installation type is console. To install IBM Cloud Manager with OpenStack using a non-default installation type, start the installer from the command line and run `-i <install type>` where `<install type>` is **`console`**, **`silent`**, or **`gui`** for graphical.

Installing IBM Cloud Manager with OpenStack on Linux

To install IBM Cloud Manager with OpenStack on the deployment server, follow the steps that are outlined in these topics.

When you install IBM Cloud Manager with OpenStack the installation process sets up the system as a deployment server. The installation process completes the following tasks:

- Installs and configures the chef-server
- Installs the IBM OpenStack RPM files in a yum repository on the deployment server
- Uploads the OpenStack cookbooks, roles, data bags and sample environments to the chef-server
- Installs an IBM extension to the chef knife command that enables IBM OpenStack topology deployments

Installing IBM Cloud Manager with OpenStack on Linux by using console installation

You can install IBM Cloud Manager with OpenStack on Linux by using a console.

Before you begin

The hostname for the deployment server must meet the following requirements:

- The host where you install the management server must have a fully qualified domain name that includes the domain suffix. For example, a fully qualified domain name is, *mydeploymentserver.ibm.com*, not *mydeploymentserver*. To verify that the deployment system hostname is a fully qualified domain name, run the following command:

  ```
  $ hostname
  ```

  If the hostname is a fully qualified domain name, it will return something like:

  ```
  $ mydeploymentserver.ibm.com
  ```

- The hostname must be resolvable. Add the hostname for the deployment server to the DNS system. To verify that a hostname is resolvable, run the following command:

  ```
  $ hostname -f
  ```

  If the hostname is resolvable it will return something like:

  ```
  $ mydeploymentserver.ibm.com
  ```

  If the hostname is not resolvable, refer to the networking documentation for the platform to get specific guidance on how to add the hostname to the DNS system.

© Copyright IBM Corp. 2012, 2014

Important: You need root authority to run the installer.

The product installer requires a functional Upstart environment. Upstart is used to manage the service state of the Chef server.

Note: This may not be the case in Docker containers or chroot environments that are common in deployment phases that use Anaconda/Kickstart.
You can verify that Upstart is running with the following command:

```
initctl list
```

About this task

To install IBM Cloud Manager with OpenStack, follow these steps:

Procedure

1. Download the following installation packages:
   - Linux: `cmwo410_xlinux_install.bin` and all `cmwo410_xlinux_install_pkgnn.tar.gz` files
   - Linux on Power: `cmwo410_plinux_install.bin` and all `cmwo410_plinux_install_pkgnn.tar.gz` files

   Copy all of the downloaded files to the same directory on the deployment system: *INSTALLER_LAUNCH_DIR*.
2. Optional: You can define a response file for the silent installation. The sample installation response file, `cmwo-install-sample.rsp` that is provided at the download site and on the installation media defines various keyword attributes that can be changed for the installation process. The sample response file includes information in the comments on the use and purpose of each keyword. It includes examples of how to specify the response file when you start the installation.
3. Navigate to the *INSTALLER_LAUNCH_DIR* directory from a terminal session and run the following commands:

Platform	Installation commands

Linux	`chmod +x ./cmwo410_xlinux_install.bin`
	
	`./cmwo410_xlinux_install.bin`

Linux on Power	`chmod +x ./cmwo410_plinux_install.bin`
	
	`./cmwo410_plinux_install.bin`

4. Follow the installation instructions.
5. The installation log, `IBM_Cloud_Manager_with_OpenStack_Install_MM_DD_YYYY_HH_mm_ss.log`, is located in the following directory: `/opt/ibm/cmwo/_installation/Logs/`

   Note: Low-level detail can be found in the following file: `/tmp/cmwo-installer.log`
6. To verify that the Chef server installed successfully, you can run the following command to check the Chef server status:

   ```
   chef-server-ctl status
   ```

   The command should return output similar to the following example:

```
[root@elvis-chef3 ~]# chef-server-ctl status
run: bookshelf: (pid 17198) 201s; run: log: (pid 17197) 201s
run: chef-expander: (pid 17145) 207s; run: log: (pid 17144) 207s
run: chef-server-webui: (pid 17357) 187s; run: log: (pid 17356) 187s
run: chef-solr: (pid 17112) 209s; run: log: (pid 17111) 209s
run: erchef: (pid 17535) 180s; run: log: (pid 17239) 199s
run: nginx: (pid 17518) 181s; run: log: (pid 17517) 181s
run: postgresql: (pid 17015) 220s; run: log: (pid 17014) 220s
run: rabbitmq: (pid 16685) 236s; run: log: (pid 16684) 236s
```

What to do next

When the installation completes, the components that are necessary for creating a cloud environment are installed.

Important: Download and install the latest fix pack for the IBM Cloud Manager with OpenStack from Fix Central. For more information, see "Getting fixes from Fix Central" on page 249.
To continue, you must create an operating system yum repository. Then, select and deploy a topology for your cloud configuration.

Installing IBM Cloud Manager with OpenStack on Linux by using graphical installation

You can install IBM Cloud Manager with OpenStack on Linux by using a graphical user interface.

Before you begin

The hostname for the deployment server must meet the following requirements:

- The host where you install the management server must have a fully qualified domain name that includes the domain suffix. For example, a fully qualified domain name is, *mydeploymentserver.ibm.com*, not *mydeploymentserver*. To verify that the deployment system hostname is a fully qualified domain name, run the following command:

  ```
  $ hostname
  ```

  If the hostname is a fully qualified domain name, it will return something like:

  ```
  $ mydeploymentserver.ibm.com
  ```

- The hostname must be resolvable. Add the hostname for the deployment server to the DNS system. To verify that a hostname is resolvable, run the following command:

  ```
  $ hostname -f
  ```

  If the hostname is resolvable it will return something like:

  ```
  $ mydeploymentserver.ibm.com
  ```

  If the hostname is not resolvable, refer to the networking documentation for the platform to get specific guidance on how to add the hostname to the DNS system.

Important: You need root authority to run the installer.

The product installer requires a functional Upstart environment. Upstart is used to manage the service state of the Chef server.

Note: This may not be the case in Docker containers or chroot environments that are common in deployment phases that use Anaconda/Kickstart.
You can verify that Upstart is running with the following command:

```
initctl list
```

About this task

To install IBM Cloud Manager with OpenStack, follow these steps:

Procedure

1. Download the following installation packages:
   - Linux: `cmwo410_xlinux_install.bin` and all `cmwo410_xlinux_install_pkgnn.tar.gz` files
   - Linux on Power: `cmwo410_plinux_install.bin` and all `cmwo410_plinux_install_pkgnn.tar.gz` files

   Copy all of the downloaded files to the same directory on the deployment system: *INSTALLER_LAUNCH_DIR*.
2. Optional: You can define a response file for the silent installation. The sample installation response file, `cmwo-install-sample.rsp` that is provided at the download site and on the installation media defines various keyword attributes that can be changed for the installation process. The sample response file includes information in the comments on the use and purpose of each keyword. It includes examples of how to specify the response file when you start the installation.
3. Navigate to the *INSTALLER_LAUNCH_DIR* directory from a terminal session and run the following commands:

Platform	Installation commands

Linux	`chmod +x ./cmwo410_xlinux_install.bin`
	
	`./cmwo410_xlinux_install.bin -i gui`

Linux on Power	`chmod +x ./cmwo410_plinux_install.bin`
	
	`./cmwo410_plinux_install.bin -i gui`

4. Follow the installation instructions.
5. The installation log, `IBM_Cloud_Manager_with_OpenStack_Install_MM_DD_YYYY_HH_mm_ss.log`, is located in the following directory: `/opt/ibm/cmwo/_installation/Logs/`

   Note: Low-level detail can be found in the following file: `/tmp/cmwo-installer.log`
6. To verify that the Chef server installed successfully, you can run the following command to check the Chef server status:

   ```
   chef-server-ctl status
   ```

   The command should return output similar to the following example:

   ```
   [root@elvis-chef3 ⌂]# chef-server-ctl status
   run: bookshelf: (pid 17198) 201s; run: log: (pid 17197) 201s
   run: chef-expander: (pid 17145) 207s; run: log: (pid 17144) 207s
   run: chef-server-webui: (pid 17357) 187s; run: log: (pid 17356) 187s
   run: chef-solr: (pid 17112) 209s; run: log: (pid 17111) 209s
   run: erchef: (pid 17535) 180s; run: log: (pid 17239) 199s
   run: nginx: (pid 17518) 181s; run: log: (pid 17517) 181s
   run: postgresql: (pid 17015) 220s; run: log: (pid 17014) 220s
   run: rabbitmq: (pid 16685) 236s; run: log: (pid 16684) 236s
   ```

What to do next

When the installation completes, the components that are necessary for creating a cloud environment are installed.

Important: Download and install the latest fix pack for the IBM Cloud Manager with OpenStack from Fix Central. For more information, see "Getting fixes from Fix Central" on page 249.
To continue, you must create an operating system yum repository. Then, select and deploy a topology for your cloud configuration.

Installing IBM Cloud Manager with OpenStack on Linux by using silent installation

You can use a response file to install IBM Cloud Manager with OpenStack on Linux silently.

Before you begin

The hostname for the deployment server must meet the following requirements:

- The host where you install the management server must have a fully qualified domain name that includes the domain suffix. For example, a fully qualified domain name is, *mydeploymentserver.ibm.com*, not *mydeploymentserver*. To verify that the deployment system hostname is a fully qualified domain name, run the following command:

  ```
  $ hostname
  ```

  If the hostname is a fully qualified domain name, it will return something like:

  ```
  $ mydeploymentserver.ibm.com
  ```

- The hostname must be resolvable. Add the hostname for the deployment server to the DNS system. To verify that a hostname is resolvable, run the following command:

  ```
  $ hostname -f
  ```

  If the hostname is resolvable it will return something like:

  ```
  $ mydeploymentserver.ibm.com
  ```

  If the hostname is not resolvable, refer to the networking documentation for the platform to get specific guidance on how to add the hostname to the DNS system.

Important: You need root authority to run the installer.

The product installer requires a functional Upstart environment. Upstart is used to manage the service state of the Chef server.

Note: This may not be the case in Docker containers or chroot environments that are common in deployment phases that use Anaconda/Kickstart.
You can verify that Upstart is running with the following command:

```
initctl list
```

About this task

Procedure

1. Download the following installation packages:
   - Linux: `cmwo410_xlinux_install.bin` and all `cmwo410_xlinux_install_pkgnn.tar.gz` files

- Linux on Power: `cmwo410_plinux_install.bin` and all `cmwo410_plinux_install_pkgnn.tar.gz` files

Copy all of the downloaded files to the same directory on the deployment system: *INSTALLER_LAUNCH_DIR*.

2. Define a response file for the silent installation. The sample installation response file, `cmwo-install-sample.rsp` that is provided at the download site and on the installation media defines various keyword attributes that can be changed for the installation process. The sample response file includes information in the comments on the use and purpose of each keyword. It includes examples of how to specify the response file when you start the installation.
3. Navigate to the *INSTALLER_LAUNCH_DIR* directory from a terminal session and run the following commands:

Platform	Silent installation command

Linux	`chmod +x ./cmwo410_xlinux_install.bin`
	
	`echo LICENSE_ACCEPTED=true > ./installer.rsp; ./cmwo410_xlinux_install.bin -i silent -f ./installer.rsp`

Linux on Power	`chmod +x ./cmwo410_plinux_install.bin`
	
	`echo LICENSE_ACCEPTED=true > ./installer.rsp; ./cmwo410_plinux_install.bin -i silent -f ./installer.rsp`

4. The installation process takes several minutes to complete.
5. Check the status of the installation by running the following command:

   ```
   echo $?
   ```

   If the return code is `0`, the installation was successful. If the return code is something other than `0`, refer to the installation logs to determine the problem.
6. The installation log, `IBM_Cloud_Manager_with_OpenStack_Install_MM_DD_YYYY_HH_mm_ss.log`, is located in the following directory: *INSTALLER_LAUNCH_DIR*.

   Note: Low-level detail can be found in the following file: `/tmp/cmwo-installer.log`
7. To verify that the Chef server installed successfully, you can run the following command to check the Chef server status:

   ```
   chef-server-ctl status
   ```

   The command should return output similar to the following example:

   ```
   [root@elvis-chef3 ~]# chef-server-ctl status
   run: bookshelf: (pid 17198) 201s; run: log: (pid 17197) 201s
   run: chef-expander: (pid 17145) 207s; run: log: (pid 17144) 207s
   run: chef-server-webui: (pid 17357) 187s; run: log: (pid 17356) 187s
   run: chef-solr: (pid 17112) 209s; run: log: (pid 17111) 209s
   run: erchef: (pid 17535) 180s; run: log: (pid 17239) 199s
   run: nginx: (pid 17518) 181s; run: log: (pid 17517) 181s
   run: postgresql: (pid 17015) 220s; run: log: (pid 17014) 220s
   run: rabbitmq: (pid 16685) 236s; run: log: (pid 16684) 236s
   ```

What to do next

When the installation completes, the components that are necessary for creating a cloud environment are installed.

Important: Download and install the latest fix pack for the IBM Cloud Manager with OpenStack from Fix Central. For more information, see "Getting fixes from Fix Central" on page 249.

To continue, you must create an operating system yum repository. Then, select and deploy a topology for your cloud configuration.

Installing IBM Cloud Manager with OpenStack self-service portal on Windows

You have the option to install IBM Cloud Manager with OpenStack self-service portal on Microsoft Windows. If you want to use the web-based self-service portal, follow the steps that are outlined in these topics.

Notes:

- Any Microsoft Windows user with administrator authority can install IBM Cloud Manager with OpenStack self-service portal.
- In previous releases, IBM Cloud Manager with OpenStack self-service portal was known as IBM SmartCloud® Entry web interface.

Graphical installation (default)

You can complete the installation through a graphical user interface.

Procedure

1. Go to the location of the Windows installer and double-click the installer icon.
2. On the first screen, select the language that you want the installer that is displayed in from the drop-down menu and click **OK**.

   Note:
   - If you see an error message similar to `Executing java permission denied`, verify that you are using the correct installer for your operating system.
   - If you see Windows error 183, you are using am unsupported 32-bit operating system.
3. On the Introduction screen, read the introduction for information about the installer and click **Next**.
4. On the License Agreement window, read the contract and accept the license. Click **Next**.
5. On the Choose Shortcut Location window, select the check boxes next to the locations where you want a IBM Cloud Manager with OpenStack link and then click **Next**.
6. To choose the location for your installation, click **Choose** on the Choose Install Folder window, specify your chosen location, and click **Next**.

   Note: The installation path cannot contain non-English characters.
7. To select the location of your property files, on the Choose Property File Install Folder window, click **Choose** and click **Next**.

   Note: The installation path cannot contain non-English characters.
8. Review your selected options on the Pre-Installation Summary window. If you are satisfied with your selections, click **Next**.
9. After the installation is finished, the Install Finished window opens.

   You have three options:
   - You can follow the installer to configure important properties.
   - You can use the installer to migrate the property files and database from a previous release.
   - You can manually migrate the configuration property files and data from a previous release at a later time.

   Click **Next** after making your selection.

Console installation

You can complete the installation by using a console.

Procedure

1. Click **Start** and then select **Run**.
2. In the run window, type `cmd` and click **OK**.
3. In the console window, navigate to the windows installer and run the following command:
   ```
   cmwo_ui410_windows_installer.exe -i console
   ```
4. On the Choose Locale screen, choose the language for the installer by typing the number corresponding to your chosen language and pressing Enter. If your preferred language is the default, press Enter.

   Note:
   - If you see an error message similar to `Executing java permission denied`, verify that you are using the correct installer for your operating system.
   - If you see Windows error 183, you are using an unsupported 32-bit operating system.
5. On the Introduction screen, read the information about the installer and press Enter.
6. On the License Agreement screen, read the contract and accept the license by entering `Y`.
7. Choose your location for the links to IBM Cloud Manager with OpenStack. To choose any combination of locations, separate each number with a comma.
8. On the Choose Install Folder screen, press Enter to accept the default location. To choose a different location for the application files, type in the full path to that folder and press Enter.

   Note: The installation path cannot contain non-English characters.
9. On the Choose Property File Install Folder screen, press Enter to accept the default location. To choose a different location for the property files, type in the full path to that folder and press Enter.

   Note: The installation path cannot contain non-English characters.
10. Review your selected options on the Pre-Installation Summary screen. If you are satisfied with your selections, press Enter.
11. When the installation is complete, the Install Finished window opens.

    You have three options:
    - You can follow the installer to configure important properties.
    - You can use the installer to migrate the property files and database from a previous release.
    - You can manually migrate the configuration property files and data from a previous release later.

Silent installation

Create a silent installation response file and use it to install IBM Cloud Manager with OpenStack self-service portal on Windows.

About this task

To create a silent installation response file, follow either “Console installation” or “Graphical installation (default)” on page 25 on Windows. On the final screen, you can create a silent installation response file. Creating a silent installation response file creates a property file that is called `installer.properties` in the installation location.

To install on Windows by using a silent installation response file, follow these steps:

Procedure

1. Open the response file and double check all of the properties. Notice that the passwords are not in the file. Notice that all paths have "\\" instead of a single "\". Make any updates and then save the file.
2. Press Start and then Run.
3. In the run prompt, type `cmd` and press Enter.
4. A console window opens. In the console window, navigate to the Windows installer.
5. Run the installer with the command `cmwo_ui410_windows_installer.exe -i silent -f <response file location>`, where `<response file location>` is the path to the response file, which defines the installation.

   Note:

   - If you see an error message similar to `Executing java permission denied`, verify that you are using the correct installer for your operating system.
   - If you see Windows error 183, you are using am unsupported 32-bit operating system.

Installing and uninstalling IBM Cloud Manager with OpenStack agents

IBM Cloud Manager with OpenStack can manage hypervisors by using OpenStack technology. To manage certain hypervisors, you must install a IBM Cloud Manager with OpenStack agent on the endpoint server. This section contains the procedures for installing and uninstalling these IBM Cloud Manager with OpenStack agents.

Installing and uninstalling the IBM Cloud Manager with OpenStack Hyper-V Agent

IBM Cloud Manager with OpenStack can manage Microsoft Hyper-V hypervisors from OpenStack technology. To manage these hypervisors, an IBM Cloud Manager with OpenStack Hyper-V Agent must be installed on the Hyper-V endpoint server.

This IBM Cloud Manager with OpenStack Hyper-V Agent contains packaging of the OpenStack technology that is required to provision to the Hyper-V server. The IBM Cloud Manager with OpenStack Hyper-V Agent can be installed on a Microsoft Hyper-V Server 2012 R2 or Microsoft Windows Server 2012 R2 with the Hyper-V role enabled. The IBM Cloud Manager with OpenStack Hyper-V Agent must be installed on all managed compute nodes. The IBM Cloud Manager with OpenStack Hyper-V Agent installation is packaged as a Microsoft Windows Installer that can be run as an installation wizard, or in silent mode. This installation installs the required OpenStack components on to the Hyper-V server and configures them to run as Microsoft Windows services.

IBM Cloud Manager with OpenStack Hyper-V Agent Installation Prerequisites

Use the following steps to prepare your environment for installation.

Preparing Your Hyper-V Server for Installation

On each Hyper-V server that is managed from IBM Cloud Manager with OpenStack, a Network Time Service (NTP) must be synchronized with the Hyper-V appliance system that is running the IBM Cloud Manager with OpenStack server. Complete the following steps to install the NTP service.

1. Access the NTP installation package, `ntp-4.2.6p5-ibm-win32-setup.exe`, in the root directory of the IBM Cloud Manager with OpenStack installation media.
2. Run the `ntp-4.2.6p5-ibm-win32-setup.exe` file to install the NTP service.
3. After the NTP package is installed, specify the NTP server IP address or host name in `C:\ntp\etc\ntp.conf`.

```
# your local system clock, could be used as a backup
# (this is only useful if you need to distribute time no
matter how good or bad it is)

server x.x.x.x

# but it should operate at a high stratum level to let the
```

4. Navigate to **Control Panel** > **System and Security** > **Administative Tools** > **Services**, and start **Network Time Protocol Daemon** service.

See the following document in the OpenStack Compute Administration Guide for more details: Hyper-V Virtualization Platform.

Note: Before you can install the IBM Cloud Manager with OpenStack Hyper-V Agent on Microsoft Windows Server 2012 R2, ensure that the Hyper-V role is enabled on the server.

Preparing the Host

The host must be a domain joined computer to support live migration. If the host is not a domain joined computer, you might see the following error display during installation:

```
Failed to modify service settings. Live migrations can be enabled only on a domain joined
computer.
```

Preparing the User

Add the user who installs the Hyper-V Agent for IBM Cloud Manager with OpenStack to the Hyper-V Administrators group.

Note: If you are creating the user profile for the first time, the Hyper-V server must be restarted before you install the IBM Cloud Manager with OpenStack Hyper-V Agent.

If the user plans to uninstall the Hyper-V Agent in the future, ensure that the user has permission to each of the installation directories on the system.

Installing the IBM Cloud Manager with OpenStack Hyper-V Agent

Follow these steps to install the IBM Cloud Manager with OpenStack Hyper-V Agent on Microsoft Windows Server 2012 R2 or Microsoft Hyper-V Server 2012 R2.

Overview of the installation

The installation completes the following steps:

- Create a product installation directory
- Create a Hyper-V external virtual switch (optional)
- Configure Hyper-V Live Migration settings for this host (optional)
- Install an independent Python environment to avoid conflicts with existing applications

  Note: This embedded Python environment is only intended for use by the IBM Cloud Manager with OpenStack Hyper-V Agent. Do not attempt to access this environment from any other application. This independent environment is designed to coexist with any preexisting Python environments that might already be installed on the system. Do not install new Python modules into the embedded Python environment.
- Install the required Python modules/packages required by the application
- Install and configure the OpenStack Nova Compute service

- Install and configure the OpenStack Hyper-V Neutron agent for networking
- Register two Windows services, which are set to auto-start by default:
  - **IBM Cloud Manager with OpenStack Network Service**
  - **IBM Cloud Manager with OpenStack Compute Service**

Important:

- If you migrate IBM SmartCloud Entry 3.2 to IBM Cloud Manager with OpenStack 4.1, the neutron password is *neutron* by default. If you install a new Hyper-V Agent, you must change the default password for neutron to *neutron* to be consistent with the controller node.
- The 3.1 Hyper-V agent installer does not prevent users from installing a previous version over the more recent Hyper-V agent.

Creating installation or uninstallation logs:

Follow these steps to create installation or uninstallation logs for use during the installation or uninstallation of IBM Cloud Manager with OpenStack Hyper-V Agent on Microsoft Windows Server 2012 R2.

About this task

Because the IBM Cloud Manager with OpenStack Hyper-V Agent installer is MSI-based, you can create an installation or uninstallation log by starting the installer with the `msiexec` command with the correct parameters. Detailed information about creating logs can be found here: How to enable Windows Installer logging.

Graphical Installation:

Follow these steps to install the IBM Cloud Manager with OpenStack Hyper-V Agent by using the graphical installation wizard.

Procedure

1. Download the latest fix for the IBM Cloud Manager with OpenStack Hyper-V Agent from Fix Central. For more information, see "Applying fixes and updates for the self-service portal on Windows" on page 39.

   Note: The IBM Cloud Manager with OpenStack Hyper-V Agent and the OpenStack controller node must be at the same level, either the GA level, or the fix level.
2. Locate the installation image, and double-click `IBM Cloud Manager with OpenStack Hyper-V Agent.msi` to start the installation wizard.
3. Follow the instructions that are provided by the installation wizard. Agree to the license terms, provide an installation destination directory, and select the type of setup you want to use.

   Note: The IBM Cloud Manager with OpenStack Hyper-V Agent must be installed to the local C: disk of the server. However, the instance directory (Instances Path) that is used to store virtual machine instance data can be on any local disk.
4. Use the **Nova Compute Configuration** window to configure the compute agent parameters. You can leave the default values provided and manually configure the `nova.conf` file, which is in the `etc\nova` folder, later. The following table shows the mappings between areas from this dialog and properties in the `nova.conf` file.

Table 16. Nova Compute Configuration fields and related properties in `nova.conf`. Mapping of field names in the installation wizard, related properties in the `nova.conf` file, and installation wizard default values

Area in dialog	Property in **nova.conf**	Installation wizard default values
Glance API Server	*glance_host*	*appliance_mgmt_ip* **Note:** Where *appliance_mgmt_ip* is the IP address of the network interface on the appliance.
Port (after Glance API Server)	*glance_port*	9292
Qpid Server	*qpid_hostname*	*appliance_mgmt_ip* **Note:** Where *appliance_mgmt_ip* is the IP address of the network interface on the appliance.
Port (after Qpid Server)	*qpid_port*	5671
Qpid User Name	*qpid_username*	qpidclient
Qpid Password	*qpid_password*	openstack1
Instances Path	*instances_path*	`C:\Program Files (x86)\IBM\Cloud Manager with OpenStack\Hyper-V Agent\`

Note: The property **qpid_hostname** is also written to the file `hyperv_neutron_agent.ini`. For more detailed descriptions about properties in the `nova.conf` file, see the List of configuration options topic in the OpenStack Compute Administration Manual.

5. Use the **Nova Compute Advanced Configuration** window to configure the advanced compute agent parameters. Select **Use Cow Images** to enable the copy on write feature and speed up deployment times. You can also manually configure the `nova.conf` file, which is in the `etc\nova` folder, later. The following table shows the mappings between areas from this dialog and properties in the `nova.conf` file.

Table 17. Nova Compute Advanced Configuration fields and related properties in `nova.conf`. Mapping of field names in installation wizard and related properties in the `nova.conf` file

Area in dialog	Property in **nova.conf**
Use Cow Images	*use_cow_images*
Verbose Logging	*verbose*
Log file	*logdir*

6. Use the **Neutron Network Configuration** window to configure network agent parameters. The installation wizard applies changes that are made to this window to properties in both the `nova.conf` and `neutron.conf` files. You can leave the default values provided and manually configure the `nova.conf` and `neutron.conf` files later. These properties files are in the `etc\nova` and `etc\neutron` folders. The following table shows the mappings between areas from this dialog and properties in the `nova.conf` and `neutron.conf` files.

Table 18. Neutron Network Configuration fields and related properties in `nova.conf` and `neutron.conf`. Mapping of field names in the installation wizard, related properties in the `nova.conf` file and `neutron.conf` file, and installation wizard default values

Area in dialog	Property in **nova.conf** and **neutron.conf**	Installation wizard default values
Neutron URL	*neutron_url*	http://*appliance_mgmt_ip*:9696 **Note:** Where *appliance_mgmt_ip* is the IP address of the network interface on the appliance.

Table 18. Neutron Network Configuration fields and related properties in `nova.conf` and `neutron.conf` (continued). Mapping of field names in the installation wizard, related properties in the `nova.conf` file and `neutron.conf` file, and installation wizard default values

Area in dialog	Property in **nova.conf** and **neutron.conf**	Installation wizard default values
Username	*neutron_admin_username*	neutron
Password	*neutron_admin_password*	openstack-network
Tenant Name	*neutron_admin_tenant_name*	service
Region name	*neutron_region_name*	RegionOne
Authentication Url	*neutron_admin_auth_url*	http://*appliance_mgmt_ip*:35357/v2.0 **Note:** Where *appliance_mgmt_ip* is the IP address of the network interface on the appliance.

Note: The file, `hyperv_neutron_agent.ini`, is also updated in the background by the installer. The following properties are updated:

- `rpc_backend=neutron.openstack.common.rpc.impl.qpid`
- `verbose=true`
- `debug=true`
- `control_exchange=neutron`
- `physical_network_vswitch_mappings = *:external`

7. Use the **Hyper-V Live Migration Settings** window to configure the Live Migration settings for the host.

   Note: IBM Cloud Manager with OpenStack with Hyper-V supports a Shared nothing live migration. To use live migration, the Hyper-V server must belong a common domain. You can also skip this step and configure the Hyper-V Live Migration setting manually later. See the following document in the OpenStack Compute Administration Guide for more details: Hyper-V Virtualization Platform

   Note: When you provide a user name in the **Nova compute service user** field, ensure that the user you select is a domain user.
8. Use the **Virtual Switch Configuration** window to configure Virtual Switch settings. For more information about Hyper-V Switches, see the following topic:Hyper-V Virtual Switch Overview. If no existing virtual switches are detected, create a new virtual switch. To add a new virtual switch, there must be at least one physical network adapter that is not bound to any existing virtual switch. To manually determine whether a physical network adapter is available, you can use the PowerShell command `get-networkadapter -physical'` to see all physical network adapters. Next, you can use the PowerShell command `get-vmswitch` to see all network adapters that are already in use. A new virtual switch must be exclusively associated with a physical network adapter. No two virtual switches can be associated with the same physical network adapter on the host. See the following document in the OpenStack Compute Administration Guide for more details on using the PowerShell command: Hyper-V Virtualization Platform

   Note: The **Shared for management** property determines whether the Hyper-V Agent can use this physical network adapter to manage network traffic.
9. After you complete the information in the installation wizard, the installation begins.

Silent installation:

Because the IBM Cloud Manager with OpenStack Hyper-V Agent is installed by using the Microsoft Installer (MSI), you can start MSI directly without using the graphical installation wizard. This process is

called silent (unattended) installation, and is useful for installing this program over a network on a remote system from a shared drive on a LAN server. Follow these steps to silently install IBM Cloud Manager with OpenStack Hyper-V Agent on Microsoft Windows Server 2012 R2.

Before you begin

If you choose to install silently, you must create an IBM Cloud Manager with OpenStack Hyper-V Agent silent installation response file and use it to drive the installation. The file is addressed in the following sections, and a sample response file is provided for reference

About this task

To install silently, you either provide installation parameters through the command line, or use an INI file (response file) to specify all the parameters in a single file. For both cases, use the `msiexec` command to start the installation. For more information about this command, see Msiexec (command-line options)

Installing silently with the command line:

Follow these steps to silently install IBM Cloud Manager with OpenStack Hyper-V Agent by using the command line.

Before you begin

If you choose to install silently, you must create an IBM Cloud Manager with OpenStack Hyper-V Agent silent installation response file and use it to drive the installation. The file is addressed in the following sections, and a sample response file is provided for reference.

Procedure

1. Download the latest fix for the IBM Cloud Manager with OpenStack Hyper-V Agent from Fix Central. For more information, see "Applying fixes and updates for the self-service portal on Windows" on page 39.

   Note: The IBM Cloud Manager with OpenStack Hyper-V Agent and the OpenStack controller node must be at the same level, either the GA level, or the fix level.
2. To start the installation through the command line directly, open a command prompt and input the following parameters, substituting the IP address and port with your own: `msiexec /i "Hyper-V-OpenStack installer.msi" /qn GLANCE_SERVER="127.0.0.1"  GLANCE_SVR_PORT = "9292"`

   Tip: The `/i` parameter means to install, and the `/qn` parameter means that the installation is done with no GUI.

   Tip: You can provide as many parameters as you like in the format `key=value`, separated by a space character at the end of the command.

   Note: The IBM Cloud Manager with OpenStack Hyper-V Agent must be installed to the local C: disk of the server. However, the instance directory (Instances Path) that is used to store virtual machine instance data can be on any local disk.

Example

The following table shows the mappings between parameters in the response file and properties in the `nova.conf` file.

Table 19. Response file parameters and related properties in `nova.conf`. This table shows the mappings between parameters in the response file and properties in the `nova.conf` file.

Parameters in the response file	Property in `nova.conf`
GLANCE_SERVER	glance_host
GLANCE_SVR_PORT	glance_port
QPID_SERVER	qpid_hostname
QPID_SVR_PORT	qpid_port
QPID_UNAME	qpid_username
QPID_PWD	qpid_password
INSTANCES1	instances_path
COW	use_cow_images
VERBOSE	verbose
NOVA_LOG_PATH	Logdir

The following table shows the mappings between parameters in the response file and properties in both the `nova.conf` and `neutron.conf` files

Table 20. Response file parameters and related properties in `nova.conf` and `neutron.conf` files. This table shows the mappings between parameters in the response file and properties in the `nova.conf` and `neutron.conf` files.

Parameters in the response file	Property in `nova.conf` and `neutron.conf`
NEUTRON_URL	neutron_url
ADMIN_USERNAME	neutron_admin_username
ADMIN_PASSWORD	neutron_admin_password
ADMIN_TENANT_NAME	neutron_admin_tenant_name
REGION_NAME	neutron_region_name
NEUTRON_AUTH_URL	neutron_admin_auth_url
ALLOW_RESIZE_TO_SAME_HOST	allow_resize_to_same_host
NEUTRON_AUTH_STRATEGY	neutron_auth_strategy

Note: The property `;AgreeToLicense` in the response file specifies your agreement to the license for the application. Its default value is set to `no`. You must specify `yes` to run the silent installation successfully.

Installing silently with a response file:

Follow these steps to run a silent installation by using a response file.

Before you begin

If you choose to install silently, you need to create an IBM Cloud Manager with OpenStack Hyper-V Agent silent installation response file and use it to drive the installation. The file is addressed in the following sections, and a sample response file is provided for reference.

Procedure

1. Download the latest fix for the IBM Cloud Manager with OpenStack Hyper-V Agent from Fix Central. For more information, see "Applying fixes and updates for the self-service portal on Windows" on page 39.

Note: The IBM Cloud Manager with OpenStack Hyper-V Agent and the OpenStack controller node must be at the same level, either the GA level, or the fix level.

2. To run the installation through the response file, you must first enter the correct parameters in your locally saved copy of the response file. See the sample response file that is provided for more details.

   Note: The IBM Cloud Manager with OpenStack Hyper-V Agent must be installed to the local C: disk of the server. However, the instance directory (Instances Path) that is used to store virtual machine instance data can be on any local disk.

3. Next, open a command prompt and input the following statement:

   ```
   msiexec /i "Hyper-V-OpenStack installer.msi" /qn USEINI="absolute path to responsefile"
   ```

Example

The sample response file provides an example INI file that can be used to drive a silent installation. This example shows all properties that are available during a graphical installation of the IBM Cloud Manager with OpenStack Hyper-V Agent.

```
[Response]
#indicate whether you agree with the liscense and its default value is "no"
AgreeToLicense=yes
GLANCE_SERVER = mySCEApplianceHostOrIP
GLANCE_SVR_PORT = 9292
QPID_SERVER = mySCEApplianceHostOrIP
QPID_SVR_PORT = 5671
QPID_UNAME = qpidclient
QPID_PWD = openstack1
NEUTRON_URL = http://9.123.106.93:9696
ADMIN_USERNAME = neutron
ADMIN_PASSWORD = neutron
ADMIN_TENANT_NAME = service
REGION_NAME = regionOne
NEUTRON_AUTH_URL = http://9.123.106.93:35357/v2.0
NEUTRON_URL_TIMEOUT = 30
ALLOW_RESIZE_TO_SAME_HOST = True
NEUTRON_AUTH_STRATEGY = keystone
INSTANCES1=C:\Program Files (x86)\IBM\Cloud Manager with OpenStack\Hyper-V Agent\instances
COW = true
ENABLELOG = 1
VERBOSE = true
NOVA_LOG_PATH = = C:\Program Files (x86)\IBM\Cloud Manager with OpenStack\Hyper-V Agent\log\nova\

#The path that IBM SCE Hyper-V Agent will be installed.
INSTALLDIR = C:\Program Files (x86)\IBM\Cloud Manager with OpenStack\Hyper-V Agent\
#The string coming after a period is the UUID of the DIM used by the installer internally.
#You can actually ignore it during the installation.
#(IntOpt)Live Migration authentication type you choose. It has two optional values, "0" and "1".
#"0" stands for "Kerberos", and "1" stands for "CredSSP".
LIVEMIGRAUTHTYPE.EDDDE39A_8D99_430B_BFF6_7644F125D2A1 = 0
NOVACOMPUTESERVICEUSER.EDDDE39A_8D99_430B_BFF6_7644F125D2A1 =
#(IntOpt)The max active virtual machine migrations.
MAXACTIVEVSMIGR.EDDDE39A_8D99_430B_BFF6_7644F125D2A1 =
(IntOpt)The max active storage migrations.
MAXACTIVESTORAGEMIGR.EDDDE39A_8D99_430B_BFF6_7644F125D2A1 =
#(IntOpt)The networks you migrate from. It has two optional values, "0" and "1".
#Set "1" means you can migrate from any network, and the following property MIGRNETWORKS will be disabled.
#Set "0" means you have to specify the network you migrate from by stating the following property.
MIGRNETWORKSANY_INTERNAL.EDDDE39A_8D99_430B_BFF6_7644F125D2A1 = 1;
#(IntOpt)Specific network you migrate from.
#This property only make sense when the MIGRNETWORKSANY_INTERNAL is set to "0".
MIGRNETWORKS.EDDDE39A_8D99_430B_BFF6_7644F125D2A1 = 10.10.10.1/32

#(IntOpt) It has two optional values, "0" and "1".
#Set to "1" means you will skip the virtual switch configuration, and the following four properties
#SKIPNOVACONF, ADDVSWITCH, VSWITCHNAME, VSWITCHNETADAPTER,NEWVSWITCHNAME, VSWITCHSHARED will be disabled.
```

```
#Set "0" means you configure the virtual switch during installation.
SKIPNOVACONF.D5E17CCE_FABA_4230_9715_2DF2AA168F6C = 1
#(IntOpt)Whether to add a virtual switch. It has two optional values, "0" and "1".
#Set "1" means a newvirtual switch will be add, and the following property VSWITCHNAME will be disabled.
#Set "0" means you will use an existing virtual switch,
#and the following property VSWITCHNETADAPTER and NEWVSWITCHNAME will be disabled.
ADDVSWITCH.D5E17CCE_FABA_4230_9715_2DF2AA168F6C = 0
#(StrOpt)The name of an existing virtual switch you choose.
VSWITCHNAME.D5E17CCE_FABA_4230_9715_2DF2AA168F6C =
#(StrOpt)The adapter you use to create a new virtual switch.
VSWITCHNETADAPTER.D5E17CCE_FABA_4230_9715_2DF2AA168F6C =
#(StrOpt)The name you use to create a new virtual switch.
NEWVSWITCHNAME.D5E17CCE_FABA_4230_9715_2DF2AA168F6C =
#(IntOpt) It has two optional values, "0" and "1".
#Set to "1" to allow management operating system toshare this network adapter. Set to "0" to disable it
VSWITCHSHARED.D5E17CCE_FABA_4230_9715_2DF2AA168F6C =

# End of the file
```

Note: The property `;AgreeToLicense` in the response file specifies your agreement to the license for the application. Its default value is set to `no`. You must specify `yes` to run the silent installation successfully.

Note: A response file must start with `[Response]`, followed by any parameters in format `key=value`.

Uninstalling the IBM Cloud Manager with OpenStack Hyper-V Agent

The IBM Cloud Manager with OpenStack Hyper-V Agent uninstaller supports uninstallation by using the Microsoft Windows Control Panel and command line.

About this task

Use the following steps to uninstall the IBM Cloud Manager with OpenStack Hyper-V Agent on Microsoft Windows:

Procedure

1. Shut down the IBM Cloud Manager with OpenStack Hyper-V Agent by using the Microsoft Windows Services panel or the appropriate command line.
2. Navigate to Start > Control Panel > Programs > Uninstall a program
3. Select **IBM OpenStack Hyper-V Agent** and click **Uninstall**.
4. Follow the instructions on the uninstallation wizard to complete the operation.

Results

After the IBM Cloud Manager with OpenStack Hyper-V Agent is uninstalled, the uninstaller will back up the following files to the `%USERPROFILE%/AppData` folder:

- `nova.conf`
- `neutron.conf`
- `hyperv_neutron_agent.ini`

Note: The uninstaller does not delete any existing instances that you started. These instances are saved in the instances folder.

Uninstalling the IBM Cloud Manager with OpenStack Hyper-V Agent on Microsoft Hyper-V Server 2012 R2

The IBM Cloud Manager with OpenStack Hyper-V Agent uninstaller supports uninstallation from the Microsoft Hyper-V Server 2012 R2 command line.

About this task

Use the following steps to uninstall The IBM Cloud Manager with OpenStack Hyper-V Agent on Microsoft Hyper-V Server 2012 R2:

Procedure

1. Shut down the IBM Cloud Manager with OpenStack Hyper-V Agent by using the Microsoft Windows Services panel or the appropriate command line.
2. Open a command-line window and enter the following command: `WMIC`.
3. Enter the following command to display a list of installed products: `product get name`.
4. Enter the command `product name call uninstall`, where `product name` is the name of the IBM Cloud Manager with OpenStack Hyper-V Agent installed product.
5. Enter `Y` to uninstall.

Results

After the IBM Cloud Manager with OpenStack Hyper-V Agent is uninstalled, the uninstaller will back up the following files to the `%USERPROFILE%/AppData` folder:

- `nova.conf`
- `neutron.conf`
- `hyperv_neutron_agent.ini`

Note: The uninstaller does not delete any existing instances that you started. These instances are saved in the instances folder.

Configuring operating system yum repositories on the deployment server

After you install IBM Cloud Manager with OpenStack on the deployment server, you must determine how your Linux nodes can access yum repositories that contain operating system packages that are required when OpenStack is deployed.

About this task

You have the following options for configuring yum repositories:

- You can use your own yum repositories. If you choose this option, the yum repositories must be configured on the nodes before deploying OpenStack. If you choose to use your own yum repositories, the following steps are unnecessary.
- You can use the following instructions to allow the deployment server to serve as a yum repository. When the deployment server is configured as a yum repository, the yum repositories are configured on the nodes automatically when OpenStack is deployed.

Note: RDO repositories are not supported because they contain packages that conflict with IBM Cloud Manager with OpenStack.

The operating system installation media for the node system contains all of the packages that are required to deploy OpenStack. You can access the operating system installation media by inserting the installation media for the node into the DVD device or by mounting the installation ISO image for the node system.

The deployment server is configured to serve these packages from the following locations:

Packages:

<*installation directory*>/yum-repo/operatingsystem/<*platform*><*platform version*>/<*architecture*>

Package updates:

<*installation directory*>/yum-repo/operatingsystem/<*platform*><*platform version*>/<*architecture*>/updates

Where the variables have the following definitions:

installation directory
: The location where IBM Cloud Manager with OpenStack is installed, typically `/opt/ibm/cmwo`.

platform
: The platform of the node system. For example, `redhat` or `ibm_powerkvm`.

platform version
: The version of the operating system on the node system. For example, `6.5` or `2.1`.

architecture
: The architecture for the node system. For example, `x86_64` or `ppc64`.

An x86 64-bit Red Hat Enterprise Linux 6.5 node system requires the packages to be in the following directories:

```
/opt/ibm/cmwo/yum-repo/operatingsystem/redhat6.5/x86_64
/opt/ibm/cmwo/yum-repo/operatingsystem/redhat6.5/x86_64/updates
```

An IBM PowerKVM 2.1 node system requires the packages and updates to be in the following directories:

```
/opt/ibm/cmwo/yum-repo/operatingsystem/ibm_powerkvm2.1/ppc64
/opt/ibm/cmwo/yum-repo/operatingsystem/ibm_powerkvm2.1/ppc64/updates
```

To configure the deployment server, complete the following steps:

Note: The following steps use the Red Hat Enterprise Linux 6.5 version and x86 64-bit architecture. However, similar steps apply for other supported options.

Procedure

1. Create the directories for packages and package updates on the deployment server as appropriate for your environment.
2. Copy the Packages and repodata directory from the installation media to the packages directory on the deployment server.

   ```
   $ cd /opt/ibm/cmwo/yum-repo/operatingsystem/redhat6.5/x86_64
   $ cp /mnt/iso/Packages .
   $ cp /mnt/iso/repodata .
   ```

   Note: Alternatively, you can create a symbolic link to the installation media. If you create the symbolic link, you do not create the *architecture* directory. Instead, make a symbolic link to the installation media.

   ```
   $ cd /opt/ibm/cmwo/yum-repo/operatingsystem/redhat6.5/x86_64
   $ ln -s /mnt/iso/Packages
   $ ln -s /mnt/iso/repodata
   ```

   If you use the symbolic link, you need to ensure that the mount point persists the deployment server is restarted.
3. Optional: If you have package updates, you can create an updates repository as follows.
   a. On the deployment server, install the `createrepo` package. This package is used to create a yum repository and is available as part of the base operating system packages.

      ```
      $ yum install createrepo
      ```

Note: Details about yum repositories and creating repositories is beyond the scope of this document. If you need more information, see your operating system documentation.

b. Acquire the updated package for your environment and place it in the updates directory. `/opt/ibm/cmwo/yum-repo/operatingsystem/redhat6.5/x86_64/updates`

c. Run the **createrepo** command to create the updates repository:

```
$ cd updates
$ createrepo .
```

4. Optional: If you are using the MySQL database, this step is required to ensure that the autoconf package is available at the time the mysql recipes require it. Set the yum repositories before running the deployment by using the following command:

```
knife os manage deploy node node name ibm-openstack-yum-server::yumrepo
-P passw0rd -E environment name
```

Where *node name* is the node where you are deploying the topology and *environment name* is the name of your environment.

Results

When the deployment server is set up as a yum repository for the operating system packages and updates, these yum repositories are automatically configured on the node systems during the deployment process. You do not need to manually add these yum repositories on the node systems.

Related reference:

"Operating system repository not found" on page 255
You might see an error in the log file that the yum repository is not available.

Applying fixes and updates

You can apply fixes and updates for the IBM Cloud Manager with OpenStack product.

Related tasks:

"Updating a deployed topology" on page 113
After you deploy a topology, you might need to apply fixes, either to the IBM Cloud Manager with OpenStack services themselves or to how the IBM Cloud Manager with OpenStack services are deployed.

"Getting fixes from Fix Central" on page 249
You can use Fix Central to find the fixes that are recommended by IBM Support for a variety of products, including IBM Cloud Manager with OpenStack. With Fix Central, you can search, select, order, and download fixes for your system with a choice of delivery options. A IBM Cloud Manager with OpenStack product fix might be available to resolve your problem.

Applying fixes and updates for IBM Cloud Manager with OpenStack

Updates for IBM Cloud Manager with OpenStack provide fixes to the product.

About this task

There are two steps to apply updates.

1. Update IBM Cloud Manager with OpenStack on the deployment server.
2. Update the topology on the client nodes.

Use the following steps to install a IBM Cloud Manager with OpenStack fix pack. The fix pack updates Chef cookbooks and other resources that are stored on the deployment server. It does not apply updates to the client nodes in your topology.

After you install the fix pack, you can deploy the fix pack to your topology.

To install a IBM Cloud Manager with OpenStack fix pack, complete the following steps:

Procedure

1. Copy the IBM Cloud Manager with OpenStack fix pack files to a temporary directory, for example `/tmp/cmwo_fixpack`, on the deployment server.
2. Extract the fix pack files.

   ```
   # cd /tmp/cmwo_fixpack
   # tar -zxf cmwo_fixpack_4.1.0.1.tar.gz
   ```

3. Run the fix pack installation script.

   ```
   # ./install_cmwo_fixpack.sh
   ```

4. The script output indicates whether the installation succeeded or failed and stores a copy of the fix pack log files in the `<product_dir>/version` directory, where `<product_dir>` is the directory where IBM Cloud Manager with OpenStack was installed. The default path is `/opt/ibm/cmwo`.

What to do next

The `<product_dir>/version` directory keeps an overall history of fix pack installs and records the fix pack level in the `product.version` file.

After you install the IBM Cloud Manager with OpenStack fix pack on the deployment server, you can deploy the fixes to your topology using the `knife os manage update` tool. The tool uses the chef server and client to update the client nodes, and can be used to update single nodes or all nodes in your topology. Update the OpenStack controller and database nodes first, then update the compute nodes.

For more information about updating the topology on the client nodes, see "Updating a deployed topology" on page 113.

Applying fixes and updates for the self-service portal on Windows

Updates for IBM Cloud Manager with OpenStack provide fixes to the self-service portal on Windows. IBM Cloud Manager with OpenStack has a Command Line Interface (CLI) that is used to update the application. The CLI is available from the OSGi console that opens when the application starts.

About this task

Depending on where you chose to download the update package, you either access a URL-based repository for remotely installing packages or local file system directory repository on the system running IBM Cloud Manager with OpenStack. Updates using remote (URL-based) or local file system (directory) repositories are the same except for the format of the repository that is specified on the CLI. URL-based repositories are of the form `http://UPDATE_REPOSITORY`, while local file system (directory) repositories have the form file: `PATH_TO_DIRECTORY`. When using a local file system repository, you must use forward slashes in the full path to the local directory.

Note: Before you install a fix pack or upgrade IBM Cloud Manager with OpenStack, back up your `skc.ini` file to a safe location. After you install the fix pack or finish upgrading, replace the `skc.ini` file with your backup version.

The following update commands are available from the CLI:

`version`
: Returns the current product version.

`showrepos`
: Returns a list of the update repositories that are associated with the product.

`addrepo [url]`
: Adds a repository to the product from the file or remote URL.

delrepo [url]
Removes a repository (* for all repositories) from the product.

checkupdates
Returns a list of the current updates available.

installupdates
Installs the available updates to the product.

updatetimestamps
Returns a list of the product update timestamps.

rollbackupdates [timestamp]
Rolls back to the timestamp. If no timestamp is given the rollback is to the previous timestamp.

To download and install fixes for IBM Cloud Manager with OpenStack self-service portal, follow these steps:

Procedure

1. Open your browser to IBM Support Fix Central at http://www-933.ibm.com/support/fixcentral/
2. Select **Select product**.
3. Select **Other Software** for the Product Group.
4. For the Product, select IBM Cloud Manager with OpenStack.
5. Select **All** for the Installed Version.
6. Select **All** for the Platform and select **Continue**.
7. Identify fixes by selecting **Browse** for fixes and select **Continue**.
8. Select the specific fix that you want and select **Continue**.
9. Authenticate to the Fix Central server to demonstrate entitlement.
10. Select the method that you want to use to download the fix and select **Continue**.
11. Store the update in a repository that is available remotely or locally on the system that is running IBM Cloud Manager with OpenStack.
12. Extract the compressed file to the disk drive of the system in a temporary directory.
13. In the IBM Cloud Manager with OpenStack OSGi command console, use the **showrepos** command to list the repositories that are associated with the IBM Cloud Manager with OpenStack. For example:

```
osgi> showrepos
Metadata repositories:
Artifacts repositories:
         file:/C:/Users/IBM_ADMIN/.eclipse/207580638/p2/org.eclipse.
equinox.p2.core/cache/
```

14. If the repository that is storing the extracted files is not available, use the **addrepo** command to add that repository.

```
osgi> addrepo file:C:/temp/myFixPack
Cloud Manager with OpenStack update repository added
```

15. Install the updates by using the **installupdates** command.

```
osgi> installupdates
Cloud Manager with OpenStack updates to install:
com.ibm.cfs.product 4.1.0.0-201405150201 ==> com.ibm.cfs.product 4.1.0.1-201405150700
Cloud Manager with OpenStack update done
```

16. When the update is complete, activate the changes by using the **close** command to end the OSGi session, then restarting IBM Cloud Manager with OpenStack.

```
osgi> close
```

17. If you want to remove the updates and return to a previous configuration of IBM Cloud Manager with OpenStack, follow these steps:

a. If you have not yet restarted IBM Cloud Manager with OpenStack after applying an update, you must close your OSGi session and restart IBM Cloud Manager with OpenStack before you can roll back the changes.
b. Determine the timestamp of the update by running the **updatetimestamp** command:

```
osgi> updatetimestamps
Update Tim-stamps:
1374633752042: Wed Jul 24 10:42:32 CST 2013: 4.1.0.1-201405131252
1374657093794: Wed Jul 24 17:11:33 CST 2013: 4.1.0.1-201407240215
1374657144808: Wed Jul 24 17:12:24 CST 2013: 4.1.0.1-201405131252
1374657179390: Wed Jul 24 17:12:59 CST 2013: 4.1.0.1-201407240215
1378173169597: Tue Sep 03 09:52:49 CST 2013: 4.1.0.1-201405131252
1378175017499: Tue Sep 03 10:23:37 CST 2013: 4.1.0.1-201407240215
```

c. Use the **rollbackupdates** command to remove the updates.

```
rollbackupdates 1378175017499
```

If you do not provide a timestamp to the **rollbackupdates** command, the last timestamp is used.
d. When the rollback is complete, activate the changes by using the **close** command to end the OSGi session, then restarting IBM Cloud Manager with OpenStack.

Applying fixes and updates for DB2

Updates for DB2 provide fixes to the DB2 database server.

About this task

These instructions assume you installed DB2 10.5, fix pack 2 and would like to install 10.5, fix pack 3.

To install a DB2 fix pack, complete the following steps:

Procedure

1. Extract the DB2 fix pack.

```
$ tar –zxvf v10.5fp3_linuxx64_server.tar.gz
```

2. Shut down DB2.
   a. View the number of instances, using the following command:**$ db2ilist**.
   b. Run the following commands on each DB2 instance.

```
$ su – db2inst1

$ db2 force applications all
DB20000I  The FORCE APPLICATION command completed successfully.
DB21024I  This command is asynchronous and may not be effective immediately.

$ db2 terminate
DB20000I  The TERMINATE command completed successfully.

$ db2stop
SQL1064N  DB2STOP processing was successful.
```

3. Stop the DB2 administrator server (DAS).

   Note: This step is not applicable if you do not have a DAS user.

```
$ su – dasusr1

$ db2admin stop
SQL4407W  The DB2 Administration Server was stopped successfully.
```

4. Remove the DB2 IPC resources for the engine and clients on each instance.

```
$ su – db2inst1

$ ipclean
Application ipclean: Removing DB2 engine and client IPC resources for db2inst1.
```

5. Upgrade to DB2 10.5, fix pack 3.

```
$ ./installFixPack -b /opt/ibm/db2/V10.5/
Do you want to choose a different installation directory for the fix pack? [yes/no]
--------------------------------------------------------------------------------
no
```

6. Verify the new version of DB2.

```
$ db2level

DB21085I  This instance or install (instance name, where applicable:
"db2inst1") uses "64" bits and DB2 code release "SQL10053" with level
identifier "0604010E".
Informational tokens are "DB2 v10.5.0.3", "s140203", "IP23551", and Fix Pack
"3".
Product is installed at "/opt/ibm/db2/V10.5".
```

Uninstalling IBM Cloud Manager with OpenStack

You can uninstall IBM Cloud Manager with OpenStack from the deployment server. You can also uninstall the optionally installed IBM Cloud Manager with OpenStack self-service portal.

Uninstalling IBM Cloud Manager with OpenStack from deployment server

The IBM Cloud Manager with OpenStack uninstaller for the deployment server, similar to the installer, supports three installation types: console, silent, and graphical installers.

About this task

The default display type is `console`.

Complete the following steps to uninstall IBM Cloud Manager with OpenStack from the deployment server:

Procedure

1. Run the uninstaller with root authority by running the following command:

   ```
   cmwo_uninstall
   ```

2. To uninstall IBM Cloud Manager with OpenStack using a non-default installation type, start the uninstaller from the command line and enter `-i <uninstall type>` where `<uninstall type>` is `silent`, or `gui`.

Results

IBM Cloud Manager with OpenStack is uninstalled from the deployment server.

Restriction: Support does not exist to uninstall a deployed topology for IBM Cloud Manager with OpenStack. The node system where a topology is deployed must be re-installed or reset back to its pre-deployment state to remove the deployed topology.

Uninstalling the self-service portal on Linux

You can disable the IBM Cloud Manager with OpenStack self-service portal.

About this task

Use the following command to uninstall the self-service portal from IBM Cloud Manager with OpenStack.

```
 $ knife os manage deploy node your-node-fqdn 'recipe[ibm-sce::uninstall]'
-E your-environment-name -x ssh-user -P ssh-password
--secret-file your-secret-key-name
```

Note: The preceding command only disables the self-service portal. It does not roll back the policy files (that the self-service portal replaced) and the *sceagent* user that is registered in OpenStack.

To roll back the OpenStack policy files, complete the following steps:

1. Log in to the controller node and go to the `"/etc/keystone/"` folder.
2. Run the following command: **`rm -f policy.json mv policy.json.openstack policy.json`**
3. Go to each of the folders below, one by one, and run the command in step 2 again from each folder.
   - `/etc/glance/`
   - `/etc/nova/`
   - `/etc/neutron/`
   - `/etc/cinder/`

Uninstalling IBM Cloud Manager with OpenStack self-service portal from Windows

The IBM Cloud Manager with OpenStack self-service portal uninstaller, similar to the installer, supports three installation types: console, silent, and swing/graphical installers.

For Windows, the default display type is `graphical`.

To uninstall IBM Cloud Manager with OpenStack self-service portal using a non-default installation type, start the uninstaller from the command line and enter `-i <uninstall type>` where `<uninstall type>` is `console` or `silent`.

Uninstalling a graphical installation of IBM Cloud Manager with OpenStack self-service portal on Windows

You can uninstall a graphical installation of IBM Cloud Manager with OpenStack self-service portal on Windows.

About this task

Use the following steps to uninstall a graphical installation of IBM Cloud Manager with OpenStack self-service portal on Windows:

Procedure

1. Shut down IBM Cloud Manager with OpenStack self-service portal.
2. Access IBM Cloud Manager with OpenStack self-service portal in the Control Panel, **Control Panel** > **Programs** > **Uninstall a program**. Select IBM Cloud Manager with OpenStack in the list and then click **Uninstall**.
3. Read the instructions on the Introduction screen. Then, click **Next**.
4. On the Choose Uninstall Features window, select **Uninstall Specific Features** to choose what to uninstall. If you want to completely uninstall IBM Cloud Manager with OpenStack self-service portal, press **Next**. Both Application and Properties product features are cleared.
5. Check the features that you want to uninstall and then press **Uninstall**.

Results

The IBM Cloud Manager with OpenStack self-service portal uninstallation is complete.

Note: If you uninstall IBM Cloud Manager with OpenStack self-service portal on a system running Microsoft Windows Server 2012 R2, you must manually delete the directory `Cloud Manager with OpenStack` from the installation path.

Uninstalling a console installation of IBM Cloud Manager with OpenStack self-service portal on Windows

You can uninstall a console installation of IBM Cloud Manager with OpenStack self-service portal.

About this task

To uninstall IBM Cloud Manager with OpenStack self-service portal that was installed using a console on Windows, follow these steps:

Procedure

1. Shut down the IBM Cloud Manager with OpenStack self-service portal executable file.
2. Open a Windows command prompt.
3. In the command prompt window, navigate to the Windows uninstaller.
4. Run the uninstaller by entering `UninstallSCE.exe -i console`
5. Read the uninstaller instructions on the Introduction window and press Enter.

Results

The IBM Cloud Manager with OpenStack self-service portal is now uninstalled.

Note: If you uninstall IBM Cloud Manager with OpenStack self-service portal on a system running Microsoft Windows Server 2012 R2, you must manually delete the directory `Cloud Manager with OpenStack` from the installation path.

Chapter 4. Migrating IBM Cloud Manager with OpenStack

This section describes different methods for migrating IBM Cloud Manager with OpenStack to a new release.

The migration steps that are required depend upon the hypervisor that is selected by the IBM SmartCloud Entry 3.2 deployment. See the following table and follow the relevant links to migrate from IBM SmartCloud Entry 3.2 to IBM Cloud Manager with OpenStack 4.1.

Table 21. Migration guidance

IBM SmartCloud Entry 3.2 deployment type	Migration steps
IBM SmartCloud Entry only (VMware only deployment)	• "Customizing for self-service portal" on page 100 • "Migrating from IBM SmartCloud Entry 3.2 web interface to 4.1 self-service portal"
IBM SmartCloud Entry and OpenStack controller node (for Hyper-V, KVM, and PowerVC)	**Attention:** Ignore the steps that instruct you to restart IBM Cloud Manager with OpenStack after the migration in section "Migrate IBM SmartCloud Entry". You must not restart until both IBM SmartCloud Entry and the OpenStack controller are migrated. 1. "Migrating cinder volume groups from the appliance to the OpenStack controller" on page 49 2. "Migrating OpenStack controller prerequisites" on page 56 3. "Deploying a new 4.1 server for migration" on page 57 4. "Migrating from IBM SmartCloud Entry 3.2 web interface to 4.1 self-service portal" 5. "Migration process" on page 58 6. "Migrating the configuration manually" on page 61
OpenStack controller (Hyper-V, KVM, and PowerVC)	1. "Migrating cinder volume groups from the appliance to the OpenStack controller" on page 49 2. "Migrating OpenStack controller from 3.2 to 4.1" on page 55
Hyper-V compute (for Hyper-V deployment)	"Upgrading the IBM Cloud Manager with OpenStack Hyper-V Agent" on page 62
KVM compute (for KVM deployment)	"Migrating the KVM Agent and CLI client" on page 65
PowerVC (for PowerVC deployment)	The PowerVC compute node upgrade relies on the PowerVC product upgrade capability.

You can also migrate from one IBM Cloud Manager with OpenStack 4.1 self-service portal server to another IBM Cloud Manager with OpenStack 4.1 self-service portal server. For more information, see "Migrating between two 4.1 self-service portals" on page 48.

Migrating from IBM SmartCloud Entry 3.2 web interface to 4.1 self-service portal

You can migrate your IBM SmartCloud Entry configuration and data when you move to IBM Cloud Manager with OpenStack 4.1.

© Copyright IBM Corp. 2012, 2014

Note: In IBM Cloud Manager with OpenStack 4.1, what was previously called the IBM SmartCloud Entry web interface is now an optional configuration that is called the IBM Cloud Manager with OpenStack self-service portal. You need to complete this process to migrate configuration and data from the IBM SmartCloud Entry web interface only if you plan to configure the IBM Cloud Manager with OpenStack self-service portal. You do not need to complete this migration if you configured the 3.2 appliance with **OpenStack Only**.

To determine the current configuration of the appliance, log in to the appliance and start sceappmgr. On the initial page of sceappmgr, the first option that is listed is **Change Appliance Usage**. The current usage is listed in parentheses after this menu option. You see one of the following configuration values listed:

- SmartCloud Entry Only
- SmartCloud Entry with OpenStack
- OpenStack Only

```
IBM-SCE> sceappmgr
====================

Manage SmartCloud Entry Appliance

1.  Change Appliance Usage (Smart Cloud Entry with OpenStack)
2.  Manage Authentication Tokens
3.  Manage Passwords
4.  Manage SmartCloud Entry Services
5.  Manage Volume Groups
6.  Support and Maintenance
7.  Manage OpenStack Configuration
8.  Manage Networking Configuration

Enter selection (Enter to cancel):
```

To migrate data and configuration files that are not migrated by the installer or to manually migrate your data and configurations, follow the instructions in these topics. When you migrate data in IBM SmartCloud Entry you must migrate sequentially; you cannot skip a version. For example, if you want to migrate IBM SmartCloud Entry version 3.1 to IBM Cloud Manager with OpenStack version 4.1, you need to migrate from version 3.1 to version 3.2 first. Before you migrate to the latest version, ensure that all available fix packs are applied. Then, migrate from IBM SmartCloud Entry version 3.2 (with fix packs) to version 4.1.

To migrate IBM Cloud Manager with OpenStack to a new release, take the following general steps:

- Stop IBM SmartCloud Entry 3.2
- Copy the `/home/sysadmin/.SCE32` folder to your new IBM Cloud Manager with OpenStack 4.1 server as the migration source folder
- Migrate your configuration
- Migrate your data (Optional)
- Manually migrate any remaining configuration files

Migrating your configuration

When you migrate to a new version of IBM Cloud Manager with OpenStack, you can choose to migrate all of your preferences from one set of configuration files to the new set.

To migrate your preferences from one set of configuration files to the new set, use the following OSGi command:

migrateConfig [*source directory*]

where *source directory* is the location of the previous IBM Cloud Manager with OpenStack release configuration files. Ensure the owner of the source directory is the same as the owner of the self-service portal service when you migrate the configuration.

Running this command migrates the following configuration files:

- `cfs.keystone`
- `ldap.xml`
- `deployment.properties`
- `authentication.properties`: `admin.username`, `admin.password`, and `admin.name` should not be updated.
- `email.properties`
- `messaging.properties`
- `metering.properties`
- `billing.properties`
- `web.properties`
- `*.jks`
- `products/*.xml`
- `server_rules.json`
- `openstack.properties`
- `server.properties`

Example

To migrate configuration files from a previous version, run the following command:

```
migrateConfig [source directory]
```

where `[source directory]` is the location of the previous IBM Cloud Manager with OpenStack release configuration files.

Migrating your data

When you are migrating to a new version of IBM Cloud Manager with OpenStack, you can choose to migrate data from your previous database.

Before you begin

To prepare for the data migration, follow these steps:

1. Ensure that the target database exists.
2. In your browser, log out and close all open self-service portal windows.

Procedure

To migrate your data, use the following OSGi command:

```
upgradeDatabase 'DB2_Path' 'DB2_User_Name' 'DB2_password'
```

where *DB2_Path* is the path of the DB2 database, *DB2_User_Name* is the name of the DB2 administrator, and *DB2_password* is the password for the DB2 administrator.

Note: *DB2_User_Name* and *DB2_password* are only needed when migrating data from a DB2 database.

Notes:

- The **upgradeDatabase** command supports only major versions of databases.
- The source database must be Derby or DB2.

- Only approved requests can be migrated, while others such as pending requests, rejected requests, and withdrawn requests cannot be migrated. If the approved requests related instances have been deleted, they cannot be migrated either.
- Before you migrate the IBM SmartCloud Entry 3.2 data, wait five minutes after starting IBM Cloud Manager with OpenStack self service portal, to ensure that the database is fully initialized.
- If any errors occur during migration, renew the target database and run the **upgradeDatabase** command again.

Results

You can see migration details in the console.

Example

- To migrate your data from a Derby database, run the following command:
  ```
  upgradeDatabase '/tmp/.SCE32/'
  ```
- To migrate your data from a DB2 database, run the following command:
  ```
  upgradeDatabase '//localhost:50001/skc''db2admin' 'db2passwd'
  ```

What to do next

You must restart the self-service portal after migration. To restart it, use the **service sce restart** command.

Migrating configuration manually

Some configuration information must be migrated manually.

This configuration information includes:

- Network configuration
- Logging.properties file

Note: When migrating configuration information for a Microsoft Windows 2008 installation, you must manually configure the cloud connection.

Migrating between two 4.1 self-service portals

You can migrate the IBM Cloud Manager with OpenStack 4.1 self-service portal to another IBM Cloud Manager with OpenStack 4.1 self-service portal.

Migrating self-service portal to another system

You can migrate your IBM Cloud Manager with OpenStack self-service portal configuration from one system to another system.

About this task

Note: Migrating a configuration from one system to another replaces the configuration on the target system.

To migrate IBM Cloud Manager with OpenStack from one system to another, follow these steps:

Procedure

1. Shut down IBM Cloud Manager with OpenStack on the target system.
2. Copy the `.SCE41` directory from the source system to the target system, overwriting the `.SCE41` directory on the target system. By default, this folder is in the following path: */var/opt/ibm*.

3. Update the directory path in the `logging.properties` file to the following:
   `java.util.logging.FileHandler.pattern=/var/opt/ibm/.SCE41/logs/skc-%g.log`
4. Update the directory path in the `server.properties` file to the following:
   `org.eclipse.equinox.http.jetty.ssl.keystore=/var/opt/ibm/.SCE41/.keystore`
5. Start IBM Cloud Manager with OpenStack on the target system.
   The configuration from the source system is now available on the target system.

Migrating a Derby database to DB2 database

Use the **migratedatabase** command to move the IBM Cloud Manager with OpenStack database from an embedded Derby database to an external DB2 database. Migrating to an external DB2 database can improve scaling, performance, and production level data store.

To prepare for the migration process, follow these steps:

1. Ensure that the DB2 database exists and is empty.
2. Configure IBM Cloud Manager with OpenStack to use DB2. For more information, see "Configuring database" on page 132.
3. In your browser, log out and close all opened IBM Cloud Manager with OpenStack windows.

To migrate the current database to the DB2 database, use the following OSGi command:

migrateDatabase *source_directory*

where *source_directory* is the property file location. You can see migration details in the console. It is not necessary to restart IBM Cloud Manager with OpenStack after migration.

Example: Migrating from a Derby database to a DB2 database

Migrate the Derby database to a DB2 database by running the following command:

```
migrateDatabase '/tmp//.SCE32\'
```

Notes:

- Only approved requests can be migrated, while others such as pending requests, rejected requests, and withdrawn requests cannot be migrated. Approved requests related instances that have been deleted cannot be migrated.
- If any errors occur during migration, renew DB2 first and try running the **migrateDatabase** command again.
- To free up the space that is occupied on IBM Cloud Manager with OpenStack server, the **migrateDatabase** command clears the Derby database following the migration. Create a backup of the Derby database before migration.
- During migration, invoice IDs might change.

Migrating cinder volume groups from the appliance to the OpenStack controller

If your IBM SmartCloud Entry 3.2 appliance was configured as a cinder-volumes server that uses the LVMISCSI cinder driver, the volumes must be copied or moved to the new OpenStack controller for IBM Cloud Manager with OpenStack.

Identifying cinder-volumes LVM volume groups

If the IBM SmartCloud Entry 3.2 appliance is running OpenStack and is using the OpenStack block storage (cinder) LVM volume driver, the associated LVM volume groups must be moved or copied to the OpenStack controller for IBM Cloud Manager with OpenStack.

About this task

Use the following steps to determine whether there are any volume groups to be migrated and to get their names:

Procedure

1. Log in to the IBM SmartCloud Entry 3.2 appliance as sysadmin.
2. Start the sceappmgr tool by using the following command:

   ```
   IBM-SCE> sceappmgr
   ```

3. The initial menu includes a **Change Appliance Usage** option that displays the current appliance configuration as part of the menu option. The following example shows that it is set to (OpenStack Only).

   ```
   IBM-SCE> sceappmgr
   ====================

   Manage SmartCloud Entry Appliance

    1.  Change Appliance Usage (OpenStack Only)
    2.  Manage Authentication Tokens
    3.  Manage Passwords
    4.  Manage SmartCloud Entry Services
    5.  Manage Volume Groups
    6.  Support and Maintenance
    7.  Manage OpenStack Configuration
    8.  Manage Networking Configuration

   Enter selection (Enter to cancel):
   ```

   If the current configuration displays **SmartCloud Entry Only**, no volume groups exist to migrate. If the current configuration is **SmartCloud Entry with OpenStack** or **OpenStack Only**, continue with these steps.

4. Select the **Manage OpenStack Configuration** option, and then the **Manage Cinder Configuration** option. The end of the menu includes options like the following example:

   ```
   Cinder configuration files:
    1.  Edit /etc/cinder/api-paste.ini
    2.  Edit /etc/cinder/cinder.conf
   . . .
   Enable or disable related services:
   13.  Disable 'openstack-cinder-volume' Service
   14.  Enable 'tgtd' Service
   ```

   If the menu includes the option to **Enable 'openstack-cinder-volume' Service**, cinder volumes is not being used and you are done with these steps. Otherwise, continue with the next step.

5. Select the **Edit /etc/cinder/cinder.conf** option.
6. Search for the **enabled_backends** option. For the vi Editor, press the Esc key, then enter `/enabled_backends` and press the Enter key. A section of the file displays that looks similar to the following example:

```
# A list of backend names to use. These backend names should
# be backed by a unique [CONFIG] group with its options (list
# value)
enabled_backends=lvmdriver
```

7. Record the list of backends: lvmdriver
8. For each backend, find that stanza in the configuration file. The stanza starts with the backend name enclosed in square brackets. For example:

```
[lvmdriver]
# Sample cinder.conf stanza for a LVM volume driver
# To enable, add the stanza name (lvmdriver) to the DEFAULT enabled_backends property
volume_group=cinder-volumes
volume_driver=cinder.volume.drivers.lvm.LVMISCSIDriver
volume_backend_name=LVM_iSCSI
```

   In vi, you can use Esc then /[*backend-name* to search for the stanza. If the stanzas are out of order, press Esc then `:0` to return to the beginning of the file.
9. If the volume_driver is **`cinder.volume.drivers.lvm.LVMISCSIDriver`**, record the name of the volume group. Sceappmgr uses volume group names that start with **`cinder-`**.
10. Exit the editor.

Identifying corresponding devices to move

With the name of a volume group, you can find the physical volumes that make up that volume group, and also determine the corresponding SCSI device name or SAN LUNs.

About this task

The following steps must be done under the root account, which can be done by using the sudo command to start a bash shell. Enter the sysadmin password if sudo requests it.

Procedure

1. Using the **`cinder-volumes`** volume group as an example, get a list of the LVM physical volumes:

```
IBM-SCE> sudo bash
[sudo] password for sysadmin:
IBM-SCE> pvs | grep cinder-volumes
  /dev/sdf1  cinder-volumes lvm2 a--  100.00g 75.00g
```

2. Find the SCSI device that contains partition sdf1. This command finds the SCSI device path that corresponds to `/dev/sdf` (the disk that contains sdf1):

   Note: The following command, must be entered on one line, even though the example shows a line break. `IBM-SCE> (for dev in /dev/disk/by-path/*; do stat $dev -c "%N"; done) | grep /sdf`

```
IBM-SCE> (for dev in /dev/disk/by-path/*; do stat $dev -c "%N"; done) |
grep /sdf
`/dev/disk/by-path/scsi-0:0:0:3' -> `../../sdf'
`/dev/disk/by-path/scsi-0:0:0:3-part1' -> `../../sdf1'
```

   The format of the names varies with hypervisor, but the example paths describe a disk that is in SCSI bus 0 slot 3. The information can be used to identify the disk at the hypervisor level.

   On a KVM virtual machine, where pvs shows device name `/dev/vdd1`, the output is similar to the following.

   Note: The following command must be entered on one line, even though the example shows a line break. `(for dev in /dev/disk/by-path/*; do stat $dev -c "%N"; done) | grep /vdd`

```
IBM-SCE> (for dev in /dev/disk/by-path/*; do stat $dev -c "%N"; done) |
grep /vdd
`/dev/disk/by-path/pci-0000:00:08.0-virtio-pci-virtio5' -> `../../vdd'
`/dev/disk/by-path/pci-0000:00:08.0-virtio-pci-virtio5-part1' -> `../../vdd1'
IBM-SCE>
```

This example describes a disk that is in SCSI bus 0, slot 8. It also appears as dev=vdd in the libvirt domain description on the KVM host:

```
[root@idpx03 ~]# virsh dumpxml sce-appl
<domain type='kvm' id='16'>
. . .
    <disk type='file' device='disk'>
      <driver name='qemu' type='raw'/>
      <source file='/var/libvirt/images/sce-appl/cinder-volumes.raw'/>
      <target dev='vdd' bus='virtio'/>
      <alias name='virtio-disk3'/>
      <address type='pci' domain='0x0000' bus='0x00' slot='0x08' function='0x0'/>
    </disk>
```

3. If you are using SAN storage with multipath configured, for a PowerVM appliance for example, you can use the output of pvs to determine the corresponding SAN LUN Worldwide Identifier (WWID):

```
# pvs
  PV                                            VG                Fmt  Attr
PSize  PFree
  /dev/mapper/360050768018086aa7800000000007816p1 user_data_vg      lvm2 a--
30.00g   0
  /dev/mapper/360050768018086aa7800000000007817p1 glance-repository lvm2 a--
10.00g   0
  /dev/mapper/360050768018086aa7800000000007818p1 cinder-volumes    lvm2 a--
100.00g   0
```

The 360050768018086aa7800000000007818 portion corresponds to the UID displayed by the SAN Volume Controller "volumes by host" view – 60050768018086AA7800000000007818.

Preparing appliance to detach devices

To detach disks so that they can be moved to another system you must first stop any virtual machines that are using storage from cinder-volumes. Then, stop OpenStack services. Then, deactivate and export the volume group. After you deactivate the volume group, detach the disks from the appliance virtual machine by using the hypervisor management tools.

About this task

The process for stopping any virtual machines that are using storage from cinder-volumes is not described here. After you stop the virtual machines, use the following instructions to prepare to detach the devices.

Procedure

1. Log in to the IBM SmartCloud Entry 3.2 appliance.
2. If your IBM SmartCloud Entry 3.2 virtual appliance manages Hyper-V agent nodes, record the volumes information. You must also detach the volumes for each Hyper-V agent node that is managed by your IBM SmartCloud Entry 3.2 virtual appliance.
   a. Record volume IDs and the corresponding devices of each instance on each Hyper-V node by using the following example. You need the volume/instance IDs and devices information to attach them back after the OpenStack controller migration.
      1) Record volume ID and corresponding instance ID for the volume that is not in available state.

```
IBM-SCE> cinder list
+--------------------------------------+-----------+--------------+------+-------------+----------+-------------+
|                  ID                  |   Status  | Display Name | Size | Volume Type | Bootable | Attached to |
+--------------------------------------+-----------+--------------+------+-------------+----------+-------------+
| 69bd80d1-6a5d-46a3-8a66-0440b481a361 | available |     None     |  1   |     None    |  false   |             |
```

```
+--------------------------------------+-----------+--------------+------+-------------+----------+-------------+
| ec721954-1b60-482b-b347-88871b2ab550 | in-use    |     None     |  1   |     None    |  false   | 434c7b80-8824-4985-9940-72c09baf347d |
+--------------------------------------+-----------+--------------+------+-------------+----------+-------------+
```

2) Record the device for the volume that is attached to the Hyper-V instance by following command.

```
IBM-SCE> cinder show $volume-id
```

For example, **IBM-SCE> cinder show ec721954-1b60-482b-b347-88871b2ab550**.

b. Run the following commands to detach volumes. Replace *$instance_id* with the instance ID in last column, and replace *$volume-id* with the volume ID in first column of the previous step. Volume state becomes available after the volume is detached.

```
# nova volume-detach $instance_id $volume-id
```

For example, **IBM-SCE> nova volume-detach 434c7b80-8824-4985-9940-72c09baf347d ec721954-1b60-482b-b347-88871b2ab550**.

c. Using command cinder list to confirm that all volumes are in available status.

```
IBM-SCE> cinder list
  +--------------------------------------+-----------+--------------+------+-------------+----------+-------------+
  |                  ID                  |   Status  | Display Name | Size | Volume Type | Bootable | Attached to |
  +--------------------------------------+-----------+--------------+------+-------------+----------+-------------+
  | 69bd80d1-6a5d-46a3-8a66-0440b481a361 | available |     None     |  1   |     None    |  false   |             |
  +--------------------------------------+-----------+--------------+------+-------------+----------+-------------+
  | ec721954-1b60-482b-b347-88871b2ab550 | available |     None     |  1   |     None    |  false   |             |
  +--------------------------------------+-----------+--------------+------+-------------+----------+-------------+
```

3. Start sceappmgr by using the following command:

```
IBM-SCE> sceappmgr
```

4. Select the option to **Manage SmartCloud Entry Services**.
5. Select **Start or Stop OpenStack Services** and stop the services.
6. Exit sceappmgr.
7. You can now deactivate and export the LVM volume groups.
   a. Stop the tgtd service by using the force-stop option:

```
IBM-SCE> sudo bash
[sudo] password for sysadmin:
IBM-SCE> sudo service tgtd force-stop
Force-stopping SCSI target daemon: [ OK ]
```

   b. Use the following example to deactivate the LVM volume groups. Repeat the **vgchange** and **vgexport** commands for each volume group that you identified previously. Replace *cinder-volumes* with your volume group name.

```
IBM-SCE>  vgchange --activate n cinder-volumes
  0 logical volume(s) in volume group "cinder-volumes" now active
IBM-SCE>  vgexport cinder-volumes
  Volume group "cinder-volumes" successfully exported
```

8. Using the management tools for your hypervisor, detach the disks that you identified earlier from the appliance virtual machine. Attach the disks to the new controller before you continue to import the volume groups there.
9. If a loop device was used, use the `losetup` command to delete the loop device. The backing file (`/var/lib/cinder/cinder-volumes.raw` in the following example) remains and can be copied to the new controller.
   a. Delete the loop device:

```
IBM-SCE> losetup -a
/dev/loop0: [fd00]:2494126 (/var/lib/cinder/cinder-volumes.raw)
IBM-SCE> losetup -d /dev/loop0
IBM-SCE>
```

   b. Copy the backing file to the new controller, and re-create the loop device there:

```
[root@xxx ⌂ ] losetup /dev/loop0 /var/lib/cinder/cinder-volumes.raw
```

10. After the disk is detached, run the following commands to be sure that the volume groups are no longer visible to the system:

```
IBM-SCE> pvscan
  PV /dev/sde1   VG glance-repository   lvm2 [10.00 GiB / 0    free]
  PV /dev/sdd1   VG user_data_vg        lvm2 [30.00 GiB / 0    free]
  Total: 2 [39.99 GiB] / in use: 2 [39.99 GiB] / in no VG: 0 [0   ]
IBM-SCE> vgs
  VG                #PV #LV #SN Attr   VSize  VFree
  glance-repository   1   1   0 wz--n- 10.00g    0
  user_data_vg        1   4   0 wz--n- 30.00g    0
```

"cinder-volumes" is not displayed in the output, indicating that the devices are removed from the IBM SmartCloud Entry 3.2 appliance.

Results

When the disks are detached, you can move the storage to the controller system that you are using for IBM Cloud Manager with OpenStack, version 4.1.

Importing volume groups on the new controller

After you move the storage to the new controller for IBM Cloud Manager with OpenStack, version 4.1, you must make the volume groups known to the new system.

Procedure

Run the following commands as the root user. Repeat the **vgimport** and **vgchange** commands for each volume group that you identified previously. Replace *cinder-volumes* with your volume group name.

```
# vgimport cinder-volumes
  Volume group "cinder-volumes" successfully imported
# vgchange --activate y cinder-volumes
  0 logical volume(s) in volume group "cinder-volumes" now active
```

Note: Loop devices do not persist when a system is restarted. If you use a loop device for the cinder volume group, you can create an initd script to invoke the losetup command. Then, register the script to be run during system startup.

Verifying cinder volume group migration

IBM Cloud Manager with OpenStack includes a tool to verify that any LVM volume groups that are used by cinder on a IBM SmartCloud Entry 3.2 appliance are moved to the new controller.

About this task

The cinder migration verification tool inspects the appliance configuration. It determines the volume groups that are used by cinder, verifies that they are no longer known to the appliance and that they are now known to the new controller.

Usage:

```
/opt/ibm/cmwo/bin/chkcindervolmigr.sh (options)
   --appl-host HOSTNAME         Hostname for SmartCloud Entry 3.2 appliance.
   --appl-pw PASSWORD           'sysadmin' password for the appliance.
   --tgt-pw PASSWORD            The ssh password for the target node
   --tgt-user USER              The ssh username for the target node.
                                Default is the current user name.
   --tgt-id-file IDENTITY_FILE  The ssh identity file for the target node.
   -v, --verbose                Prints verbose output for debugging.
If a password or identify file is not provided, the password will be prompted.
```

Procedure

Use the following example command to verify that the cinder migration was successful.

Note: The following command must be run on one line, even though the example shows a line break. # /opt/ibm/cmwo/bin/chkcindervolmigr.sh --appl-host vs256.mycompany.com --tgt-host vs412.mycompany.com

```
# /opt/ibm/cmwo/bin/chkcindervolmigr.sh --appl-host vs256.mycompany.com
--tgt-host vs412.mycompany.com
Enter sysadmin password for vs256.mycompany.com:
Enter root password for vs412.mycompany.com:

Checking cinder volume groups on appliance...
Backend 'lvmdriver' uses volume_group 'cinder-volumes'
Volume group 'cinder-volumes' has been removed from the appliance.
Source appliance vs256.mycompany.com has completed cinder volume
preparation.

Checking cinder volume groups on target system...
Volume group 'cinder-volumes' was found on the target system.
Cinder volume group migration appears to be complete.
```

Results

The following examples display results from an incomplete cinder migration:

```
# /opt/ibm/cmwo/bin/chkcindervolmigr.sh --appl-host vs256.mycompany.com
--tgt-host vs412.mycompany.com --tgt-id-file vs412.id_rsa
Enter sysadmin password for vs256.mycompany.com:

Checking cinder volume groups on appliance...
Backend 'lvmdriver' uses volume_group 'cinder-volumes'
ERROR: Volume group 'cinder-volumes' is still attached to the appliance.
Source appliance vs256.mycompany.com has not completed cinder volume
preparation.
```

```
# /opt/ibm/cmwo/bin/chkcindervolmigr.sh --appl-host vs256.mycompany.com
--tgt-host vs412.mycompany.com --tgt-id-file vs412.id_rsa
Enter sysadmin password for vs256.mycompany.com:

Checking cinder volume groups on appliance...
Backend 'lvmdriver' uses volume_group 'cinder-volumes'
Volume group 'cinder-volumes' has been removed from the appliance.
Source appliance vs256.mycompany.com has completed cinder volume
preparation.

Checking cinder volume groups on target system...
ERROR: Volume group 'cinder-volumes' was not found on the target system.
One or more cinder volume groups were not found on the target system
vs412.mycompany.com.
```

Migrating OpenStack controller from 3.2 to 4.1

You can migrate IBM SmartCloud Entry 3.2 to IBM Cloud Manager with OpenStack 4.1, including image repository, OpenStack configurations, database, user, and password for QPID and DB2.

About this task

In addition, the network configuration of IBM SmartCloud Entry 3.2 is also migrated. However, it is assumed that the network configuration and topology are the same.

There are three servers involved in the migration process: the source, IBM SmartCloud Entry 3.2 virtual appliance; the target, IBM Cloud Manager with OpenStack 4.1 server; and the deployment server to assist the migration process.

When the migration process is completed, you can take the 3.2 appliance offline. The IBM Cloud Manager with OpenStack 4.1 has the same network topology as the previous IBM SmartCloud Entry 3.2 appliance.

Migrating OpenStack controller prerequisites

Complete or verify the following prerequisites before you migrate to IBM Cloud Manager with OpenStack 4.1.

Before you begin

Note: The following information applies to migrating OpenStack data only. If you want to also migrate data for IBM SmartCloud Entry 3.2, you must first copy the `/home/sysadmin/.SCE32` folder to your new 4.1 server that you plan to install the self-service portal to. You can copy it to a temporary folder, and remember this temporary folder because the IBM SmartCloud Entry 3.2 server will be shut down after the OpenStack data is migrated. After you migrate the IBM SmartCloud Entry database data, you must finish migrating the OpenStack data before you restart the IBM SmartCloud Entry service. For database data migration, if your database is IBM DB2 and runs on the 3.2 appliance, then you must migrate your database before migrating the OpenStack data.

1. The hardware prerequisite for IBM Cloud Manager with OpenStack 4.1 controller node is: 4 CPUs, 16 GB physical memory, 6 GB free disk space for installed code and empty databases.
2. Ensure that the 3.2 virtual appliance has the most recent, generally available fix pack. For more information, see the IBM SmartCloud Entry 3.2 administrator guide, chapter, titled "Applying fixes and updates for IBM SmartCloud Entry".
3. Ensure that the 3.2 virtual appliance, deployment server, and IBM Cloud Manager with OpenStack 4.1 can resolve the fully qualified domain names of one another.
4. Ensure that the Network Manager service is stopped in the IBM Cloud Manager with OpenStack 4.1 host, if this service is installed. To stop the service, run **`service NetworkManager stop`**.
5. Ensure that the amount, number, function, and connectivity of network interface cards for the 3.2 virtual appliance and IBM Cloud Manager with OpenStack 4.1 are the same. For example, if the 3.2 virtual appliance has three network interface cards (`eth0`, `eth1` and `eth2`) and `eth0` is working as the management network then, IBM Cloud Manager with OpenStack 4.1 must be the same. So, IBM Cloud Manager with OpenStack 4.1 must also have `eth0`, `eth1` and `eth2`, with `eth0` as the management network. For *eth0*, they must connect to the same management network and be able to access each other. For more information, see “Deploying the cloud environment” on page 76.
6. Do not store any new data on the OpenStack controller node database since that data is replaced by the appliance 3.2 data after migration.
7. Ensure that there is sufficient space for `/mount/glance-repository` on 4.1 to store the images that are migrated from the 3.2 appliance. To verify, you can use `df -h /mount/glance-repository` on theIBM SmartCloud Entry 3.2 controller node. Check the **Used** size, for example:

   ```
   IBM-SCE> df -h /mount/glance-repository/
   Filesystem            Size  Used Avail Use% Mounted on
   /dev/mapper/glance--repository-lv0
                         9.9G  7.1G  2.3G  76% /mount/glance-repository
   ```

   For the previous example, you must ensure that there is more than 7.1G available in `/mount/glance-repository`. If `/mount/glance-repository` does not exist on the IBM Cloud Manager with OpenStack 4.1 controller node, you need to ensure there is more than 7.1G available in `/mount`.

   Notes:

   - If there is already a `/mount/glance-repository` directory on IBM Cloud Manager with OpenStack 4.1, it is retained. If not, it is created during the migration process.
   - You might want to use shared storage, attached volume, or others as the image repository for IBM Cloud Manager with OpenStack 4.1. They are supported; however, ensure that the directory `/mount/glance-repository` exists on the storage that you intend to use.

Deploying a new 4.1 server for migration

About this task

1. Log in to the deployment system as the root user. This is the system where IBM Cloud Manager with OpenStack was installed.
2. Create a directory to store the files for the topology that you deploy. Change your-deployment-name to the name for your deployment.

   ```
   $ mkdir your-deployment-name
   $ chmod 600 your-deployment-name
   $ cd your-deployment-name
   ```
3. Copy the example environment for the topology that you deploy. Change `your-environment-name` to the name for your environment.

   ```
   $ knife environment show example-ibm-os-single-controller-n-compute -d -Fjson > your-environment-name.json
   ```
4. Change the following JSON attributes in your environment file, `your-environment-name.json`:
   - **`Name`**: Set to your environment name: *your-environment-name*.
   - **`Description`**: Set to the description for your environment.
   - **`openstack.endpoints.host, openstack.endpoints.bind-host and openstack.endpoints.mq.host`**: Change from *127.0.0.1* to the IP address of the node system for the topology.
   - **`openstack.network.core_plugin:`** Set to *neutron.plugins.openvswitch.ovs_neutron_plugin.OVSNeutronPluginV2*.
   - **`openstack.network.service_plugins`**: Set to *[]*.
   - **`openstack.block-storage.volume.create_volume_group`**: Set to *false*.

     Note: If this attribute is set to the default value, *true*, then the cinder-volumes volume group is created on the controller node with a size that is determined by **`openstack.block-storage.volume.volume_group_size`**. Because the volume group is migrated manually, you do not need to create the volume group.
   - **`openstack.block-storage.storwize.storwize_svc_vol_rsize`**: Set to *2*.
5. If you do not have a PowerVC compute node, copy the following example topology to a file, `your-topology-name.json`. Change `your-topology-name` to the name for your topology.

   ```
   {
    "name":"CHANGEME",
    "description":"CHANGEME",
    "environment":"CHANGEME",
    "secret_file":"CHANGEME",
    "run_sequentially":false,
    "nodes": [
      {
        "fqdn":"CHANGEME",
        "password":"CHANGEME",
        "identity_file":"CHANGEME",
        "quit_on_error":true,
        "run_order_number":1,
        "runlist": [
          "role[ibm-os-single-controller-node]"
        ]
      }
    ]
   }
   ```
6. If you do have a PowerVC compute node, copy the following example topology to a file, `your-topology-name.json`. Change `your-topology-name` to the name for your topology. Here is an example topology to manage to PowerVC.

   ```
   {
     "name":"CHANGEME",
     "description":"CHANGEME",
     "environment":"CHANGEME",
   ```

```
  "secret_file":"CHANGEME",
  "run_sequentially":false,
  "nodes": [
    {
      "fqdn":"CHANGEME",
      "password":"CHANGEME",
      "identity_file":"CHANGEME",
      "quit_on_error":true,
      "run_order_number":1,
      "runlist": [
        "role[ibm-os-single-controller-node]",
        "role[ibm-os-powervc-driver-node]"
      ]
    }
  ]
}
```

7. (Optional) If you want to deploy the self-service portal to this controller node, see "Customizing for self-service portal" on page 100 to finish the self-service portal configuration.
8. Customize the topology file. Change the following JSON attributes in your topology file, `your-topology-name.json`:
   - **`Name`**: Set to your topology name: *your-topology-name*.
   - **`Description`**: Set to the description for your topology.
   - **`Environment`**: Set to the environment for your topology: *your-environment-name*.
   - **`nodes.fqdn`**: For each node, you must set to the fully qualified domain name of the node system. The deployment system must be able to ssh using the fully qualified domain name. You can also set to the public IP address, private IP address, or host name.
   - **`nodes.password`** or **`nodes.identity_file`**: Set to the appropriate SSH root user authentication for the node system. Either a password and an SSH identity file can be used for authentication. Remove the unused attribute.
9. Customize the passwords and secrets before deploying. For instructions, see "Customizing passwords and secrets" on page 102.
10. Deploy your topology.
    - Upload the environment for your deployment.
      ```
      $ knife environment from file your-environment-name.json
      ```
    - Deploy the topology.
      ```
      $ knife os manage deploy topology your-topology-name.json
      ```

If IBM SmartCloud Entry 3.2 virtual appliance manages KVM Agent nodes, please record the following information for each KVM Agent node managed by your IBM SmartCloud Entry 3.2 virtual appliance. These nodes are migrated after the appliance is migrated.

1. Record the fully qualified domain names for each node.
2. Record the root user login information (either password or SSH identity file) for each node.
3. Record the values for the following options in the `/etc/nova/nova.conf` file for each node: `state_path`, `lock_path`, and `instances_path`

   Note: Ignore lines beginning with #.
4. Record whether or not the CLI client was installed on each node.
5. Record the number and name of each network interface card on each node. The network configuration and hypervisor type limit the type of networks that can be defined. For more information about network considerations, see "Network considerations" on page 16.

Migration process

The migration process can take time to run, depending on the configuration of the IBM SmartCloud Entry 3.2 appliance.

About this task

During the migration, both IBM SmartCloud Entry 3.2 and IBM Cloud Manager with OpenStack 4.1 will exist in the environment, and do not use both IBM SmartCloud Entry 3.2 and IBM Cloud Manager with OpenStack 4.1 before the migration is finished.

To migrate to IBM Cloud Manager with OpenStack 4.1, you need to install IBM Cloud Manager with OpenStack 4.1 and migrate the settings from the IBM SmartCloud Entry 3.2 (previous version) appliance. When the migration is completed, the 3.2 appliance can be offline, the IBM Cloud Manager with OpenStack 4.1 assumes the network setting will be the same as the 3.2 appliance.

Procedure

1. Run the following knife command on the deployment server to start the migration.

   ```
   knife os manage appliance migrate node NODE_FQDN (options)
   ```

   where **NODE_FQDN** is the fully qualified domain name of IBM Cloud Manager with OpenStack 4.1, IP address. Common options include:

   - **sysadmin_pass**: The password of user **sysadmin** to log in to the IBM SmartCloud Entry 3.2 operating system. The default value is *passw0rd*.
   - **db2inst1_pass**: The password of **db2inst1** in IBM SmartCloud Entry 3.2. The default value is *passw0rd*.
   - **openstack_pass**: The password for the user that is identified in the `/home/sysadmin/.keystonerc` file for OpenStack in IBM SmartCloud Entry 3.2.
   - **os_endpoints_host**: The management IP of the IBM SmartCloud Entry 3.2 OpenStack controller. The value of which replaces the value of **openstack.endpoints.host** in your Chef environment that you use to deploy the IBM Cloud Manager with OpenStack 4.1 controller.
   - **os_endpoints_bind_host**: The bind host IP of IBM SmartCloud Entry 3.2 OpenStack controller. The value of which replaces the value of **openstack.endpoints.bind-host** in your Chef environment that you use to deploy the IBM Cloud Manager with OpenStack 4.1 controller.
   - **os_endpoints_mq_host**: The message queue IP of IBM SmartCloud Entry 3.2 OpenStack controller, the value of which replaces the value of **openstack.endpoints.mq.host** in your Chef environment that you use to deploy the IBM Cloud Manager with OpenStack 4.1 controller.
   - **sce32**: The fully qualified domain name of the IBM SmartCloud Entry 3.2. This value can be specified in the environment that was used to deploy the topology. The attribute name is **['ibm-openstack']['app-migration']['sce32']**
   - **nic**: The network interface card that is used to connect to IBM SmartCloud Entry 3.2.
   - **secret-file**: A file containing the secret key to use to encrypt data bag item values. If not specified, the default secret file in deployment server is used. This secret file must be identical as the file you use to deploy the topology.
   - **ssh-username**: The username that is used to run the chef cookbooks in OpenStack controller node of IBM Cloud Manager with OpenStack 4.1.
   - **ssh-password**: The password of ssh-username to log in to the OpenStack controller node of IBM Cloud Manager with OpenStack 4.1.

   For more information about the options, enter **knife os manage appliance migrate node NODE_FQDN --help**.

   The following example shows the command using some options:

   ```
   $ knife os manage appliance migrate node cmwo41.mycompany.com
   --sce32 sceappliance32.mycompany.com --sysadmin_pass passw0rd
   --db2inst1_pass passw0rd --openstack_pass passw0rd --nic eth0
   --ssh-user root --ssh-password Passw0rd --secret-file <your_special_secret_file_path>
   --os_endpoints_host <SmartCloud Entry 3.2 controller management ip>
   --os_endpoints_bind_host <SmartCloud Entry 3.2 controller service bind-host ip>
   --os_endpoints_mq_host <SmartCloud Entry 3.2 controller message queue ip>
   ```

The <*your_special_secret_file_path*> indicates the path to the secret file you used to deploy the IBM Cloud Manager with OpenStack 4.1 new topology.

2. You are prompted to restart. Before you continue, ensure that the applications on the 4.1 server are closed properly and other services (not OpenStack related services) or jobs are saved. Then, press **Enter**.

   Note: After the restart, you might receive the following error messages. You can ignore these errors.

   - You receive the `ERROR: No route to host – connect(2` message before you receive a message stating `The node has been reconnected successfully...`.
   - You receive the `Chef encountered an error...` message before you receive a message stating `The node has been bootstrapped successfully`.

What to do next

When the migration is completed, the IBM SmartCloud Entry 3.2 appliance is shut down and IBM Cloud Manager with OpenStack 4.1 restarts to complete the migration process. During this process, neither are available and should not be used. If errors occur during the migration process, correct the source of the error, and restart the migration process as needed. After migration is complete, the deployment server will manage the 4.1 controller.

After migration, when you want to update the topology `<topology.json>` cli, with the `topology.json` file that was used to deploy the IBM Cloud Manager with OpenStack 4.1 controller, you must update the fully qualified domain name attribute to point to the new fully qualified domain name of the controller node since the previous fully qualified domain name has been changed to IBM SmartCloud Entry 3.2 fully qualified domain name.

For example, the original topology to deploy the OpenStack control node before migration is:

```
{
  "name":"CHANGEME",
  "description":"CHANGEME",
  "environment":"CHANGEME",
  "secret_file":"CHANGEME",
  "run_sequentially":false,
  "nodes": [
    {
      "fqdn":"CMWO41.example.com",
      "password":"CHANGEME",
      "identity_file":"CHANGEME",
      "quit_on_error":true,
      "run_order_number":1,
      "runlist": [
        "role[ibm-os-single-controller-node]"
      ]
    }
  ]
}
```

The fully qualified domain name of the migration source node is `SCE32.example.com`, then after migration, the fully qualified domain name of the "migrated to" node is changed from *CMWO41.example.com* to *SCE32.example.com*. So, when you want to update the topology, including the OpenStack controller node, you must modify the topology as follows, otherwise the update fails.

```
{
  "name":"CHANGEME",
  "description":"CHANGEME",
  "environment":"CHANGEME",
  "secret_file":"CHANGEME",
  "run_sequentially":false,
  "nodes": [
    {
```

```
      "fqdn":"SCE32.example.com",
      "password":"CHANGEME",
      "identity_file":"CHANGEME",
      "quit_on_error":true,
      "run_order_number":1,
      "runlist": [
        "role[ibm-os-single-controller-node]"
      ]
    }
  ]
}
```

Migrating the configuration manually

Procedure

1. (Optional) Access the backup configuration files. The OpenStack configuration files of the 3.2 appliance are kept as a backup to IBM Cloud Manager with OpenStack 4.1. To access the backup, append `'.sce32'` to the original path. For example, the configuration location for Nova is `'/etc/nova'`, then the backup location would be `'/etc/nova.sce32'`.
2. (Optional) Migrate changed attributes in different files manually. Files with the suffix `'.diff'` are generated to contain the attributes that you modified using **sceappmgr** on the 3.2 appliance, but were not merged automatically during migration. The `'.diff'` file is in the same folder as the corresponding configuration file.

   If you want to preserve the configuration for the attributes in different files, you must migrate them manually. Refer to "Mapping attributes to services" on page 230 to add or update the corresponding environment attributes.

   For example, if you updated `keystone.conf` on the 3.2 appliance to enable debug, then a different file is created with the name, `keystone.conf.diff`, under the directory `/etc/keystone` of IBM Cloud Manager with OpenStack 4.1 after migration. The `keystone.conf.diff` file contains the lines below since the debug attribute is out of scope during migration. You must migrate the debug attribute manually, and update your environment.

   ```
   [DEFAULT]
   debug = True
   ```

   Note: For more information about mapping attributes, see "Mapping attributes to services" on page 230.
3. For cinder migration, this command supports only one backend and must be configured in the [DEFAULT] section with **volume_driver** specified. In most cases, if you used sceappmgr to configure cinder, **enabled_backends** was used to configure the volume driver with multiple backends. If you have configured multiple backends, you must manually migrate the configuration. The old configuration can be found in `/etc/cinder.sce32/cinder.conf`. Specifically, **enabled_backends** and related sections must be added to the current `/etc/cinder/cinder.conf` file manually.

   For more information about configuring multiple block storage backends, see "Configuring multiple block storage backends" on page 122.

   Important: The migration process merges the host configuration in the [DEFAULT] section in `cinder.conf` automatically, but there is no matching Chef attribute for this configuration. Therefore, if you run the **knife os manage update topology xxx.json** command to update the topology, you must manually merge the host configuration in [DEFAULT] section in `cinder.conf` after updating, in order for the cinder service to run properly. The value of the host configuration comes from `/etc/cinder.sce32/cinder.conf`. For example:

   ```
   [DEFAULT]
   ....
   host = sce32host
   ```
4. If you changed passwords and secrets of the 3.2 appliance (which are not in the second column of following table), you must update the corresponding data bags manually.

Table 22. Data bag and migrated items

Data bag	Migrated Items
db_passwords	All items are consistent with the 3.2 appliance after migration .
secrets	• openstack_identity_bootstrap_token • openstack_simple_token
service_passwords	All items are consistent with the 3.2 appliance after migration .
user_passwords	• pvcadmin • pvcqpid • qpidclient

Note: For more information about data bags, see "Data bags" on page 229.

5. After migration, Chef updates the environment according to the settings in the 3.2 appliance. However, the environment file (for example, `your-environment-name.json`) that you created has not been updated, so you must update the `JSON` file manually. To update the new environment `.json` file, use the following command:

   ```
   knife environment show your-environment-name -d -F json > your-environment-name.json
   ```

   Modify `your-environment-name` to the environment name that is used before migration. Choose the one used to deploy before migration and modify `your-environment-name.json` to the `JSON` file name that you want to be updated.

   Note: Use the **`knife environment list`** command to view all the environments in the Chef server.

   For future environment updates (knife environment from file `xxx.json`), you must use the `.json` file that is updated in the previous step; otherwise, the environment settings do not maintain the 3.2 appliance settings.

6. If your IBM SmartCloud Entry 3.2 virtual appliance manages Hyper-V agent nodes, run the following commands to attach volumes back to the instances on each Hyper-V agent node that is managed by your controller node. Replace *$instance_id*, *$volume-id* and *$device* with your records in "Preparing appliance to detach devices" on page 52.

   ```
   # nova volume-attach $instance_id $volume-id $device
   ```

   For example, **`# nova volume-attach 434c7b80-8824-4985-9940-72c09baf347d ec721954-1b60-482b-b347-88871b2ab550 /dev/vdd`**.

Upgrading the IBM Cloud Manager with OpenStack Hyper-V Agent

You can upgrade the IBM Cloud Manager with OpenStack Hyper-V Agent by using the graphical installation wizard or silently.

Note: You must upgrade the Hyper-V Agent one version at a time. For example, if you have Hyper-V Agent version 3.1, FP2 installed, you must upgrade to version 3.2 and then upgrade to version 3.2, FP1. However, if you want to upgrade the Hyper-V agent to IBM Cloud Manager with OpenStack version 4.1, you must upgrade directly from IBM SmartCloud Entry version 3.2, FP2 to IBM Cloud Manager with OpenStack version 4.1, FP1. If you upgrade the Hyper-V agent from IBM SmartCloud Entry version 3.2, FP2 to IBM Cloud Manager with OpenStack version 4.1 without installing FP1 for version 4.1, you are likely to encounter a problem.

Attention: During a graphical installation of the Hyper-V agent, if you provide a user name in the `Nova compute service user` field on the "Hyper-V Live Migration Settings" window, the IBM Cloud Manager with OpenStack Compute Service runs using the domain user. In this case, after you upgrade, you must complete the following steps to ensure that the instances can support live migration successfully on the Hyper-V compute node.

1. Open **Control Panel** -> **Administrative Tools** -> **Services**.
2. Locate the service that is named `IBM Cloud Manager with OpenStack Compute Service`, right-click, and select the **properties** menu.
3. Switch to the **Log On** tab, select the second option for `Log on as` and provide the domain user account and password.
4. Click **Apply** and **OK**.
5. Locate the service that is named `IBM Cloud Manager with OpenStack Compute Service`, right-click, and select **Restart** to restart the service.

Graphical Upgrade

To upgrade the IBM Cloud Manager with OpenStack Hyper-V Agent by using the graphical installation wizard, follow these steps:

Procedure

1. Download the latest fix for the IBM Cloud Manager with OpenStack Hyper-V Agent from Fix Central. For more information, see “Applying fixes and updates for the self-service portal on Windows” on page 39.

   Note: The IBM Cloud Manager with OpenStack Hyper-V Agent and IBM Cloud Manager with OpenStack must be at the same level, either the GA level, or the fix level.
2. Locate the installation image, and double-click **IBM Cloud Manager with OpenStack Hyper-V Agent.msi** to start the installation wizard.
3. Follow the instructions that are provided by the installation wizard. For example, you must agree to the license terms and specify configuration information, such as target directory and basic nova and virtual switch configuration information.
4. After you complete the information in the installation wizard, the upgrade begins.

Silent upgrade

To upgrade IBM Cloud Manager with OpenStack Hyper-V Agent silently, follow these steps:

Before you begin

Procedure

1. Download the latest fix for the IBM Cloud Manager with OpenStack Hyper-V Agent from Fix Central. For more information, see “Applying fixes and updates for the self-service portal on Windows” on page 39.

   Note: The IBM Cloud Manager with OpenStack Hyper-V Agent and IBM Cloud Manager with OpenStack must be at the same level, either the GA level, or the fix level.
2. To run the installation through the response file, you must first enter the correct parameters in your locally saved copy of the response file. For detailed information, see the sample response file.

   The sample response file provides an example INI file that can be used to drive a silent installation. This example shows all properties that are available during a graphical installation of the IBM Cloud Manager with OpenStack Hyper-V Agent.

   ```
   [Response]
   #indicate whether you agree with the license and its default value is "no"
   AgreeToLicense=yes
   ```

Note: The property ;AgreeToLicense in the response file specifies your agreement to the license for the application. Its default value is set to no. You must specify yes to run the silent installation successfully.

3. Next, open a command prompt and input the following statement:

   Note: The following statement must be entered on a single line, even though the example shows a line break for formatting purposes.

   ```
   msiexec /i "IBM Cloud Manager with OpenStack Hyper-V Agent.msi"
   /qn USEINI="absolute path to responsefile"
   ```

   The following properties in the nova.conf file are copied:

   ```
   glance_host;
   glance_port;
   qpid_hostname;
   qpid_port;
   qpid_username;
   qpid_password;
   instances_path;
   neutron_url;
   neutron_admin_username;
   neutron_admin_password;
   neutron_admin_tenant_name;
   neutron_region_name;
   neutron_admin_auth_url;
   auth_strategy with default value "keystone";
   flat_injected with default value "true";
   rpc_thread_pool_size with default value "128";
   rpc_conn_pool_size with default value "60";
   rpc_response_timeout with default value "600";
   use_cow_images=" ";
   vswitch_name=" ";
   use_ipv6 with default value "true";
   verbose with default value "false";
   log_dir=" ";
   ```

   These following properties in neutron.conf file are copied:

   ```
   qpid_hostname;
   qpid_port;
   qpid_username;
   qpid_password;
   allow_overlapping_ips;
   rpc_thread_pool_size with default value "128";
   rpc_conn_pool_size with default value "60";
   rpc_response_timeout with default value "600";
   ```

   The following properties in hyperv_neutron_agent.ini file are copied:

   ```
   qpid_hostname;
   qpid_username;
   qpid_password;
   physical_network_vswitch_mappings;
   rpc_backend;
   verbose;
   debug;
   control_exchange;
   tenant_network_type;
   network_vlan_ranges;
   ```

   Note: If you manually modified properties that are not shown, you must manually modify those properties after the upgrade is complete.

Migrating the KVM Agent and CLI client

If you had the KVM Agent and CLI client installed on endpoints in your IBM SmartCloud Entry 3.2 environment, use the following information to migrate to the IBM Cloud Manager with OpenStack 4.1 compute nodes. You must complete the preparing to migrate steps, the upgrade steps, and the completion steps.

Preparing to migrate the KVM Agent and CLI client

If you had the KVM Agent and CLI client installed on endpoints in your IBM SmartCloud Entry 3.2 environment, prepare to migrate the KVM nodes to IBM Cloud Manager with OpenStack 4.1.

About this task

Complete the pre-migration setup steps for the KVM compute nodes:

Procedure

1. Log in to the deployment system as the root user. This is the system where you installed IBM Cloud Manager with OpenStack.
2. Change to the directory where you stored the files for the deployment of the OpenStack controller node topology. In the following example, *your-deployment-name* is the name of your deployment.

   ```
   $ cd your-deployment-name
   ```
3. Download the current environment for the topology. The name of the environment can be found in the topology file.

   ```
   $ knife environment show your-environment-name -d -F json > your-environment-name.json
   ```
4. Optional: If live migration was enabled on any of the KVM Agent nodes, update the environment by using the following steps.
   a. Change the **openstack.compute.misc_nova** JSON attribute in your environment file to add the following item to the run list:

      ```
      live_migration_flag=VIR_MIGRATE_UNDEFINE_SOURCE,VIR_MIGRATE_PEER2PEER,VIR_MIGRATE_LIVE
      ```
   b. Upload the updated environment.

      ```
      $ knife environment from file your-environment-name.json
      ```
5. Optional: Create node-specific attribute files.

   Important: This step is only required when one or more of the KVM Agent nodes require different attribute values from those that are defined in your environment file.

   A separate node-specific attribute file must be created for each node that requires different attribute values. Use the information recorded in the prerequisite steps to determine if any node-specific attribute files are required. For example, a node might not use the default values for the **openstack.compute.state_path**, **openstack.compute.lock_path**, and **openstack.compute.instances_path**. The default values for these attributes, if they are not defined in the environment file, are as follows:

   - **openstack.compute.state_path**: `/var/lib/nova`
   - **openstack.compute.lock_path**: `/var/lib/nova/lock`
   - **openstack.compute.instances_path**: `/var/lib/nova/instances`

   Also, a node may not use `eth0` as the management network, `eth1` as the virtual machine data network, or `br-eth1:eth1` as the data network bridge. These values are the default values for some of the attributes in the environment file.

   The following example node attribute file illustrates different paths, management network interface, and data network bridge for a KVM Agent node:

```
{
  "openstack": {
    "endpoints": {
      "network-openvswitch": {
        "bind_interface": "eth2"
      },
      "compute-vnc-bind": {
         "bind_interface": null,
         "host": "0.0.0.0"
      }
    },
    "network": {
      "openvswitch": {
        "bridge_mappings": "default:br-ex",
        "bridge_mapping_interface": ""
      }
    },
    "compute": {
      "state_path": "/opt/stack/data/nova/",
      "lock_path": "/opt/stack/data/nova/",
      "instances_path": "/opt/stack/data/nova/instances"
    }
  }
}
```

6. Update your topology file to include the KVM Agent nodes.
   a. The first node in your topology file is your single controller node. Copy the following compute node section into the topology file for each KVM Agent node. Ensure that the compute node sections are comma-separated.

```
{
      "fqdn":"CHANGEME",
      "password":"CHANGEME",
      "identity_file":"CHANGEME",
      "quit_on_error":true,
      "run_order_number":2,
      "runlist": [
        "role[ibm-os-compute-node-kvm]"
      ]
    }
```

   b. Change or add the following JSON attributes in the topology file for each new compute node section:
      - **nodes.fqdn**: Set to the fully qualified domain name of the node system. You can also set to the public IP address, private IP address, or hostname.
      - **nodes.password** or **nodes.identity_file**: Set to the appropriate SSH root user authentication for the node system. Either a password and an SSH identity file can be used for authentication. Remove the unused attribute.
      - **nodes.runlist**: Add `role[ibm-os-client-node]` to the run list only if the CLI client was installed on the node. If the CLI client was not installed on the IBM SmartCloud Entry endpoint, then no changes are required.
      - **nodes.attribute_file**: Add this attribute only if the KVM Agent node requires a node-specific attribute file. Set to the node-specific attribute file name.

Upgrading the IBM Cloud Manager with OpenStack KVM Agent and CLI Client

You must upgrade each KVM Agent compute node by following the instructions in either the console mode or silent mode section.

Overview of the upgrade

The upgrade completes the following steps:

For KVM Agent

- Stop OpenStack Nova Compute and Neutron Open vSwitch agent
- Upgrade the OpenStack Nova Compute and Neutron Open vSwitch agent
- Restart OpenStack Nova Compute and Neutron Open vSwitch agent

For CLI Client

- Upgrade the following clients:
  - Nova
  - Neutron
  - Cinder
  - Keystone
  - Glance
- Configure the `openrc` file environment variable for accessing these Clients
- Get certification from the IaaS gateway server

Upgrading by using console mode

Follow these steps to upgrade the IBM Cloud Manager with OpenStack KVM Agent and CLI Client by using the console upgrade installation wizard.

Procedure

1. Download the latest fix for the IBM Cloud Manager with OpenStack KVM Agent and CLI Client from Fix Central. For more information, see "Applying fixes and updates for the self-service portal on Windows" on page 39.

   Note: The IBM Cloud Manager with OpenStack KVM Agent and CLI Client and the IBM Cloud Manager with OpenStack appliance must be at the same level, either the GA level, or the fix level.
2. Locate the installation image, and run the following command to start the installation wizard:
   ```
   IBM_Cloud_Manager_with_OpenStack_KVM_Agent.bin
   ```
3. Follow the instructions that are provided by the installation wizard. Choose locale language and agree to the license terms.
4. On the Choose Upgrade step, the installation wizard indicates that IBM Cloud Manager with OpenStack is installed, and prompts you to choose whether to upgrade IBM Cloud Manager with OpenStack.
5. On the Install Folder Confirmation step, the installation wizard displays the default path to the installed product folder. Press **Enter** to use this default path, or specify the path to the installed product folder. The installation wizard performs the upgrade.
6. For information about the upgrade, view the logs in the following directory:
   ```
   $USER_INSTALL_DIR/_$PRODUCT_NAME_installation/Logs/$PRODUCT_NAME_yum_upgrade.log
   ```
7. To verify whether the services installed correctly, follow these steps:
   a. Run the following commands to check whether these two processes are running:
      - **ps -ef|grep nova-compute**
      - **ps -ef|grep neutron-openvswitch-agent**
   b. Log onto the IBM Cloud Manager with OpenStack Appliance and run the following command to check whether the IBM Cloud Manager with OpenStack KVM agent is displayed in the resulting list:
      ```
      nova-manage service list
      ```

Note: The upgrade does not change the Nova and Neutron configurations or status.

What to do next

After the upgrade, set the SELinux state to enforcing mode by using the following command.

```
$ setenforce 1
```

For more information about SELinux, see the SELinux Project Wiki.

Also, save backup copies of the `/etc/nova/nova.conf` and `/etc/neutron/neutron.conf` files in case you need to refer to them in the future.

```
$ cp /etc/nova/nova.conf /etc/nova/nova.conf.cmwo_backup
$ cp /etc/neutron/neutron.conf /etc/neutron/neutron.conf.cmwo_backup
```

Upgrading by using silent mode

You can upgrade the IBM Cloud Manager with OpenStack Linux Kernel-based Virtual Machine (KVM) Agent and CLI Client by using silent mode.

About this task

To upgrade silently, you either provide upgrade parameters through the command line, or use a properties file (response file) to specify all the parameters in a single file.

Upgrading silently by using the command line:

To silently install the IBM Cloud Manager with OpenStack Linux Kernel-based Virtual Machine (KVM) Agent and CLI Client by using the command line, follow these steps:

Procedure

1. Download the latest fix for the IBM Cloud Manager with OpenStack KVM Agent and CLI Client from Fix Central. For more information, see "Applying fixes and updates for the self-service portal on Windows" on page 39.

   Note: The IBM Cloud Manager with OpenStack KVM Agent and CLI Client and the IBM Cloud Manager with OpenStack KVM appliance must be at the same level, either the GA level, or the fix level.

2. Open a command prompt and run the following command:

   ```
   /IBM_Cloud_Manager_with_OpenStack_KVM_Agent.bin -i silent -DLICENSE_ACCEPTED=TRUE
          -DUSER_INSTALL_DIR=/opt/other/other_product
          -DPRODUCT_ALIAS=other_product -DIS_UPGRADE=YES
   ```

   Tips:

   a. The **-i silent** parameter indicates silent (unattended) upgrade.

   b. The **Dkey = value** parameter contains a key and value pair to pass to the upgrade. You can specify as many of these key and value pair parameters as you want in the format `Dkey=value`, separated by a space character.

What to do next

After the upgrade, set the SELinux state to enforcing mode by using the following command.

```
$ setenforce 1
```

For more information about SELinux, see the SELinux Project Wiki.

Also, save backup copies of the /etc/nova/nova.conf and /etc/neutron/neutron.conf files in case you need to refer to them in the future.

```
$ cp /etc/nova/nova.conf /etc/nova/nova.conf.cmwo_backup
$ cp /etc/neutron/neutron.conf /etc/neutron/neutron.conf.cmwo_backup
```

Upgrading silently by using a response file:

Follow these steps to run a silent upgrade by using a response file.

Before you begin

If you choose to upgrade silently, you must create an IBM Cloud Manager with OpenStack Linux Kernel-based Virtual Machine (KVM) Agent and CLI Client silent installation response file and use the file for the upgrade.

Procedure

1. Download the latest fix for the IBM Cloud Manager with OpenStack KVM Agent and CLI Client from Fix Central. For more information, see "Applying fixes and updates for the self-service portal on Windows" on page 39.

   Note: The IBM Cloud Manager with OpenStack KVM Agent and CLI Client and the IBM Cloud Manager with OpenStack KVM appliance must be at the same level, either the GA level, or the fix level.
2. To run the upgrade by using a response file, you must specify the correct parameters in a locally saved copy of the response file. The response file properties are used to in the silent upgrade.

   The following example shows all the properties that are available during a console upgrade of the IBM Cloud Manager with OpenStack KVM Agent and CLI Client for the KVM Agent. You can modify these properties for a silent upgrade.

   ```
   #Indicate whether the license agreement been accepted
   #-----------------------------------------------------
   LICENSE_ACCEPTED=TRUE
   #Choose Install Folder
   #---------------------
   USER_INSTALL_DIR=/opt/other/other_product
   #product alias used as install folder name
   #--------------------
   PRODUCT_ALIAS=other_product
   #Choose to upgrade or not
   #------------------
   IS_UPGRADE=YES
   ```

   Notes:

   a. The property ;LICENSE_ACCEPTED in the response file specifies your agreement to the license for the application. You must specify TRUE to successfully run the silent upgrade.
   b. If you did not install the product into the default installation path, you must use the USER_INSTALL_DIR property to specify the installation path for the product.
3. When the response file is ready, open a command prompt and run the following command:

   ```
   /IBM_Cloud_Manager_with_OpenStack_KVM_Agent.bin -i silent -f <path to response file>
   ```

   where <*path to response file*> is the full path to the response file.

What to do next

After the upgrade, set the SELinux state to enforcing mode by using the following command.

```
$ setenforce 1
```

For more information about SELinux, see the SELinux Project Wiki.

Also, save backup copies of the /etc/nova/nova.conf and /etc/neutron/neutron.conf files in case you need to refer to them in the future.

```
$ cp /etc/nova/nova.conf /etc/nova/nova.conf.cmwo_backup
$ cp /etc/neutron/neutron.conf /etc/neutron/neutron.conf.cmwo_backup
```

Uninstalling the IBM Cloud Manager with OpenStack Linux Kernel-based Virtual Machine (KVM) Agent and CLI Client

If you need to uninstall IBM Cloud Manager with OpenStack Linux Kernel-based Virtual Machine (KVM) Agent and CLI Client in the future, the uninstaller supports uninstallation by using the command line.

About this task

Use the following steps to uninstall the IBM Cloud Manager with OpenStack KVM Agent and CLI Client on Red Hat Enterprise Linux:

Note: When you uninstall the IBM Cloud Manager with OpenStack KVM Agent, you can choose not to uninstall Open vSwitch and Libvirt to preserve the virtual server and virtual switches on the host. Then if you install the IBM Cloud Manager with OpenStack KVM Agent again, you can reuse these virtual resources.

Procedure

1. (KVM Agent only) Run the following commands to shut down the IBM Cloud Manager with OpenStack Management Services:

   ```
   service openstack-nova-compute stop
   service neutron-openvswitch-agent stop
   ```

2. Navigate to the $USER_SHORTCUTS$/bin folder, which contains the uninstallation program.
3. Run the following command:

   ```
   ./Uninstall_IBM_SmartCloud_Entry_KVM_Agent
   ```

4. Follow the instructions in the uninstallation wizard to complete the operation.
5. (KVM Agent only) Choose whether to uninstall Open vSwitch. If you choose to uninstall Open vSwitch,the virtual switches that were created by Open vSwitch are removed.

   Note: While the installation mode is silent, the uninstallation is in silent mode automatically, and Open vSwitch virtual switches are not removed. To remove Open vSwitch virtual switches, append the following option to the uninstallation command:

   ```
   -DIS_OVS_UNINSTALL=YES
   ```

6. (KVM Agent only) Choose whether to uninstall Libvirt.

   Note: While the installation mode is silent, the uninstallation is in silent mode automatically, and Libvirt is not removed. To remove Libvirt , append the following option to the uninstallation command:

   ```
   -DIS_OVS_UNINSTALL=YES
   ```

7. Complete the uninstallation.

Results

After the IBM Cloud Manager with OpenStack KVM Agent and CLI Client is uninstalled, the uninstaller backs up the following files by renaming them to *.rpmsave:

- nova.conf
- neutron.conf

- `ovs_neutron_plugin.ini`

Completing the migration process for KVM agent or CLI client

Complete the migration process for your KVM Agent nodes to ensure that the nodes can be managed as part of your topology.

Procedure

1. Log in to the deployment system as the root user. The deployment system is the system where IBM Cloud Manager with OpenStack is installed.
2. Change to the directory where you stored the files for your topology deployment. Change *your-deployment-name* to the name of the deployment.

   ```
   $ cd your-deployment-name
   ```
3. Complete the migration of all the KVM Agent nodes. This step ensures that the nodes can be managed as part of your topology.

   ```
   $ knife os manage deploy topology your-topology-name.json
   ```

   Note: Console access to instances from the OpenStack Dashboard might fail for existing instances until after the instance is rebooted from OpenStack.

Chapter 5. Deploying IBM Cloud Manager with OpenStack

After IBM Cloud Manager with OpenStack is installed, you must follow these instructions to complete the installation.

To continue, select a topology and deploy the components that are necessary to create your cloud environment.

1. Select the topology that you want to deploy. Review the "Topology overview" for more information.
2. Deploy the components that are necessary to create your cloud environment. See "Deploying the cloud environment" on page 76 and follow the steps for the topology you select to deploy.

Topology overview

The IBM Cloud Manager with OpenStack solution provides some predefined topologies.

The following topologies are the supported configurations for IBM Cloud Manager with OpenStack.

Table 23. Supported topologies

Topology	Description
Minimal	For product evaluation purposes. This topology is the simplest topology and does not require any customization.
Controller +*n* compute	For smaller production environments. This topology provides a single controller node, plus any number of compute nodes. You can configure this topology for your specific needs. For example, configure networking, resource scheduler, and other advanced customization.
Distributed database	For larger production environments. This topology is similar to the controller +*n* compute topology; however, it allows the IBM Cloud Manager with OpenStack database service to run on a separate compute node. It also supports advanced customization.

The following terminology is used by the topologies:

Deployment server
The deployment server is the system where you install the IBM Cloud Manager with OpenStack components. The deployment server runs the Chef server and maintains the IBM Cloud Manager with OpenStack repository.

IBM Cloud Manager with OpenStack repository
The repository is the location that where all of the IBM Cloud Manager with OpenStack components are stored. The repository includes Chef resources, RPMs, commands, licenses, dependencies, and more. As an administrator, you manage your topology using this repository.

The repository, along with the associated OpenStack Chef cookbooks and dependencies provide the basic building blocks that you can use to deploy IBM Cloud Manager with OpenStack.

The topologies support the following operating systems, database, message queue, and OpenStack components.

- OpenStack Components: Identity, Image, Network, Compute, Orchestration, Block Storage, Telemetry, and Dashboard

© Copyright IBM Corp. 2012, 2014

- OpenStack Networking: Neutron
- OpenStack Compute Scheduler: Compute scheduler filters (default), IBM Platform Resource Scheduler (PRS)
- Operating System for OpenStack controller: Red Hat Enterprise Linux 6.5 (default) and 6.4
- Database: IBM DB2 (default) , MySQL
- Message Queue: Qpid
- Hypervisor types: Linux Kernel-based Virtual Machine (KVM), PowerKVM, QEMU, z/VM, and Hyper-V
- Virtualization Manager: PowerVC
- Self-service portal (optional)

You can use a single IBM Cloud Manager with OpenStack setup to manage multiple topologies and multiple environments for those topologies. To manage, use the IBM Cloud Manager with OpenStack commands that simplify administration.

Minimal deployment

This topology is for product evaluation or proof of concept deployments.

With the minimal topology (example-ibm-os-allinone environment), all IBM Cloud Manager with OpenStack services run from one system. The IBM Cloud Manager with OpenStack node runs basic services, including the KVM/QEMU Nova compute service.

The deployment server is the system to which you install IBM Cloud Manager with OpenStack. The solution delivery repository is the repository that contains all the IBM Cloud Manager with OpenStack components. As an administrator, you manage the topology using this repository. The graphic below shows an example of a minimum topology.

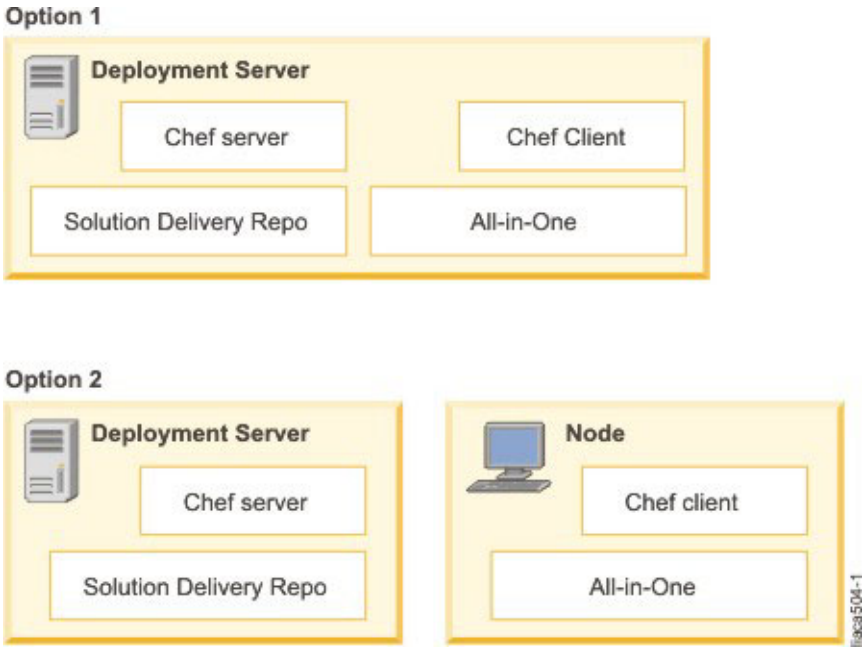

Related tasks:

"Deploying minimal with customizations" on page 78
Use the following procedure to deploy the minimal topology with customizations.

Controller +*n* compute deployment

This topology is for smaller deployments.

The controller *+n* compute deployment topology (example-ibm-os-single-controller-n-compute environment) is similar in size to the topology used by IBM SmartCloud Entry 3.1 and 3.2. The IBM Cloud Manager with OpenStack single controller node provides the basic IBM OpenStack services, excluding the OpenStack compute node services. The OpenStack compute node services are provided by the IBM Cloud Manager with OpenStack compute nodes. The IBM Cloud Manager with OpenStack single controller node can run PowerVC driver services. This topology supports all of the IBM Cloud Manager with OpenStack supported compute hypervisors. The specific compute hypervisors supported include:

- KVM/QEMU
- PowerKVM
- z/VM
- Hyper-V

The deployment server can double as the IBM Cloud Manager with OpenStack single controller node, as shown in option 1.

Note: The Chef client does not run on Hyper-V compute nodes since the Hyper-V Agent installer is used.

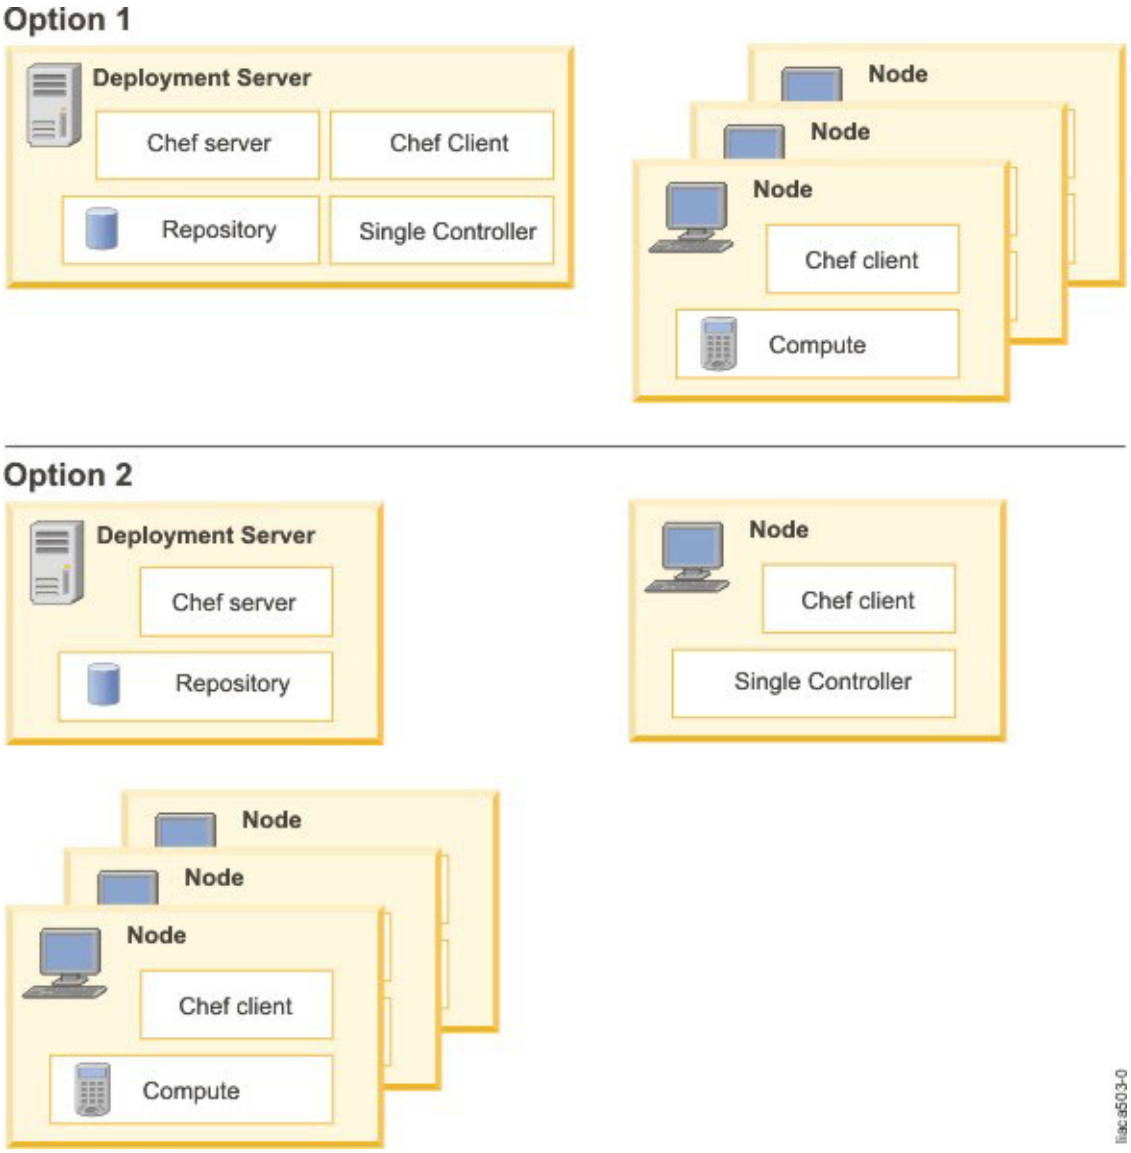

Related tasks:

“Deploying a controller +*n* compute or distributed database topology” on page 80
Deploy the components that are necessary to create a cloud environment with more than one node. Use the instructions that apply to your specific hypervisor for each compute node.

Distributed database deployment

This topology is for larger deployments that use a distributed database.

The distributed database deployment topology (`example-ibm-os-single-controller-n-compute environment`) is similar to the controller +*n* compute topology; however, it allows the IBM Cloud Manager with OpenStack database service to run on a separate node. The IBM Cloud Manager with OpenStack single controller node provides the basic IBM OpenStack services, excluding the OpenStack compute node services. The OpenStack compute node services are provided by the IBM Cloud Manager with OpenStack compute nodes. The IBM Cloud Manager with OpenStack single controller node can run PowerVC driver services. This topology supports all of the IBM Cloud Manager with OpenStack supported compute hypervisors. The specific compute hypervisors supported include:

- KVM/QEMU hypervisor
- PowerKVM
- z/VM
- Hyper-V

This image shows the IBM Cloud Manager with OpenStack database service running on a separate node.

Note: The Chef client does not run on Hyper-V compute nodes since the Hyper-V Agent installer is used.

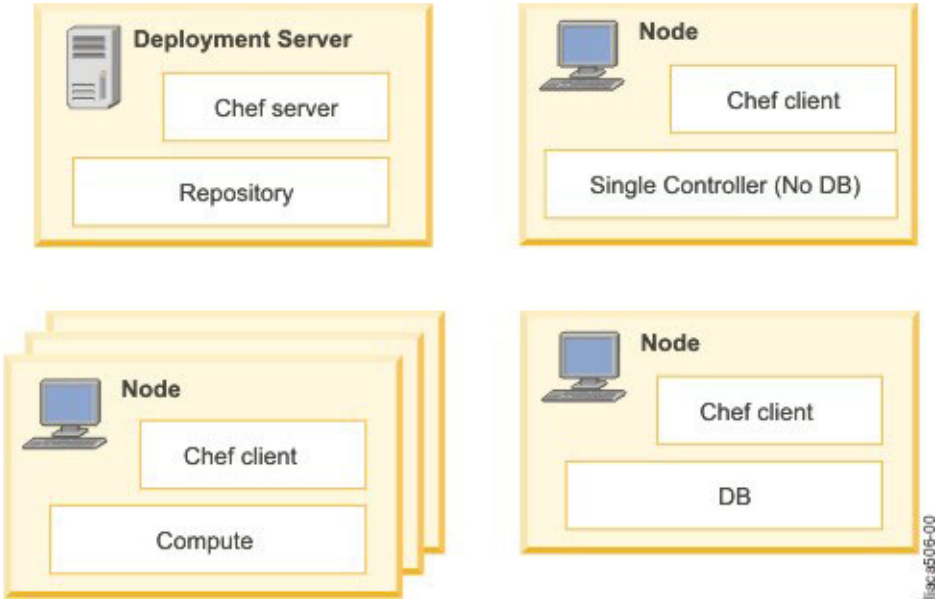

Deploying the cloud environment

Deploy the components that are necessary to create your cloud environment.

Complete the prerequisites, then follow the steps for the topology you want to deploy.

Deploying prerequisites

Before deploying an IBM Cloud Manager with OpenStack topology, you need to complete the prerequisites.

Procedure

1. Install the latest fix packs for IBM Cloud Manager with OpenStack before you proceed. For information about fix packs, see “Applying fixes and updates” on page 38.
2. Identify the node systems for the topology.

   Note: You can use the deployment server as a node.

   a. Verify that the time on the node systems is within 15 minutes of the time that is shown on the deployment server. Consider synchronizing the system clock of the deployment server and node systems with a network time protocol (NTP) server. For more information, see “Installing a network time service” on page 122.
   b. Record the IP addresses for each node.
   c. Record the fully qualified domain names for each node.
   d. Record the root user login information (either password or SSH identity file) for each node.
   e. Record the number and name of each network interface card on each node. The example environment defaults to *eth0* for the management network and a dedicated *eth1* for the virtual machine data network. If the deployment nodes do not have an *eth0* or *eth1* interface or if the management and dedicated data network are on different interfaces besides *eth0* and *eth1*, then you must change the default environment settings when you configure the deployment. The example environment assumes that the network configuration is identical across all of the deployment nodes.
   f. The network configuration and hypervisor type limit the type of networks that can be defined. For more information about network considerations, see “Network considerations” on page 16.
3. Verify the OpenStack controller node system meets the following criteria:
   - Operating System: Red Hat Enterprise Linux 6.4 or 6.5
   - Architecture: *x86_64* or *ppc64*
4. If applicable, verify the KVM, QEMU compute node system meets the following criteria:
   - See KVM prerequisites for details.
   - To use the KVM hypervisor type on a node system, the node must support KVM acceleration. For details, see the OpenStack documentation.
5. If applicable, verify that the z/VM compute node system meets the following criteria:
   - See “z/VM prerequisites” on page 13 for details.
   - To use the z/VM hypervisor, one or more x86_64 Red Hat Enterprise Linux system nodes should be used to install the compute and network driver to manage the z/VM hypervisor. One x86_64 Red Hat Enterprise Linux system node is supported for each z/VM node. For more information on configuring the z/VM hypervisor, see the z/VM OpenStack user manual.
6. If applicable, verify the PowerKVM compute node system meets the following criteria: See “PowerKVM prerequisites” on page 11 for details.
7. If applicable, verify that the Hyper-V compute node system meets the following criteria: See “Installing and uninstalling the IBM Cloud Manager with OpenStack Hyper-V Agent” on page 27 for details.
8. Linux node systems must have access to a yum repository, which contains base operating system packages for your node systems. Many of the OpenStack components depend on operating system packages that are automatically installed on the node system during the deployment process. If you do not have a yum repository that is configured on the node system, which contains these packages, the Deployment Server can be configured as a yum repository for operating system packages. For configuration information, see “Configuring operating system yum repositories on the deployment server” on page 36.
9. Consider this restriction before deploying IBM Cloud Manager with OpenStack, in case you must undo a deployment.

- Support does not exist to uninstall a deployed topology for IBM Cloud Manager with OpenStack. You must reinstall or reset the node back to its pre-deployment state. You cannot attempt to redeploy to the same managed system without first setting the node back to its pre-deployment state. Back up your node system before deployment by using existing snapshot or capture capabilities that are provided by the underlying virtualization manager or other backup methods.
- The node must also be deleted from the chef server. For more information, see "Redeploying a node cleanup" on page 113.

Deploying a minimal topology

Deploy the components that are necessary to create a very basic cloud environment. The deployment can be done with or without customizations.

Deploying evaluation

Use the following procedure to deploy the minimal topology without customizations.

Procedure

1. Log in to the deployment system as the root user. This is the system where IBM Cloud Manager with OpenStack was installed.
2. Deploy your topology. Change **node.fqdn.com** to the fully qualified domain name of the node.

   ```
   $ knife os manage deploy evaluation node.fqdn.com
   ```
3. After the deployment completes, the IBM Cloud Manager with OpenStack services are ready to use. Log into the node system. The administrator user name and password can be found in the `openrc` file under the root user home directory. The IBM Cloud Manager with OpenStack dashboard is available at https://node.fqdn.com/, where **node.fqdn.com** is the fully qualified domain name of the node. For information about managing IBM Cloud Manager with OpenStack services, see "Managing IBM Cloud Manager with OpenStack services" on page 169.

Results

You are ready to start using IBM Cloud Manager with OpenStack services.

Deploying minimal with customizations

Use the following procedure to deploy the minimal topology with customizations.

Procedure

1. Log in to the deployment system as the root user. This is the system where IBM Cloud Manager with OpenStack was installed.
2. Create a directory to store the files for the topology that you deploy. Change **your-deployment-name** to the name for your deployment.

   ```
   $ mkdir your-deployment-name
   $ chmod 600 your-deployment-name
   $ cd your-deployment-name
   ```
3. Copy the example environment for the topology that you deploy. Change **your-environment-name** to the name for your environment.

   ```
   $ knife environment show example-ibm-os-allinone -d -Fjson > your-environment-name.json
   ```
4. Change the following JSON attributes in your environment file, `your-environment-name.json`
   - **Name**: Set to your environment name: *your-environment-name*.
   - **openstack.endpoints.host, openstack.endpoints.bind-host and openstack.endpoints.mq.host**: Change from *127.0.0.1* to the IP address of the node system for the topology.
5. Copy the following example topology to a file, `your-topology-name.json`. Change **your-topology-name** to the name for your topology.

```
{
  "name":"CHANGEME",
  "description":"CHANGEME",
  "environment":"CHANGEME",
  "run_sequentially":true,
  "nodes": [
    {
      "fqdn":"CHANGEME",
      "password":"CHANGEME",
      "identity_file":"CHANGEME",
      "quit_on_error":true,
      "runlist": [
        "role[ibm-os-allinone-kvm]"
      ]
    }
  ]
}
```

6. Change the following JSON attributes in your topology file, `your-topology-name.json`:
   - **`Name`**: Set to your topology name: *your-topology-name*.
   - **`Description`**: Set to the description for your topology.
   - **`Environment`**: Set to the environment for your topology: *your-environment-name*.
   - **`nodes.fqdn`**: For each node, you must set to the fully qualified domain name of the node system. The deployment system must be able to ssh using the fully qualified domain name. You can also set to the public IP address, private IP address, or host name.
   - **`nodes.password`** or **`nodes.identity_file`**: Set to the appropriate SSH root user authentication for the node system. Either a password and an SSH identity file can be used for authentication. Remove the unused attribute.
7. (Optional) Complete any optional customizations.
   - "Customizing for self-service portal" on page 100
   - "Customizing Platform Resource Scheduler" on page 102
   - "Customizing basic options" on page 104
   - "Customizing advanced options" on page 105
8. Deploy your topology.
   a. Upload the environment for your deployment.
      ```
      $ knife environment from file your-environment-name.json
      ```
   b. Deploy the topology.
      ```
      $ knife os manage deploy topology your-topology-name.json
      ```
9. After the deployment is complete, the IBM Cloud Manager with OpenStack services are ready to use. Log in to the node system. The administrator user name and password can be found in the `openrc` file under the root user home directory. The IBM Cloud Manager with OpenStack dashboard is available at `https://node.fqdn.com/`, where **`node.fqdn.com`** is the fully qualified domain name of the node.
   a. If you also included the **`ibm-sce-node`** role in your topology, you can check the IBM Cloud Manager with OpenStack self-service portal status on your node by using **`service sce status`**. The web interface for IBM Cloud Manager with OpenStack self-service portal is available at `http://node.fqdn.com:18080/cloud/web/login.html`. You can log in using admin user with *admin* as the default password. After you are logged in, you can add the IaaS gateway and an OpenStack cloud using the *sceagent* user with *sceagent* as the default password.
   b. For more information about managing IBM Cloud Manager with OpenStack services, see "Managing IBM Cloud Manager with OpenStack services" on page 169.

What to do next

You are ready to start using IBM Cloud Manager with OpenStack services.

Related concepts:

"Minimal deployment" on page 74
This topology is for product evaluation or proof of concept deployments.

Related reference:

Platform Resource Scheduler online product documentation

Deploying a controller +*n* compute or distributed database topology

Deploy the components that are necessary to create a cloud environment with more than one node. Use the instructions that apply to your specific hypervisor for each compute node.

Related concepts:

"Controller +*n* compute deployment" on page 74
This topology is for smaller deployments.

Related reference:

"Customizing advanced options" on page 105
You can update the JSON attributes in your environment file, `your-environment-name.json` to customize advanced options such as SELinux, iptables, and message queue performance.

Deploying with Hyper-V compute nodes

Deploy the components that are necessary to create a cloud environment with Hyper-V compute nodes.

Before you begin

Before you begin, ensure you completed the "Deploying prerequisites" on page 76 steps.

About this task

Use the following procedure to deploy the topology to your node systems.

Procedure

1. Log in to the deployment system as the *root* user. This is the system where IBM Cloud Manager with OpenStack was installed.
2. Create a directory to store the files for the topology that you deploy. Change **your-deployment-name** to the name for your deployment.

   ```
   $ mkdir your-deployment-name
   $ chmod 600 your-deployment-name
   $ cd your-deployment-name
   ```
3. Copy the example environment for the topology that you deploy. Change **your-environment-name** to the name for your environment.

   Note: The following command must be entered on a single line, even though the example shows a line break for formatting purposes.

   ```
   $ knife environment show example-ibm-os-single-controller-n-compute -d -Fjson
   > your-environment-name.json
   ```
4. Change the following JSON attributes in your environment file, `your-environment-name.json`:
   a. **Name**: Set to your environment name: your-environment-name.
   b. **Description**: Set to the description for your environment.
   c. **openstack.endpoints.host, openstack.endpoints.bind-host, and openstack.endpoints.mq.host**: Change from *127.0.0.1* to the IP address of the controller node system for the topology.
   d. **openstack.network.ml2.mechanism_drivers**: *hyperv*
   e. **openstack.network.ml2.tenant_network_types**: *vlan* or *flat*, depending on your environment. The default value is *vlan*.
   f. **openstack.network.ml2.flat_networks**: Add your Hyper-V flat networks. For example, *physnet1*.

g. **openstack.network.ml2.network_vlan_ranges**: Add your Hyper-V network VLAN range. For example, *physnet1:1000:2999*.

5. Copy the following example topology to a file, `your-topology-name.json`. Change **your-topology-name** to the name for your topology. Here is an example topology with Hyper-V compute nodes.

```
{
  "name":"CHANGEME",
  "description":"CHANGEME",
  "environment":"CHANGEME",
  "secret_file":"CHANGEME",
  "run_sequentially":false,
  "nodes": [
    {
      "fqdn":"CHANGEME",
      "password":"CHANGEME",
      "identity_file":"CHANGEME",
      "quit_on_error":true,
      "run_order_number":1,
      "runlist": [
        "role[ibm-os-single-controller-node]"
      ]
    }
  ]
}
```

6. Customize the topology file. The node in your topology file is your single controller node. The Hyper-V compute nodes will be deployed after the single controller node. Change the following JSON attributes in your topology file, `your-topology-name.json`:
   a. **Name**: Set to your topology name: *your-topology-name*.
   b. **Description**: Set to the description for your topology.
   c. **Environment**: Set to the environment for your topology: *your-environment-name*.
   d. **nodes.fqdn**: For each node, you must set to the fully qualified domain name of the node system. The deployment system must be able to ssh using the fully qualified domain name. You can also set to the public IP address, private IP address, or host name.
   e. **nodes.password** or **nodes.identity_file**: Set to the appropriate SSH root user authentication for the node system. Either a password and an SSH identity file can be used for authentication. Remove the unused attribute.
7. Customize the passwords and secrets before deploying. For instructions, see "Customizing passwords and secrets" on page 102.
8. Configure the OpenStack block storage (cinder) driver. By default, the environment is configured to use the LVM iSCSI cinder driver. You can change the following JSON attributes in your environment file, `your-environment-name.json`, to customize the LVM iSCSI cinder driver configuration.
   a. **openstack.block-storage.volume.create_volume_group**: If set to *true* (default) then the cinder-volumes volume group is created on the controller node with a size determined by **openstack.block-storage.volume.volume_group_size**. Set to *false* if you want to create the volume group manually.
   b. **openstack.block-storage.volume.volume_group_size**: The amount of storage you use must be smaller than the size available. If necessary, you can set the value to your free disk size. The default value is 40 GB.
   c. **openstack.block-storage.volume.iscsi_ip_address**: Change from *127.0.0.1* to the management IP address of the volume service node.

   To customize your environment for the IBM Storwize Cinder driver, see"Configuring IBM Storwize Cinder driver" on page 107.
9. (Optional) Complete any optional customizations.
   - "Customizing for a distributed database topology" on page 100
   - "Customizing for self-service portal" on page 100

- "Customizing basic options" on page 104
- "Customizing advanced options" on page 105

10. Deploy your topology.
    a. Upload the environment for your deployment.
       ```
       $ knife environment from file your-environment-name.json
       ```
    b. Deploy the topology.
       ```
       $ knife os manage deploy topology your-topology-name.json
       ```
11. Install the Hyper-V agent on your Hyper-V compute nodes. For instructions, see "Installing and uninstalling the IBM Cloud Manager with OpenStack Hyper-V Agent" on page 27.
12. After the deployment is complete, the IBM Cloud Manager with OpenStack services are ready to use. Log in to the controller node system. The administrator user name and password can be found in the `openrc` file under the root user home directory. The IBM Cloud Manager with OpenStack dashboard is available at `https://controller.fqdn.com/`, where **`controller.fqdn.com`** is the fully qualified domain name of the controller node in your topology.
    a. If you also include the **`ibm-sce-node`** role in your topology, you can check the IBM Cloud Manager with OpenStack self-service portal status on your controller node by using **`service sce status`**. The web interface for IBM Cloud Manager with OpenStack self-service portal is available at `http://contoller.fqdn.com:18080/cloud/web/login.html`. You can log in using admin user with *openstack1* as the default password (unless customized). After you are logged in, you can add the IaaS gateway and an OpenStack cloud. For more information, see "Configuring an OpenStack cloud" on page 198.
    b. For more information about managing IBM Cloud Manager with OpenStack services, see "Managing IBM Cloud Manager with OpenStack services" on page 169.

Results

You are ready to start using IBM Cloud Manager with OpenStack services.

Related reference:

"Troubleshooting errors when deploying topologies" on page 253
If a topology deployment fails, review the log output from the deployment command for more information. Identify errors, starting at the beginning of the log output. The first error found is typically the most useful. Errors found at the end of the log output might be the result of earlier errors. The following topics might provide additional information to solve the problem.

Deploying with KVM or QEMU compute nodes

Deploy the components that are necessary to create a cloud environment with KVM or QEMU compute nodes.

Before you begin

- Before you begin, ensure you completed the "Deploying prerequisites" on page 76 steps.
- For the node systems, the network configuration and hypervisor type limit the type of networks that can be defined. For example:

Table 24. Supported network configuration

Number of network interface cards (NICs) per node	Network type	Hypervisor type
1 - *eth0* is dedicated as the management network	Local or GRE	KVM only **Note:** If one or more of the compute hypervisors is not KVM, then GRE cannot be used.
2 or more • *eth0* is dedicated as the management network. • *eth1* is dedicated as the virtual machine data network. **Note:** Recommended configuration	• Local • Flat • VLAN • GRE	KVM only **Note:** If one or more of the compute hypervisors is not KVM, then GRE cannot be used.

Note:

– The virtual machine data network must be on a dedicated interface. If there is an IP address on this data network interface, it will no longer be ping-able after the node is deployed. Communication to the node must be done through the management network or another interface on the node.
– The `local` network type can be configured; however, the network traffic is limited to the current node. The minimum topology uses the `local` network option, by default.

For more information about network considerations, see "Network considerations" on page 16.

Use the following procedure to deploy the topology to your node systems.

Procedure

1. Log in to the deployment system as the root user. This is the system where IBM Cloud Manager with OpenStack was installed.
2. Create a directory to store the files for the topology that you deploy. Change **your-deployment-name** to the name for your deployment.

   ```
   $ mkdir your-deployment-name
   $ chmod 600 your-deployment-name
   $ cd your-deployment-name
   ```
3. Copy the example environment for the topology that you deploy. Change **your-environment-name** to the name for your environment.

   Note: The following command must be entered on a single line, even though the example shows a line break for formatting purposes.

   ```
   $ knife environment show example-ibm-os-single-controller-n-compute -d -Fjson
   > your-environment-name.json
   ```
4. Change the following JSON attributes in your environment file, `your-environment-name.json`:
   - **Name**: Set to your environment name: *your-environment-name*.
   - **Description**: Set to the description for your environment.
   - **openstack.endpoints.host**, **openstack.endpoints.bind-host**, and **openstack.endpoints.mq.host**: Change from *127.0.0.1* to the IP address of the controller node system for the topology.
   - **openstack.compute.libvirt.virt_type**: Set to the hypervisor type, *kvm or qemu*, for the topology.
   - **(Single network interface card or no virtual machine data network)**: If you are using a GRE network with a single network interface card on the nodes (or no virtual machine data network), you must change the following default values in the environment:

     ```
     openstack.network.openvswitch.tenant_network_type = "gre"
     openstack.network.openvswitch.bridge_mappings = ""
     openstack.network.openvswitch.network_vlan_ranges = ""
     openstack.network.openvswitch.bridge_mapping_interface = ""
     openstack.network.ml2.tenant_network_types = "gre"
     openstack.network.ml2.network_vlan_ranges = ""
     openstack.network.ml2.flat_networks = ""
     ```

     If the management network interface of the nodes is not *eth0*, then update all occurrences of *eth0* in the environment file to match your network configuration on the nodes.

     (Recommended network configuration) If the management network interface of the nodes is not *eth0*, if the virtual machine data network interface is not *eth1*, or both apply, then update all occurrences of *eth0*, *eth1*, or both in the environment file to match your network configuration. The following list displays some of the networking properties and their default values (from the example environment) that you might need to change. In most cases, these default values should be sufficient and do not need to be changed.

     – **openstack.network.core_plugin**: *neutron.plugins.ml2.plugin.Ml2Plugin*.
     The **core_plugin** property controls which OpenStack plug-in is used by the Neutron server. The two recommended choices for this property include:
     "neutron.plugins.openvswitch.ovs_neutron_plugin.OVSNeutronPluginV2" and

"neutron.plugins.ml2.plugin.Ml2Plugin". If you are migrating from an earlier version of the IBM SmartCloud Entry product, it is recommended that you use that OVS plug-in.

In the example environment, the **openstack.network.core_plugin** property is set to *"neutron.plugins.ml2.plugin.ML2Plugin"* and the **openstack.network.ml2.mechanism_drivers** property is set to *"openvswitch"*. When the **core_plugin** property is set to ML2, then the ML2 specific properties (properties that begin with **openstack.network.ml2.***) must be consistent and kept in sync with the properties of the mechanism driver (**openstack.network.openvswitch.***). The example environment is set up with these two sets of properties in sync.

If you change the **openstack.network.core_plugin** property from *"neutron.plugins.ml2.plugin.Ml2Plugin"* to *"neutron.plugins.openvswitch.ovs_neutron_plugin.OVSNeutronPluginV2"*, you MUST remove the **openstack.network.service_plugins** property or set the value of the property to *[ ]*. In the example environment, the **openstack.network.service_plugins** property defaults to *"neutron.services.l3_router.l3_router_plugin.L3RouterPlugin"*. If the **core_plugin** is set to OVS, the ML2 specific properties in the environment are ignored.

- **openstack.network.openvswitch.bridge_mapping_interface**: *"br-eth1:eth1"*. The **bridge_mapping_interface** property is used to control the creation of the data network OVS bridge on the nodes. If Open vSwitch is installed and the data network bridge is already configured on the nodes, this property is not necessary and you can set the variable to *" "*. If a specific network configuration is needed for the data network (for example, bonding), you must set this variable to "" and complete set up manually before or after the node is converged.
- **openstack.network.openvswitch.bridge_mappings**: *"default:br-eth1"*.
  The **bridge_mappings** property controls which OVS bridge is used for flat and VLAN network traffic from the node. If this OVS bridge does not exist, the Open vSwitch agent does not start. This bridge can be automatically created by setting the **bridge_mapping_interface** property.
- **openstack.network.openvswitch.network_vlan_ranges** and **openstack.network.ml2.network_vlan_ranges**: These two properties define the default vlan range used when creating a tenant network. Both of these properties default to *default:1:4094*. This vlan range might need to be adjusted based on the vlan configuration of the physical switches and hypervisors in your environment. If the **core_plugin** value is set to *ML2*, the values of both properties must be the same.

5. Copy the following example topology to a file, `your-topology-name.json`. Change **your-topology-name** to the name for your topology. Here is an example topology with KVM or QEMU compute nodes.

```
{
  "name":"CHANGEME",
  "description":"CHANGEME",
  "environment":"CHANGEME",
  "secret_file":"CHANGEME",
  "run_sequentially":false,
  "nodes": [
    {
      "fqdn":"CHANGEME",
      "password":"CHANGEME",
      "identity_file":"CHANGEME",
      "quit_on_error":true,
      "run_order_number":1,
      "runlist": [
        "role[ibm-os-single-controller-node]"
      ]
    },
    {
      "fqdn":"CHANGEME",
      "password":"CHANGEME",
      "identity_file":"CHANGEME",
      "quit_on_error":true,
      "run_order_number":2,
      "runlist": [
```

```
        "role[ibm-os-compute-node-kvm]"
      ]
    }
  ]
}
```

6. Customize the topology file.
   a. The first node in your topology file is your single controller node. The second node in your topology file is for a compute node. If your topology requires extra compute nodes, copy the compute node section as many times as needed. Ensure that additional compute node sections are comma-separated.
   b. Change the following JSON attributes in your topology file, `your-topology-name.json`:
      - **`Name`**: Set to your topology name: *your-topology-name*.
      - **`Description`**: Set to the description for your topology.
      - **`Environment`**: Set to the environment for your topology: *your-environment-name*.
      - **`nodes.fqdn`**: For each node, you must set to the fully qualified domain name of the node system. The deployment system must be able to ssh using the fully qualified domain name. You can also set to the public IP address, private IP address, or host name.
      - **`nodes.password`** or **`nodes.identity_file`**: For each node, set to the appropriate SSH root user authentication for the node system. Either a password and an SSH identity file can be used for authentication. Remove the unused attribute for each node.
   c. (Optional) Create node specific attribute files. This step is only required when one or more nodes in your topology require different attributes from those defined in your environment file `your-environment-name.json`.
      1) Create a node-specific attribute file that is similar to the following format. For example, a node may not have an *eth0* network interface, which is the default value for some attributes. Below is an example node attribute file that can be used to change the *eth0* default network on a compute node over to use *eth2*.

```
{
  "openstack": {
    "endpoints": {
      "network-openvswitch": {
        "bind_interface": "eth2"
      },
      "compute-vnc-bind": {
        "bind_interface": "eth2"
      }
    }
}
```

      After creating the node specific attribute files, add the **`nodes.attribute_file`** JSON attributes in your topology file, `your-topology-name.json`:
      - **`nodes.attribute_file`**: For each node, set to the attribute JSON file which overrides the attributes in the **`default_attributes`** section of the environment file.
7. Customize the passwords and secrets before deploying. For instructions, see "Customizing passwords and secrets" on page 102.
8. Configure the OpenStack block storage (cinder) driver. By default, the environment is configured to use the LVM iSCSI cinder driver. You can change the following JSON attributes in your environment file, `your-environment-name.json`, to customize the LVM iSCSI cinder driver configuration.
   a. **`openstack.block-storage.volume.create_volume_group`**: If set to *true* (default) then the cinder-volumes volume group is created on the controller node with a size determined by **`openstack.block-storage.volume.volume_group_size`**. Set to *false* if you want to create the volume group manually.
   b. **`openstack.block-storage.volume.volume_group_size`**: The amount of storage you use must be smaller than the size available. If necessary, you can set the value to your free disk size. The default value is 40 GB.

c. **openstack.block-storage.volume.iscsi_ip_address**: Change from *127.0.0.1* to the management IP address of the volume service node.

To customize your environment for a different cinder driver, see "Configuring Cinder drivers" on page 107.

9. (Optional) Complete any optional customizations.
   - "Customizing for a distributed database topology" on page 100
   - "Customizing for self-service portal" on page 100
   - "Customizing Platform Resource Scheduler" on page 102
   - "Customizing basic options" on page 104
   - "Customizing advanced options" on page 105
10. Deploy your topology.
    a. Upload the environment for your deployment.
       ```
       $ knife environment from file your-environment-name.json
       ```
    b. Deploy the topology.
       ```
       $ knife os manage deploy topology your-topology-name.json
       ```
11. After the deployment is complete, the IBM Cloud Manager with OpenStack services are ready to use. Log in to the controller node system. The administrator user name and password can be found in the `openrc` file under the root user home directory. The IBM Cloud Manager with OpenStack dashboard is available at `https://controller.fqdn.com/`, where **controller.fqdn.com** is the fully qualified domain name of the controller node in your topology.
    a. If you also include the **ibm-sce-node** role in your topology, you can check the IBM Cloud Manager with OpenStack self-service portal status on your controller node by using **service sce status**. The web interface for IBM Cloud Manager with OpenStack self-service portal is available at `http://contoller.fqdn.com:18080/cloud/web/login.html`. You can log in using admin user with *openstack1* as the default password (unless customized). After you are logged in, you can add the IaaS gateway and an OpenStack cloud. For more information, see "Configuring an OpenStack cloud" on page 198.
    b. For more information about managing IBM Cloud Manager with OpenStack services, see "Managing IBM Cloud Manager with OpenStack services" on page 169.
12. (Optional) Verify the Open vSwitch (OVS) configuration for your network. See "Verifying Open vSwitch configuration" on page 110.

What to do next

You are ready to start using IBM Cloud Manager with OpenStack services.

Related reference:

"Troubleshooting errors when deploying topologies" on page 253
If a topology deployment fails, review the log output from the deployment command for more information. Identify errors, starting at the beginning of the log output. The first error found is typically the most useful. Errors found at the end of the log output might be the result of earlier errors. The following topics might provide additional information to solve the problem.

Platform Resource Scheduler online product documentation

Configure the Image Service

"Creating initial networks" on page 117
After you deploy the components for creating a cloud environment, you can create several different types of networks.

Deploying with PowerKVM compute nodes

Deploy the components that are necessary to create a cloud environment with IBM PowerKVM compute nodes.

Before you begin

Before you begin, ensure you completed the "Deploying prerequisites" on page 76 steps.

About this task

Use the following procedure to deploy the topology to your node systems.

Procedure

1. Log in to the deployment system as the root user. This is the system where IBM Cloud Manager with OpenStack was installed.
2. Create a directory to store the files for the topology that you deploy. Change **your-deployment-name** to the name for your deployment.

   ```
   $ mkdir your-deployment-name
   $ chmod 600 your-deployment-name
   $ cd your-deployment-name
   ```

3. Copy the example environment for the topology that you deploy. Change **your-environment-name** to the name for your environment.

   Note: The following command must be entered on a single line, even though the example shows a line break for formatting purposes.

   ```
   $ knife environment show example-ibm-os-single-controller-n-compute -d -Fjson
   > your-environment-name.json
   ```

4. Change the following JSON attributes in your environment file, `your-environment-name.json`:
   - **Name**: Set to your environment name: *your-environment-name*.
   - **Description**: Set to the description for your environment.
   - **openstack.endpoints.host**, **openstack.endpoints.bind-host**, and **openstack.endpoints.mq.host**: Change from *127.0.0.1* to the IP address of the controller node system for the topology.
   - **openstack.block-storage.volume.create_volume_group**: If set to *true* (default) then the cinder-volumes volume group is created on the controller node with a size determined by **openstack.block-storage.volume.volume_group_size**. Set to *false* if you want to create the volume group manually.
   - **openstack.block-storage.volume.volume_group_size**: The amount of storage you use must be smaller than the size available. If necessary, you can set the value to your free disk size. The default value is 40 GB.
   - **openstack.block-storage.volume.iscsi_ip_address**: Change from *127.0.0.1* to the management IP address of the volume service node.
   - **(Recommended network configuration)** If the management network interface of the nodes is not *eth0*, if the virtual machine data network interface is not *eth1*, or both apply, then update all occurrences of *eth0*, *eth1*, or both in the environment file to match your network configuration. The following list displays some of the networking properties and their default values (from the example environment) that you might need to change. In most cases, these default values should be sufficient and do not need to be changed.
     - **openstack.network.core_plugin**: *neutron.plugins.ml2.plugin.Ml2Plugin*.
       The **core_plugin** property controls which OpenStack plug-in is used by the Neutron server. The two recommended choices for this property include: *"neutron.plugins.openvswitch.ovs_neutron_plugin.OVSNeutronPluginV2"* and *"neutron.plugins.ml2.plugin.Ml2Plugin"*. If you are migrating from an earlier version of the IBM SmartCloud Entry product, it is recommended that you use that OVS plug-in.

       In the example environment, the **openstack.network.core_plugin** property is set to *"neutron.plugins.ml2.plugin.ML2Plugin"* and the **openstack.network.ml2.mechanism_drivers** property is set to *"openvswitch"*. When the **core_plugin** property is set to ML2, then the ML2 specific properties (properties that begin with **openstack.network.ml2.***) must be consistent and

kept in sync with the properties of the mechanism driver (**openstack.network.openvswitch.***). The example environment is set up with these two sets of properties in sync.

If you change the **openstack.network.core_plugin** property from *"neutron.plugins.ml2.plugin.Ml2Plugin"* to *"neutron.plugins.openvswitch.ovs_neutron_plugin.OVSNeutronPluginV2"*, you MUST remove the **openstack.network.service_plugins** property or set the value of the property to *[ ]*. In the example environment, the **openstack.network.service_plugins** property defaults to *"neutron.services.l3_router.l3_router_plugin.L3RouterPlugin"*. If the **core_plugin** is set to OVS, the ML2 specific properties in the environment are ignored.

- **openstack.network.openvswitch.bridge_mapping_interface**: *"br-eth1:eth1"*. The **bridge_mapping_interface** property is used to control the creation of the data network OVS bridge on the nodes. If Open vSwitch is installed and the data network bridge is already configured on the nodes, this property is not necessary and you can set the variable to *" "*. If a specific network configuration is needed for the data network (for example, bonding), you must set this variable to "" and complete set up manually before or after the node is converged.
- **openstack.network.openvswitch.bridge_mappings**: *"default:br-eth1"*.
  The **bridge_mappings** property controls which OVS bridge is used for flat and VLAN network traffic from the node. If this OVS bridge does not exist, the Open vSwitch agent does not start. This bridge can be automatically created by setting the **bridge_mapping_interface** property.
- **openstack.network.openvswitch.network_vlan_ranges** and **openstack.network.ml2.network_vlan_ranges**: These two properties define the default vlan range used when creating a tenant network. Both of these properties default to *default:1:4094*. This vlan range might need to be adjusted based on the vlan configuration of the physical switches and hypervisors in your environment. If the **core_plugin** value is set to *ML2*, the values of both properties must be the same.

5. Copy the following example topology to a file, `your-topology-name.json`. Change **your-topology-name** to the name for your topology. Here is an example topology with PowerKVM compute nodes.

```
{
  "name":"CHANGEME",
  "description":"CHANGEME",
  "environment":"CHANGEME",
  "secret_file":"CHANGEME",
  "run_sequentially":false,
  "nodes": [
    {
      "fqdn":"CHANGEME",
      "password":"CHANGEME",
      "identity_file":"CHANGEME",
      "quit_on_error":true,
      "run_order_number":1,
      "runlist": [
        "role[ibm-os-single-controller-node]"
      ]
    },
    {
      "fqdn":"CHANGEME",
      "password":"CHANGEME",
      "identity_file":"CHANGEME",
      "quit_on_error":true,
      "run_order_number":2,
      "runlist": [
        "role[ibm-os-compute-node-powerkvm]"
      ]
    }
  ]
}
```

6. Customize the topology file.

a. The first node in your topology file is your single controller node. The second node in your topology file is for a compute node. If your topology requires extra compute nodes, copy the compute node section as many times as needed. Ensure that additional compute node sections are comma-separated.

b. Change the following JSON attributes in your topology file, `your-topology-name.json`:

- **`Name`**: Set to your topology name: *your-topology-name*.
- **`Description`**: Set to the description for your topology.
- **`Environment`**: Set to the environment for your topology: *your-environment-name*.
- **`nodes.fqdn`**: For each node, you must set to the fully qualified domain name of the node system. The deployment system must be able to ssh using the fully qualified domain name. You can also set to the public IP address, private IP address, or host name.
- **`nodes.password`** or **`nodes.identity_file`**: For each node, set to the appropriate SSH root user authentication for the node system. Either a password and an SSH identity file can be used for authentication. Remove the unused attribute for each node.

c. (Optional) Create node specific attribute files. This step is only required when one or more nodes in your topology require different attributes from those defined in your environment file `your-environment-name.json`.

1) Create node specific attribute files. This step is only required when one or more nodes in your topology require different attributes from those defined in your environment file `your-environment-name.json`. For example, a node may not have an *eth0* network interface, which is the default value for some attributes. Below is an example node attribute file that can be used to change the *eth0* default network.

```
{
  "openstack": {
    "endpoints": {
      "network-openvswitch": {
        "bind_interface": "CHANGEME"  (Run the ifconfig command to identify
         the name of the interface, for example, enP3p9s0f0.)
      },
      "compute-vnc-bind": {
        "bind_interface": "CHANGEME"  (Run the ifconfig command to identify
         the name of the interface, for example, enP3p9s0f0.)
      }
    }
}
```

After creating the node specific attribute files, add the **`nodes.attribute_file`** JSON attributes in your topology file, `your-topology-name.json`:

- **`nodes.attribute_file`**: For each node, set to the attribute JSON file which overrides the attributes in the **`default_attributes`** section of the environment file.

7. Customize the passwords and secrets before deploying. For instructions, see "Customizing passwords and secrets" on page 102.
8. Configure the OpenStack block storage (cinder) driver. By default, the environment is configured to use the LVM iSCSI cinder driver. You can change the following JSON attributes in your environment file, `your-environment-name.json`, to customize the LVM iSCSI cinder driver configuration.

a. **`openstack.block-storage.volume.create_volume_group`**: If set to *true* (default) then the cinder-volumes volume group is created on the controller node with a size determined by **`openstack.block-storage.volume.volume_group_size`**. Set to *false* if you want to create the volume group manually.

b. **`openstack.block-storage.volume.volume_group_size`**: The amount of storage you use must be smaller than the size available. If necessary, you can set the value to your free disk size. The default value is 40 GB.

c. **`openstack.block-storage.volume.iscsi_ip_address`**: Change from *127.0.0.1* to the management IP address of the volume service node.

To customize your environment for a different cinder driver, see "Configuring Cinder drivers" on page 107.

9. (Optional) Complete any optional customizations.
   - "Customizing for a distributed database topology" on page 100
   - "Customizing for self-service portal" on page 100
   - "Customizing Platform Resource Scheduler" on page 102
   - "Customizing basic options" on page 104
   - "Customizing advanced options" on page 105
10. Deploy your topology.
    a. Upload the environment for your deployment.
       ```
       $ knife environment from file your-environment-name.json
       ```
    b. Deploy the topology.
       ```
       $ knife os manage deploy topology your-topology-name.json
       ```
11. After the deployment is complete, the IBM Cloud Manager with OpenStack services are ready to use. Log in to the controller node system. The administrator user name and password can be found in the `openrc` file under the root user home directory. The IBM Cloud Manager with OpenStack dashboard is available at `https://controller.fqdn.com/`, where **`controller.fqdn.com`** is the fully qualified domain name of the controller node in your topology.
    a. If you also include the **`ibm-sce-node`** role in your topology, you can check the IBM Cloud Manager with OpenStack self-service portal status on your controller node by using **`service sce status`**. The web interface for IBM Cloud Manager with OpenStack self-service portal is available at `http://contoller.fqdn.com:18080/cloud/web/login.html`. You can log in using admin user with *openstack1* as the default password (unless customized). After you are logged in, you can add the IaaS gateway and an OpenStack cloud. For more information, see "Configuring an OpenStack cloud" on page 198.
    b. For more information about managing IBM Cloud Manager with OpenStack services, see "Managing IBM Cloud Manager with OpenStack services" on page 169.

What to do next

You are ready to start using IBM Cloud Manager with OpenStack services.

Related reference:

"Troubleshooting errors when deploying topologies" on page 253
If a topology deployment fails, review the log output from the deployment command for more information. Identify errors, starting at the beginning of the log output. The first error found is typically the most useful. Errors found at the end of the log output might be the result of earlier errors. The following topics might provide additional information to solve the problem.

Platform Resource Scheduler online product documentation

Deploying with z/VM compute nodes

Deploy the components that are necessary to create a cloud environment with z/VM compute nodes.

Before you begin

- Ensure you completed the "Deploying prerequisites" on page 76 steps.
- In order to use the z/VM hypervisor, one or more x86_64 Red Hat Enterprise Linux system nodes should be used to install the compute and network driver to manage the z/VM hypervisor. One x86_64 Red Hat Enterprise Linux system node is supported for each z/VM node. For more information on configuring the z/VM hypervisor, see the z/VM OpenStack user manual.

About this task

Use the following procedure to deploy the topology to your node systems.

Procedure

1. Log in to the deployment system as the root user. This is the system where IBM Cloud Manager with OpenStack was installed.
2. Create a directory to store the files for the topology that you deploy. Change **your-deployment-name** to the name for your deployment.

   ```
   $ mkdir your-deployment-name
   $ chmod 600 your-deployment-name
   $ cd your-deployment-name
   ```

3. Copy the example environment for the topology that you deploy. Change **your-environment-name** to the name for your environment.

   Note: The following command must be entered on a single line, even though the example shows a line break for formatting purposes.

   ```
   $ knife environment show example-ibm-os-single-controller-n-compute -d -Fjson
   > your-environment-name.json
   ```

4. Change the following JSON attributes in your environment file, `your-environment-name.json`:

   a.
   - **Name**: Set to your environment name: *your-environment-name*.
   - **Description**: Set to the description for your environment.
   - **openstack.endpoints.host**, **openstack.endpoints.bind-host**, and **openstack.endpoints.mq.host**: Change from *127.0.0.1* to the IP address of the controller node system for the topology.
   - **openstack.block-storage.volume.create_volume_group**: Set to *false*.

   b.
   - **(Recommended network configuration)** If the management network interface of the nodes is not *eth0*, if the virtual machine data network interface is not *eth1*, or both apply, then update all occurrences of *eth0*, *eth1*, or both in the environment file to match your network configuration.
   - Update the following attributes for the controller node in your environment file.
     - **openstack.network.ml2.mechanism_drivers**: Add *zvm*, using a comma-separated list.
     - **openstack.network.ml2.tenant_network_types**: Set to *vlan* or *flat*, depending on your environment.
     - **openstack.network.ml2.flat_networks**: Add your z/VM flat networks. For example, *xcatvswflat*.
     - **openstack.network.ml2.network_vlan_ranges**: Add your z/VM network VLAN range. For example, *xcatvswvlan:10:100*.

   c. z/VM compute nodes configurations: You must update the following attributes for the z/VM compute nodes. If an attribute is not the same across all compute nodes then you will be able to customize it in step 6.
   - **ibm-openstack.zvm-driver.xcat.server**: Specify the xCAT MN IP address or host name.
   - **ibm-openstack.zvm-driver.xcat.username**: Specify the xCAT REST API user name.
   - **ibm-openstack.zvm-driver.xcat.zhcp_nodename**: Specify the zHCP node name in xCAT.
   - **ibm-openstack.zvm-driver.xcat.master**: Specify the xCAT master node (the node name in the xCAT definition).
   - **ibm-openstack.zvm-driver.xcat.mgt_ip**: Specify the first IP address of the management network.

Note: Remember the xCAT management interface IP address. xCAT uses this IP address to connect a newly-deployed instance server.

- **ibm-openstack.zvm-driver.xcat.mgt_mask**: Specify the network mask of the xCAT management network. For example: 255.255.255.0.
- **ibm-openstack.zvm-driver.diskpool**: Specify the disk pool name from where xCAT allocates disks for new servers. The disk pool name is the name of the storage 'group' defined in the Directory Manager.
- **ibm-openstack.zvm-driver.diskpool_type**: Specify the disk pool type, either FBA or ECKD™.
- **ibm-openstack.zvm-driver.zvm_host**: Specify the xCAT node name of the z/VM hypervisor.
- **ibm-openstack.zvm-driver.host**: Specify the host that is used to distinguish different nova compute hosts. It can be the same as **zvm_host**.
- **ibm-openstack.zvm-driver.user_profile**: Specify the default template of the user directory for new servers. Do not use *lnxdflt* but define your own profile.
- **ibm-openstack.zvm-driver.config_drive.inject_password**: Define whether to place the password in the config drive. If **inject_password** is set to *False*, the default os root password of the new booted virtual machine is the password in data bag **user_passwords.zlinuxroot**. If **inject_password** is set to *True*, the default os root password can be set using nova user-data. If you do not specify the password in nova user-data, the default os root password is a random value of *adminPass* property that is shown in the output of the virtual machine boot console.
- **ibm-openstack.zvm-driver.scsi_pool**: Specify the name of the xCAT SCSI pool. You can specify any name. xCAT creates and manages it.
- **ibm-openstack.zvm-driver.fcp_list**: Specify the list of FCPs used by instances. Each instance needs one FCP in order to attach a volume to itself. Those FCPs should be well planned and made online before OpenStack can use them. OpenStack does not check their status, so if they are not ready, you might receive errors. The format of this variable should look like *'min1-max1;min2-max2;min3-max3'*. Contact your z/VM system manager if you do not know what FCPs you can use.
- **ibm-openstack.zvm-driver.zhcp_fcp_list**: Specify the list of FCPs used only by xCAT HCP node. It must be different from *zvm_fcp_list* or you receive errors. The format of this variable should look like *'min1-max1;min2-max2;min3-max3'*. Only specify one FCP for HCP to avoid wasting resources. Contact your z/VM system manager if you do not know what FCPs you can use.
- **ibm-openstack.zvm-driver.external_vswitch_mappings**: Set the OSA configuration for each of the virtual switches. These configurations are required if the virtual switch connects outside of z/VM. The format of this variable is *'xcatvsw2:6243,6245;xcatvsw3:6343'*. Where *xcatvsw2* and *xcatvsw3* are the virtual switches and *6243, 6245, 6343* are RDEV addresses of the OSA cards that are connected to the virtual switch.

5. Copy the following example topology to a file, `your-topology-name.json`. Change **your-topology-name** to the name for your topology. Here is an example topology with z/VM compute nodes.

```
{
  "name":"CHANGEME",
  "description":"CHANGEME",
  "environment":"CHANGEME",
  "secret_file":"CHANGEME",
  "run_sequentially":false,
  "nodes": [
    {
      "fqdn":"CHANGEME",
      "password":"CHANGEME",
      "identity_file":"CHANGEME",
      "quit_on_error":true,
      "run_order_number":1,
      "runlist": [
        "role[ibm-os-single-controller-node]"
```

```
      ]
    },
    {
      "fqdn":"CHANGEME",
      "password":"CHANGEME",
      "identity_file":"CHANGEME",
      "quit_on_error":true,
      "run_order_number":2,
      "runlist": [
        "role[ibm-os-zvm-driver-node]"
      ]
    }
  ]
}
```

6. Customize the topology file.
   a. The first node in your topology file is your single controller node. The second node in your topology file is for a compute node. If your topology requires extra compute nodes, copy the compute node section as many times as needed. Ensure that additional compute node sections are comma-separated.
   b. Change the following JSON attributes in your topology file, `your-topology-name.json`:
      - **`Name`**: Set to your topology name: *your-topology-name*.
      - **`Description`**: Set to the description for your topology.
      - **`Environment`**: Set to the environment for your topology: *your-environment-name*.
      - **`nodes.fqdn`**: For each node, you must set to the fully qualified domain name of the node system. The deployment system must be able to ssh using the fully qualified domain name. You can also set to the public IP address, private IP address, or host name.
      - **`nodes.password`** or **`nodes.identity_file`**: For each node, set to the appropriate SSH root user authentication for the node system. Either a password and an SSH identity file can be used for authentication. Remove the unused attribute for each node.
   c. (Optional) Create node specific attribute files. This step is only required when one or more nodes in your topology require different attributes from those defined in your environment file `your-environment-name.json`. For example, one or more z/VM compute nodes may have different configurations than those set in step 4. Example node specific attributes for a z/VM compute node.

```
{
  "ibm-openstack": {
    "zvm-driver": {
      "xcat": {
        "username": "CHANGEME",
        "server": "CHANGEME",
        "zhcp_nodename": "CHANGEME",
        "master": "CHANGEME",
        "mgt_ip": "CHANGEME",
        "mgt_mask": "CHANGEME"
      },
      "diskpool": "CHANGEME",
      "diskpool_type": "CHANGEME",
      "zvm_host": "CHANGEME",
      "host": "CHANGEME",
      "user_profile": "CHANGEME",
      "scsi_pool": "CHANGEME",
      "fcp_list": "CHANGEME",
      "zhcp_fcp_list": "CHANGEME",
      "external_vswitch_mappings": "",
      "config_drive": {
        "inject_password": "CHANGEME"
      }
    }
  }
}
```

After creating the node specific attribute files, add the **nodes.attribute_file** JSON attributes in your topology file, `your-topology-name.json`:

- `nodes.attribute_file`: For each node, set to the attribute JSON file which overrides the attributes in the **default_attributes** section of the environment file.

7. Customize the passwords and secrets before deploying. For instructions, see "Customizing passwords and secrets" on page 102.
8. (Optional) Complete any optional customizations.
   - "Customizing for a distributed database topology" on page 100
   - "Customizing for self-service portal" on page 100
   - "Customizing basic options" on page 104
   - "Customizing advanced options" on page 105
9. Deploy your topology.
   a. Upload the environment for your deployment.
      ```
      $ knife environment from file your-environment-name.json
      ```
   b. Deploy the topology.
      ```
      $ knife os manage deploy topology your-topology-name.json
      ```
10. After the deployment is complete, the IBM Cloud Manager with OpenStack services are ready to use. Log in to the controller node system. The administrator user name and password can be found in the `openrc` file under the root user home directory. The IBM Cloud Manager with OpenStack dashboard is available at `https://controller.fqdn.com/`, where **controller.fqdn.com** is the fully qualified domain name of the controller node in your topology.
    a. If you also include the **ibm-sce-node** role in your topology, you can check the IBM Cloud Manager with OpenStack self-service portal status on your controller node by using **service sce status**. The web interface for IBM Cloud Manager with OpenStack self-service portal is available at `http://contoller.fqdn.com:18080/cloud/web/login.html`. You can log in using admin user with *openstack1* as the default password (unless customized). After you are logged in, you can add the IaaS gateway and an OpenStack cloud. For more information, see "Configuring an OpenStack cloud" on page 198.
    b. For more information about managing IBM Cloud Manager with OpenStack services, see "Managing IBM Cloud Manager with OpenStack services" on page 169.
11. Now that the z/VM compute nodes are deployed successfully, you must exchange ssh-key between xCAT and Nova. For more information, see Configuration for SSH Between xCAT and Nova in the z/VM OpenStack user manual.

What to do next

You are ready to start using IBM Cloud Manager with OpenStack services.

Related reference:

"Troubleshooting errors when deploying topologies" on page 253
If a topology deployment fails, review the log output from the deployment command for more information. Identify errors, starting at the beginning of the log output. The first error found is typically the most useful. Errors found at the end of the log output might be the result of earlier errors. The following topics might provide additional information to solve the problem.

Configure the Image Service

"Creating initial networks" on page 117
After you deploy the components for creating a cloud environment, you can create several different types of networks.

Deploying to manage to PowerVC

Deploy the components that are necessary to create a cloud environment to manage to PowerVC.

Before you begin, ensure you completed the "Deploying prerequisites" on page 76 steps.

Use the following procedure to deploy the topology to your node systems.

1. Log in to the deployment system as the root user. This is the system where IBM Cloud Manager with OpenStack was installed.
2. Create a directory to store the files for the topology that you deploy. Change **your-deployment-name** to the name for your deployment.

   ```
   $ mkdir your-deployment-name
   $ chmod 600 your-deployment-name
   $ cd your-deployment-name
   ```

3. Copy the example environment for the topology that you deploy. Change **your-environment-name** to the name for your environment.

   Note: The following command must be entered on a single line, even though the example shows a line break for formatting purposes.

   ```
   $ knife environment show example-ibm-os-single-controller-n-compute -d -Fjson
   > your-environment-name.json
   ```

4. Change the following JSON attributes in your environment file, `your-environment-name.json`:
   - **Name**: Set to your environment name: *your-environment-name*.
   - **Description**: Set to the description for your environment.
   - **openstack.endpoints.host**, **openstack.endpoints.bind-host**, and **openstack.endpoints.mq.host**: Change from *127.0.0.1* to the IP address of the controller node system for the topology.
   - **ibm-openstack.powervc-driver.powervc.admin_user**: Specify the administrator for PowerVC.
   - **ibm-openstack.powervc-driver.powervc.auth_url**: Specify the authentication url for PowerVC.
   - **ibm-openstack.powervc-driver.powervc.qpid.host**: Specify the host name for the PowerVC.
   - **ibm-openstack.powervc-driver.db_create**: Set to *true* for PowerVC driver.
   - **openstack.block-storage.volume.create_volume_group**: Set to *false*.

   Note: Depending on your configuration, change the following values:
   - If you configured Neutron to use ML2 as the core plug-in (default configuration), set the following values:
     - **openstack.network.ml2.tenant_network_types** list *"vlan"* in the value list.
     - **openstack.network.ml2.network_vlan_ranges** list *"default:1:4094"* in the value list.
   - If you configured Neutron to use Open vSwitch as the core plug-in, set the following values:
     - **openstack.network.openvswitch.tenant_network_type**: *"vlan"*
     - **openstack.network.openvswitch.network_vlan_ranges** list *"default:1:4094"* in the value list.
   - If you configured Neutron to use Linux bridge as the core plug-in, set the following values:
     - **openstack.network.linuxbridge.tenant_network_type**: *"vlan"*
     - **openstack.network.linuxbridge.network_vlan_ranges** list *"default:1:4094"* in the value list.
   - In a PowerVC only environment, where no other compute node is registered to a controller node that is running the PowerVC driver, the default configuration can be simplified. You do not need to set up a dedicated virtual machine data network. In the sample environment file, the GRE network type is enabled by the default settings. However, since PowerVC does not support the GRE network type, set the following attributes to disable GRE and specify a valid network type.
     - **openstack.network.openvswitch.bridge_mappings** = *""*
     - **openstack.network.openvswitch.bridge_mapping_interface** = *""*
     - **openstack.network.openvswitch.enable_tunneling** =*"False"*

- **openstack.network.openvswitch.tunnel_type** = *""*
- **openstack.network.openvswitch.tunnel_id_ranges** = *""*

In a PowerVC only environment, you must also add an attribute to automatically confirm when an instance is resized.

- **openstack.compute.misc_nova**: *['api_rate_limit=false', 'flat_injected=true', 'use_ipv6=true', "resize_confirm_window=1"]* You must add *resize_confirm_window=1* to the list of values for the attribute. This setting applies to the PowerVC driver environment only.

  Note: If you are using a mixed hypervisor environment, then you must manually set this option in the controller node of the PowerVC driver (`/etc/nova/nova.conf`).

5. Copy the following example topology to a file, `your-topology-name.json`. Change **your-topology-name** to the name for your topology. Here is an example topology to manage to PowerVC.

```
{
  "name":"CHANGEME",
  "description":"CHANGEME",
  "environment":"CHANGEME",
  "secret_file":"CHANGEME",
  "run_sequentially":false,
  "nodes": [
    {
      "fqdn":"CHANGEME",
      "password":"CHANGEME",
      "identity_file":"CHANGEME",
      "quit_on_error":true,
      "run_order_number":1,
      "runlist": [
        "role[ibm-os-single-controller-node]",
        "role[ibm-os-powervc-driver-node]"
      ]
    }
  ]
}
```

6. Customize the topology file. The node in your topology file is your single controller node that includes the PowerVC driver. Change the following JSON attributes in your topology file, `your-topology-name.json`:
   - **Name**: Set to your topology name: *your-topology-name*.
   - **Description**: Set to the description for your topology.
   - **Environment**: Set to the environment for your topology: *your-environment-name*.
   - **nodes.fqdn**: For each node, you must set to the fully qualified domain name of the node system. The deployment system must be able to ssh using the fully qualified domain name. You can also set to the public IP address, private IP address, or host name.
   - **nodes.password** or **nodes.identity_file**: Set to the appropriate SSH root user authentication for the node system. Either a password and an SSH identity file can be used for authentication. Remove the unused attribute.
7. You must change the passwords in the **user_passwords** data bag to the actual password for the PowerVC administrator and PowerVC QPID *qpidclient* user. To change the passwords in *pvcadmin* and *pvcqpid* data bags, see Customizing passwords and secrets.
8. (Optional) Complete any optional customizations.
   - "Customizing for a distributed database topology" on page 100
   - "Customizing for self-service portal" on page 100
   - "Customizing basic options" on page 104
   - "Customizing advanced options" on page 105
9. Deploy your topology.

a. Upload the environment for your deployment.

```
$ knife environment from file your-environment-name.json
```

b. Deploy the topology.

```
$ knife os manage deploy topology your-topology-name.json
```

10. After the deployment is complete, the IBM Cloud Manager with OpenStack services are ready to use. Log in to the controller node system. The administrator user name and password can be found in the `openrc` file under the root user home directory. The IBM Cloud Manager with OpenStack dashboard is available at `https://controller.fqdn.com/`, where **`controller.fqdn.com`** is the fully qualified domain name of the controller node in your topology.
    a. If you also include the **`ibm-sce-node`** role in your topology, you can check the IBM Cloud Manager with OpenStack self-service portal status on your controller node by using **`service sce status`**. The web interface for IBM Cloud Manager with OpenStack self-service portal is available at `http://contoller.fqdn.com:18080/cloud/web/login.html`. You can log in using admin user with *openstack1* as the default password (unless customized). After you are logged in, you can add the IaaS gateway and an OpenStack cloud. For more information, see "Configuring an OpenStack cloud" on page 198.
    b. For more information about managing IBM Cloud Manager with OpenStack services, see "Managing IBM Cloud Manager with OpenStack services" on page 169.

You are ready to start using IBM Cloud Manager with OpenStack services.

Deploying multi-region support

In a multi-region configuration, you can set up two different controllers that use the same keystone server, but use different regions and different hypervisors.

About this task

Use the following instructions to build your own multi-region environment.

Note: This example uses two controller nodes as an example for multiple regions; however, you could also have several controller nodes added into the configuration for multiple regions.

Procedure

1. Create two environment files, which are copied from `example-ibm-os-single-controller-n-compute`. One file is used for region one and the second file is used for region two.
2. Create two topology files. One file is used for region one and the second file is used for region two.
3. Update the environment files for each region. Set the following attributes. The two unique updates are the region attributes and identity endpoint attributes. Any remaining updates are normal updates for a stand-alone deployment.
    a. For a specific region, update the **`openstack.region`** attribute for the applicable region.
    b. For all regions in multiple region configurations, update the **`openstack.endpoint.identity-api.host`** and **`openstack.endpoints.identity-admin.host`** attributes to specify the Keystone node that is used for the multiple region configuration. The node can be any controller node, within the multiple region configuration, which is used as the Keystone node. The following example json snippet is added to the environment file inside **`override_attributes`** > **`openstack`** > **`endpoints`**:

```
"identity-api": {
   "host": "X.X.X.X"
 },
 "identity-admin": {
   "host": "X.X.X.X"
 }
```

    Where X.X.X.X is the management IP for the region one controller where the shared Keystone server resides.

c. For all regions that are part of a multiple region configuration, update the following environmental attributes to reference the management IP for the controller in its region: **openstack.endpoint.host**, **openstack.endpoint.bind-host.host**, **openstack.endpoint.mq.host**, and **openstack.endpoint.db.host** The following example json snippet is added to the environment file inside **override_attributes** > **openstack** > **endpoints**:

```
"host": "Y.Y.Y.Y",
"bind-host": "Y.Y.Y.Y",
"mq": {
  "host": "Y.Y.Y.Y",
  "port": "5671"
},
"db": {
  "host": "Y.Y.Y.Y",
  "port": "50000"
},
```

Where Y.Y.Y.Y is the management IP of the controller in its region.

4. Update the topology for each region. They are similar but have different environments and can have different secrets and data bags.

   If the passwords are different between regions, then you must create the **user_passwords** data bag for each region and specify the data bag name in the environment file. For more information about customizing secrets and data bags, see "Customizing passwords and secrets" on page 102.

   Note:
   - The data bags for OpenStack regular and service users must have the same passwords in all regions since there is a shared keystone.
   - The data bag for **openstack_identity_bootstrap_token** must also be the same between the two regions.
5. Deploy the region one topology: **knife os manage deploy topology topology_region_one.json**. Deploy topology one first, as it has the shared keystone information.
6. Deploy the region two topology: **knife os manage deploy topology topology_region_two.json**.
7. (Self-service portal): If you have the self-service portal installed, restart the IaaS gateway server on the Keystone node using multi-region support. Run the **service openstack-iaasgateway restart** command to restart the services.

   Note: Only one self-service portal needs to be installed to manage the multiple regions.

Deploying multiple hypervisor types

Use this information to deploy the components that are necessary to create a cloud environment with multiple hypervisor types.

About this task

It is possible to configure different types of hypervisors in the same OpenStack environment.

The following example shows how to build a dual hypervisor type environment using IBM Cloud Manager with OpenStack. The example uses the following system setup:

- System 1: This is the deployment system where IBM Cloud Manager with OpenStack is installed.
- System 2: The OpenStack single controller node.
- System 3: The OpenStack x86 Linux Kernel-based Virtual Machine (KVM) compute node.
- System 4: The OpenStack PowerKVM compute node.

Important: This procedure is for example purposes only.

Procedure

1. Install IBM Cloud Manager with OpenStack on system 1. For instructions, see "Installing IBM Cloud Manager with OpenStack on Linux" on page 19.
2. Follow the deployment prerequisites for all systems. See "Deploying prerequisites" on page 76 for details.
3. From system 1, complete the following steps.
   a. Create a directory to store the files for the topology that you deploy. Change **your-deployment-name** to the name for your deployment.

      ```
      $ mkdir your-deployment-name
      $ chmod 600 your-deployment-name
      $ cd your-deployment-name
      ```

   b. Copy the example environment for the topology that you deploy. Change **your-environment-name** to the name for your environment.

      Note: The following command must be entered on a single line, even though the example shows a line break for formatting purposes.

      ```
      $ knife environment show example-ibm-os-single-controller-n-compute -d -Fjson
      > your-environment-name.json
      ```

   c. Create a topology file that is based on the following example topology file, `your-topology-name.json`.

      ```
      {
        "name":"CHANGEME",
        "description":"CHANGEME",
        "environment":"CHANGEME",
        "secret_file":"CHANGEME",
        "run_sequentially":false,
        "nodes": [
          {
            "fqdn":"CHANGEME_TO_SYSTEM2",
            "password":"CHANGEME",
            "quit_on_error":true,
            "run_order_number":1,
            "runlist": [
              "role[ibm-os-single-controller-node]"
            ]
          },
          {
            "fqdn":"CHANGEME_TO_SYSTEM3",
            "password":"CHANGEME",
            "quit_on_error":true,
            "run_order_number":2,
            "runlist": [
              "role[ibm-os-compute-node-kvm]"
            ]
          },
          {
            "fqdn":"CHANGEME_TO_SYSTEM4",
            "password":"CHANGEME",
            "quit_on_error":true,
            "run_order_number":2,
            "runlist": [
              "role[ibm-os-compute-node-powerkvm]"
            ]
          }
        ]
      }
      ```

4. From system 1, modify the environment and topology files for the deployment. See "Deploying a controller +*n* compute or distributed database topology" on page 80 for details.
5. From system 1, deploy your topology.
   a. Upload the environment for your deployment.

```
$ knife environment from file your-environment-name.json
```

b. Deploy the topology.

```
knife os manage deploy topology your-topology-name.json
```

6. After the deployment completes, the IBM Cloud Manager with OpenStack services are ready to use. Log in to the controller node system. The administrator user name and password can be found in the `openrc` file under the root user home directory. The IBM Cloud Manager with OpenStack dashboard is available at `https://controller.fqdn.com/`, where **`controller.fqdn.com`** is the fully qualified domain name of the controller node in your topology.

Related reference:

"Troubleshooting errors when deploying topologies" on page 253
If a topology deployment fails, review the log output from the deployment command for more information. Identify errors, starting at the beginning of the log output. The first error found is typically the most useful. Errors found at the end of the log output might be the result of earlier errors. The following topics might provide additional information to solve the problem.

"Roles" on page 237
The following roles are provided in support of the reference topologies.

Deployment customization options

Configure these advanced customization options before deploying IBM Cloud Manager with OpenStack.

Note: The minimal topology does not require any customization to deploy.

Customizing for a distributed database topology

To customize your deployment for a distributed database topology then you must update your environment and topology files.

Procedure

1. Change the following JSON attribute in your environment file, `your-environment-name.json`:
   - **`openstack.endpoints.db.host`**: Change from *127.0.0.1* to the IP address of the database node system for the topology.
2. Add the following database node system as the first node in your topology file, `your-topology-name.json`.

```
{
  "fqdn":"CHANGEME",
  "password":"CHANGEME",
  "identity_file":"CHANGEME",
  "quit_on_error":true,
  "run_order_number":0,
  "runlist": [
    "role[ibm-os-database-server-node]"
  ]
},
```

3. Modify the role for the controller node system in your topology file, `your-topology-name.json`. Change the **`"role[ibm-os-single-controller-node]"`** role to **`"role[ibm-os-single-controller-distributed-database-node]"`**.

Customizing for self-service portal

If you want to use the self-service portal to manage IBM Cloud Manager with OpenStack, you must manually include the self-service portal in a topology.

About this task

The IBM Cloud Manager with OpenStack self-service portal is not included in any topology. To deploy the self-service portal, you can manually include the self-service portal in a topology by updating the environment and topology file. The self-service portal can be included in all topologies, including the

Minimal topology and the **Controller +*n* compute** topology. You can install the self-service portal using the `ibm-sce-node` role when you are deploying to Red Hat Enterprise Linux (x86 and Power) nodes.

Note:

- If you want to install the self-service portal to Windows nodes, see "Installing IBM Cloud Manager with OpenStack self-service portal on Windows" on page 25 for installation instructions.
- If you deploy the self-service portal to a Linux node, it is recommended to deploy to the OpenStack controller node.

To include the self-service portal in the topology, complete the following steps.

Procedure

1. Update the following JSON attributes in your environment file, `your-environment-name.json`.
   - **`ibm-sce.service.enabled`**: *true*. This attribute results in the following actions to allow self-service portal users access to the OpenStack cloud:
     – Replaces the default OpenStack policy files with the updated self-service portal policy files.
     – Registers the *sceagent* user with the *admin* role in OpenStack.
2. Add the **`ibm-sce-node`** role in your topology file, `your-topology-name.json`.
   - Append **`"role[ibm-sce-node]"`** in the **`run_list`** for the node. For example, in a **Minimal** topology file, you would append the *ibm-sce-node* role, as shown.

```
"runlist":[
  "role[ibm-os-allinone-kvm]",
  "role[ibm-sce-node]"
]
```

3. To apply a fix pack or JRE update for the IBM Cloud Manager with OpenStack self-service portal, use the following steps to update your environments.
   a. Enter the self-service portal installer repository, `<CMWO ROOT>/file-repo/SCE`
   b. Find the fix pack file named `4.1.0.1-IBM-SCE-FP001-201405150700.zip`.
   c. Update the attribute value for **`ibm-sce.package.fixpack.file`** with its name in the environment file. For example:

```
"ibm-sce": {
         "package": {
           "fixpack": {
             "file": "4.1.0.1-IBM-SCE-FP001-201405150700.zip"
           }
         }
       }
```

   d. For a JRE update, find the applicable JRE file name, either `sce410_jre_linux_installer.bin` (Linux) or `sce410_jre_lop_installer.bin` (Power).
   e. Update the attribute value for **`ibm-sce.package.jre.update.file`** with the file name in your environment file. For example:

```
"ibm-sce": {
         "package": {
           "update": {
             "jre": {
               "file": "sce410_jre_linux_installer.bin"
             }
           }
         }
       }
```

Related reference:

"Managing clouds" on page 197
You can add, configure, edit, and remove clouds.

“Roles” on page 237
The following roles are provided in support of the reference topologies.

Customizing Platform Resource Scheduler

If you want to use IBM Platform Resource Scheduler (PRS) instead of the default OpenStack compute scheduler filters, then you need to complete more customization tasks.

Before you begin

If you plan to deploy a topology with PRS enabled, ensure that Java™ Runtime Environment (JRE), version 1.7 or later, is installed on controller node.

About this task

The PRS customization is only supported for the KVM, QEMU, and PowerKVM hypervisors. To use PRS as the resource scheduler, you must install the PRS components on the OpenStack controller and all compute nodes. For the Minimal topology, all PRS components are installed on the same node.

Procedure

1. Append the following PRS roles to the run list for the single controller and all compute nodes, as follows:

```
# PRS roles to append to the OpenStack single controller's run list

"role[ibm-os-prs-ego-master]",
"role[ibm-os-prs-controller-node]"

# PRS roles to append to an OpenStack compute node's run list

"role[ibm-os-prs-compute-node]"
```

2. If you want to use the optimization feature of PRS to run live migrations of virtual machines, ensure that you enable live migration for IBM Cloud Manager with OpenStack. For more information, see Configuring migration for a KVM node with NFS shared storage.

Related reference:

“Roles” on page 237
The following roles are provided in support of the reference topologies.

Platform Resource Scheduler online product documentation

Customizing passwords and secrets

You can customize the passwords and secrets that are used during the deployment process.

If you are deploying a **Minimal topology**, then all passwords and secrets are the same as their associated user name (for example, admin/admin). Customization is not enabled for the topology.

If you are deploying a **Controller +*n* compute or Distributed database** topology, then by default, all passwords, and secrets are obtained through encrypted data bags. An example secret key (**/opt/ibm/cmwo/chef-repo/data_bags/example_data_bag_secret**) and data bags (**/opt/ibm/cmwo/chef-repo/data_bags/**) are provided by the IBM Cloud Manager with OpenStack installation. The following steps guide you through customizing these examples for your deployment:

1. Create a secret key for encrypting the data bags. A secret key can be created in many ways. The following example uses the **OpenSSL rand** command to generate pseudo-random bytes that can be used as the secret key.

```
$ openssl rand -base64 512 | tr -d '\r\n' > your-secret-key-name
$ chmod 600 your-secret-key-name
```

2. Copy the example data bags. Change **your_env_db_passwords**, **your_env_secrets**, **your_env_service_passwords**, **your_env_user_passwords**, and **your_env_appliance32_passwords** to the

name of your environment data bags. If you are not migrating existing data from IBM SmartCloud Entry then the **appliance32_passwords** data bag is not required.

Note: Each of the following lines must be entered on a single line, even though the example shows a line break for formatting purposes.

```
$ mkdir data_bags
$ cp -r /opt/ibm/cmwo/chef-repo/data_bags/db_passwords data_bags/your_env_db_passwords
$ cp -r /opt/ibm/cmwo/chef-repo/data_bags/secrets data_bags/your_env_secrets
$ cp -r /opt/ibm/cmwo/chef-repo/data_bags/service_passwords data_bags/
  your_env_service_passwords
$ cp -r /opt/ibm/cmwo/chef-repo/data_bags/user_passwords data_bags/your_env_user_passwords
$ cp -r /opt/ibm/cmwo/chef-repo/data_bags/appliance32_passwords data_bags/
  your_env_appliance32_passwords
$ chmod -R 600 data_bags/
```

Note: These steps assume the default IBM Cloud Manager with OpenStack installation path on the deployment server (`/opt/ibm/cmwo`).

3. Update the passwords and secrets in your environment data bags. This is done by changing the value for the data bag item's name. For example, using the **your_env_user_passwords/admin.json** data bag item, change the value at *CHANGEME* to the password.

```
{
  "id": "admin",
  "admin": "CHANGEME"
}
```

Note:

- z/VM related data bag item: **xcat** and **zlinuxroot**

4. Upload the data bag items for your deployment. Run the following commands from the parent of the `data_bags` directory that you created in step 2. If you are not migrating existing data from IBM SmartCloud Entry 3.2 then the **appliance32_passwords** data bag is not required.

```
$ knife data bag create your_env_db_passwords
$ knife data bag create your_env_secrets
$ knife data bag create your_env_service_passwords
$ knife data bag create your_env_user_passwords
$ knife data bag create your_env_appliance32_passwords
$ knife data bag from file your_env_db_passwords --all --secret-file your-secret-key-name
$ knife data bag from file your_env_secrets --all --secret-file your-secret-key-name
$ knife data bag from file your_env_service_passwords --all --secret-file your-secret-key-name
$ knife data bag from file your_env_user_passwords --all --secret-file your-secret-key-name
$ knife data bag from file your_env_appliance32_passwords --all --secret-file your-secret-key-name
```

5. Change the **openstack.secret.*_data_bag** JSON attributes in your environment file, `your-environment-name.json`, to match the names of the data bags upload in the previous steps. For example, set **openstack.secret.user_passwords_data_bag** to **your_env_user_passwords**.
6. Change the **secret_file** JSON attribute in your topology file, `your-topology-name.json`, to point to the secret key used to encrypt the data bags.
7. Remove the local data bag items since they are no longer needed.

```
$ rm -rf data_bags/
```

Related reference:

"Data bag item not found" on page 253
You might see an error in the log that a data bag item is missing.

"Data bags" on page 229
The following example data bags are provided in support of the IBM Cloud Manager with OpenStack topologies.

Customizing basic options

You can update the JSON attributes in your environment file, `your-environment-name.json` to customize basic options for NTP, OpenStack and database.

Customizing NTP attributes:

You can change the NTP servers attribute in your environment file to modify the NTP servers used by the node systems throughout your topology deployment.

Procedure

Update the following JSON attributes in your environment file, `your-environment-name.json`.

- **`ntp.servers`**: Set to the NTP servers accessible to your deployment.

Customizing OpenStack attributes:

You can change the OpenStack attributes in your environment file to control logging, quotas and uploading images.

Procedure

Update the following JSON attributes in your environment file, `your-environment-name.json`.

- **`openstack.compute.config.quota_*`** and **`openstack.network.quota.*`**: The environment file provides default quota settings for the IBM Cloud Manager with OpenStack compute and network components. Set the quotas as needed for your deployment. For more information, see the OpenStack manage quotas documentation.
- **`openstack.image.upload_image`** and **`openstack.image.upload_images`**: These attributes allow an image to be uploaded to OpenStack as part of your deployment. **`openstack.image.upload_image`** is a set of key-value pairs where the key is the image name and the value is the URL location of the image. **`openstack.image.upload_images`** is a list of keys from **`openstack.image.upload_image`** that you want to upload. For example:

```
"openstack": {
  "image": {
    "upload_image": {
      "myimage": "https://myimagesource.com/myimage"
    },
    "upload_images": [
      "myimage"
    ]
  }
}
```

- **`openstack.identity.debug`** and **`openstack.identity.verbose`**: Set to *true* or *false* for the desired logging level for the OpenStack identity service.
- **`openstack.image.debug`** and **`openstack.image.verbose`**: Set to *true* or *false* for the desired logging level for the OpenStack image service.
- **`openstack.network.debug`** and **`openstack.network.verbose`**: Set to *true* or *false* for the desired logging level for the OpenStack network service.
- **`openstack.compute.debug`** and **`openstack.compute.verbose`**: Set to *true* or *false* for the desired logging level for the OpenStack compute service.
- **`openstack.block-storage.debug`** and **`openstack.block-storage.verbose`**: Set to *true* or *false* for the desired logging level for the OpenStack block storage service.
- **`openstack.telemetry.debug`** and **`openstack.telemetry.verbose`**: Set to *true* or *false* for the desired logging level for the OpenStack telemetry service.
- **`openstack.orchestration.debug`** and **`openstack.orchestration.verbose`**: Set to *true* or *false* for the desired logging level for the OpenStack orchestration service.

- **ibm-openstack.iaas-gateway.debug** and **ibm-openstack.iaas-gateway.verbose**: Set to *true* or *false* for the desired logging level for the IBM OpenStack IaaS gateway service.

Customizing database attributes:

You can change the database used by OpenStack from DB2 (default) to MySQL.

About this task

Update the following JSON attributes in your environment file, `your-environment-name.json`.

- **openstack.db.service_type**: Set to *mysql*.
- **openstack.db.telemetry.nosql.used**: Set to *false*.
- **mysql.allow_remote_root**: Set to *true*.
- **mysql.server_root_password**: Set to the root password.
- **openstack.endpoints.db.port**: Remove this attribute from the environment file.

MySQL requires the MySQL ruby gem. During deployment this gem, by default, is downloaded and installed from the internet location https://rubygems.org. If your node does not have access to the internet, you can configure IBM Cloud Manager with OpenStack to download and install the MySQL gem from the deployment server by completing the following steps:

Procedure

1. Add the attribute **ibm-openstack.remote_files.mysql_gem_download_enabled**: true to your environment file.
2. Acquire the mysql ruby gem file and add the file into the following path in the installation directory for your deployment server, installation directory/file-repo/gems. The default value for installation directory is `/opt/ibm/cmwo`.
3. By default, IBM Cloud Manager with OpenStack is configured to download and install the gem with a file name of `mysql-2.9.1.gem`. If the gem file name is different in your environment, you can use a different file name by setting the following attribute in your environment file: **ibm-openstack.remote_files.mysql_gem_filename**: gem file name. For example: **ibm-openstack.remote_files.mysql_gem_filename:mysql-2.9.1.gem**.

Customizing Qpid messaging service attributes:

You can change the Qpid messaging service binding interface in your environment file to make it bind to the desired IP address.

About this task

Update the following JSON attributes in your environment file, `your-environment-name.json`. By default, the Qpid messaging service binds to *0.0.0.0*.

- **qpid.broker.interface**: Set to the IP address that is applicable. In most cases, this is the IP address of the controller node system.

Customizing advanced options

You can update the JSON attributes in your environment file, `your-environment-name.json` to customize advanced options such as SELinux, iptables, and message queue performance.

Related concepts:

"Deploying a controller +*n* compute or distributed database topology" on page 80
Deploy the components that are necessary to create a cloud environment with more than one node. Use the instructions that apply to your specific hypervisor for each compute node.

Customizing SELinux attributes:

You can change the SELinux attribute in your environment file to modify the SELinux state on the node systems to which you deploy a topology.

Procedure

Update the following JSON attributes in your environment file, `your-environment-name.json`.

`selinux.state`
The default value for this attribute is *nothing*. With the default value, the SELinux state is unchanged for the node systems in your topology. If you need to change the SELinux state on each node system, the attribute can be set to *enforcing*, *permissive* or *disabled*. For more information about SELinux, see the SELinux Project Wiki.

Customizing iptables attributes:

You can change the iptables attributes in your environment file to control iptables rules customization, access to IBM Cloud Manager with OpenStack services, and SSH access on the nodes in your topology.

Procedure

You can update the following JSON attributes in your environment file, `your-environment-name.json`, to customize the **`ibm-openstack.iptables.*`** attributes.

`ibm-openstack.iptables.status`
The default value for this attribute is *enabled*. With the default value, the iptables rules can be customized for the node systems in your topology using the other **`ibm-openstack.iptables.*`** attributes. If you want to manually configure the iptables for each node system, you must set this attribute to *unchanged*.

`ibm-openstack.iptables.use_default_rules`
When this attribute is set to *true* (default value), iptables rules are created to allow access to the IBM Cloud Manager with OpenStack services that are running on each node system.

`ibm-openstack.iptables.include_ssh_default_rule`
When this attribute is set to *true* (default value), iptables rules are created to allow SSH access to the node systems.

`ibm-openstack.iptables.custom_rules`
You can use this attribute to create custom iptables rules.

Note: If you want to manually configure the iptables for each node system, you must set the **`ibm-openstack.iptables.status`** attribute to *unchanged*. The *unchanged* value will leave the iptables unchanged on each node system and allow your manual configuration to be maintained.

Customizing message queue attributes:

You can change the message queue performance attributes in your environment file to customize settings such as qpid brokers, compute settings, and network settings.

About this task

The message queue services are used extensively by the IBM Cloud Manager with OpenStack services to provide process communication and event notifications. The default values for the message queue-related customization values are sufficient for most topology deployments. However, you might need to customize the values for heavily used deployments.

Procedure

You can update the following JSON attributes in your environment file, `your-environment-name.json`, to customize the message queue performance attributes.

qpid.broker.max-connections, qpid.broker.connection-backlog, qpid.broker.worker-threads
The default value for **qpid.broker.max-connections** is *1000*. The default value for **qpid.broker.connection-backlog** is *10*. The default value for **qpid.broker.worker-threads** is *20*. For more information about customizing these values, see the Qpid broker man page.

openstack.compute.rpc_thread_pool_size, openstack.compute.rpc_conn_pool_size, openstack.compute.rpc_response_timeout
The default values for **openstack.compute.rpc_thread_pool_size** is *2048*. The default value for **openstack.compute.rpc_conn_pool_size** is *60*. The default value for **openstack.compute.rpc_response_timeout** is *960*. For more information about customizing these values, see the nova.conf configuration options.

openstack.network.rpc_thread_pool_size, openstack.network.rpc_conn_pool_size, openstack.network.rpc_response_timeout
The default value for **openstack.network.rpc_thread_pool_size** is *128*. The default value for **openstack.network.rpc_conn_pool_size** is *60*. The default value for **openstack.network.rpc_response_timeout** is *600*. For more information about customizing these values, see the neutron.conf configuration options.

openstack.compute.osapi_compute_workers, openstack.compute.metadata_workers, openstack.compute.conductor.workers
The default values are equal to the number of CPUs on the node system (with a maximum value of 8). Changes to the default values should be based on the number of CPUs that are available on the single controller node system. For more information about customizing these values, see the nova.conf configuration options.

openstack.sysctl.net.core.somaxconn
You can run the **sysctl net.core.somaxconn** command on the node systems in your topology to determine the default value for each node system. Heavily used deployments might require a minimum value of 512.

Note: The attribute name for **net.core.somaxconn** includes the period (**.**) in the name.

You might need to increase the maximum file descriptor limit on the single controller node system in your topology. The message queue service runs on the single controller node. You can increase the maximum file descriptor limit by adding the following lines to the `/etc/security/limits.conf` file. After the update, new processes will use the new limit. Existing processes continue to use their current limit.

```
* soft nofile 65536
* hard nofile 65536
```

Configuring Cinder drivers

Use this information to configure Cinder storage drivers appropriately for IBM Cloud Manager with OpenStack, version 4.1.

Configuring IBM Storwize Cinder driver:

You can configure IBM Cloud Manager with OpenStack to use the IBM Storwize Cinder driver.

About this task

Complete the following steps to configure IBM Cloud Manager with OpenStack to use the IBM Storwize Cinder driver.

Procedure

1. Add the following JSON attribute to the default_attributes section of your environment file, `your-environment-name.json`.
   - **`openstack.block-storage.volume.driver`**: Set to *cinder.volume.drivers.ibm.storwize_svc.StorwizeSVCDriver*
2. Add the following JSON attributes as needed to the default_attributes section of your environment file. The values for these attributes must be based on your IBM Storwize Cinder configuration.
   - **`openstack.block-storage.san.san_ip`**: IP address of SAN controller. The default value is *127.0.0.1*.
   - **`openstack.block-storage.san.san_login`**: User name for SAN controller. The default value is *admin*.
   - **`openstack.block-storage.san.san_private_key`**: File name of the private key to use for SSH authentication. The default is */v7000_rsa*.
   - **`openstack.block-storage.storwize.storwize_svc_volpool_name`**: Storage system storage pool for volumes. The default is *volpool*.
   - **`openstack.block-storage.storwize.storwize_svc_vol_rsize`**: Storage system space-efficiency parameter for volumes. The default value is *2*.
   - **`openstack.block-storage.storwize.storwize_svc_vol_warning`**: Storage system threshold for volume capacity warnings. The default value is *0*.
   - **`openstack.block-storage.storwize.storwize_svc_vol_autoexpand`**: Storage system auto-expand parameter for volumes. The default is *true*.
   - **`openstack.block-storage.storwize.storwize_svc_vol_grainsize`**: Storage system grain size parameter for volumes. The default is *256*.
   - **`openstack.block-storage.storwize.storwize_svc_vol_compression`**: Storage system compression option for volumes. The default value is *false*.
   - **`openstack.block-storage.storwize.storwize_svc_vol_easytier`**: Enable Easy Tier® for volumes. The default value is *true*.
   - **`openstack.block-storage.storwize.storwize_svc_flashcopy_timeout`**: Maximum number of seconds to wait for FlashCopy® to be prepared. The default value is *120*.
   - **`openstack.block-storage.storwize.storwize_svc_vol_iogrp`**: The I/O group in which to allocate volumes. The default value is *0*.
   - **`openstack.block-storage.storwize.storwize_svc_connection_protocol`**: Connection protocol (iSCSI or Fibre Channel). The default value is *iSCSI*.

     Note: Only iSCSI is supported by Hyper-V compute nodes.
   - **`openstack.block-storage.storwize.storwize_svc_iscsi_chap_enabled`**: Configure CHAP authentication for iSCSI connections. The default value is *true*.

     Note: If the compute node is Hyper-V, specify the value in the environment file as *false*.
   - **`openstack.block-storage.storwize.storwize_svc_multipath_enabled`**: Connect with multipath (Fibre Channel only; iSCSI multipath is controlled by Nova). The default value is *false*.
   - **`openstack.block-storage.storwize.storwize_svc_multihostmap_enabled`**: Allows vdisk to multi-host mapping. The default value is *true*.

Configuring IBM Storwize V7000 Unified and IBM SONAS Cinder driver:

You can configure IBM Cloud Manager with OpenStack to use the IBM Storwize V7000 Unified and IBM Scale Out Network Attached Storage Cinder driver.

About this task

Complete the following steps to configure IBM Cloud Manager with OpenStack to use the IBM Storwize V7000 Unified and IBM SONAS Cinder driver.

Procedure

1. Add the following JSON attribute to the default_attributes section of your environment file, `your-environment-name.json`.
   - **openstack.block-storage.volume.driver**: Set to *cinder.volume.drivers.ibm.ibmnas.IBMNAS_NFSDriver*.
2. Add the following JSON attributes as needed to the default_attributes section of your environment file. The values for these attributes must be based on your IBM Storwize V7000 Unified and IBM SONAS Cinder driver configuration.
   - **openstack.block-storage.ibmnas.nas_ip**: Management IP address of IBM SONAS or IBM Storwize V7000 Unified storage. The default value is *127.0.0.1*.
   - **openstack.block-storage.ibmnas.nas_login**: User name to authenticate for IBM SONAS or IBM Storwize V7000 Unified storage. The default value is *admin*.
   - **openstack.block-storage.ibmnas.nas_access_ip**: Host name or public IP address to access shares. There is no default value.
   - **openstack.block-storage.ibmnas.nas_ssh_port**: SSH port to access shares. There is no default value.
   - **openstack.block-storage.ibmnas.shares_config**: File that contains a list of IBM SONAS or IBM Storwize V7000 Unified shares. The default value is */etc/cinder/nfs_shares.conf*.
   - **openstack.block-storage.ibmnas.mount_point_base**: Storage system mount point path for volume creation. The default value is */mnt/cinder-volumes*.
   - **openstack.block-storage.ibmnas.nfs_sparsed_volumes**: Storage system volume creation method. The default value is *true*.
   - **openstack.block-storage.ibmnas.export**: Storage system shares / export path parameter. There is no default value.
3. Create a data bag item in the user passwords data bag for your environment. This data bag item will contain the password to authenticate to the storage system.
   a. Create a directory for the user passwords data bag. Change *your_env_user_passwords* to the name of the user passwords data bag for your environment. The **openstack.secret.user_passwords_data_bag** attribute in your environment file contains the data bag name to use.

      ```
      $ mkdir -p data_bags/your_env_user_passwords
      $ chmod -R 600 data_bags/
      ```

   b. Create a data bag item with the following contents. The data bag item name must match the value of the **openstack.block-storage.ibmnas.nas_login** attribute. The following example assumes that the value is *your_data_bag_item_name*.

      ```
      $ cat data_bags/your_env_user_passwords/your_data_bag_item_name.json
      {
        "id": "your_data_bag_item_name",
        "your_data_bag_item_name": "CHANGE_TO_PASSWORD"
      }
      ```

   c. Upload the data bag item that you created in the previous step. Change *your-secret-key-name* to the secret key for your topology. The **secret_file** JSON attribute in your topology file contains the secret file to use.

      ```
      $ knife data bag from file your_env_user_passwords
      your_data_bag_item_name.json --secret-file your-secret-key-name
      ```

   d. Remove the local data bag item since it is no longer needed.

      ```
      $ rm -rf data_bags/
      ```

Configuring IBM GPFS Cinder driver:

You can configure IBM Cloud Manager with OpenStack to use the IBM GPFS™ Cinder driver.

About this task

Complete the following steps to configure IBM Cloud Manager with OpenStack to use the IBM GPFS Cinder driver.

Procedure

1. Add the following JSON attribute to the default_attributes section of your environment file, `your-environment-name.json`.
   - **openstack.block-storage.volume.driver**: Set to *cinder.volume.drivers.gpfs.GPFSDriver*.
2. Add the following JSON attributes as needed to the default_attributes section of your environment file. The values for these attributes must be based on your IBM GPFS Cinder driver configuration.
   - **openstack.block-storage.gpfs.gpfs_mount_point_base**: Path to directory in GPFS file system where volume files are located. There is no default value.
   - **openstack.block-storage.gpfs.gpfs_images_dir**: Path to directory in GPFS file system where images are located. There is no default value.
   - **openstack.block-storage.gpfs.gpfs_images_share_mode**: Path to directory in GPFS file system where Glance images are located. The default value is *copy_on_write*.
   - **openstack.block-storage.gpfs.gpfs_sparse_volumes**: Create volumes as sparse or fully allocated files. The default value is *true*.
   - **openstack.block-storage.gpfs.gpfs_max_clone_depth**: Maximum clone indirections that are allowed when volume file snapshots clones are created. The default value is *8*.
   - **openstack.block-storage.gpfs.gpfs_storage_pool**: GPFS storage pool that volumes are assigned to. The default value is *system*.

   For more information about configuration parameters, see OpenStack information to Enable the GPFS driver.

Verifying Open vSwitch configuration

Verify the Open vSwitch configuration for the applicable network that you configured.

Procedure

1. If you configured a GRE network, verify that there is at least one port beginning with `"gre-"` created on the *br-tun* bridge. Also, verify that the local and remote IP addresses that are listed in the options field match the IP addresses of the controller and the compute node.

```
> ovs-vsctl show

    Bridge br-tun
        Port br-tun
            Interface br-tun
                type: internal
        Port "gre-2"
            Interface "gre-2"
                type: gre
                options: {in_key=flow, local_ip="192.168.0.1",
out_key=flow, remote_ip="192.168.0.2"}
        Port patch-int
            Interface patch-int
                type: patch
                options: {peer=patch-tun}
```

2. If you configured a flat or VLAN network, verify you see the following output:
   a. Verify that a bridge was created matching the name that you specified in the **openstack.network.openvswitch.bridge_mappings** property. Typically the bridge is called *br-ethX*.
   b. Verify that port *ethX* was added to bridge *br-ethX* as an interface.
   c. Verify that port *phy-br-ethX* was added to bridge *br-ethX* as an interface.

d. Verify that port *int-br-ethX* was added to bridge *br-int* as an interface.

```
> ovs-vsctl show

    Bridge br-int
        Port br-int
            Interface br-int
                type: internal
        Port "int-br-eth1"
            Interface "int-br-eth1"
    Bridge "br-eth1"
        Port "br-eth1"
            Interface "br-eth1"
                type: internal
        Port "eth1"
            Interface "eth1"
        Port "phy-br-eth1"
            Interface "phy-br-eth1"
```

Note: The previous example was truncated to display only what you must verify. If you do not see *"phy-br-ethX'* and *'int-br-ethX'* ports in the `ovs-vsctl` show output (c and d from the previous list), verify that **neutron-openvswitch-agent** is active (**service neutron-openvswitch-agent status**) and check the error log in `/var/log/neutron/openvswitch-agent.log`.

Adding a compute node to a deployed topology

After deploying your topology, you can add a compute node system to your deployment. This does not apply to the Minimal topology.

About this task

Use the following procedure to add a Linux Kernel-based Virtual Machine (KVM), QEMU, PowerKVM or z/VM compute node to your deployment. To add another Hyper-V compute node, see "Installing and uninstalling the IBM Cloud Manager with OpenStack Hyper-V Agent" on page 27.

Procedure

1. Log in to the deployment system as the root user. This is the system where IBM Cloud Manager with OpenStack was installed.
2. Locate the topology file for the deployment to determine the secret file and environment for the topology along with the run list used for the compute nodes.
3. Create a node specific attribute file. This step is only required when the new compute node requires different attributes from those defined in the environment for the topology. For more information, see the "Deploying a controller +*n* compute or distributed database topology" on page 80 for the compute node.
4. Deploy the new compute node. Change the following in the command below.
   - **node-fqdn**: Set to the fully qualified domain name of the node system. You can also set to the public IP address, private IP address, or hostname.
   - **node-run-list**: Set to the run list determined in step 2.
   - **topology-environment**: Set to the environment determined in step 2.
   - **node-ssh-password**: SSH root user password for the node. To use an SSH identity file instead of a password, use the -i node-ssh-identity option, where node-ssh-identity is the SSH identity file for the node.
   - **topology-secret-file**: Set to the secret file determined in step 2.
   - **node-attribute-file**: Set to the node specific attribute file created in step 3. Remove this option if not required.

```
knife os manage deploy node node-fqdn 'node-run-list' -E topology-environment -x root
-P node-ssh-password --secret-file topology-secret-file --attribute-file node-attribute-file
```

5. After the deployment completes, you can check the status of the IBM Cloud Manager with OpenStack services, see “Checking status of OpenStack services” on page 171.
6. Update your topology file to include the new node. This step ensures that your topology file is accurate if you need to update the deployed topology in the future.

Results

You are ready to start using IBM Cloud Manager with OpenStack services on the new compute node system.

Changing passwords and secrets

You can change the passwords and secrets that were used during the deployment process.

About this task

This task applies to the **Controller +*n* compute or Distributed database** topologies. Customizing passwords and secrets is not done for the **Minimal** topology.

Note: This procedure only applies to passwords and secrets that are found in the “Data bags” on page 229 topic.

Procedure

1. Change the passwords using the appropriate management interface. For a list of users and passwords that are managed by IBM Cloud Manager with OpenStack deployments, see “User names and passwords” on page 223 and “Data bags” on page 229. For example, the database user passwords are changed using DB2 or MySQL and the OpenStack service user passwords are changed using OpenStack keystone (for example, **keystone user-password-update**). Other management interfaces for changing user passwords include the IBM Cloud Manager with OpenStack self-service portal, Qpid, PowerVC and z/VM.
2. You need to refresh your topology deployment to apply the password changes. You also change the secrets as part of this step.
   a. Log in to the deployment system as the root user. This is the system where IBM Cloud Manager with OpenStack was installed.
   b. Change to the directory where you stored the files for the topology that you deployed. Change **your-deployment-name** to the name for your deployment.
      ```
      $ cd your-deployment-name
      ```
   c. Copy the example data bags. Change **your_env_db_passwords**, **your_env_secrets**, **your_env_service_passwords**, and **your_env_user_passwords** to the name of the data bags for your environment. The **openstack.secret.*_data_bag** JSON attributes in your environment file contain the data bag names to use.
      ```
      $ mkdir data_bags
      $ cp -r /opt/ibm/cmwo/chef-repo/data_bags/db_passwords data_bags/your_env_db_passwords
      $ cp -r /opt/ibm/cmwo/chef-repo/data_bags/secrets data_bags/your_env_secrets
      $ cp -r /opt/ibm/cmwo/chef-repo/data_bags/service_passwords data_bags/your_env_service_passwords
      $ cp -r /opt/ibm/cmwo/chef-repo/data_bags/user_passwords data_bags/your_env_user_passwords
      $ chmod -R 600 data_bags/
      ```
   d. Change the passwords and secrets in your environment data bags. The password changes must be based on any changes made earlier. This is done by changing the value for the data bag item's name. For example, using the **your_env_user_passwords/admin.json** data bag item, change the value at *CHANGEME* to the password.
      ```
      {
        "id": "admin",
        "admin": "CHANGEME"
      }
      ```

e. Upload the data bag items for the passwords and secrets that you changed in the previous step. For each change, run the following command from the parent of the `data_bags` directory that you created. Change **your_data_bag_name** and **your_data_bag_item_name.json** to the data bag item changed and change **your-secret-key-name** to the secret key for your topology. The **secret_file** JSON attribute in your topology file contains the secret file to use.

   Note: The following command must be entered on a single line, even though the example shows a line break for formatting purposes.

   ```
   $ knife data bag from file your_data_bag_name your_data_bag_item_name.json
   --secret-file your-secret-key-name
   ```

f. Remove the local data bag items since they are no longer needed.

   ```
   $ rm -rf data_bags/
   ```

g. Refresh the topology deployment with the changed passwords and secrets.

   ```
   $ knife os manage update topology your-topology-name.json
   ```

Redeploying a node cleanup

If you need to deploy a topology to a node where a topology has previously been deployed, ensure that the node is in a clean state and that any corresponding node data on the deployment server has been removed.

About this task

Before you redeploy a topology, complete the following steps to prepare your node and deployment server.

Procedure

1. Revert your node system where a topology was previously deployed to a snapshot or backup version of the node. The node must be in the state that it was in prior to when the topology was first deployed.
2. On the deployment server, you must delete the data that is stored about a previously deployed node. The deployment server maintains state and other data about a node system in the deployment server database. If you are redeploying, or if you replace a node with the same fully qualified domain name, run the following commands on the deployment server to remove the data that has been stored about the node:

   ```
   $ knife node delete node.fqdn -y
   $ knife client delete node.fqdn -y
   ```

   These commands reset the node to a clean state for the Chef server.
3. Depending on how your node is configured, you might also need to delete the SSH fingerprint on your deployment server. The deployment knife commands use SSH to communicate with the node system to deploy OpenStack. Use one of the following methods to delete SSH fingerprints from the `.ssh/known_hosts` file on the deployment server:
   - Run the following command:

     ```
     ssh-keygen -R node fqdn
     ```

   - Edit the `.ssh/known_hosts` file in a text editor and delete the line that contains your node system.

Updating a deployed topology

After you deploy a topology, you might need to apply fixes, either to the IBM Cloud Manager with OpenStack services themselves or to how the IBM Cloud Manager with OpenStack services are deployed.

About this task

In addition, you might need to modify the customizations for the topology. To update a deployed topology, complete the following steps:

Procedure

1. The environment file `your-environment-name.json` for your deployment has constraints on the cookbook versions used. By default, these constraints allow cookbook patches only (that is, bug fixes) to be applied when you refresh the topology deployment. If you need minor cookbook fixes, then you must update the cookbook version constraints in your environment file using the instructions in the IBM Cloud Manager with OpenStack fix pack.
2. If you included the *ibm-sce-node* role in your topology and you want to apply a fix pack or JRE update for the IBM Cloud Manager with OpenStack self-service portal, use the following steps to update your environments.
   a. Enter the self-service portal installer repository, `<CMWO ROOT>/file-repo/SCE`
   b. Find the fix pack file. For example, `4.1.0.1-IBM-SCE-FP001-201405150700.zip`.
   c. Update the attribute value for **`ibm-sce.package.fixpack.file`** with its name in the environment file. For example:

```
"ibm-sce": {
         "package": {
           "fixpack": {
             "file": "4.1.0.1-IBM-SCE-FP001-201405150700.zip"
           }
         }
       }
```

   d. For a JRE update, find the applicable JRE file name. For example, either `sce410_jre_linux_installer.bin` (Linux) or `sce410_jre_lop_installer.bin` (Power).
   e. Update the attribute value for **`ibm-sce.package.jre.update.file`** with the file name in your environment file. For example:

```
"ibm-sce": {
         "package": {
           "update": {
             "jre": {
               "file": "sce410_jre_linux_installer.bin"
             }
           }
         }
       }
```

3. Optional: You can also modify some of the customizations for the IBM Cloud Manager with OpenStack services. You can change configuration files for IBM Cloud Manager with OpenStack services or customize other basic node system settings. Modify the customized settings in the environment file `your-environment-name.json`.

   Note:

   - Major changes to the topology, such as switching the database from IBM DB2 to MySQL, for example, cannot be completed as an update to the topology.
   - Manually modifying IBM Cloud Manager with OpenStack configuration files that are managed by the environment file for your topology gets overwritten when you update the deployment. Such configuration files have a banner similar to the following at the beginning of the file:

     ```
     # This file autogenerated by Chef
     # Do not edit, changes will be overwritten
     ```

     If you must manually modify such configuration files then the modification will need to be redone after the update.

4. If your environment was updated as part of a previous steps, upload the environment for your deployment.

   ```
   $ knife environment from file your-environment-name.json
   ```

   Note: If you are using the self-service portal, you must restart it after the update as well. To restart the self-service portal, run the **service sce restart** command.

5. Refresh the topology deployment.

   ```
   $ knife os manage update topology your-topology-name.json
   ```

   All of the nodes that are defined in the topology are refreshed. The following topology JSON items, which are required during a topology deployment, are optional during a topology update:

   Topology Keys: environment
   The environment to use with all nodes in the topology. The environment must exist on the Chef server.

   Node Keys: runlist
   The runlist for the node. The runlist is a list of roles and cookbook recipes to be applied to the node. The roles and cookbook recipes must exist on the Chef server.

   If the environment, or runlist values are specified in the topology JSON, they are set again for each node in the topology. If the environment is not specified in the topology JSON, it is not set on all nodes. If the runlist is not specified for a particular node in the topology JSON, it is not set on that node.

   If there are nodes in the topology that you do not want to refresh, there are two options to prevent them from being updated (while still leaving them in the topology).

   - The **knife os manage update topology** command can be run in interactive mode by specifying the `--interactive` command-line option:

     ```
     $ knife os manage update topology your-topology-name.json --interactive
     ```

     When run in interactive mode, and choosing not to update all nodes, the command asks whether to update each node. You can then skip any nodes in the topology that you do not want to refresh.
   - If you do not want to refresh a node in the topology, you can set a key in the topology `JSON` file that is used to specify whether to update it.

     Node Keys: allow_update
     A Boolean value, which indicates whether to update the node. Set to true to update the node. Set to false if you do not want to update the node. The default is true.

6. Optionally, refresh a single node instead of a whole topology. To refresh a single node, instead of the whole topology, run the following command:

   ```
   $ knife os manage update node your-node-fqdn -x ssh-user -P ssh-password
   ```

   You can set the node environment and run list, but it is not required. If you want to change the environment or run list of the node, enter the appropriate command-line options:

   Note: The following command must be entered on a single line, even though the example shows a line break for formatting purposes.

   ```
   $ knife os manage update node your-node-fqdn 'role[ibm-os-compute-node-kvm]'
   -E your-environment-name -x ssh-user -P ssh-password
   ```

Related concepts:

"Applying fixes and updates" on page 38
You can apply fixes and updates for the IBM Cloud Manager with OpenStack product.

Chapter 6. Configuring web interface properties

You can configure many features through a web user interface, but you can configure all features by modifying configuration property files that are located in the home directory. This section contains information about these configuration property files and the values that you can modify.

Configuring IBM Cloud Manager with OpenStack properties

Various configuration tasks for IBM Cloud Manager with OpenStack include changing the configuration of a topology after you deploy it, creating initial networks, and configuring SSH and ICMP. You can also configure live migration for a KVM node with shared storage, install a network time service, and configure multiple storage backends.

Changing IBM Cloud Manager with OpenStack properties

You can change the configuration of the IBM Cloud Manager with OpenStack topology after you deploy it.

About this task

To change the configuration, update attributes in the environment or node attribute files, then redeploy the configuration. If you change configuration files in the topology and redeploy later (for example, to deploy a fix pack), the configuration files are overwritten.

You can make temporary changes, for example to enable debug logging, by editing the files directly.

Procedure

1. Update the chef environment. If the change is common to the entire topology, change the environment.
   a. If you do not have the original environment file, you can download it from the Chef server.
      ```
      # knife environment show your-environment-name -d -F json > your-environment-name.json
      ```
   b. To see a list of environments, use the following command.
      ```
      # knife environment list
      ```
   c. Edit the environment file.
   a. Upload the modified environment to the chef server.
      ```
      # knife environment from file your-environment-file.json
      ```
2. Edit the node attribute files. If the change is specific to a single node, edit or create a node attribute file.
3. Deploy the updated topology.
   a. To update all the nodes in the topology, use the following command.
      ```
      # knife os manage update topology your-topology-name.json
      ```

      All the topology nodes are refreshed.
   b. To update a single node in the topology, use the following command.
      ```
      # knife os manage update node node-fqdn -E your-environment-name -x ssh-user -P ssh-password
      ```

Creating initial networks

After you deploy the components for creating a cloud environment, you can create several different types of networks.

© Copyright IBM Corp. 2012, 2014

About this task

The type of network that you must create depends on the type of network connectivity and hypervisor that you are using. For more information, see “Network considerations” on page 16.

Ensure that you understand the components of OpenStack networking before you create your initial networks. For more information, see Networking API, Tenant and provider networks, and Terminology.

The type of network is not specified for a tenant network because the type of network is determined by one of the following OpenStack properties:

- **openstack.network.openvswitch.tenant_network_type**
- **openstack.network.ml2_tenant_network_types**

When you create a network, you can specify the network type when you specify the provider network extension. You must have the administrator role to use the provider network extensions. For more information about provider network extensions, see Provider attributes.

You can create and manage your networks by using any of the following methods:

OpenStack CLI
For details, see Basic networking operations, Advanced Networking operations, and Provider extension API operations.

OpenStack dashboard
You can create a network in either of the following ways:

- To create both a network and a subnet - Log in to the dashboard, select **Project** > **Network** > **Networks**. Then, click **Create Network**. You cannot specify the network provider settings by using this method. For more information, see Create and manage networks and refer to the section entitled, "Create a network."
- To create a network and specify the network provider settings - Log in to the dashboard, select **Admin** > **System Panel** > **Networks**. Then, click **Create Network**. You cannot create a subnet by using this method. You can create the subnet after the network is created. For more information, see Configuring IBM Cloud Manager with OpenStack dashboard and refer to the section entitled "Provider Network Settings."

IBM Cloud Manager with OpenStack web interface
For details, see “Managing network configurations” on page 201.

Related tasks:

“Deploying with KVM or QEMU compute nodes” on page 82
Deploy the components that are necessary to create a cloud environment with KVM or QEMU compute nodes.

“Deploying with z/VM compute nodes” on page 90
Deploy the components that are necessary to create a cloud environment with z/VM compute nodes.

Configuring SSH and ICMP access to deployed virtual machines

After you deploy the components for creating a cloud environment, you can configure SSH and ICMP for accessing your virtual machines.

About this task

To configure SSH and ICMP access to your virtual machines by using the OpenStack dashboard, see Configure access and security for instances and refer to the section entitled, "Add a rule to the default security group."

For information about using commands to configure SSH and ICMP to access your virtual machines, see Enable ping and SSH on VMs (security groups). Using these commands result in the following error:

```
Multiple security_group matches found for name 'default', use an ID to be more specific.
```

Each tenant has a security group called `default`. To determine the security group ID, you can take one of the following actions:

- Deploy a virtual machine and look at the Neutron port settings of the IP address that is assigned to that virtual machine. The security group ID is a property of the Neutron port.
- Use the keystone tenant-list to determine the tenant ID and then dump each Neutron security group to determine which one has the matching tenant ID.

After you identify the security group ID, replace `default` with the security group ID when you run the original **`neutron security-group-rule-create`** commands.

OpenStack provides the following default security rules.

Table 25. Default security rules provided by OpenStack

Direction	EtherType	IP protocol	Remote	Description
Ingress	IPv4 or IPv6	Any	default	Allows all inbound network traffic from any virtual machine that is using this security group. Allows virtual machines that are deployed with this security group to receive data from each other. To prevent or limit communication between deployed virtual machines, delete this rule and add a rule that is more restrictive.
Egress	IPv4 or IPv6	Any	• 0.0.0.0/0 (CIDR) for IPv4 • ::/0 (CIDR) for IPv6	Allows all outbound network traffic to be sent to all destinations (0.0.0.0). Does not restrict any outbound network traffic from the virtual machines. If you want to limit outbound traffic from a deployed virtual machine, delete this rule and add a rule that is more restrictive.

Configuring migration for a KVM node with NFS shared storage

You can migrate running instances from one IBM Cloud Manager with OpenStack compute node to another compute node.

Before you begin

The following prerequisites must be true:

- IBM Cloud Manager with OpenStack was deployed using the controller +*n* compute or distributed database topology with two or more KVM compute nodes.
- You must have NFS shared storage configured.

About this task

Complete the following steps to enable migration capabilities with IBM Cloud Manager with OpenStack on KVM over NFS shared storage.

Procedure

1. Change the following JSON attributes in your environment file:
   a. Under **misc_nova** add the following value:

```
"misc_nova" : [
      " live_migration_flag=VIR_MIGRATE_UNDEFINE_SOURCE,
        VIR_MIGRATE_PEER2PEER,VIR_MIGRATE_LIVE ",
      ...
],
```

   b. Under **selinux.booleans.virt_use_nfs**, add the following value.

```
"selinux" : {
   "booleans" : {
     "virt_use_nfs" : "on"
  }
},
```

2. Change the following node-specific JSON attributes in your topology file for the controller node:

```
"ibm-openstack": {
   "iptables": {
      "custom_rules": "-A FWR -m state --state NEW -m tcp -p tcp --dport 2049 -j
 ACCEPT -m comment --comment \"nfs\"\n-A FWR -m multiport -m state --state NEW -m
 udp -p udp --dports 111,875,32803,32769,892,662,2020,20049 -j
 ACCEPT -m comment --comment \"nfs\""
  }
},
```

3. Update the topology. For instructions, see "Updating a deployed topology" on page 113.
4. Verify the following after the IBM Cloud Manager with OpenStack deployment is completed. Update as needed.
   a. Verify that **virt_use_nfs** Boolean is set to *ON*.

```
# getsebool -a|grep virt_use_nfs
virt_use_nfs --> on

Note: If needed, run the following command to set virt_use_nfs Boolean is set to ON:
# setsebool -P virt_use_nfs=on
```

   b. In the `/etc/nova/nova.conf` path within the controller and compute nodes, verify the following values.

```
live_migration_flag=VIR_MIGRATE_UNDEFINE_SOURCE,VIR_MIGRATE_PEER2PEER,VIR_MIGRATE_LIVE
```

   c. Ensure that the UID and GID of user **'nova'** are identical between each of your nodes.

```
# id nova
uid=162(nova) gid=162(nova) groups=162(nova)
```

   - To change a node, run the following commands:

```
# usermod -u 103 nova (to change UID to 103)
# groupmod -g 103 nova (to change GID to 103)
```

   - Run the following commands to change file ownership:

```
#service nova-api stop
#service libvirt-bin stop
#find / -uid 162 -exec chown nova {} \;  (UID 162 is the old nova UID)
#find / -gid 162 -exec chgrp nova {} \;  (GID 162 is the old nova GID)
#service nova-api restart
#service libvirt-bin restart
```

5. Configure NFS.
   a. Ensure that the controller and compute nodes are configured with DNS or `/etc/hosts`. Make sure that all nodes can perform name resolution with each other. You might use the ping command to ping each host from one another.
   b. Ensure that the UID and GID of your Compute user (Nova) are identical among all of your nodes. This ensures that the permissions on the NFS mount work correctly.
   c. Export `INSTANCES_PATH` (that is `/var/lib/nova/instances`) from the controller node, and have it readable and writable by the Compute user on compute nodes.
   d. Configure the NFS server at controller node.
      1) Add the following line to the `/etc/exports` file:

         ```
         INSTANCES_PATH Controller_IP/MASK(rw,sync,fsid=0,no_root_squash)
         ```

      2) Set the subnet mask to the appropriate value to include the IP addresses of all compute nodes. Then, configure and start services:

         ```
         # exportfs –r
         # chkconfig nfs on
         # service nfs restart
         ```

   e. Set the appropriate permissions on your shared directory. Run the following on all nodes.

      ```
      # chmod 775 INSTANCES_PATH
      ```

   f. Configure NFS mounts on compute nodes.
      1) Add the following line to the `/etc/fstab` file:

         ```
         Controller_IP:/ /INSTANCES_PATH nfs4 defaults 0 0
         ```

      2) Ensure that you can mount the exported directory can be mounted:

         ```
         # mount -av
         ```

      3) Verify that mount was successful:

         ```
         $ df -k
         Filesystem      1K-blocks   Used    Available  Use%  Mounted on
         ...
         Controller_IP:  921515008  101921792  772783104  12%
           /var/lib/nova/instances  ( <-- verify this)
         ```

      4) Verify that the `INSTANCES_PATH` directory is accessible and writable by the **'nova'** user:

         ```
         # ls -ld INSTANCES_PATH
         drwxrwsr-x. 5 nova nova 4096 Apr 20 18:25 /var/lib/nova/instances/

         # su – nova
         # touch INSTANCES_PATH/test.txt (this file should be successfully created)
         # rm -f INSTANCES_PATH/test.txt (should be successfully removed)
         # exit
         ```

      Note: For more information, see https://help.ubuntu.com/community/SettingUpNFSHowTo

6. Use the Nova live-migration command to migrate the instances:

   ```
   # nova live-migration server host_name
   ```

   The *server* variable can be either the ID of the server or the name. The *host_name* variable is the new compute host as shown in the Nova `service-list`.

What to do next

After IBM Cloud Manager with OpenStack is deployed and NFS shares a mount, all instances that are created from this point forward can be migrated between compute nodes.

Important: You must configure the firewall (IP tables) to allow **libvirt** to communicate between nodes. By default, **libvirt** listens on TCP port 16509, and an ephemeral TCP range 49152 - 49261 is used for the KVM communications. As a quick test, you can shut down the **IPtables** service to eliminate any firewall issues. When you change firewall rules, be careful choosing what ports you open and understand who has access. For information about ports that are used with **libvirt**, see http://libvirt.org/remote.html#Remote_libvirtd_configuration.

For more information about migration, see

- Migrate instances
- Configure migrations

Installing a network time service

Consider synchronizing the system clock of the deployment server and node systems with a network time server.

Before you begin

Note: This task is for Windows NTP server setup only. For Red Hat Enterprise Linux NTP server setup, see your Red Hat Enterprise Linux documentation.

About this task

Complete the following steps to install the Network Time Protocol (NTP) service:

1. Access the NTP installation package, `ntp-4.2.6p5-ibm-win32-setup.exe`, in the root directory of the IBM Cloud Manager with OpenStack installation media.
2. Run the `ntp-4.2.6p5-ibm-win32-setup.exe` file to install the NTP service.
3. After the NTP package is installed, specify the NTP server IP address or host name in `C:\ntp\etc\ntp.conf`.

```
# your local system clock, could be used as a backup
# (this is only useful if you need to distribute time no
matter how good or bad it is)

server x.x.x.x

# but it should operate at a high stratum level to let the
```

4. Navigate to **Control Panel** > **System and Security** > **Administative Tools** > **Services**, and start **Network Time Protocol Daemon** service.

For more information about how to synchronize the system clock with an NTP server, see the Network Time Protocol setup documentation ⌂for the platform.

Configuring multiple block storage backends

When you deploy an IBM Cloud Manager with OpenStack topology, you can configure a single block storage backend (cinder-volume) on the controller. To use multiple block storage backends, more configuration is required.

About this task

Important: When you use the following instructions to configure multiple storage backends, the changes are made outside of the Chef environment. These changes are overwritten if you use the Chef client to make updates to the node in the future. Plan to make a copy of the changes you make to the `cinder.conf` file.

To configure multiple block storage backends, complete the following steps:

Procedure

1. On the node where the cinder-volume service is running (typically the controller node), edit the `/etc/cinder/cinder.conf` file by using the information in OpenStack Configuration Reference: Block Storage. For detailed information, see i'Configure a multiple-storage back-end.
2. Add the **`enabled_backends`** property to the [DEFAULT] section of the file.
3. Add a section for each storage backend with the necessary properties. The following example shows configurations for several specific drivers. The examples do not include all properties. Copy the related properties into the new section and remove them from the DEFAULT section.

```
[DEFAULT]
. . .
enabled_backends = lvm_driver, powervc_driver, storwize_svc_driver
. . .

[lvm_driver]
# Sample cinder.conf stanza for a LVM volume driver
# To enable, add the stanza name (lvm_driver) to the DEFAULT enabled_backends
# property
volume_driver=cinder.volume.drivers.lvm.LVMISCSIDriver
volume_backend_name=LVM_iSCSI
volume_group=cinder-volumes
iscsi_ip_address=ip-addr

[powervc_driver]
# Sample cinder.conf stanza for a PowerVC volume driver
# To enable, add the stanza name (powervc_driver) to the DEFAULT
# enabled_backends property
volume_driver=powervc.volume.driver.powervc.PowerVCDriver
volume_backend_name=PowerVC

[storwize_svc_driver]
# Sample cinder.conf stanza for an IBM Storwize SVC volume driver
# To enable, add the stanza name (storwize_svc_driver) to the DEFAULT
# enabled_backends property
volume_driver=cinder.volume.drivers.ibm.storwize_svc.StorwizeSVCDriver
volume_backend_name=StorwizeSVC
san_ip=san-ip-addr
san_login=san-user-name
san_private_key=path-to-private-key
storwize_svc_volpool_name=X86_mdiskgrp1
storwize_svc_connection_protocol=iscsi
storwize_svc_iscsi_chap_enabled=false
```

4. Make a copy of the changes you made to the `cinder.conf` file. If the Chef client is run on the node in the future, by running the **`knife os manage update`** command on the deployment server, you must reapply the changes to the `cinder.conf` file. Edit the `/etc/cinder/cinder.conf` file to reapply the changes and restart the openstack-cinder-volume service.

Configuring IBM Cloud Manager with OpenStack self-service portal properties

You can configure many self-service portal features through a web user interface. However, you can configure all self-service portal features by modifying configuration property files that are in the self-service portal home directory. This section contains information about these configuration property files and the values that you can modify.

The `/.SCE41` directory is the self-service portal home directory and contains all of the self-service portal configuration property files and log files.

The `/.SCE41` directory is created at the following default locations, which are based on platform:

- On Windows platforms, the default location is the user home directory.
- On Linux platforms, the default location is the `/var/opt/ibm` directory.

The `authentication.properties` file settings are required for the self-service portal to interface with your cloud manager. The other features provide flexibility and convenience.

To configure basic properties and features of the self-service portal, you must modify properties that are defined within the configuration files in the home directory. These configuration files include the following properties:

`authentication.properties`
Specifies the user authentication type to use for the self-service portal.

`database.properties`
Specifies the database type and path.

`email.properties`
Specifies email notification capabilities.

`cloud.properties`
Specifies the common configuration of the self-service portal and the URL for the User Guide documentation.

`deployment.properties`
Specifies properties for configurations of virtual appliances to simplify deployments for cloud users.

`logging.properties`
Specifies self-service portal log file properties.

`networkConfiguration.properties`
Specifies network configurations for deployments.

`billing.properties`
Enables charging for cloud resources upon deployment.

`metering.properties`
Enables metering for cloud resources upon deployment.

`web.properties`
Specifies properties for the user interface configuration.

`server.properties`
Specifies properties for enabling the optional Secure Sockets Layer (SSL) configuration.

Notes:

- If you do not know your user home directory in a Windows OS, enter `%HOMEPATH%` in the address bar of a Windows Explorer window.

- If you modify these properties while the self-service portal server is active, you must stop and restart the self-service portal for the new property values to take effect.

Configuring secure shell (SSH) communication

The IBM Cloud Manager with OpenStack self-service portal ships a self-signed certificate for SSL communication between a client machine, such as a web browser, and the IBM Cloud Manager with OpenStack self-service portal node. This certificate is stored in the *<home directory>*/.keystore file.

This self-signed certificate is shipped for testing purposes only. It is not associated with a qualified host and domain name. Additionally, it is self-signed so a security warning is displayed when you access the host using https. To use SSL configuration in production, create a different self-signed or CA issued certificate that is designated specifically for the qualified host. Additionally, the keystore password must be changed or another keystore must be used to contain this certificate with a secure password. The new passwords would then be used in the following server.properties file configuration example.

To export a certificate to be used by clients, run the following command from the .SCE41 directory:

```
"<jre path>/keytool" -export
-v -alias SKC -file SKC.cer -keystore .keystore -storepass password
```

where *password* is the password you specify.

Notes:

- In order for this command to run properly, the Java/bin directory must be added to the system %PATH% variable.
- **keytool** is a key and certificate management utility that is included with Java SE 6.0.

After this certificate is imported into a client, the client can communicate with the IBM Cloud Manager with OpenStack self-service portal by using the trusted certificate with no additional user intervention required. If the import is not done, the client, such as a browser, might prompt the user to verify it and confirm that the certificate is trusted. After you confirm that you accept the risk of the certificate, you will be able to use SSL.

Note: When you use Internet Explorer to install a self-signed certificate, ensure that the certificate issuer name exactly matches the domain name of the URL that you are using it for. For example, if the URL is https://*ip_address*/cloud/web/login.html, where *ip_address* is your IP address, the **CN** setting must be CN=*ip_address* and the command is as follows:

```
keytool -genkey -dname "CN=ip_address, OU=Cloud, O=IBM, L=RTP, S=NC, C=US" -alias SKC
         -keystore .keystore -keyalg RSA -keysize 1024
```

If you still cannot install the certificate using Internet Explorer, it might be necessary to modify the system date time to synchronize with the IBM Cloud Manager with OpenStack self-service portal time. Also, ensure that you shut down and restart all instances of Internet Explorer after you install the certificate.

SSL is enabled on the server by configuring the server.properties file in the self-service portal home directory as follows:

```
# HTTP server port
org.osgi.service.http.port=18080

# Flag to enable/disable HTTP. If it is necessary for the protocol to be only SSL,
# set this flag to false.
org.eclipse.equinox.http.jetty.http.enabled=true

# Flag to enable/disable HTTPS
org.eclipse.equinox.http.jetty.https.enabled=true
```

```
# HTTPS port
org.eclipse.equinox.http.jetty.https.port=18443

# SSL password
org.eclipse.equinox.http.jetty.ssl.password=password

# Keystore password
org.eclipse.equinox.http.jetty.ssl.keypassword=password

# The full path location of the keystore
org.eclipse.equinox.http.jetty.ssl.keystore=home directory/.keystore

# The SSL protocol
org.eclipse.equinox.http.jetty.ssl.protocol=SSL_TLS
```

Note: The **org.eclipse.equinox.http.jetty.ssl.protocol** property is *SSL_TLS* if running on an IBM JRE. The property is *TLS* if running on a Sun or Oracle JRE.

Restart the server after you change the `server.properties` file. With the server running, point your client to `https://system:18443/cloud/api/users` to test it. Depending on whether you imported the certificate from above, you might be prompted to accept the certificate.

Creating a new certificate for your host

You can use the **keytool** tool to create a self-signed certificate for the host you are deploying IBM Cloud Manager with OpenStack self-service portal on or to create a certificate signing request (CSR) to send to a certificate authority (CA) to get a CA-issued certificate that is trusted by clients automatically.

For example, to generate a new keystore with specific customer information, use the following command:

```
keytool -genkey -dname "CN=cloud.ibm.com, OU=Cloud Services, O=IBM, L=RTP, S=NC, C=US"
-alias SKC -keystore .keystore -keyalg RSA -keysize 1024
```

CN Specifies the customers domain.

OU Specifies the organization within the customer's company.

O Specifies the company.

L Specifies the city of the company location.

S Specifies the state where that city resides.

CB Specifies the country.

To generate a certificate signing request from the keystore, run the following command:

```
keytool -certreq -alias SKC -keystore .keystore -file NewCertSignRequest.csr
```

To import the trusted certificate (.cer) file, run this command:

```
keytool -import -trustcacerts -alias SKC -file ./TrustedCertificate.cer -keystore .keystore
```

See the **keytool** documentation for your JRE for instructions. For the IBM JRE, the instructions are available at http://www.ibm.com/developerworks/java/jdk/security/60/secguides/keytoolDocs/keytool.html.

Note: When the CA is not trusted by clients automatically and you are attempting to access the self-service portal using https protocol, an exception is encountered that says the connection is untrusted. You must confirm that the risks are understood and must add an exception to continue. Even with a trusted certificate, when using Internet Explorer, you are likely to run into a similar exception.

Connecting using SSH

To further minimize security risks when connecting using OpenSSH, change the OpenSSH daemon configuration file so that the line containing `Protocol` is changed to `2`. Anything less than 2 is more susceptible to attack. The OpenSSH daemon implemented under the IBM Cloud Manager with OpenStack self-service portal uses port 22 as a default for communication.

Configuring user registry authentication

The IBM Cloud Manager with OpenStack self-service portal supports both Local and Lightweight Directory Access Protocol (LDAP) user registry authentication mechanisms. Authentication is performed by using the self-service portal local user registry.

Local user registries are intended for small-scale usage, such as proof-of-concept scenarios, demonstrations, or environments with up to 30 users and projects.

LDAP user registries provide the highest level of security and scalability for production environments and enable the self-service portal to share a central user registry for existing users with other applications that support LDAP.

For more information about user registry authentication, see "Supported IBM Cloud Manager with OpenStack user registries" on page 9.

Configuring Keystone authentication for the self-service portal

By default, IBM Cloud Manager with OpenStack self-service portal uses the local authentication mode. If you are using OpenStack and you want the highest level of security, configure Keystone authentication.

About this task

After successfully installing IBM Cloud Manager with OpenStack, there is a directory called `/.SCE41`. This directory is the IBM Cloud Manager with OpenStack home directory and contains all of the IBM Cloud Manager with OpenStack configuration property files and log files.

Procedure

1. Open the `authentication.properties` file in the home directory.
2. Configure the **`authentication.type`** property to use Keystone authentication, as follows:
   `authentication.type=KEYSTONE`
3. Log in to the Keystone and IaaS gateway server.
4. Open the `adapters.json` file in the `/etc/iaasgateway/` directory and change the `adapters.json` file as follows:

   Note: Each of the following lines must be entered on a single line, even though the example shows a line break for formatting purposes.

```
{"/v2.0/tokens/{token_id:.*}/endpoints": "iaasgateway.adapter.OSPKITokenAdapter"},
{"/v2.0/tokens/{token_id:.*}": "iaasgateway.adapter.OSPKITokenAdapter"},
{
        "authentication": {
                ".*": [
                    {"/{path_info:.*}": "iaasgateway.adapter.OSAuthAdapter"}
                ]
        },
        "authorization": {

        },
        "request": {
                ".*": [
                        {"/v2.0/tokens": "iaasgateway.adapter.OSPostTokensAdapter"},
```

```
                        {"/v3/auth/tokens": "iaasgateway.adapter.OSV3TokensAdapter"},
                        {"/v2.0/tokens/{token_id:.*}/endpoints": "iaasgateway.adapter.
                          OSPKITokenAdapter"},
                        {"/v2.0/tokens/{token_id:.*}": "iaasgateway.adapter.
                          OSPKITokenAdapter"},
                    {"/{path_info:.*}": "iaasgateway.adapter.HTTPPassThroughAdapter"}
                ]
        }
}
```

5. Open the `keystone.properties` file in the home directory.
6. In the Keystone properties file, configure the Keystone configuration information, as follows:

```
#host address of keystone/IaaS Gateway (the IaaS GW listening agent)
host= x.x.x.x
#listening port of the keystone/IaaS Gateway
port=9973
#user name for connecting the keystone/IaaS Gateway
user_name=sceagent
#user password for connecting the keystone/IaaS Gateway
password=openstack1
#if the keystone is ssl enabled.
isSSLEnabled=true
#default tenant
tenant=Public
```

   Note: You must be able to connect to Keystone during this step.
7. Restart the IBM Cloud Manager with OpenStack self-service portal by using the following command:

```
sudo ./cmwo_services.sh 'sce' restart
```

   For more information about the `cmwo_services.sh` script, see Restarting IBM Cloud Manager with OpenStack services.

LDAP authentication

The IBM Cloud Manager with OpenStack includes an LDAP authenticator that is designed to authenticate users in a wide range of environments, whether the authentication process is a simple bind or a multistage process.

LDAP authentication is performed by defining a series of steps for the authentication process, defining the inputs and outputs of each step, and then running them in sequence. If all the steps run successfully, the user is authenticated.

Configuring LDAP authentication manually

Beginning in IBM Cloud Manager with OpenStack 3.1 the web interface is the primary means of configuring LDAP. If you have a migrated LDAP configuration from a previous release, or if you want to enable user name case sensitivity, you can edit the `ldap.xml` configuration file.

For more information about configuring LDAP by using the web interface, see “Configuring LDAP authentication using the web interface” on page 175.

Properties of an LDAP authentication step

Host A string host name for the LDAP host. This property is required.

Search Context
If an LDAP lookup is to be performed, a search context must be provided. This property is required only if a search filter is provided.

Search Filter
If an LDAP lookup is to be performed, a search filter format string must be provided. This string tells the authenticator how to create a search filter that ensures that only one result is returned, as LDAP does not guarantee ordering of results if there are more than one. Additionally, the string

FILTER is a special value in the search filter. This string is replaced with the user ID entered during login. If you do not use the string *FILTER* in your configuration file, there is no replacement during authentication. If the strings that are defined in your configuration file are static, and a search context is provided, the search filter property is required.

Authentication DN

This property specifies the distinguished name (DN) used for authenticating to the LDAP server. If you are using this property to perform a search, you can specify the property as:

```
</authDN password="password">dnname</authDN>
```

If the property is specifying the DN to use for authentication, define the DN in one of the following ways:

- The DN can be constructed from the user ID. For example, the DN for a user logging in as *joe* can be constructed by using the following:

  ```
  <authDN>uid={FILTER},ou=people,dc=site</authDN>
  ```

  This example creates the DN `uid=joe,cn=users,ou=people,dc=site`
- The DN of the LDAP user entry that is returned by a previous search step is represented using the special value {PERSON_DN}, as shown in this example:

  ```
  <authDN>{PERSON_DN}</authDN>
  ```

In both cases, the password that the user entered to log in is also used to authenticate to the LDAP server.

To perform an anonymous search, do not specify the authentication DN property.

Admin Users

This attribute specifies a list of LDAP user accounts to be given administrator privileges:

```
<adminUsers>admin@company.com,me@company.com</adminUsers>
```

User name case sensitive

This attribute specifies whether the LDAP server defines the user name as case sensitive or not.

```
<userNameCaseSensitive>true</userNameCaseSensitive>
```

Secure connection enablement

This attribute specifies whether to enable a secure connection for LDAP authentication. Some LDAP servers enable the **StartTLS** operation by default, and other LDAP server do not. As an administrator, you can turn off the secure connection, if the LDAP server does not support **StartTLS** operation. The possible values for this attribute are **true** or **false**. To enable a secure connection, specify this property in the config element:

```
<enableSecureConn>true</enableSecureConn>
```

Outputs

This value indicates what information is needed during the current step for the next step and how to pass that information along. The Outputs value is essentially a list of instructions that gets an attribute value, for example *foo*, and passes it along as *bar*. This value is optional.

- **User account name**: An output can be flagged as the name for a user account by adding `attribute="fullname"` to the output tag. This value is retrieved and used as the user name by IBM Cloud Manager with OpenStack. If you do not specify this value, the user ID is used for the user display name.
- **User e-mail address**: An output can be flagged as the email for a user account by adding `attribute="email"` to the output tag. This value is retrieved and used as the user email address by IBM Cloud Manager with OpenStack.

Example of an ldap.xml file

In this example of an `ldap.xml` file, an authenticated secure search is performed to find the directory entry where the mail attribute matches the value that is passed into the username field.

```
<?xml version="1.0"?>
<config>
 <host>ldap://ldap.company.com</host>
  <adminUsers>admin@company.com,me@company.com</adminUsers>
  <enableSecureConn>false</enableSecureConn>
  <userNameCaseSensitive>true</userNameCaseSensitive>
  <step>
      <authDN password="password">cn=ldapAdmin,ou=directory,o=company.com</authDN>
      <searchFilter>(|(mail={FILTER}))</searchFilter>
      <searchContext>ou=directory,o=company.com</searchContext>
      <outputs>
         <output attribute="fullname">
            <get>cn</get>
         </output>
      </outputs>
    </step>
    <step>
        <authDN>{PERSON_DN}</authDN>
    </step>
</config>
```

Changing authentication mode

You can change the IBM Cloud Manager with OpenStack to LDAP authentication mode by editing the `authentication.properties` file.

About this task

Note: Beginning in IBM Cloud Manager with OpenStack 3.1, the web interface is the primary means of configuring LDAP. If you use the web interface to configure LDAP, the steps in this task are not required. For more information about configuring LDAP by using the web interface, see "Configuring LDAP authentication using the web interface" on page 175.

To change IBM Cloud Manager with OpenStack to LDAP authentication mode, complete the following steps:

Procedure

1. Open the `authentication.properties` file in the home directory.
2. Set the `authentication.type` property to *LDAP* as shown in the following example:

   ```
   authentication.type=LDAP
   ```
3. Open the `ldap.xml` file in the home directory.
4. Configure the LDAP steps as described in the "Configuring LDAP authentication manually" on page 128.
5. Restart the IBM Cloud Manager with OpenStack server.

   You can change the authentication mode back to local by setting the `authentication.type` property back to *LOCAL*.

Configuring local authentication

By default IBM Cloud Manager with OpenStack is set up to use local authentication mode. Local authentication is intended for small-scale usage, such as proof-of-concept scenarios, demonstrations, or environments with up to 30 users and projects. For large-scale production environments, configure LDAP to ensure the highest level of security.

About this task

Validate the configuration by following these steps:

Procedure

1. Open the `authentication.properties` file in the home directory.
2. Configure the `authentication.type` property to use local authentication, such as:

   ```
   authentication.type=Local
   ```

3. Configure the default administrator user name, name, and password, similar to the following example:

   ```
   admin.username=admin
   admin.name=SCE Administrator
   admin.password=<password>
   ```

   Notes:

   a. These fields might already be populated or configured during installation.

   b. The values of the `admin.username`, `admin.name`, and `admin.password` that are shown are examples. You should update these values according to your business or security guidelines.

   c. To prevent too many invalid login attempts, a user can attempt to login to IBM Cloud Manager with OpenStack three times within a 24-hour period. After three failed login attempts, both the user and administrator roles are locked out. However, the administrator can unlock the user record.

   To configure this limitation, enable the `com.ibm.cfs.failedlogincheck.enabled` property as follows:

   ```
   com.ibm.cfs.failedlogincheck.enabled=false
   ```

   This property is disabled by default.

   d. The account 'Locked' field associated with the user record in IBM Cloud Manager with OpenStack, is valid only when the account is locked in IBM Cloud Manager with OpenStack using LOCAL authentication, rather than LDAP authentication. If you are using the LDAP authentication and have the 'account locking' feature enabled on your LDAP server, do not enable it on the IBM Cloud Manager with OpenStack server. In this case, set the `com.ibm.cfs.failedlogincheck.enabled` property to `false` in the `authentication.properties` file on the IBM Cloud Manager with OpenStack server.

Configuring REST API authentication

You can configure IBM Cloud Manager with OpenStack to require authentication when it calls to the IBM Cloud Manager with OpenStack REST APIs.

About this task

IBM Cloud Manager with OpenStack supports the following authentication methods:

- Basic HTTP authentication for a user login and REST API-based validation
- Encrypted token-based authentication for REST API calls

The basic strategy for using encrypted tokens is as follows:

- HTTP/REST agents (browser or REST client) initially use the login authentication REST API to authenticate their user ID and password credentials.
- The user ID and password are validated against the LOCAL or LDAP repository (depending if LOCAL or LDAP is configured).
- Upon successful login authentication, an encrypted token and its expiration are returned in the login response.

- The agent can use (as an HTTP header cookie) the encrypted token for subsequent REST API calls to identity themselves until the token expires.
- After the authentication token expires, the agent must use the login REST API again to validate their user ID and password. When complete, the agent obtains a new authentication token.

Note: The system that is running the IBM Cloud Manager with OpenStack web interface or REST client must have the date, time, and time zone that is correctly configured for its physical location.

To require authentication when IBM Cloud Manager with OpenStack calls to the Rest APIs, complete the following configuration steps:

Procedure

1. Open the `authentication.properties` file in the home directory.
2. Set the authentication.secure property to `true` as shown in the following example:

   ```
   authentication.secure=true
   ```

   When the property is set to `true`, the caller is prompted for credentials before it processes the API request. The credentials are validated against the user registry that is configured, such as Local or LDAP.
3. If IBM Cloud Manager with OpenStack is configured to use Single Sign On with other SmartCloud products, you must set the shared secret key. Use the same shared secret key in all applications by using Single Sign On. If IBM Cloud Manager with OpenStack is not using Single Sign On, leave the property unset and the application generates and save a new secret key when it first starts.

   ```
   com.ibm.cfs.token.key=The Secret Encryption Key
   ```
4. Optional: Set the name of the HTTP header cookie. The cookie is used to transport the encrypted authentication token. This property specifies the name of the HTTP header cookie, which is used in HTTP REST API requests to transport the encrypted authentication token. The default value is *simpletoken*.

   ```
   com.ibm.cfs.token.header.field.name=simpletoken
   ```
5. Optional: Set the time duration for authentication tokens (in seconds). This time duration determines how long an authentication token is valid. After a token expires, the agent must obtain a new token by using the login authentication REST API.

   ```
   com.ibm.cfs.token.duration=14400
   ```
6. Optional: Disable automatic renewal for the authentication token. When enabled, authentication tokens renew (by using the specified duration period) each time they are successfully used on an API call. If this option is disabled, the only way to renew an authentication token is to obtain a new token by using the login authentication REST API.

   ```
   com.ibm.cfs.token.autorenew.enabled=false
   ```
7. Restart your IBM Cloud Manager with OpenStack server for the changes to take effect.

Configuring database

By default, the IBM Cloud Manager with OpenStack self-service portal uses an internal Derby database that is created inside the home directory. For larger environments, you might want to use an external database. The self-service portal supports the use of an external DB2 database.

About this task

For more information about DB2 support, see “Supported IBM Cloud Manager with OpenStack databases” on page 9.

The self-service portal also supports initial use of a Derby database and migration to a DB2 database at a future point. For details on the migration process, see “Migrating a Derby database to DB2 database” on page 49.

To change the self-service portal database configuration to use DB2, complete the following steps:

Procedure

1. Open the `database.properties` file in the home directory.
2. Set the database.type property to DB2, as shown in the following example:

   ```
   database.type=db2
   ```
3. Set the database.username property to the user ID defined for the database, as shown in the following example:

   ```
   database.username=<db2user>
   ```
4. Set the database.password property to the password ID defined for the database, as shown in the following example:

   ```
   database.password=<db2passwd>
   ```

   Note: The clear text password is replaced with an encrypted password after the self-service portal starts the first time.
5. Set the database.db2.path property to the location of the DB2 database, as shown in the following example:

   ```
   database.db2.path=//localhost:50000/cfs:
   ```

   Note:

   - One or more connection directives can be appended to the database path, and they must be separated by a semicolon. For example:

     ```
     database.db2.path=//localhost:50000/cfs:retrieveMessagesFromServerOnGetMessage=true;
     ```
   - Replace `localhost` with a full IP address (it can be a remote host) and verify the port number. Here are a few links to help you find the port number:

     UNIX: http://publib.boulder.ibm.com/infocenter/cmgmt/v8r3m0/index.jsp?topic=%2Fcom.ibm.sysadmin.hlp%2Fcsa10010.htm or as a potential shortcut, use the `grep i db2 /etc/services` command.

     Windows: http://publib.boulder.ibm.com/infocenter/cmgmt/v8r3m0/index.jsp?topic=%2Fcom.ibm.sysadmin.hlp%2Fcsa10010.htm or as a potential shortcut, look for the DB2 entries in the services file at `C:\WINDOWS\system32\drivers\etc\services`.

Configuring email notifications

To receive email notifications when various actions are taken, you must configure the notification properties in the home directory.

About this task

The following examples are user and administrator events that can be configured to trigger an email notification.

- Account is created or deleted, reaches its balance threshold, is delinquent, receives a payment, has a bill created
- Instances are deployed, fail, expire, have snapshots complete
- Project access is requested or granted
- User is created, requests access, has a password reset
- Virtual machine is backed up, restored, has a snapshot created, or reverts to a snapshot
- Volumes are created, deleted, detached and resized

To set up notification for the IBM Cloud Manager with OpenStack self-service portal, follow these steps:

Procedure

1. Open the `email.properties` file in the home directory.
2. Set the `com.ibm.cfs.email.relay.host` property to the host name of the relay host that the self-service portal uses for outgoing SMTP emails.
3. Optionally, you can set the email subject prefix for all self-service portal emails and the "from" name, by setting the following properties:

   ```
   com.ibm.cfs.email.subject.prefix
   com.ibm.cfs.email.from.name
   com.ibm.cfs.email.from.address
   ```

4. Save the `email.properties` file and restart the self-service portal service.

   You can globally disable email notifications in the self-service portal by setting the `com.ibm.cfs.email.default.notifications` property in the `email.properties` file to `false`.

   Individual users can disable notifications through the self-service portal.

   Note: Ensure that the self-service portal administrator has an email address that is configured to receive notifications. The notifications email for volume operations can be sent only to the self-service portal Administrator.

Configuring common cloud properties

Common cloud properties are configured by providing information such as the refresh interval and online help configuration in the `cloud.properties` file.

About this task

To configure your cloud manager, complete the following steps:

Procedure

1. Open the `cloud.properties` file in the home directory.
2. Edit the properties by providing values for each configuration property.
3. Save the `cloud.properties` file and restart the IBM Cloud Manager with OpenStack self-service portal server.

Cloud refresh interval

IBM Cloud Manager with OpenStack checks for new images and instances in the cloud and synchronizes the properties for these images and instances.

By default, IBM Cloud Manager with OpenStack receives messages from the cloud to synchronize with the cloud manager. The frequency of this synchronization is determined by the `com.ibm.cfs.cloud.refresh.interval` property in the `cloud.properties` file. If the property is not set, a default of 30 seconds is used.

IBM Cloud Manager with OpenStack scans the cloud for updates on instances as often as the refresh interval property specifies. However, you can change the synchronization mode so that IBM Cloud Manager with OpenStack periodically checks with the cloud without waiting for messages.

For more information about setting the synchronization mode, see “Configuring cloud synchronization mode” on page 137.

Cloud online help configuration

The IBM Cloud Manager with OpenStack has a configurable documentation property that enables IBM Cloud Manager with OpenStack to open the User Guide when the Help link is selected by the user.

About this task

To configure the URL for the Help documentation, follow these steps:

Procedure

1. Open the `cloud.properties` file in the home directory.
2. Configure the property `com.ibm.cfs.cloud.documentation` to be set to the URL for the User Guide location. By default, this property is set to the IBM Cloud Manager with OpenStack User Guide. Using the default property setting, the user can access the User Guide on the IBM Cloud Manager with OpenStack wiki in any of the supported languages by selecting the link for the language of choice. If something other than this default behavior is wanted, the property can be changed to any URL where the User Guide document is located.

Configuring global image deployment

Image deployment customization properties that apply equally to all images in the IBM Cloud Manager with OpenStack self-service portal image library can be configured through the `deployment.properties` configuration file in the home directory.

To simplify the image deployment process for self-service portal users, configure all of the images before you make the self-service portal available to users. Image property customization often requires knowing the low-level details of the environment or having advanced knowledge of the data center.

You can configure image deployment customization properties through the self-service portal for individual image settings or through the `deployment.properties` file in the home directory for global image settings. For more information about configuring individual image settings, see "Configuring image deployment properties" on page 183.

The contents of the `deployment.properties` configuration file depend on what is expected by the cloud manager software, either VMware or OpenStack and what hardware is available.

Note: Global configurations are refreshed only when manually reset or when the deployment target changes.

VMware

When you use VMware as the cloud manager, the following properties are supported for Linux and Windows images:

```
vmware.linux.dns1=9.8.8.8
vmware.linux.dns2=9.8.8.7
vmware.linux.hostname=myhost
vmware.linux.domainname=cloud.company.com

vmware.windows.computername=WINDOWS
vmware.windows.workgroup=WORKGROUP
vmware.windows.timezone=20
vmware.windows.username=John Doe
vmware.windows.organization=Cloud Company
vmware.windows.productkey=xxxx-xxxx-xxxx-xxxx-xxxx
vmware.windows.password=Default_password_for_windows_deployments

vmware.dnssuffixlist=cloud.company.com,company.com

vmware.networkdevice.Network adapter 1.network=VM Network
vmware.networkdevice.Network adapter 1.usedhcp=false
vmware.networkdevice.Network adapter 1.ipaddress=
vmware.networkdevice.Network adapter 1.netmask=255.255.255.0
vmware.networkdevice.Network adapter 1.gateway1=9.9.9.9
vmware.networkdevice.Network adapter 1.gateway2=
```

```
vmware.networkdevice.Network adapter 1.dns1=9.8.8.8
vmware.networkdevice.Network adapter 1.dns2=9.8.8.7
vmware.networkdevice.Network adapter 1.primaryWINS=9.8.8.10
vmware.networkdevice.Network adapter 1.secondaryWINS=9.8.8.11
```

For VMware, you can also find these properties in the `deployment.properties` file.

OpenStack

When you use OpenStack as the cloud manager, the following properties are supported for AIX®, Linux, and Windows images in the `deployment.properties` file:

```
openstack.openstack.flavors
openstack.openstack.keypairs
openstack.openstack.server.personality.source.[1-5]
openstack.openstack.server.personality.target.[1-5]
openstack.openstack.server.customizations
openstack.networkdevice.Network adapters.networks
openstack.config.drive
```

More deployment properties are available for images that have a configuration strategy. For more information about configuration strategies, see "Configuring images with OpenStack" on page 146.

Configuring a deployment target

By default, IBM Cloud Manager with OpenStack deploys images to any known pool or host in the cloud, where "pool" refers to a resource pool when you are using VMware. OpenStack only supports the cloud as the deployment target. For VMware, you can set a different default global deployment target.

About this task

To change this default target selection strategy, follow these steps:

Procedure

1. Open the `deployment.properties` file in the home directory.
2. Set the value of the `com.ibm.cfs.deployments.target.strategy` property to any of the following target selection strategies:

   byName
   Use the target with the given name. The name might refer to a host, system pool, resource pool, or cluster depending on the type of cloud adapter that is being used. Set the property value to `byName:{targetName}`, where `{targetName}` is the actual name of the desired system.

   byID Use the system with the specified ID. Set the property value to `byID:{targetOID}`, where `{targetOID}` is the actual OID of the desired system.

   anyPool
   Use any system that is a pool.

   anyHost
   Use any system that is a physical host.

   anyPoolOrHost
   Use any pool or physical host.

   anyCluster
   Use any cluster (applicable to VMware only).

   anyTarget
   Use any pool or host or cluster for VMware.
3. Save the `deployment.properties` file and restart the IBM Cloud Manager with OpenStack server.

Configuring cloud synchronization mode

IBM Cloud Manager with OpenStack scans the cloud for updates on instances as often as the refresh interval property specifies. However, you can change the synchronization mode so that IBM Cloud Manager with OpenStack periodically checks with the cloud without waiting for messages.

To change the sync mode, open the `deployment.properties` file and modify the following settings:

```
com.ibm.cfs.cloud.sync=push
```

To enable IBM Cloud Manager with OpenStack to synchronize with the cloud by using periodic checking, set this property to `poll`. Enable the configuration by ensuring that you uncomment the **`com.ibm.cfs.cloud.sync`** line (remove any preceding # symbol).

For more information about setting the refresh interval, see "Cloud refresh interval" on page 134.

Configuring a staging project

By default, IBM Cloud Manager with OpenStack scans the cloud for new images periodically. When IBM Cloud Manager with OpenStack finds a new image or instance, it places it in the Public project where all users have access to it. However, you can configure a staging project to store newly discovered images or instances, allowing administrators to configure images before making them available to other users.

For more information about newly discovered images, see Cloud refresh interval.

To configure this staging project, add or uncomment this line in the `deployment.properties` file:

```
com.ibm.cfs.staging.project=Staging
```

Save the `deployment.properties` file and restart the IBM Cloud Manager with OpenStack server. The property takes effect after the server is restarted.

Note: When using the VMware adapter, virtual servers that are defined as templates on the vCenter server are automatically discovered and displayed on the IBM Cloud Manager with OpenStack Images area. The IBM Cloud Manager with OpenStack administrator defines which images belong to which user profiles and therefore defines which VMware virtual server templates a user can access. IBM Cloud Manager with OpenStack discovers all virtual server templates regardless of which datacenter they reside in. There is currently no option to limit IBM Cloud Manager with OpenStack to specific datacenters.

Configuring global priority of an instance when relocating

IBM Cloud Manager with OpenStack enables you to choose whether you want your users to update the global priority of an instance when relocating the instance from host to host. *Instance priority* is the priority for relocating instances from one host to another host, when the instance is deployed in a pool. Depending on your site administration policies, you might not want users to change the priority of instances.

About this task

To configure the ability of updating instance priority, follow these steps:

Procedure

1. Open the `deployment.properties` file in the home directory.
2. To disable the ability to update instance priority, set the `com.ibm.cfs.deployments.update.priority` property to *false*. The default value of this property is *false*. If this property does not exist in the `deployment.properties` file, add it to the file.
3. Save the `deployment.properties` file and restart the IBM Cloud Manager with OpenStack server.

Configuring access to advanced deployment form

IBM Cloud Manager with OpenStack allows you to choose whether you want your users to see the advanced deployment form or just the basic deployment form.

About this task

The advanced deployment form allows a user or administrator to access all of the image properties that can be configured when an image is deployed. The basic deployment form contains only a subset of the image properties. Depending on your site administration policies, you may or may not want to show this advanced panel to the user.

To configure the visibility of the advanced deployment form, follow these steps:

Procedure

1. Open the `deployment.properties` file in the home directory.
2. Set the `deployment.advanced.form.enabled` property to *true*. This value enables the advanced deployment form so that it is displayed to all users. The default value of this property is *false*; users, by default, do not see the advanced deployment form.
3. Save the `deployment.properties` file and restart the IBM Cloud Manager with OpenStack server.

Results

Note: Administrators can also configure which advanced properties are shown on the basic deployment form. Use the web interface to configure those values for an image. When the `deployment.advanced.form.enabled` property is set to *true*, project owners can also configure which advanced properties are shown on the basic deployment form

Configuring the number and maximum size of additional storage

IBM Cloud Manager with OpenStack interface allows a user to add additional secondary disks to a virtual image using the Add Storage property. Adding additional secondary disks is also supported for VMware when you are deploying an image. IBM Cloud Manager with OpenStack provides a configurable property to set the maximum number of disks that can be attached to a virtual machine. This property applies during and after deployment of the virtual machine.

About this task

Note: This feature is not supported if the following statements are true:

- The virtual machine is deployed in the Shared Storage Pool.
- The instance that is being deployed is an IBM i instance.

To configure the secondary disk properties, follow these steps:

Procedure

1. Open the `deployment.properties` file in the home directory.
   - To set the maximum number of disks to use, edit the `com.ibm.cfs.vs.max.disks` property. The default value of this property is 3.
   - To set the maximum size in megabytes, edit the `com.ibm.cfs.vs.max.disksize` property. The default value of this property is `2048000`.
2. Save the `deployment.properties` file and restart the IBM Cloud Manager with OpenStack server.

Configuring images with VMware

This section describes some additional setup and configuration considerations when using the VMware cloud manager.

VMware considerations when deploying an image:

- IBM Cloud Manager with OpenStack requires the vCenter virtual switches and distributed virtual switch port groups to be defined before deploying instances. IBM Cloud Manager with OpenStack users and administrators are allowed to choose which virtual network the instance uses. IBM Cloud Manager with OpenStack supports either standard VMware vSwitches or distributed virtual switch port groups. IBM Cloud Manager with OpenStack supports the IBM DVS 5000V and the IBM SDN VE distributed virtual switches. If you are using a distributed virtual switch other than those, check the type in vCenter to ensure that it is supported.

  To check the type in vCenter, follow these steps:

  1. Browse to `https://<your vCenter>/mob`.
  2. Log in with an administrator account.
  3. Select the content link.
  4. Select the root folder.
  5. Select the data center that contains the third party distributed virtual switch.
  6. Select the network folder.
  7. Select the distributed switch that you want (the id starts with **dvs-**).

  The top of the page shows the managed object type of the switch. If the switch type is *VmwareDistributedVirtualSwitch* or *DistributedVirtualSwitch* then it is supported. If the distributed switch type is something other than the types listed, it is not supported; you receive an error when you deploy to a port group that uses that distributed switch.
- IBM Cloud Manager with OpenStack connects to the vCenter server by using a single user ID and password. It is recommended that this user has the vCenter administrator role. If you choose to use a different user ID, that user must have sufficient permissions to perform the operations on the virtual machines. The user ID must also have access to various virtual machine resources, such as networks and datastores.
- Do not change a virtual machines UUID in the vSphere Infrastructure Client. In some cases, such as manually moving a virtual machine, the vSphere Infrastructure Client asks if you want the UUID changed. When an instance is deployed, IBM Cloud Manager with OpenStack keeps the UUID of the virtual machine in its database and uses that UUID to find the virtual machine in vCenter, therefore you should not change the UUID.
- At some point, you might decide to migrate from IBM Cloud Manager with OpenStack to the IBM Tivoli Service Automation Manager (TSAM) product or other IBM cloud product. To ease the transition, it is highly recommended that you set up your Windows and Linux guest images as required by TSAM. Even if you have no plans to migrate, see Creating operating system image templates for VMware in the IBM Tivoli Service Automation Manager information center at http://pic.dhe.ibm.com/infocenter/tivihelp/v10r1/topic/com.ibm.tsam_7.2.4.1.doc/rdp/c_supported_os_vmware.html for more information about configuring your guest images.
- It is recommended that you install VMware tools on your guest operating systems before you make the virtual machine a template.

  Note: The VMware tools must be current with your version of VMware.
- If you are using the VMware Converter tool to import your virtual machines into vCenter, you should fully start the virtual server by using the VMware vSphere Client before you make it a template. This action allows vCenter to discover whether VMware tools are installed into the guest.
- All virtual machines in vCenter are shown as instances in IBM Cloud Manager with OpenStack. If you convert a virtual machine to a template, the status of that instance in IBM Cloud Manager with OpenStack changes to Unknown since the virtual machine is now a template. Do not delete this instance from IBM Cloud Manager with OpenStack since that deletes the underlying template on vCenter. If you want to make this instance not available for non-administrative users, use the Hide option instead.

- IBM Cloud Manager with OpenStack does not allow you to import vApps. If you want to enable users to deploy a single virtual server vApp, follow these steps:
  - Import the vApp by using vCenter
  - Modify the vApp properties as required by the application.
  - Convert the vApp to a template
- vApps with more than one virtual server are not supported.
- During the deployment of an image, IBM Cloud Manager with OpenStack uses the VMware API to apply customizations to the new virtual machine image. Therefore the VMware template image must be customizable by VMware. While VMware vCenter enables you to deploy a template without customizations, this option is not available in IBM Cloud Manager with OpenStack.
- When you deploy an image, IBM Cloud Manager with OpenStack instructs VMware to power on the newly deployed virtual machine. In some cases, depending on server loads, VMware might not power on the virtual machine. In this case, IBM Cloud Manager with OpenStack displays, by default, the instance in STOPPED state. Since the instance has been successfully cloned and configured, IBM Cloud Manager with OpenStack does not display an error and the instance can be started by the user. You can change this default behavior by using a deployment property that is described in the Waiting for a deployed virtual machine (VMware) topic.
- By default, IBM Cloud Manager with OpenStack reports a newly deployed instance in OK state as soon as VMware powers on the virtual machine. Depending on the guest operating system and image definition, it might be a few minutes before the virtual machine is completely up and running and can be used. In some cases, VMware might restart the virtual machine more than once during the customization process. You can change this default behavior by using a deployment property that is described in the Waiting for a deployed virtual machine (VMware) topic.
- IBM Cloud Manager with OpenStack allows the target of a deployment to be either a specific host, a host resource pool that is associated with a host, a cluster, or a resource pool that is associated with a cluster. If the cluster is a DRS enabled cluster, VMware chooses the appropriate host and therefore you cannot choose an individual host in that cluster. By default, IBM Cloud Manager with OpenStack is configured to randomly choose a host that is not associated with a cluster. If no hosts are available, you get an error when you deploy. You can change the default behavior by modifying the `deployment.properties` file as described in section “Configuring a deployment target” on page 136. However, it is recommended that an administrator configure each images target. For more information, see “Configuring global image deployment” on page 135.
- IPv6 is not supported when deploying to a VMware cloud. In order for deployments to not obtain an IPv6 address, the administrator needs to disable IPv6 in the VM template that is used for deploys.

Saving, restoring, and deleting virtual servers:

The IBM Cloud Manager with OpenStack allows users to save back up copies of their virtual server disks and configuration files. These copies can be restored later.

In addition, IBM Cloud Manager with OpenStack provides functions to allow users to view and delete their saved images. IBM Cloud Manager with OpenStack allows users to keep an administrator configured number of saved images. When the limit is reached, the system automatically deletes the oldest saved image. For information about how users perform save, restore, view, and delete operations, see the IBM Cloud Manager with OpenStack User Guide.

Deleting a virtual server

When a virtual server is deleted, all the saved images are deleted at the same time. There is no option to keep images beyond the life of the virtual server.

Approvals

The save and restore functions can be configured for approval control. This requires an IBM Cloud Manager with OpenStack administrator to first approve any save or restore request.

Note: If the virtual server is powered on, the save function powers down the virtual server before starting the save operation. If the approval process is enabled, the virtual server remains powered on until the administrator approves the save or restore request. There is no approval process for deleting a saved virtual server image. To enable the approval process, see "Managing approval policies" on page 190.

Authorization

Only the creator of an instance, an IBM Cloud Manager with OpenStack administrator, or the project owner is allowed to save, restore, or delete a virtual server image. Users within the same project are not allowed to perform save, restore, or delete operations on other user images within a project.

Notifications

The save, restore, and delete images functions log events to the IBM Cloud Manager with OpenStack event log. In addition, save image and restore image operations send email notifications, if the user configuration is enabled to receive email notifications.

For more information about email notifications, see "Configuring email notifications" on page 133.

Setting saved image limit:

By default, IBM Cloud Manager with OpenStack allows you to keep up to three saved virtual server images.

About this task

To change this limit, follow these steps:

Procedure

1. Open the `deployment.properties` file.
2. Update the `com.ibm.cfs.vs.max.backups` property.
   For example, to keep 10 saved virtual server images, change the property to the following setting:
   `com.ibm.cfs.vs.max.backups=10`
3. Save the `deployment.properties` file.
4. Restart the IBM Cloud Manager with OpenStack server.

VMware datastore assignment during deployment:

There is a new deployment property that you can set to select the target storage to be used when you deploy a virtual image. You can set the **`Target Storage`** property value to *datastores* or *datastore clusters* that are attached to the selected deployment target. If the selected deployment target is changed, the target storage value is updated to match what is available on the newly selected deployment target. IBM Cloud Manager with OpenStack always sets the default target storage value to use the default storage selection algorithm.

The default storage selection algorithm retrieves the list of datastores and datastore clusters that are associated with the deployment target. It then selects one of the datastores or datastore clusters that is

backed by a block device that has enough free space to contain the virtual machine disk sizes. If a block storage cannot be selected, then the appropriate NFS file storage, with the largest free space is chosen. If there is no available storage, the deployment fails.

You can specify a set of datastores and datastore clusters for the selection algorithm to exclude or include or both. Set the `com.ibm.cfs.cloud.vmware.enable.clone.template.properties` property in the `vmware.properties` file to `true`. To exclude datastores and datastore clusters from being selected, edit the `com.ibm.cfs.cloud.vmware.datastore.exclude.list` property and add a comma-separated list of datastore and datastore cluster names. To set the datastores and datastore clusters that can be selected, edit the `com.ibm.cfs.cloud.vmware.datastore.include.list` property and add a comma-separated list of datastore and datastore cluster names.

Note: By default, the selection algorithm includes all of the datastores and datastore clusters that are associated with the deployment target.

For example,

```
com.ibm.cfs.cloud.vmware.enable.clone.template.properties=true
com.ibm.cfs.cloud.vmware.datastore.exclude.list=dscluster3,localdisk4
com.ibm.cfs.cloud.vmware.datastore.include.list=localdisk1,dscluster2
```

To disable the algorithm, specify the datastore that you want to use in the `com.ibm.cfs.cloud.vmware.target.datastore.names` property.

For example:

```
com.ibm.cfs.cloud.vmware.enable.clone.template.properties=true
com.ibm.cfs.cloud.vmware.target.datastore.names=san,san,san
```

Each datastore that is listed is for a different virtual disk. The first entry is the datastore where the virtual server configuration files are located. The subsequent datastores in the list are for the virtual system hard disk devices. For example, if your virtual server has three disks you must specify four datastores in this list. These datastores can all be the same datastore or a combination of different datastores. If the first entry is for a datastore cluster, then the remaining entries are ignored. The datastore cluster is used for both the configuration files and disks. Datastore clusters are ignored when specified for subsequent entries in the list.

These `vmware.properties` file changes apply globally to IBM Cloud Manager with OpenStack and therefore to all deployment targets. Make sure that you specify datastores and datastore clusters that are available to all hosts that are the potential targets for every image. You cannot use the include list and datastore names properties, if the following is true:

- You have multiple cluster targets, each with its own set of storage, that is managed by the same vCenter.
- You want IBM Cloud Manager with OpenStack to target all the clusters.

In the datastore names property, you can specify different datastores for different disks if your virtual server templates have more than one disk.

Note:

- Datastore clusters are only available when you are using vSphere 5 Enterprise Plus edition. For more information about configuring and using datastore clusters, see the vSphere documentation.
- When you create the `vmware.properties` file, you must restart IBM Cloud Manager with OpenStack server. However, changes made to the property file after you restart the server are automatically updated.

Setting VMware user data during deployment:

When you deploy an image, you can set user data for the deployed virtual machine by using instance customization. The user data customization enables you to pass your own configuration information to the deployed virtual machine. If this customization fails during deployment, IBM Cloud Manager with OpenStack displays the instance in FAILED state.

About this task

IBM Cloud Manager with OpenStack does not require a specific content or format for the user data. Other than base64 decoding, IBM Cloud Manager with OpenStack passes the data unchanged to the guest. It is therefore the responsibility of the guest provider to process the user data. This usually requires writing and configuring a startup script to read this data from the CD device. IBM Cloud Manager with OpenStack does not provide nor recommend any particular script.

IBM Cloud Manager with OpenStack supports base64 encoded and plain text user data. The default value is set to base64 encoded. If decoding the user data fails, the user data is treated as plain text. IBM Cloud Manager with OpenStack passes the decoded user data to the deployed virtual machine through a CD backed by an ISO file. This gives the virtual machine access to the user data through the `user-data` file on one of its CD devices.

IBM Cloud Manager with OpenStack must be configured properly to create an ISO file that contains the user data. When IBM Cloud Manager with OpenStack is installed on Linux, the the **`mkisofs`** or **`genisoimage`** binary must exist in the `/usr/bin` directory.

In addition, you can also set the following properties in the `vmware.properties` file to control the user data:

com.ibm.cfs.cloud.vmware.user.data.file.name
The name of the file on the ISO that contains the actual user data. The default value is *user-data*.

com.ibm.cfs.cloud.vmware.user.data.iso.temp.path
The name of the path that is used to temporarily store ISO files on the IBM Cloud Manager with OpenStack server. The default value is the IBM Cloud Manager with OpenStack home directory. This path must end with a path separator, such as '/' on Linux.

Note:

- The ISO files that are created are managed by IBM Cloud Manager with OpenStack and not VMware. As a result, when you delete a virtual machine outside of IBM Cloud Manager with OpenStack, such as through the vSphere client interface, the ISO file that is created for the virtual machine is not removed.
- The ISO file of the virtual machine that contains the user data is not preserved with any saved virtual server images or capture instances that are created from the virtual machine.

Set secure access during deployment:

When you deploy a Linux image, you can set secure access to the deployed virtual machine by using instance customizations.

Using these optional customizations, you can set the password or SSH public key or both for a user on the guest operating system. If these customizations fail during deployment, IBM Cloud Manager with OpenStack displays the instance in FAILED state. You can set secure access for a root or a non-root user.

IBM Cloud Manager with OpenStack must be provided with the current root user name and password for the guest operating system to set secure access during deployment. In addition, IBM Cloud Manager

with OpenStack uses VMware guest operations to complete the customizations and there are special requirements for performing VMware guest operations. For more information, see “Requirements for VMware guest operation” on page 145.

When you set a password for a user, access the IBM Cloud Manager with OpenStack server by using a secure connection. This ensures that the password is encrypted when sent to the IBM Cloud Manager with OpenStack server. For more information, see “Configuring secure shell (SSH) communication” on page 125.

When you set an SSH public key for a user, the guest operating system must have OpenSSH installed and configured to take advantage of this customization. In addition, the SSH public key must be specified according to the OpenSSH authorized_keys file format. For more information about OpenSSH, see http://www.openssh.org/.

Resetting secure access during capture:

When you capture a Linux instance, the SSH public keys for a user are removed from the guest operating system if they were set by IBM Cloud Manager with OpenStack when the instance was deployed.

Note: For more information, see “Set secure access during deployment” on page 143.

Removing the SSH public keys prevents the keys from being available to instances deployed from the captured image. If the SSH public keys are unable to be removed during capture, IBM Cloud Manager with OpenStack displays warning messages in the image logs. In such cases, you must manually remove the SSH public keys for the user.

IBM Cloud Manager with OpenStack needs the current root user name and password for the guest operating system to reset secure access during capture. IBM Cloud Manager with OpenStack obtains this information from the virtual server credentials, which are initially set based on the instance customizations. The virtual server credentials can be reset later if changed in the guest operating system after deployment. For more information, see the *GET* and *PUT /instances/{id}/virtualServers/{id}/credentials* REST APIs in the IBM Cloud Manager with OpenStack Software Development Kit (SDK) Reference guide.

IBM Cloud Manager with OpenStack uses VMware guest operations to reset secure access during capture and there are special requirements for performing VMware guest operations.

Note: For more information, see “Requirements for VMware guest operation” on page 145.

Waiting for a deployed virtual machine (VMware):

IBM Cloud Manager with OpenStack provides a deployment property that enables users to specify whether to wait for a deployed virtual machine to be started and ready for use before it reports the newly deployed instance in OK state.

About this task

This deployment property option is labeled “Wait for the deployed virtual machine to be started and ready for use.” When this option is enabled, IBM Cloud Manager with OpenStack displays the instance in a FAILED state, if the deployed virtual machine is not started and ready for use in the allotted time. A failure can be caused by problems during the customization process (for example, specifying an incorrect Windows product key) or if the virtual machine cannot be powered on.

If VMware Tools is installed on the guest operating system, the virtual machine is considered started and ready for use when the virtual machine is powered on and the necessary network customizations are completed for it. If VMware Tools is not installed, then only the powered on state is checked.

You can configure the amount of time to wait for a deployed virtual machine to be started and ready for use by setting the *com.ibm.cfs.cloud.vmware.deployed.vm.start.wait.time* in the `vmware.properties` file. The time is in seconds and defaults to 2,700 (or 45 minutes). For example, *com.ibm.cfs.cloud.vmware.deployed.vm.start.wait.time=1800*

You can also configure the default value for this deployment property by setting *com.ibm.cfs.cloud.vmware.default.deployed.vm.start.wait* in the `vmware.properties` file. The default value is *false*, which disables the option so that IBM Cloud Manager with OpenStack reports a newly deployed instance in OK state as soon as VMware powers on the virtual machine. You can override the *com.ibm.cfs.cloud.vmware.default.deployed.vm.start.wait* setting when you configure or deploy an image. For example, *com.ibm.cfs.cloud.vmware.default.deployed.vm.start.wait=true*.

Note:

- The wait time starts after IBM Cloud Manager with OpenStack attempts to power on the deployed virtual machine.
- This deployment property applies globally to IBM Cloud Manager with OpenStack and therefore to all VMware deployments.
- When you create the `vmware.properties` file, you must restart IBM Cloud Manager with OpenStack server; however changes made to the property file after that are picked up automatically.
- During deployment, some optional instance customizations require the deployed virtual machine to be started and ready for use before the customizations can be completed. In such cases,IBM Cloud Manager with OpenStack waits for the deployed virtual machine regardless of the value that is specified for this deployment property.

  Note: For example, see “Set secure access during deployment” on page 143.

Requirements for VMware guest operation:

VMware guest operations are used to perform certain optional customizations of instances.

When you request such customizations, your request must meet the following requirements to successfully perform the VMware guest operations:

1. vCenter version 5.0 or later is required.
2. The vSphere host machine that is used for the instance must be at version 5.0 or later and IBM Cloud Manager with OpenStack must have network connectivity to it.
3. VMware tools must be installed and current on the guest operating system for the virtual machine.
4. VMware guest operations must be enabled for both the virtual machine and the host machine. They are enabled by default, but can be disabled.

Note: If it is necessary to connect to the host machine to complete the VMware guest operations, IBM Cloud Manager with OpenStack automatically accepts the security certificate for the host machine. The security certificate is stored in the `<host machine>.jks` file in your IBM Cloud Manager with OpenStack home directory.

Configuring shutdown of VMware instances:

In previous versions of IBM Cloud Manager with OpenStack stopping an active instance of VMware instantly powered off the running instance. The virtual machine was given no delay to allow it to perform its own shutdown process. IBM Cloud Manager with OpenStack now provides a 90-second delay for the system to complete the shutdown process. If by the end of 90 seconds the system is not shut down, IBM Cloud Manager with OpenStack forces the VMware instance to power down immediately.

You can configure the behavior of how VMware instances are shut down by modifying the following statements in `vmware.properties`:

com.ibm.vmware.client.shutdown.delay.in.milliseconds=90000
This property allows the VMware instance time to shut down before a power off is called. The default is 90000 milliseconds if the property is not specified. Setting this property to `0` (zero) prevents a shutdown from being called.

com.ibm.vmware.client.disable.save.image.shutdown=false
This property disables shutdown when save image is called. The default value is set to `false`, which allows shutdown to be called before save image. Specifying a value of `true` prevents a shutdown from being called on save image operations.

VMware limitations:

The following limitations apply to VMware and IBM Cloud Manager with OpenStack.

- The saved images are managed by IBM Cloud Manager with OpenStack and not VMware. Deleting an image outside of IBM Cloud Manager with OpenStack, such as through the vSphere client interface, does not remove the saved images.
- Properties that are defined in the `deployment.properties` file and the `vmware.properties` file are global to all users, instances, and images. There is no option to configure these options on a more granular level.
- If you have an image that is defined in IBM Cloud Manager with OpenStack and you rename the associated virtual machine template by using the vSphere Client, the name of the image in IBM Cloud Manager with OpenStack does not change. The IBM Cloud Manager with OpenStack image is still associated with the renamed template and can continue to be used. The image details page displays the renamed template name in the `Original name` field. You can manually change the name of the image by clicking the name of the image on the Images details page and pressing **Save**.
- If you have an image that is defined in IBM Cloud Manager with OpenStack and you convert it to a virtual machine by using the vSphere Client, the image in IBM Cloud Manager with OpenStack shows a state of `unknown`. This state is displayed because it is no longer a template on the VMware server; the conversion made it a virtual machine, which shows up as an IBM Cloud Manager with OpenStack instance. If the unknown IBM Cloud Manager with OpenStack image is no longer needed, it can be deleted. A deletion of the IBM Cloud Manager with OpenStack image does not affect the IBM Cloud Manager with OpenStack instance or virtual machine on the VMware server.
- In some cases, when you use special characters in names of VMware objects such as port group names and cluster names, the VMware API encodes these specials characters in a URL encoding scheme. For example a / character is encoded as a `%2f`. When the names are displayed in IBM Cloud Manager with OpenStack, the characters are not decoded. IBM Cloud Manager with OpenStack displays the encoded name. For example, if you have a cluster named `DRS/Cluster` it is displayed as `DRS%2f%Cluster`.
- IBM Cloud Manager with OpenStack creates and manages custom fields for internal use when using VMware virtual machines. The custom fields have a `"SKC_"` prefix and you should not modify or remove them using the vSphere client.

Configuring images with OpenStack

The following topics cover information to configure images in an OpenStack environment.

Customizing an OpenStack instance:

When an image is deployed, you might have to customize the resulting instance on startup to apply network configurations, login information, application settings, and so on, before the instance is ready for use.

About this task

IBM Cloud Manager with OpenStack provides a deployment property to enable OpenStack config drive support when deploying an image. This support is used to pass customizations (for example: server metadata, user data, personality files, and SSH keys) to an instance. The config drive can be accessed by

any guest operating system capable of mounting an ISO9960 file system. Images that are built with a recent version of the cloud-init software package, or similar software package such as IBM SmartCloud init, can automatically access and apply the supported customizations that are passed to the instance by the config drive.

Note: For more information about the cloud-init software package and the customizations that it supports, see CloudInit.

IBM Cloud Manager with OpenStack also supports the Pluggable Configuration Strategy feature added by IBM to OpenStack. This feature is similar to the config drive support in that it provides an instance with the necessary customizations. Like config drive, the image must be built with the correct software package for the configuration strategy to automatically access and apply the customizations. In particular, this feature provides support for Open Virtualization Format (OVF) or Microsoft Windows System Preparation (Sysprep) configuration. For more information, see the following resources:

- For information about OVF configuration, see Open Virtualization Format (OVF).
- For information about Sysprep configuration, see Sysprep Technical Reference.

Cloud-init software package:

IBM Cloud Manager with OpenStack supports the cloud-init software package.

When you deploy an image with config drive enabled, IBM Cloud Manager with OpenStack makes the instance customizations available to cloud-init using the OpenStack config drive support. The cloud-init software package can then access the config drive and apply the customizations to the instance. The following are some of the customizations made available through the config drive:

1. User data
2. Personality files
3. SSH key pair
4. Network adapter information (for static networks)
5. Password of the default user (the `password` option)

Using IBM Cloud Manager with OpenStack, you can enter the contents of the user data and personality files by using the deployment properties. You can enter the contents when you configure or deploy an image. The contents can be either base64 encoded or plain text. There are also deployment properties for the SSH key pair and network adapters that are based on the SSH key pairs and networks available. You can set the network adapters when you configure or deploy an image. However, the SSH key pair should not be set when you configure an image because OpenStack SSH key pairs are scoped to a user. Instead, the user who deploys the image should select an appropriate SSH key pair.

Note:

- Config drive is ignored when deploying PowerVC images. As a result, the instance customizations that are passed by the config drive are also ignored.
- The cloud-init `password` option sets the password for the default user, which is usually not the root user.
- For more information about the cloud-init software package and the customizations that it supports, see CloudInit.
- For more information about the OpenStack config drive support, see Config drive.

Configuration strategies:

IBM Cloud Manager with OpenStack supports the OVF and Sysprep types of pluggable configuration strategies.

When an image is deployed that has one of these configuration strategies, the OpenStack Pluggable Configuration Strategy feature determines the customizations made available to the instance and how they are made available. The appropriate software package (for the configuration strategy type) on the image is expected to access and apply the customizations. The customizations that are provided by OpenStack come from the following sources:

1. Server metadata that is provided by OpenStack itself.
2. Server metadata that is provided by the user deploying the image.

The following server metadata is provided by OpenStack:

```
server.admin_password               Random administrator password generated by OpenStack.
server.hostname                     The hostname for the instance.
server.domainname                   Domain name from the dhcp_domain configuration option.
server.dns-client.pri_dns           Primary DNS server IP address.
server.dns-client.sec_dns           Secondary DNS server IP address.
server.dns-client.dns_list          Space separated list of DNS server IPs.
server.network.[n].mac              Mac address for network interface number n.
server.network.[n].mac_alt          Mac address formatted with '-' rather than ':'.
server.network.[n].slotnumber       Slot number for network interface number n.
                                    Defined as the decimal value of the last two digits
                                    of the mac address.
server.network.[n].[v4|v6].address  IPv4 or IPv6 address for network interface number n.
server.network.[n].[v4|v6].netmask  IPv4 or IPv6 netmask for network interface number n.
server.network.[n].[v4|v6].cidr     IPv4 or IPv6 address and netmask in CIDR notation
                                    for network interface number n.
server.network.[n].[v4|v6].gateway  IPv4 or IPv6 gateway for network interface number n.
server.network.[n].v4.use_dhcp      'true' if the network uses DHCP.
```

Server metadata that is provided by the user during deployment are prefixed with `'server.metadata'`.

Creating a configuration strategy:

A complete and accurate OVF or Sysprep configuration strategy is important to ensure that an image can be deployed and customized properly. A poor configuration strategy can cause the deployment to fail or prevent the instance from being customized.

About this task

A configuration strategy consists of the following parts:

- Type
- Template
- Mapping
- User metadata

For information about how to add, update or delete the configuration strategy of an image, see "Updating an image configuration strategy (OpenStack only)" on page 182.

Type The type is required and can either be ovf or sysprep.

Template
The template is required. When you are using an ovf configuration strategy type, this contains the OVF descriptor for the image. When you are using a sysprep configuration strategy type, this contains the template `unattend.xml` file for the image.

Mapping
The mapping is required. It defines how to map the server metadata that is provided by both OpenStack and the user deploying the image to the appropriate elements/parts of the template. The mapping is a JavaScript Object Notation (JSON) array of objects, where each object has a source representing the server metadata to map to the target element/part in the template. For example:

```
[
  {
    "source": "server.network.1.v4.address",
    "target": "com.ibm.vsae.2_1.network-interface.ipaddr"
  },

  {
    "source": "server.network.1.v4.netmask",
    "target": "com.ibm.vsae.2_1.network-interface.netmask"
  },

  {
    "source": "server.network.1.v4.gateway",
    "target": "com.ibm.vsae.2_1.network-interface.gateway"
  },

  {
    "source": "server.hostname",
    "target": "com.ibm.vsae.2_1.system-host.hostname"
  },

  {
    "source": "server.domainname",
    "target": "com.ibm.vsae.2_1.system-host.domainname"
  },

  {
    "source": "server.dns-client.pri_dns",
    "target": "com.ibm.vsae.2_1.dns-client.pri_dns"
  },

  {
    "source": "server.metadata.username",
    "target": "com.ibm.vsae.2_1.system-user.username"
  },

  {
    "source": "server.metadata.system.password",
    "target": "com.ibm.vsae.2_1.system-user.password"
  }
]
```

IBM Cloud Manager with OpenStack uses the mapping to create additional deployment properties for the image. Every object in the mapping with a source prefix of `'server.metadata.'` is added to the configurable deployment properties for the image. Doing so allows such properties to be customized by the user when the image is deployed. For more information about defining the mapping, see "OVF configuration strategy" on page 151 and "Sysprep Configuration Strategy" on page 152 topics.

Note:

- The same source can be mapped to multiple targets. To do this, you must define a separate source/target object in the JSON array for each mapping.
- An empty mapping (for example, []) must be used only for testing purposes since all deploys will use the same template and thus have the same customizations applied.
- When you define a source mapping name with the `'server.metadata.'` prefix, avoid using `'.'` in the suffix portion of the name.

User Metadata

The user metadata is optional. It determines how IBM Cloud Manager with OpenStack defines, displays, and processes the configurable deployment properties created based on the mapping. If no user metadata is provided for a mapping, a basic string deployment property is used.

Defining detailed user metadata helps users properly configure and deploy the image. The user metadata is a JSON array of objects where each object might contain the following:

1. name
2. type
3. subtype
4. description
5. required
6. min
7. max
8. allowed_values
9. default_value

For example:

```
[
 {
  "name": "system.username",
  "type": "STRING",
  "description": "System user name",
  "required": "true"
 },
 {
  "name": "system.password",
  "type": "STRING",
  "subtype": "PASSWORD",
  "description": "System user password hash",
  "required": "true"
 }
]
```

The name string is required. The name corresponds to the mapping source without the `'server.metadata.'` prefix.

The type string is optional. It is the native type of the deployment property (INT, LONG, FLOAT, BOOLEAN, or STRING). The default is *STRING*.

The subtype string is optional. It is a more descriptive type to allow for early validation of the deployment property. A STRING type can have the following subtypes: *IPV4_ADDRESS, DOMAIN_NAME, DOMAIN_NAMES, IPV4_SUBNET_MASK, HOST_NAME and PASSWORD.* A BOOLEAN type can have the following subtypes: *DHCP_FLAG* and *DNS_FLAG*. The default is no specific subtype.

The description string is an optional description of the deployment property. If no description is provided, the name is used for the description.

The required flag is optional. It is a flag indicating whether the deployment property is required when deploying the image. The default is false.

The min and max strings are optional. They provide minimum and maximum boundaries for *INT, LONG, FLOAT* and *STRING* type deployment properties. The default is no boundaries.

The *allowed_values* string is optional. It is a comma-separated list of allowed values for the deployment property. When you specify a list of allowed values, also provide the *default_value* and ensure that the allowed values are valid for the type. The default is any allowed values corresponding to the type.

The *default_value* string is optional. It is the default value for the deployment property. The default value should be valid for the type. If no default value is provided, a value must be explicitly set by the user in order for the deployment property to be used when deploying the image.

OVF configuration strategy:

The OVF configuration strategy is an example of an OpenStack Pluggable Configuration Strategy. It is designed for use with OVF configuration, which is a way to package and provide configuration options for an image.

About this task

The OVF configuration strategy supports OVF version 1.1. For more information about OVF standards, see Open Virtualization Format (OVF).

The OVF configuration strategy passes the configuration options in an `ovf-env.xml` file in a disk that is presented to the guest system. It is expected that an activation engine, such as IBM VSAE, embedded in the image mounts the drive, read the `ovf-env.xml`, and apply the customizations when an instance deployed from the image starts.

The `ovf-env.xml` file is created based on the default values in the OVF descriptor (that is, the template in the configuration strategy) and the configuration options that are mapped using the mapping that is specified in the configuration strategy.

To know what mappings to specify in the configuration strategy, you must know the properties that the image expects in the `ovf-env.xml` file. The properties that the image expects in the `ovf-env.xml` are specified in the OVF descriptor's ProductSection elements, as documented in the OVF 1.1 specification, section 9.5. Here is an example ProductSection from an OVF descriptor:

```
<ovf:ProductSection ovf:class="com.ibm.vsae.2_1.network-interface">
 <ovf:Info>System network interface configuration</ovf:Info>
 <ovf:Property ovf:key="ipaddr" ovf:type="string" ovf:userConfigurable="true"
  ovf:value="192.168.71.129">
  <ovf:Description/>
  <ovf:Label>IP address</ovf:Label>
 </ovf:Property>
</ovf:ProductSection>
```

Using the previous example, the image can have a property *com.ibm.vsae.2_1.network-interface.ipaddr* that defaults to *192.168.71.129*. You might want to have the IP address set to the value that OpenStack assigns to it, which is given in the *server.network.1.v4.address* server metadata. To do this, you would create the following mapping:

```
{
 "source": "server.network.1.v4.address",
 "target": "com.ibm.vsae.2_1.network-interface.ipaddr"
}
```

Here is another example ProductSection:

```
<ovf:ProductSection ovf:class="com.ibm.vsae.2_1.ntp-client">
 <ovf:Info>activates the openntp client</ovf:Info>
 <ovf:Property ovf:key="ntp-server" ovf:type="string" ovf:userConfigurable="true"
  ovf:value="0.pool.ntp.org">
  <ovf:Description>Ntp server</ovf:Description>
  <ovf:Label>Ntp server</ovf:Label>
 </ovf:Property>
</ovf:ProductSection>
```

Using the previous example, there is no OpenStack provided server metadata that contains the NTP server's IP address. Therefore, if you want users to be able to override the value when they deploy the image, you would create the following `'server.metadata.'` mapping:

```
{
 "source": "server.metadata.ntp-server",
 "target": "com.ibm.vsae.2_1.ntp-client.ntp-server"
}
```

The OVF configuration strategy also supports using the wildcard character * to map multiple server metadata items that are provided by OpenStack using a single source/target mapping. When the wildcard character is used, the system matches existing configuration properties against the source to generate targets for the wildcard matches. This support is useful when you need to dynamically add network adapters at the time you deploy an image. Here is an example mapping that uses the wildcard character:

```
{
  "source": "server.network.*.v4.address",
  "target": "com.ibm.vsae.2_1.network-interface.ipaddr.*"
}

If the following server metadata items were provided by OpenStack when deploying an image:
server.network.1.v4.address = 192.168.1.101
server.network.2.v4.address = 192.168.1.102
server.network.3.v4.address = 192.168.1.103

Then the following mapping would be generated when deploying the image:
[
  {
    "source": "server.network.1.v4.address,
    "target": "com.ibm.vsae.2_1.network-interface.ipaddr.1"
  },
  {
    "source": "server.network.2.v4.address,
    "target": "com.ibm.vsae.2_1.network-interface.ipaddr.2"
  },
  {
    "source": "server.network.3.v4.address,
    "target": "com.ibm.vsae.2_1.network-interface.ipaddr.3"
  }
]
```

The wildcard character support has the following restrictions:

1. The wildcard character replaces a string of decimal digits only that translates internally to regular expression *'d+'*.
2. The mapping source must have only one wildcard character.
3. The mapping target must have at least one wildcard character. If more than one wildcard character is used, they are all replaced.

Note: The OVF configuration strategy has the following limitations:

- The OVF configuration strategy support is only for image activation (that is, ProductSection elements) and does not support actions such as adding disks to the image.
- After activation is complete, OpenStack does not automatically detach the disk drive that contains the `ovf-env.xml` file.
- Extensions to IBM VSAE might be required to support network configurations with both IPv4 and IPv6 addresses.

Sysprep Configuration Strategy:

The Sysprep configuration strategy is an example of an OpenStack Pluggable Configuration Strategy. It is designed for use with Microsoft Windows System Preparation (Sysprep) configuration, which allows customizing many aspects of a Windows system as it starts. For more information about Sysprep, see Sysprep Technical Reference.

The Sysprep configuration strategy passes the image configuration options in an unattend.xml file in a CD-ROM device that is presented to the guest system. Before adding the Sysprep configuration strategy to an image, it is expected that the image is ready for Sysprep, and that it runs Sysprep to read the unattend.xml file, and apply the customizations when starting an instance deployed from the image.

The unattend.xml file is created based on the default values in the template unattend.xml file (that is, the template in the configuration strategy) and the configuration options that are mapped using the mapping that is specified in the configuration strategy. To know what mappings to specify in the configuration strategy, you must know the properties that the image expects in the unattend.xml file.

The format of the target values in the configuration strategy mapping is as follows:

- To identify an element, the format is an XPATH. In this case, the contents of the element are replaced with the value of the source configuration property.
- To identify an attribute, the format is like XPATH@attribute-name. In this case, the attribute in the element is set to the value of the source configuration property.

For documentation on the XPATH format, see the python documentation. If the path identifies more than one element, only one of the elements are the target element. If the path does not identify an element in the template the boot of the instance fails with an error.

Example template unattend.xml file:

```
<?xml version="1.0" encoding="utf-8"?>
<unattend xmlns="urn:schemas-microsoft-com:unattend">
 <settings pass="oobeSystem">
  <component name="Microsoft-Windows-Shell-Setup" processorArchitecture="amd64"
 xmlns:wcm="http://schemas.microsoft.com/WMIConfig/2002/State">
   <Display>
    <ColorDepth>16</ColorDepth>
    <HorizontalResolution>1024</HorizontalResolution>
    <RefreshRate>60</RefreshRate>
    <VerticalResolution>768</VerticalResolution>
   </Display>
   <RegisteredOrganization>OpenStack</RegisteredOrganization>
   <OOBE>
    <HideEULAPage>true</HideEULAPage>
    <NetworkLocation>Cluster</NetworkLocation>
    <ProtectYourPC>1</ProtectYourPC>
    <SkipMachineOOBE>true</SkipMachineOOBE>
    <SkipUserOOBE>true</SkipUserOOBE>
   </OOBE>
  </component>
 </settings>
</unattend>
```

Using the above example, the following mapping target would identify the ColorDepth element: *.//{urn:schemas-microsoft-com:unattend}Display/{urn:schemas-microsoft-com:unattend}ColorDepth*

Using the above example, the following mapping target would identify the processorArchitecture attribute in the component element: *.//{urn:schemas-microsoft-com:unattend}component[@name='Microsoft-Windows-Shell Setup']@processorArchitecture*

Considerations for capturing an OpenStack instance:

Consider the following information before capturing an OpenStack instance.

When you capture an instance that was deployed by using a pluggable configuration strategy (OVF or Sysprep), the configuration strategy is copied from the source image for the instance to the new image. If the source image for the instance was deleted, the configuration strategy cannot be copied to the new image. As a result, you need to manually add the configuration strategy to the new image.

After an instance is deployed using a configuration strategy, the activation software (for example, Sysprep or IBM VSAE) has run and applied the configuration. Therefore, you might have to perform more actions when you capture the instance. If you want the activation software to run again when you deploy the

image that was created, the activation software must be reset. If you do not have to run the activation software again, you can delete the configuration strategy from the image that was created.

For more information about the reset requirements and other capture prerequisites, see the documentation for the applicable activation software. For more information about pluggable configuration strategies, see “Configuration strategies” on page 147.

Configuring PowerVC images

Use the following information to configure PowerVC images with IBM Cloud Manager with OpenStack.

About this task

Placement policy

IBM Cloud Manager with OpenStack provides a deployment property to control the placement of a deployed PowerVC virtual machine. The placement of the virtual machine can either be handled by the OpenStack scheduler or the PowerVC scheduler. This deployment property option is labeled **Use the PowerVC scheduler to place the deployment**. Select the appropriate option for your environment.

- When this option is enabled, the PowerVC scheduler is used.
- When this option is disabled, the OpenStack scheduler is used.

Note: This option is enabled by default.

Network adapter configuration

OpenStack images to be deployed only to PowerVC support VLAN type networks with a physical network name equal to "default". Such networks must also exist in the PowerVC environment. In addition, if the network supports both IPv4 and IPv6 addresses, only the IPv4 addresses are used by PowerVC.

Flavor storage size

PowerVC does not support resizing during image deployment. However, you can resize the virtual machine after it is successfully deployed. As a result, regardless of the OpenStack flavor that is selected, PowerVC uses the size of the image as the storage size of the deployed PowerVC virtual machine. For more information about PowerVC flavors, see the PowerVC Information Center.

Config drive and virtual machine customizations

PowerVC does not support config drive and the related virtual machine customizations during image deployment. If specified, these configurations are ignored. For more information about the cloud-init software package and the customizations that it supports, see “Cloud-init software package” on page 147.

Boot Volume Storage Template

By default, when you deploy an image, the boot volume is created using the default storage template, as configured on the PowerVC server. If you would like the boot volume to use a different storage template, you need to add the boot volume type property to the extra specifications of the flavor that is used to deploy the image. Edit the flavor and add the following extra specifications.

- Set the property key *powervm:boot_volume_type*
- Set the value of the property to the ID of the volume type you want to use.

You can find the volume type id by running the following OpenStack command: `cinder type-list`

Note: The volume type ID that is specified must be a PowerVC volume type. When you run a command in the IBM Cloud Manager with OpenStack appliance, PowerVC volume type names are prefixed with a *pvc:*. For more information about flavors, see Flavors.

Related reference:

PowerVC Standard Edition

Nova placement policy

Configuring instance resize timeout

IBM Cloud Manager with OpenStack allows the administrator to configure a timeout for an instance resize action. This is optional, but can be configured in the event that the cloud manager does not respond in a reasonable timeframe.

About this task

To configure the instance resize timeout, follow these steps:

Procedure

1. Open the `deployment.properties` file in the home directory.
2. Set the `com.ibm.cfs.deployments.resize.timeout` property to the time in minutes to wait for an instance resize action to complete in the cloud. The default value is 10 minutes.
3. Save the `deployment.properties` file and restart the IBM Cloud Manager with OpenStack server.

Identifying expired instances

IBM Cloud Manager with OpenStack has a configurable task that allows administrators to identify expired instances and how long to keep them after they have expired.

About this task

To set how often expired instances are identified and how long they are kept, follow these steps:

Procedure

1. Open the `deployment.properties` from your home directory.
2. Configure each property in the file.
   a. `com.ibm.cfs.expired.deployments.finder.interval=10`

      When the expiration date is reached, the deployment property is put into an `EXPIRED` state and the virtual machine is stopped in the deployment. The `com.ibm.cfs.expired.deployments.finder.interval` defines how often IBM Cloud Manager with OpenStack identifies expired deployments. This interval is set in seconds.

   b. `com.ibm.cfs.deployments.expired.delete.interval=1`

      This property defines the interval for deleting instances that have been identified as expired. The unit is hour. If not set or valid, the default value is 24 hours.

      This property is also defined by the length of the grace period set when you create an expiration policy for a cloud or a project. For more information, see “Managing expiration policies” on page 194.

Configuring multiple instances for a single deployment

You can deploy multiple instances through a single deployment.

About this task

If you enable the multiple instance deployment function, a user can deploy multiple instances through a single deployment. The deployed instances use the deployment name as the prefix of each single instance. The new names also use **-x** as the suffix, where **x** is the index of that instance.

Procedure

1. To enable or disable this feature, set the following property value within the `deployment.properties` file that is in the `.SCE41` directory:
   - **`com.ibm.cfs.deployments.multi.enabled=true`** to enable the function.
   - **`com.ibm.cfs.deployments.multi.enabled=false`** to disable the function.

   Note: By default, this feature is enabled.
2. To control the maximum number of instances that a user is allowed to deploy at one time, set the following property value within the `deployment.properties` file:

   `com.ibm.cfs.deployments.multi.max.value=5`

   Note:
   - The default value is 5.
   - If this value is set too high, it might overload the connected cloud.

Configuring logging

Log files are automatically saved by the IBM Cloud Manager with OpenStack self-service portal. You can configure the default number of log files that are saved and the types of messages that are logged.

About this task

By default, the self-service portal saves 9 log files of 50 MB each. These defaults can be modified in the `logging.properties` file in the home directory.

To change the default logging options:

Procedure

Open the `logging.properties` file in the home directory.

- To change the number of log files that are saved, set the `java.util.logging.FileHandler.count` property to the number of log files that you to save. The default is `9`.
- To change the types of messages that are saved, set the `java.util.logging.FileHandler.level` property to the level of messages that you want to receive. The default is `INFO`.

  The types of messages that are logged in the log file are informational, warning, and error messages. Use the debugging messages only for troubleshooting and debugging purposes because performance can be impacted by excessive logging.

What to do next

Modifying the `logging.properties` file requires restarting the self-service portal server to pick up the changes.

For more information about logging, see Chapter 10, “Troubleshooting and support for IBM Cloud Manager with OpenStack,” on page 243.

Configuring a network

The IBM Cloud Manager with OpenStack self-service portal provides a convenient way to manage and apply network settings by using network configurations. Network configurations are a group of network settings for a particular environment, typically a virtual network. These settings can be managed as a single entity and applied to image configurations or instance deployment settings.

For example, suppose that a cloud environment contains two virtual networks applicable to instance deployment: a public and a private virtual network. In this case, an administrator might create two network configurations, one for the public and one for the private. In the public configuration, the administrator would specify all the public network settings such as primary DNS, secondary DNS, and primary gateway. The same would be done for the private network configuration. After the configurations are created, the administrator can configure the images to use the appropriate network configuration. This action saves time by not requiring the administrator to specify each network setting in each image. It also allows an easier way to manage the network settings on a virtual network.

While the actual settings specified in a configuration are tailored to a specific environment, the network configurations themselves are a superset of all network settings regardless of image, operating system, or cloud management system. Therefore, all settings that are specified in a configuration are applicable. For example, the primary and secondary WINS settings of a network configuration are only applicable to Windows based images. So when you create a configuration for an image that is not using Windows, these values are not needed and can be left blank.

Note: With the self-service portal, you can specify the network configuration for a cloud. The self-service portal displays only the fields that are applicable for that cloud. Before you can create an OpenStack network configuration, you must select an existing OpenStack cloud.

When network configuration settings are applied to either an image configuration or during an advanced instance deployment, their individual settings can be overridden or manually specified, if wanted.

Note: You cannot override or manually specify OpenStack network configuration settings.

You can modify your network connections through the self-service portal or though the property files in the home directory. For more information about modifying your network connections through the self-service portal, see “Managing network configurations” on page 201.

Note: You cannot use property files to specify OpenStack network configuration settings. You must use the self-service portal.

To modify your network connections from the home directory, create a `.properties` file and save it to your home directory. The name of these files should be prefixed with `networkConfiguration` followed by an arbitrary suffix and the `.properties` file extension, similar to `networkConfiguration.properties`, `networkConfiguration-vlan1.properties`, or `networkConfiguration3.properties`.

Each property file contains a group of network setting. For example, assume that there is a file that is named `networkConfiguration.properties` in the home directory, which containing the following settings:

```
name=VLAN1
dns1=9.10.244.100
dns2=9.10.244.200
gateway1=9.5.40.1
gateway2=9.5.40.2
domain=mydomain.company.com
subnet=255.255.252.0
networkId=[Network 1]=hostVnet:ETHERNET0/1
useDHCP=false
hostnamePrefix=sce
computerNamePrefix=sce
workgroup=workgroup
description=default network configuration
9.5.42.250
9.5.42.251
9.5.43.23
```

Note: When you use a brocade switch, you must configure a host name prefix in the `networkConfiguration.properties` file: `hostnamePrefix=sce`.

When the self-service portal starts, the network configuration named "VLAN1" is added to the network configuration list.

(VMware only) In the VWware environment, the value of the *Network ID* field is the name of a VMware standard switch network, port group name, or distributed port group. A typical VMware network ID is *VM Network*. This value is used to assign the virtual network adapter to the VMware network during a deployment. The rest of the values in the network configuration should be appropriate for that network. The network configuration displays all available port groups and distributed port groups. Not all port groups or distributed port groups might be available on all target hosts. Validation of this field occurs only at deployment time when the actual deployment target is known. If the selected port group or distributed switch is not available on the selected target host, then an error occurs and the instance deployment fails.

Configuring billing

The IBM Cloud Manager with OpenStack self-service portal has a configurable billing and accounting interface. The interface allows the self-service portal to monitor resource use and create subsequent billing to the self-service portal user accounts for the usage.

For more information about accounts, see “Managing accounts” on page 213.

Configuring billing

To enable billing, edit the `billing.properties` file and define what action to take when an account becomes delinquent. Also, set the time intervals to determine accounts that are delinquent or at their account balance threshold.

About this task

Important: For billing to work, you must also enable metering. Account bills are generated based on metering results.

To configure billing, follow these steps:

Procedure

1. Open the `billing.properties` file in the home directory.
2. Configure each property in the file.

   `com.ibm.cfs.billing.enabled=true`
   Defines whether to enable the billing and accounting functionality in IBM Cloud Manager with OpenStack. True enables and false disables billing and accounting.

   `com.ibm.cfs.billing.delinquency.policy= com.ibm.cfs.services.billing.policies.shutdown`
   Determines the action IBM Cloud Manager with OpenStack takes against existing instances and volumes when an account becomes delinquent. Possible values are as follows:

Table 26.

Value	Description
com.ibm.cfs.services.billing.policies.destroy	This value destroy all deployments and detaches all volumes for the delinquent account.
com.ibm.cfs.services.billing.policies.shutdown	This value suspends all deployments for the delinquent account.
com.ibm.cfs.services.billing.policies.do.nothing	This value ignores delinquent accounts.
com.ibm.cfs.services.billing.policies.detach	This value detaches all volumes for the delinquent account.

Table 26. (continued)

Value	Description
com.ibm.cfs.services.billing.policies.delete	This value deletes all volumes for the delinquent account.
com.ibm.cfs.services.billing.policies.shutdowndetach	This value suspends all deployments and detaches all volumes for the delinquent account.
com.ibm.cfs.services.billing.policies.shutdowndelete	This value suspends all deployments and deletes all volumes for the delinquent account.
com.ibm.cfs.services.billing.policies.destroydelete	This value destroy all deployments and deletes all volumes for the delinquent account.

com.ibm.cfs.billing.delinquency.finder.interval=120
This property represents the number of seconds to wait before running a job that examines each account to determine whether the account is delinquent.

com.ibm.cfs.billing.account.balance.threshold.interval= 24
This property represents the number of hours to wait before running a job to find accounts that are at their account balance threshold. The default value of this property is 24 hours or 1 day.

Note:

- The `billing.properties` file is not configurable through the web user interface.
- For PowerVC support, the PowerVC driver file, `/etc/powervc/powervc.conf`, has staging user (staging_user) and project (staging_project_name) properties. These properties control which OpenStack user and project owns instances that are synchronized from PowerVC. The default configuration uses the IBM Cloud Manager with OpenStack administrator and the public project. An instance must be owned by a user in OpenStack, so when you enable billing, the default owner is also billed for those instances in the PowerVC cloud.

What to do next

After you enable billing, ensure that you also enable metering. For more information, see "Configuring metering" on page 162.

Configuring billing details

IBM Cloud Manager with OpenStack can produce charges that are billed back to users when using a specific cloud resource, such as an instance.

About this task

IBM Cloud Manager with OpenStack currently has the following configurable products:

Table 27. Configurable products by cloud type

OpenStack configurable products	VMware configurable products
• Processor • Memory • Disks • Volume	• Processor • Memory • Disks

A cloud product might be something similar to processor by the hour, 1 GB of RAM per day, a fixed rate charge for a running VM, or 20 GB of active disks per day. IBM Cloud Manager with OpenStack loads those cloud products into a product catalog. System events, such as deploying an instance, can cause the

creation of a bill with one or more charges from one or more cloud products. IBM Cloud Manager with OpenStack automatically deducts money from the account to which the instance owner belongs.

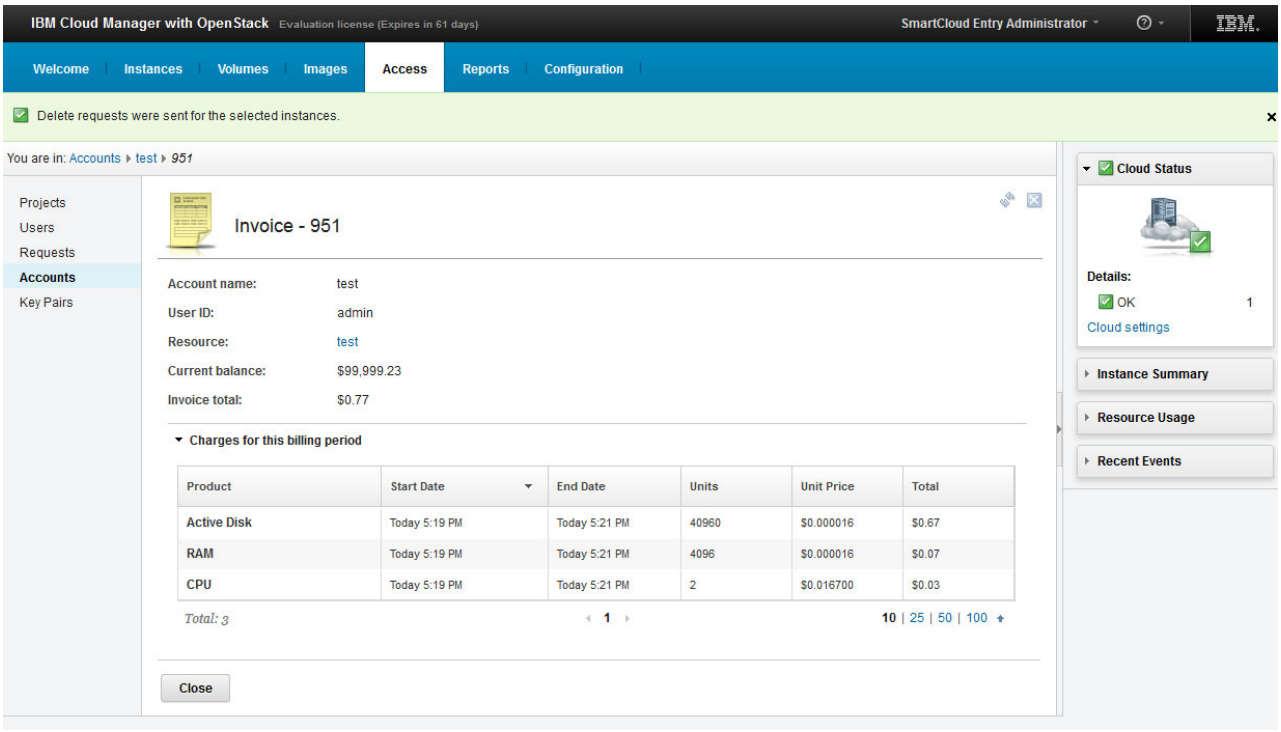

Figure 1. Sample billing account summary

The settings for product price per interval time are configurable. To configure product pricing information, follow these steps:

Procedure

1. Open the `products` directory in the home directory to locate the product configuration files. The cloud type is added to the file name as a prefix as shown in the following examples:

   For OpenStack:

   - `Openstack_CPU.xml`
   - `Openstack_RAM.xml`
   - `Openstack_Disk.xml`.
   - `Openstack_Volume.xml`

   For VMware:

   - `VMware_CPU.xml`
   - `VMware_RAM.xml`
   - `VMware_Disk.xml`.

   The file that you configure depends on the cloud and product type.
2. Configure processor price in the *`cloud_type`*`_cpu.xml`.

   ```
   <pricing currency="USD" interval="3600" price="1.000"/>
   ```

   This property specifies that the default IBM Cloud Manager with OpenStack collector collects charges on virtual servers by using the number of processors that are assigned to them at a rate of $1.00 per hour. Collecting at an interval less than the actual described rate (for example, hours instead of days)

enables users to get a more accurate approximation of their actual charges. Having an accurate look at the charges might be important for accounting purposes or in situations where account credit is limited.

3. Configure Memory price in *cloud_type*_ram.xml.

   ```
   <pricing currency="USD" interval="3600" price="0.000976563"/>
   ```

   This property specifies that the default IBM Cloud Manager with OpenStack collector collects charges on virtual machines by using the number of bytes of RAM assigned to them at a rate of $0.000976563 per MB per hour, which is about $1.00 per hour per GB.

4. Configure disks in *cloud_type*_disk.xml.

   ```
   <cloudProduct id="com.ibm.cfs.cloud.vmc.products.storage">
   <name>Active Disk</name>
   <description>The amount of total disk storage in MB used in a workload
   per minute.</description>
   <!-- $1.627604167E-5 per megabyte per minute -->
   <pricing currency="USD" interval="60" price="1.627604167E-5"/>
   </cloudProduct>
   ```

   These properties specify that the default IBM Cloud Manager with OpenStack collector collects charges on virtual machines by using the disks that are assigned to them at a rate of $0.000976563 per MB per hour, which is about $1.00 per hour per GB.

   The <name> and <description> can also be overridden from the IBM Cloud Manager with OpenStack defaults by specifying different values.

   Note: The currency for all configurable products must be consistent. For example, set US dollar (USD) for both or Chinese Yuan (CNY) for both. Using inconsistent currencies causes incorrect product charges.

5. For OpenStack clouds, configure volume price in Openstack_Volume.xml.

   ```
   <cloudProduct id="com.ibm.cfs.cloud.openstack.products.volume"
     cloudType="Openstack" objectType="Volume">
   <name>Volume</name>
   <description>The amount of a volume in MB per minute.</description>
   <!-- $1.627604167E-5 per megabyte per minute -->
   <pricing currency="USD" interval="60" price="1.627604167E-5"/>
   </cloudProduct>
   ```

   These properties specify that the default IBM Cloud Manager with OpenStack collector collects charges on volumes at a rate of $1.627604167E-5 per MB per hour, which is about $1.00 per hour per GB. The <name> and <description> can also be overridden from the IBM Cloud Manager with OpenStack defaults by specifying different values.

   Note: The currency for all configurable products must be consistent. For example, set US dollar (USD) for both or Chinese Yuan (CNY) for both. Using inconsistent currencies causes incorrect product charges.

Results

After you configure account billing, you can view account billing information in the IBM Cloud Manager with OpenStack interface.

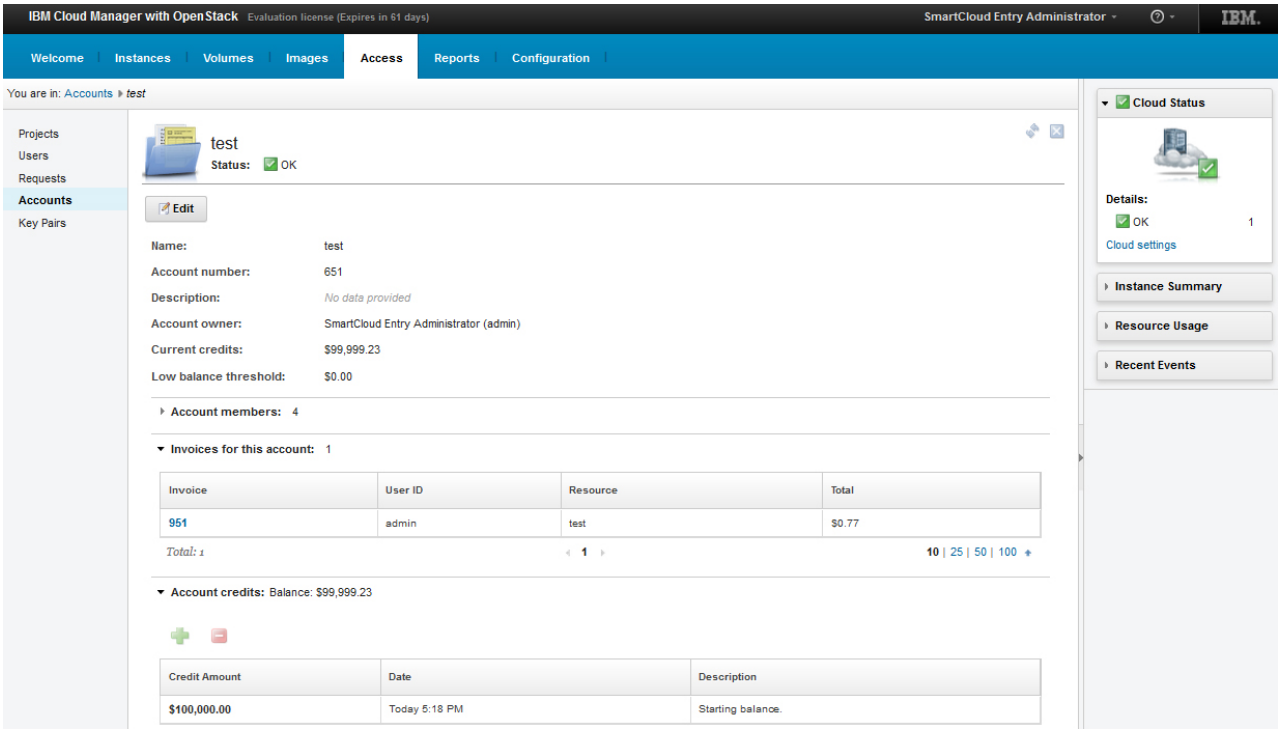

Figure 2. Sample billing account settings

Configuring metering

The IBM Cloud Manager with OpenStack self-service portal has a configurable metering framework that enables the self-service portal to record and present metering data.

About this task

You can download metering data files through the metering data API. To enable metering with the self-service portal, configure the following properties:

Procedure

1. Open the `metering.properties` file in the home directory.
2. Configure the property `com.ibm.cfs.metering.enabled=true` to enable the metering function within the self-service portal. The default value for this property is false.
3. Configure the property `com.ibm.cfs.metering.interval=<time in minutes>` where *<time in minutes>* is the time in minutes between each metering record synchronization. The default value is 1441, or every day. If you want a more frequent synchronization, you can decrease this value.
4. Configure the property `com.ibm.cfs.metering.data.path = cfshome/metricsdata/`. This property allows the administrator to configure the storage system where the metrics data is located. The default location is the self-service portal `home directory/metricsdata/` if not specified.
5. Configure the property `com.ibm.cfs.metering.data.export.interval = <interval time, hour as unit>`. This property is used for how often to export the metering data to file. The default value is 1 hour.
6. Configure the property `com.ibm.cfs.metering.data.expired.days = <day as unit>`. This property is used to set the number of days that the metering data is expired. The default value is 370 days.
7. Configure the property `com.ibm.cfs.statistics.interval = <interval time, seconds as unit>` This property is used to set the frequency of the synchronization of the statistics resource usage from the

cloud. By default, the self-service portal retrieves resource usage from the cloud. These statistics include processors in core, memory, and storage in megabytes. If the property is not set, a default of 60 seconds is used.

Results

After you configure usage metering, you can monitor the cloud resource usage from the self-service portal by selecting **Reports** > **Usage Metering**. View details about a specific virtual server by selecting the virtual server from the Usage metering grid.

For more information about using Usage Metering, see the IBM Cloud Manager with OpenStack User Guide.

Configuring capacity and overcommit rates

The capacity in the IBM Cloud Manager with OpenStack self-service portal indicates the total and used (or allocated) resources, including processors, memory, and storage.

The capacity view is enabled for all the supported virtualization environments. All x86 clouds support resource overcommit by the hypervisors. Therefore, the total amount of virtual resources that can be allocated are larger than the total physical resources available. Therefore, the self-service portal supports the ability to set an overcommit rate to limit the resource use.

Note: Overcommit rate is not supported for PowerVC virtualization environments.

The overcommit rate is represented in the following fields:

- *totalCpu* - Represents the physical processor cores that are multiplied by the processor overcommit ratio
- *totalMem* - Represents the physical memory size that is multiplied by the memory overcommit ratio

Note: Overcommitted storage is not allowed.

The capacity API tells administrators the amount of physical resources, the amount of resources (after overcommitted), and the resources that are allocated. The user interface shows the physical resource and the overcommit rate only. It excludes the amount of resources after they are overcommitted.

To configure the overcommit rate, open the `cloud.properties` file and set the following properties:

```
# The cpu overcommit rate in OpenStack cloud
com.ibm.cfs.openstack.overcommit.cpu=16
# The memory overcommit rate in OpenStack cloud
com.ibm.cfs.openstack.overcommit.memory=1.5

# The cpu overcommit rate in VMware cloud
com.ibm.cfs.vmware.overcommit.cpu=10
# The memory overcommit rate in VMware cloud
com.ibm.cfs.vmware.overcommit.memory=1.5

com.ibm.cfs.openstack.overcommit.memory is invalid for hyper-v hypervisor.
```

The self-service portal checks whether there are sufficient available resources when deploying. If the available resource is less than the requested resource, the self-service portal stops the deployment process. This feature can be enabled or disabled by using the following property in the `deployment.properties` file:

```
#Enable/Disable the capacitiy check against the selected target while deploying a workload.
com.ibm.cfs.resource.check.enable=true
```

Configuring web user interface

Configuring user interface widgets

The widgets in the web user interface of IBM Cloud Manager with OpenStack and the properties of the widgets are configurable. Using configuration settings, you can control which widgets appear and in what order they appear.

About this task

To configure user interface widgets for IBM Cloud Manager with OpenStack, perform the following steps:

Procedure

1. Open the `web.properties` file in the home directory.
2. Set the `com.ibm.cfs.web.pods.order` property to the names of widgets that are to be shown in the IBM Cloud Manager with OpenStack user interface, in the order you want them displayed. The names are not case-sensitive and must be separated by a comma. Possible names include the following names:
   - **CloudStatus**
   - **WorkloadsStatus**
   - **ResourceUsageStatus**
   - **RecentEvents**
3. Set the properties of each widget. The following example shows a widget property configuration example using the CloudStatus widget.
   a. `com.ibm.cfs.web.pods.cloudstatus.enabled=true`

      If the value is *true*, the CloudStatus widget is displayed in the IBM Cloud Manager with OpenStack web user interface. If the value is *false*, the property is not specified in the file, or you specify an incorrect value (*truue*) then the CloudStatus widget is not displayed.

   b. `com.ibm.cfs.web.pods.cloudstatus.closed`

      If the value is true, the CloudStatus widget is initially displayed in a collapsed form. Otherwise, the CloudStatus widget is initially expanded in the IBM Cloud Manager with OpenStack web user interface.

   c. `com.ibm.cfs.web.pods.cloudstatus.refresh.interval=30`

      The value of this property indicates how often the CloudStatus widget is refreshed. The value is specified in seconds and must be an integer of 1 or higher.

      Repeat these substeps for each additional named widget to be configured, including WorkloadsStatus, ResourceUsageStatus, and RecentEvents. The following properties can be set:

   WorkloadsStatus:
   - `com.ibm.cfs.web.pods.workloadsstatus.enabled`
   - `com.ibm.cfs.web.pods.workloadsstatus.closed`
   - `com.ibm.cfs.web.pods.workloadsstatus.refresh.interval`

   ResourceUsageStatus
   - `com.ibm.cfs.web.pods.resourceusagestatus.enabled`
   - `com.ibm.cfs.web.pods.resourceusagestatus.closed`
   - `com.ibm.cfs.web.pods.resourceusagestatus.refresh.interval`

   RecentEvents
   - `com.ibm.cfs.web.pods.recentevents.enabled`
   - `com.ibm.cfs.web.pods.recentevents.closed`
   - `com.ibm.cfs.web.pods.recentevents.refresh.interval`

4. Save the web.properties file and restart the IBM Cloud Manager with OpenStack server. The properties of each widget take effect after the server is restarted.

   Note:
   - If a widget is not listed in com.ibm.cfs.web.pods.order and its property com.ibm.cfs.web.pods.name.enabled is set to true, it is displayed in the IBM Cloud Manager with OpenStack user interface after all the widgets specified in the com.ibm.cfs.web.pods.order property.
   - If the web.properties file does not exist, all user interface widgets show by default.

Configuring session timeout

You can configure how long a web interface session for an IBM Cloud Manager with OpenStack user can remain inactive before the session times out.

About this task

To configure the timeout value, follow these steps:

Procedure

1. Open the web.properties file in the home directory.
2. Set the com.ibm.cfs.client.idle.timeout property to the number of minutes for which the session is allowed to be inactive. The number must be a positive number and greater than one. After the specified amount of time passes, the user session with IBM Cloud Manager with OpenStack expires.

   If the property is set to -1, the user session with IBM Cloud Manager with OpenStack never expires.
3. Save the web.properties file and restart the IBM Cloud Manager with OpenStack server. The property takes effect after the server is restarted.

   Note: If com.ibm.cfs.client.idle.timeout property is not present or is set to an invalid value, a default value of 30 minutes is used.

Configuring the Welcome page

You can configure IBM Cloud Manager with OpenStack to display the welcome page for all users.

Procedure

1. Open the web.properties file in the home directory.
2. To display the welcome page for all users, set the com.ibm.cfs.web.welcomepage.enabled property to true.
3. Save the web.properties file and restart the IBM Cloud Manager with OpenStack server. The property takes effect after the server is restarted.

   Note: If com.ibm.cfs.web.welcomepage.enabled property is not present or is set to an invalid value, the welcome page is displayed.

Configuring the default instance name

You can configure IBM Cloud Manager with OpenStack to use a default instance name when deploying an image. If you set this property to true, a default instance name based on the image name that is being deployed is generated; otherwise no default is used.

Procedure

1. Open the web.properties file in the home directory.
2. To set the default instance name to the image name being deployed, set the com.ibm.cfs.web.workloadname.default.enabled property to true.

3. Save the `web.properties` file and restart the IBM Cloud Manager with OpenStack server. The property takes effect after the server is restarted.

   Note: If `com.ibm.cfs.web.workloadname.default.enabled` property is not present or is set to an invalid value, the default name is set.

Configuring IBM Cloud Manager with OpenStack dashboard properties

You can configure various web interface properties for the IBM Cloud Manager with OpenStack dashboard.

After you successfully install IBM Cloud Manager with OpenStack, you can configure unique properties and features if you use the IBM Cloud Manager with OpenStack dashboard to manage your cloud.

The following properties are unique to the IBM Cloud Manager with OpenStack solution. For information about customizing the dashboard settings and configuration, see Horizon Settings and Configuration Documentation.

Custom Properties (IMAGE_RESERVED_CUSTOM_PROPERTIES)

In the IBM Cloud Manager with OpenStack dashboard, an administrator can use the Image Custom Properties pane to add, update, or remove custom properties for an image. You can use the dashboard setting `IMAGE_RESERVED_CUSTOM_PROPERTIES` to specify which image custom properties (if any) should not be displayed on that pane. For example, you might create a different pane for a particular property in the IBM Cloud Manager with OpenStack dashboard and, as a result, do not want to display that property on the Image Custom Properties pane. To add the value, modify the following property.

```
IMAGE_RESERVED_CUSTOM_PROPERTIES = ['property_1','property_2']
```

where *property_1* and *property_2* represent the image custom properties you want to suppress.

Extra Filtering Capabilities

The IBM Cloud Manager with OpenStack dashboard provides extra API filters to some of the tables that can be paged. When the filters are applied, the filter is included in the API request and not performed in the dashboard. This behavior allows the filters to be applied across pages of data, if you are using pagination.

- From the **Images** tab, you can apply the following filters:

  ```
  Name:  value must match image name
  Status:  value must match current image status
  Container Format:  value must match image container format
  Disk Format:  value must match image disk format
  Min. Size (MB):  value must be a number and is the minimum image size in MB (for example, 8.5)
  Max. Size (MB):  value must be a number and is the maximum image size in MB (for example, 8.5)
  ```

- From the **Instances** tab, you can apply the following filters:

  ```
  Project:  value must match the project name
  Name:  value is a regular expression (for example, server.*)
  Status:  value must match current instance status
  Image ID:  value must match the image ID
  Flavor ID:  value must match the flavor ID
  Hostname:  value must match the hostname to which the instance is assigned
  IPv4 Address:  value is a regular expression for matching the IP address
  IPv6 Address:  value is a regular expression for matching the IPv6 address
  ```

  Note: The **Project** field is unique, in that, it is not an API filter and does not get applied across pages.

- From the **Projects** tab, select **Instances** to apply the following filters, as needed:

```
Name:  value is a regular expression (for example, server.*)
Status:  value must match current instance status
Image ID:  value must match the image ID
Flavor ID:  value must match the flavor ID
```

Provider Network Settings

Using the IBM Cloud Manager with OpenStack dashboard, you can set the provider network settings when you create a network. These settings require the Neutron provider extension.

- From the **Networks** tab, select **Create Network**. In the **Provider Network Type** field, choose the applicable network type.
  - *Local:*Local is the default value.
  - *Flat*: If selected, the **Physical Network** field displays and you need to enter the name of the physical network.
  - *VLAN*: If selected, the **Physical Network and Segmentation ID** fields display. In this case, the **Segmentation ID** is the VLAN virtual ID (VID) on the physical network that routes to the virtual network. Valid values are 1 - 4094.
  - *GRE*: If selected, the **Segmentation ID** field displays. In this case, the Segmentation ID is the tunnel ID. Valid values can include any unsigned 32-bit integer.

Note: For more information, see Networking in the OpenStack documentation.

Rebranding the dashboard

IBM Cloud Manager with OpenStack provides a recipe to customize the branding that is applied to the dashboard.

The branding information must be contained in a compressed file. For example, `openstack-branding.tar.gz`. The compressed file content is a folder that contains up to three items:

- A file that is named `product_name.txt` that contains the dashboard name on the first line.
- A folder that is named `icons` that contains custom icons that override the default icons.
- A folder that is named `styles` that contains custom style sheets to be applied to the dashboard.

Here is an example branding folder structure:

```
branding
 |-- product_name.txt
 |-- icons
 |    |-- favicon.png
 |    |-- logo.png
 |    |-- logo-splash.png
 |-- styles
 |    |-- custom_1.css
 |    |-- custom_2.less
```

To create the tar file, put the branding files into a local directory structure as noted above.

1. Change the active directory to the branding directory.
2. Run **cd branding**.
3. Run the tar command: **tar -czf ../openstack-branding.tar.gz ***.
4. Make tar file available to Chef server.
5. Add the **ibm-os-dashboard-branding** role to the end of the run list of the controller node which is running Horizon and is to be re-branded
6. Update the node. See steps in the “Updating a deployed topology” on page 113 topic.

Changing the dashboard name

You can change the dashboard name by providing a `product_name.txt` file. The first line of this file is used in the page title and displayed on the Login page and in the page header.

Changing Icons

You can change the icons that are displayed in the dashboard by providing custom icons in the icons folder. The content of this folder is copied into the dashboard icons directory, overwriting any existing icons. To override an icon, provide an icon of the same name in the icons folder. You might change the following icons for rebranding purposes:

- `favicon.png`: The favicon image that is displayed in the browser.

  Note: The `favicon.ico` file is not used.
- `logo.png`: The logo that is displayed in the page header.
- `logo-splash.png`: The logo that is displayed on the Login page

The dashboard icons are in `/usr/share/openstack-dashboard/openstack_dashboard/static/dashboard/img`.

Providing Custom Style sheets

You can provide custom style sheets by placing them in the styles folder. The supported style sheet types include Less (.less extension) and CSS (.css extension). Files of any other extension type are ignored. Multiple style sheets are applied in alphanumeric order. For example, a style sheet that is named `custom1.css` is applied before a style sheet named `custom2.css`.

Changing the Custom Branding

You can change or remove any custom branding that is applied to the dashboard by running the dashboard branding recipe again using different branding information. To remove a previously applied customization, omit the customization information from the provided branding folder. For example, to return the icons back to the default, run the branding recipe using a branding folder that does not contain an icons folder. By using an empty branding folder, all customizations are removed.

Chapter 7. Managing IBM Cloud Manager with OpenStack as an Administrator

Use the following information for working with IBM Cloud Manager with OpenStack. The management tasks vary depending on the IBM Cloud Manager with OpenStack role and access level you are using. Select the content that applies below, depending on whether you log in as an *Administrator* or a *User*.

Configuring the license key

The IBM Cloud Manager with OpenStack product includes a default trial license, with a 90-day trial period. You can use this license to investigate IBM Cloud Manager with OpenStack.

About this task

To update the license, follow these steps:

Procedure

1. Copy the `cmwo_4.1.lic` file from the IBM Cloud Manager with OpenStack installation media to the deployment server system.
2. To apply the permanent license, run the following command:
   ```
   /opt/ibm/cmwo/bin/cmwo_lic.sh -a /product_dir/cmwo_4.1.lic
   ```
   where *product_dir* is the directory to which you copied the `cmwo_4.1.lic` file.

Results

If the license is valid, you receive a message that states the update was successful. If it is not valid, you receive a message that states the license is invalid.

What to do next

Note: If you have installed IBM Cloud Manager with OpenStack self-service portal to run on Windows, you must also update the license for the self-service portal. For more information, see Configuring the license key for the self-service portal running on Windows.

Related reference:

"License information" on page 6
Learn about license information for the IBM Cloud Manager with OpenStack product.

Managing IBM Cloud Manager with OpenStack services

Use this information to see which services run in your IBM Cloud Manager with OpenStack topology deployment. Additionally, learn to check the status of services and restart services.

IBM Cloud Manager with OpenStack services

The IBM Cloud Manager with OpenStack services provided throughout your topology deployment are determined by the roles that are defined for each node.

Overview

The following lists detail the services that are provided by each role. The services are listed in their start order. If you stop services, you must stop the services in reverse order. If you need to restart database

© Copyright IBM Corp. 2012, 2014

services, then all services on the OpenStack controller node must be stopped before the database services can be restarted. In addition, if you need to restart the Qpid messaging service, then the services running on the OpenStack compute nodes must also be restarted.

Note: Nodes that have the Hyper-V agent installed are not assigned a role. For more information about managing the IBM Cloud Manager with OpenStack services that are running on Hyper-V agent nodes, see Installing and uninstalling the IBM Cloud Manager with OpenStack Hyper-V Agent.

The following roles run the same services as other roles that are described in detail below.

ibm-os-allinone-kvm
: Runs the same services as the **ibm-os-single-controller-node** role followed by the **ibm-os-compute-node-kvm** role.

ibm-os-single-controller-node
: Runs the same services as the **ibm-os-database-server-node** role followed by the **ibm-os-single-controller-distributed-database-node** role.

ibm-os-single-controller-powervc-driver
: Runs the same services as the **ibm-os-single-controller-node** role followed by the **ibm-os-powervc-driver-node** role.

ibm-os-single-controller-distributed-database-node

```
qpidd
memcached
openstack-keystone
openstack-iaasgateway
openstack-glance-api
openstack-glance-registry
openvswitch
neutron-server
neutron-openvswitch-agent
neutron-dhcp-agent
openstack-nova-conductor
openstack-nova-api
openstack-nova-scheduler
openstack-nova-cert
openstack-nova-novncproxy
openstack-nova-consoleauth
openstack-cinder-volume
openstack-cinder-api
openstack-cinder-scheduler
openstack-heat-api
openstack-heat-api-cfn
openstack-heat-api-cloudwatch
openstack-heat-engine
openstack-ceilometer-api
openstack-ceilometer-central
openstack-ceilometer-collector
openstack-ceilometer-notification
openstack-ceilometer-alarm-evaluator
openstack-ceilometer-alarm-notifier
httpd
```

ibm-os-compute-node-kvm

```
openvswitch
neutron-openvswitch-agent
openstack-nova-compute
openstack-ceilometer-compute
```

ibm-os-compute-node-powerkvm

Runs the same services as the **ibm-os-compute-node-kvm** role.

ibm-os-zvm-driver-node

```
neutron-zvm-agent
openstack-nova-compute
```

ibm-os-powervc-driver-node

```
openstack-glance-powervc
openstack-cinder-powervc
openstack-neutron-powervc
openstack-nova-powervc
```

ibm-os-database-server-node

```
# When DB2 database is used.
db2.service
db2.nosql.service

# When MySQL database is used:
mysqld
```

ibm-sce-node

```
sce
```

ibm-os-prs-ego-master

```
ego
```

ibm-os-prs-controller-node

```
openstack-nova-ibm-notification
openstack-nova-ibm-ego-resource-optimization
openstack-nova-ibm-ego-ha-service
```

Related reference:

"Roles" on page 237
The following roles are provided in support of the reference topologies.

Checking status of OpenStack services

After a IBM Cloud Manager with OpenStack topology is deployed, you can check the status of the IBM Cloud Manager with OpenStack services. You can use the OpenStack Dashboard or create a script to check more detailed status.

About this task

To check the services status with OpenStack Dashboard, log in and select **Admin** > **System Panel**.

If you want a more detailed status check of the CMWO services, complete the following steps for each node in your topology deployment.

Procedure

1. Determine the roles that are applied to the node. You can find this information in topology file for your deployment, `your-topology-name.json`.
2. Log in to the node.
3. Create a `cmwo_services.sh` script with the following contents.

```
#!/bin/bash

# Example:
# sudo ./cmwo_services.sh 'CHANGEME' status

SERVICE_LIST=$1
SERVICE_OPTION=$2

# Handle the services
for SERVICE in $SERVICE_LIST
do
  if [[ $SERVICE == "db2.service" ]]; then
    if [[ $SERVICE_OPTION == "stop" ]]; then
      sleep 10  # Allow database connections to close.
    fi
  fi
  service $SERVICE $SERVICE_OPTION
  if [[ $SERVICE_OPTION == "start" ]]; then
    sleep 5  # Allow service to initialize.
  fi
done
```

4. After you create the script, give the owner execute permission, `chmod +x ./cmwo_services.sh`
5. Run the `cmwo_services.sh` script as follows, where *CHANGEME* is set to the IBM Cloud Manager with OpenStack services that are available based on the roles of the node. For more information about the list of services that are provided by each role, see "IBM Cloud Manager with OpenStack services" on page 169.

   ```
   sudo ./cmwo_services.sh CHANGEME status
   ```

Restarting IBM Cloud Manager with OpenStack services

If you need to restart services for IBM Cloud Manager with OpenStack, follow a specified order for restarting the services successfully.

About this task

For more information about the order of stopping and starting services, see "IBM Cloud Manager with OpenStack services" on page 169.

Procedure

1. Log in to the node.
2. Create a `cmwo_services.sh` script with the following contents.

```
#!/bin/bash

# Example:
# sudo ./cmwo_services.sh 'CHANGEME' status

SERVICE_LIST=$1
SERVICE_OPTION=$2

# Handle the services
for SERVICE in $SERVICE_LIST
do
  if [[ $SERVICE == "db2.service" ]]; then
    if [[ $SERVICE_OPTION == "stop" ]]; then
      sleep 10  # Allow database connections to close.
    fi
  fi
  service $SERVICE $SERVICE_OPTION
  if [[ $SERVICE_OPTION == "start" ]]; then
    sleep 5  # Allow service to initialize.
  fi
done
```

3. After you create the script, give the owner execute permission, `chmod +x ./cmwo_services.sh`
4. Run the `cmwo_services.sh` script as follows, where *CHANGEME_** is set to the services that are available based on the roles that are defined for the node. For more information about the order of stopping and starting services, see "IBM Cloud Manager with OpenStack services" on page 169.

```
sudo ./cmwo_services.sh 'CHANGEME_STOP_ORDER' stop
sudo ./cmwo_services.sh 'CHANGEME_START_ORDER' start
sudo ./cmwo_services.sh 'CHANGEME_START_ORDER' status
```

Managing the IaaS gateway

The IaaS gateway is a part of IBM Cloud Manager with OpenStack.

About this task

It is a light-weight proxy middleware container capable of providing pluggable adaptation to normalize interactions across multiple IaaS cloud provider vendors.

Procedure

1. Customize the `iaasgateway` attributes. You can update the following JSON attributes in your environment file, `your-environment-name.json`, to customize the **`ibm-openstack.iaasgateway.*`** attributes.
   - **`ibm-openstack.iaas-gateway.iaasgateway_ip`**: The default value for this attribute binds to the management IP address specified by the **`openstack.endpoints.host`** value. If you want this value to be different from the **`openstack.endpoints.host`** value, set this attribute individually.
   - **`ibm-openstack.iaas-gateway.listen_port`**: The default value for this attribute is *9973*. With the default value, the **`iaasgateway`** listens to port 9973.
   - **`ibm-openstack.iaas-gateway.logging.enabled`**: The default value for this attribute is *true*. With the default value, when you start the **`iaasgateway`**, it logs basic information.
   - **`ibm-openstack.iaas-gateway.logging.debug`**: The default value for this attribute is *false*. With the default value, only the informational level (*INFO*) messages are logged. If you want more detailed messages, set this value to *true*.
2. Manage the license. After `iaasgateway` is installed, there are some default licenses installed in `/etc/iaasgateway/ssl` that expire in 10 years. If you want to provide you own certifications, replace them with the same name or use the environment to set them before using IBM Cloud Manager with OpenStack or deploying the controller node. If you use the environment to set the values, you must add the following attributes into the environment. For more information about how to create a license and its expiration date, see http://www.openssl.org/docs/apps/openssl.html.
   - **`ibm-openstack.iaas-gateway.ssl.certfile_url`**: The default value for this attribute is *nil*. The value must be a url where you can download the certification file. If you set this value, both of the following values must be set, otherwise this setting does not work.
   - **`ibm-openstack.iaas-gateway.ssl.keyfile_url`**: The default value for this attribute is *nil*. The value must be a url where you can download the `key` file. If this value is not *nil*, then **`certfile_url`** and **`ca_certs_url`** cannot be nil, or this setting does not work.
   - **`ibm-openstack.iaas-gateway.ssl.ca_certs_url`**: The default value for this attribute is *nil*. The value must be a url where can download the `ca` file. If this value is not nil, then **`certfile_url`** and **`keyfile_url`** cannot be nil, or this setting does not work.

Managing with the IBM Cloud Manager with OpenStack self-service portal (Administrator access)

With the Administrator role in the IBM Cloud Manager with OpenStack self-service portal, you perform tasks such as configuring clouds, creating projects, and managing images, instances, and requests.

For information about using the self-service portal as a non-administrative user, see the **Managing as a User (self-service portal)** section.

Starting and stopping IBM Cloud Manager with OpenStack self-service portal

The following steps are required for starting IBM Cloud Manager with OpenStack self-service portal.

Note: When you start or restart the self-service portal on a high scale cloud, the synchronization between the self-service portal and the cloud might take longer than expected. This resynchronization might cause operations such as deploying, deleting, or resizing an instance to be delayed or even fail. Wait for the synchronization to complete before you attempt these actions.

Starting and stopping IBM Cloud Manager with OpenStack self-service portal on Windows

To start the self-service portal, go to **All Programs** > **IBM Cloud Manager with OpenStack 4.1** > **Cloud Manager with OpenStack**. Alternatively, you can double-click the **Cloud Manager with OpenStack 4.1** icon that might be installed on your desktop. In 10 to 20 seconds, the server is available and you can access the self-service portal by opening `http://localhost:18080/cloud/web/index.html` in a supported browser.

Note: The host name *localhost* and port *18080* are the default host and port names. Substitute the appropriate values for your environment if necessary.

To stop the self-service portal, shut down the instance by closing the window or by using Ctrl-C to end it.

If the self-service portal instance was started with an OSGi console running in the background, you might have to telnet to it with the correct port to access the console. Refer to OSGi console for more details. From the OSGi console, type `shutdown` to stop the self-service portal and then `exit` to exit immediately. Exit the self-service portal before you restart it.

Starting and stopping IBM Cloud Manager with OpenStack self-service portal on Linux

The IBM Cloud Manager with OpenStack self-service portal installation on Linux can be started by the root user or by users who are members of the sce group.

After the self-service portal on Linux is installed, the self-service portal service is started. You can check the self-service portal status on your controller node with the `service sce status` command.

Use the following commands to start, stop, and restart the self-service portal service on Linux:

- Start the self-service portal: `service sce start`
- Stop the self-service portal: `service sce stop`
- Restart the self-service portal: `service sce restart`

Access the self-service portal user interface by opening `http://localhost:18080/cloud/web/index.html` in a supported browser.

Note: The host name *localhost* and port *18080* are the default host and port names. Substitute the appropriate values for your environment if necessary.

Configuring the default administrator user account

The default administrator account is created the first time the IBM Cloud Manager with OpenStack self-service portal is started. As administrator, configure the default administrator user account to receive email and notification from users.

About this task

To modify the default administrator user account, follow these steps:

Procedure

1. In the self-service portal, log in as the cloud administrator.
2. Select **Cloud Administrator** in the upper right title bar of the screen, and click **Show user preferences**.
3. On the User Profile dialog, enter the administrator email address.
4. Check **Send notifications about instances and other events**.
5. Verify the **Timezone** and **Language** for the administrator.
6. To change the Cloud Administrator password, click **Change Password**.
7. Click **Update.**

Configuring LDAP authentication using the web interface

Use the web interface to configure IBM Cloud Manager with OpenStack self-service portal as an LDAP client.

Procedure

1. Log in to the self-service portal as an administrator.
2. Click the **Configuration** tab and select **LDAP** in the navigation pane.
3. Click **Edit** and enter the configuration settings to specify how to connect to the LDAP host.

   LDAP provider hostname
   The fully qualified server host name or IP address of the LDAP host.

   Port The port number of the LDAP service on the host for either transaction level security (TLS) or for no security protocol. The default port number is 389.

   Security Protocol
   The self-service portal allows transaction level security (TLS) to be used.

   Certificate
   If transaction level security is used, you must provide the certificate (public key) used by the LDAP server for securing the connection. For information about obtaining a certificate, see your LDAP server documentation.

   LDAP search DN
   This is the distinguished name that should be used to connect to the LDAP host to perform a directory search, for example cn=Administrator,cn=users,dc=cfs1,dc=us

   Note: This field might be required based on the configuration of the LDAP server. For example, if the LDAP server does not support anonymous bind, or if you specify transaction level security (TLS) for the Security Protocol, this field is required. If the LDAP search DN is required, the ID must have read authority on the LDAP server.

   Password
   This is the password that is associated with the LDAP search DN.

   Note: This field is required if the LDAP search DN is required.

Search filter
This is the filter that is used to authenticate users when they log in. Include the special value {FILTER} in the filter to specify where the user ID that is provided during the login should be substituted. For example, (|(userPrincipalName={FILTER}))

Search context
The search context for providing the LDAP lookup.

User ID attribute
The name of the LDAP field to use as the user ID in the self-service portal.

User name attribute
The name of the LDAP field to use as the user name in the self-service portal.

Email address attribute
The name of the LDAP field to use as the email address in the self-service portal.

4. Click **Save**.
5. Restart the self-service portal service for the settings to take effect.

Example

The following examples show settings for different connections:

- **Example 1: Non secure connection (transaction level security is disabled)**

  LDAP provider hostname
  your.host.com.

  Port 389

  Security Protocol
  None

  LDAP search DN
  cn=Manager,dc=sce,dc=com

  Password
  password

  Search filter
  (|(cn={FILTER}))

  Search context
  ou=People,dc=sce-svt,dc=com

  User ID attribute
  uid

  Username attribute
  cn

  Email addressattribute
  mail

- **Example 2: Transaction level security is enabled**

  LDAP provider hostname
  your.host.com.

  Port 389

  Security Protocol
  Transaction level security (TLS)

  Certificate
  certificate_file.cer

LDAP search DN
cn=Manager,dc=sce,dc=com

Password
password

Search filter
(|(cn={FILTER}))

Search context
ou=People,dc=sce-svt,dc=com

User ID attribute
uid

Username attribute
cn

Email addressattribute
mail

What to do next

Notes:

- The self-service portal cannot be returned to use local authentication (non-LDAP authentication) through the web interface. If it is necessary to restore local authentication, see Configuring local authentication for more information. Local authentication is intended only for non-production environments such as for a proof of concept or for performing a demo.
- If you want to enable user name case sensitivity, you must update the `ldap.xml` file after setting the initial LDAP configuration in the web interface. For more information, see "Configuring LDAP authentication manually" on page 128 for more information.

Configuring the license key for the self-service portal running on Windows

The IBM Cloud Manager with OpenStack product includes a default trial license, with a 90-day trial period. You can use this license to investigate IBM Cloud Manager with OpenStack.

About this task

To update the license for the self-service portal when it is running on Windows, follow these steps:

Procedure

1. In the IBM Cloud Manager with OpenStack self-service portal, log in as the administrator.
2. Click the **Configuration** tab and select **License** in the navigation pane.
3. Click **Update License**.
4. Click **Browse...** and choose the license file that you want to use.
5. If you want to update the IaaS gateway license, select the **Update IaaS Gateway license** option.
6. Click **Apply** to complete the update.

Results

If the license is valid, you receive a message that states the update was successful. If it is not valid, you receive a message that states the license is invalid.

Managing images

In the **Images** tab, you can manage and configure the images that are available for deployment. You can view image properties and deploy images.

In IBM Cloud Manager with OpenStack, each image has a status that is associated with it. If the status is *OK,* then the image is ready to be deployed. Click the refresh arrow to update the status.

To view the properties of an image, click the name of the image.

If the list of images does not contain the image that you want, ensure that the current cloud, project, and architecture filters are set correctly.

Building images

Building images manually is a complex and error-prone process. By pre-building images from specific software bundles for reuse by others, administrators can streamline this process. There are several different ways of building images.

Building images with VMware Studio

VMware Studio and the OVF Toolkit simplify the process of image creation. The images that are created in VMware Studio can be imported and deployed by using vSphere Client. For more information about using VMware Studio, see VMware Studio Documentation at http://www.vmware.com/support/developer/studio/.

Building images manually

You can choose to build images manually using open source tools. This method requires significant virtualization and image configuration (for example, OVF, Sysprep or cloud-init) experience.

For more information about creating OpenStack compatible images, see Create images manually in the OpenStack documentation.

Importing images (OpenStack only)

Using IBM Cloud Manager with OpenStack, you can import images to an OpenStack cloud.

About this task

You can use existing OpenStack compatible images. For more information about creating OpenStack compatible images, see Create images manually in the OpenStack documentation.

IBM Cloud Manager with OpenStack supports OpenStack images with the following disk formats supported by different hypervisor types.

Table 28. Supported disk formats by hypervisor

Disk Format	Hyper-V	KVM	PowerKVM	PowerVC	Details
VHD	Y	N	N	N	Microsoft virtual hard disk format
RAW	N	Y	Y	Y	Raw virtual machine disk format
QCOW2	N	Y	Y	N	QEMU disk format
VMDK	N	Y	N	N	VMware virtual machine disk format

Table 28. Supported disk formats by hypervisor (continued)

Disk Format	Hyper-V	KVM	PowerKVM	PowerVC	Details
AMI/AKI/ARI	N	Y	N	N	Amazon machine/kernel/ ramdisk disk format
ISO	N	N	N	N	Disk format with an optical disk file system
VDI	N	N	N	N	Virtual desktop infrastructure disk format

IBM Cloud Manager with OpenStack supports OpenStack images for one of the following guest operating systems:

Table 29. Supported guest operating systems by hypervisor

Guest operating system	Hyper-V	KVM	PowerKVM	PowerVC	VMware	z/VM
Windows • Windows 8 • Windows 7 • Windows Server 2012 R2	Y	Y	N	N	Y	N
Linux • Red Hat Enterprise Linux 6.5 • SUSE Linux Enterprise Server 11.2	Y	Y	N	N	Y	N
Linux on Power • Red Hat Enterprise Linux 5.9 • Red Hat Enterprise Linux 6.5 • SUSE Linux Enterprise Server 11 SP3	N	N	Y	Y	N	N
Linux on IBM System z • Red Hat Enterprise Linux 6.5 • SUSE Linux Enterprise Server 11.2	N	N	N	N	N	Y

Table 29. Supported guest operating systems by hypervisor (continued)

Guest operating system	Hyper-V	KVM	PowerKVM	PowerVC	VMware	z/VM
AIX • AIX 6.1, TL 9 • AIX 7.1, TL 3	N	N	N	Y	N	N
IBM i	N	N	N	N	N	N
Ubuntu • Ubuntu 12.10 • Ubuntu 13.10 • Ubuntu 14.04 LE	Y (Ubuntu 12.10 only)	Y (Ubuntu 12.10 only)	Y (Ubuntu 14.04 LE only)	N	Y (Ubuntu 13.10 only)	N

Notes:

- The imported image file is stored in the OpenStack cloud and not in the IBM Cloud Manager with OpenStack database.
- An image with the VMDK disk format must have its disk that is contained in a single VMDK file.
- IBM Cloud Manager with OpenStack does not support directly deploying images with the AKI and ARI disk formats. Such images are deployed with an AMI disk formatted image. As a result, images with the AKI and ARI disk formats have an Undeployable state in IBM Cloud Manager with OpenStack.
- IBM Cloud Manager with OpenStack does not support importing images for deployment to PowerVC. Only the PowerVC images that were synchronized into the OpenStack cloud and made available to IBM Cloud Manager with OpenStack can be deployed to PowerVC. For information about working with PowerVC images, see the IBM Power Virtualization Center Standard information center:

  http://www.ibm.com/support/knowledgecenter/SSXK2N_1.2.1/com.ibm.powervc.standard.help.doc/powervc_images_hmc.html.

Procedure

1. In the IBM Cloud Manager with OpenStack interface, click the **Images** tab.
2. Click **More** and choose **Import image...** from the menu to open the Import Image window.
3. Update the cloud, project, image name, disk format, and container format for the image being imported.

   Note: When the AMI disk format is selected, you can select the associated AKI (kernel image) and ARI (ramdisk image) disk formatted images. If these images have not been imported, then you can edit these image properties later.
4. Update the hypervisor type for the image being imported. The hypervisor type might have a default value set based on the disk format selected. If the image does not have a specific hypervisor type requirement, then select the “Not Specified” option. If a specific hypervisor type is selected and there is no hypervisor in the OpenStack cloud that has a matching type, the image is not deployable.
5. Optional: If the image has a minimum memory or storage size requirement, update the **minimum memory (MB)** and **minimum storage (GB)** fields.

Note: By default, OpenStack uses the size of the image as the minimum storage size requirement when deploying the image. This default is often sufficient. However, if the image uses a compressed disk format, such as QCOW2, then the minimum storage size requirement should be set to the decompressed image size.

6. Optional: If the image has specific OpenStack architecture requirements, update the architecture fields.
7. Click **Import**.

Note: When you upload an image file using some older browser versions, space for the image file is required in the server `temp` directory. This temporary file is deleted when the upload completes. If the upload does not complete successfully, it is possible that the temporary file is not deleted automatically. If you must use an older browser, place the image file in a location where it can be imported using a URL.

Results

After the image is successfully imported, you can edit a subset of the image properties. For more information on editing OpenStack image properties, see Editing image properties (OpenStack only).

If an image is not deployable after it is imported, check the log entries for the image for more information. You might must modify the image properties to make the image deployable.

Related reference:

Building images
This topic contains more information about building images.

Related information:

Getting virtual machine images
This site contains example images that are compatible with OpenStack.

Editing image properties (OpenStack only)

IBM Cloud Manager with OpenStack supports viewing and editing a subset of the properties that are stored with an image in the OpenStack cloud.

Procedure

1. In the IBM Cloud Manager with OpenStack web interface, click an image to view or edit the properties of that image.
2. The following basic OpenStack image properties can be edited: name, disk format, container format, minimum memory (MB), and minimum storage (GB).
3. Optional: You can also view, create, update, and delete additional OpenStack image properties. The additional properties might include architecture, hypervisor_type, kernel_id, ramdisk_id, os_version or os_distro. For more information about OpenStack image properties, see the OpenStack Compute Administration Guide at http://docs.openstack.org/admin-guide-cloud/content/section_image-mgmt.html.

   Some of the additional OpenStack image properties are specific to certain OpenStack hypervisor types. For example, the KVM hypervisor supports the hw_vif_model, hw_disk_bus and hw_cdrom_bus properties. For more information about these properties, see Libvirt Custom Hardware Configuration at https://wiki.openstack.org/wiki/LibvirtCustomHardware.

   Notes:

   - An image configuration strategy is part of the additional OpenStack image properties. However, the configuration strategy cannot be viewed or edited with this task.
   - To deploy an image to an OpenStack KVM hypervisor, the **qemu** value is used for the hypervisor_type property of the image.

Updating an image configuration strategy (OpenStack only)

IBM Cloud Manager with OpenStack supports adding, updating, and deleting the configuration strategy for an image in an OpenStack cloud.

About this task

The configuration strategy is stored with the image in the OpenStack cloud.

Procedure

1. Select the image that you want to update.
2. Click **More** and choose **Configuration Strategy...** from the menu.
3. Click **Edit**.

   Note: If a configuration strategy exists, a Delete button is provided to delete the existing configuration strategy. The Edit button can be used to add or update the configuration strategy.
4. Update the configuration strategy type, template, user metadata, and mapping for the image.
5. Click **Save**.

   Note: After you update the configuration strategy, reset the image configuration in order for the updated configuration strategy to be applied when you configure or deploy the image.

Related tasks:

"Configuring images with OpenStack" on page 146
The following topics cover information to configure images in an OpenStack environment.

Creating a VMware linked virtual machine

Linked virtual machines can be created from a snapshot or from the current running point. A linked clone is a virtual machine whose disks are linked with the template that it was cloned from. Duplicated data is shared between the linked virtual machine and the template. Linked clones deploy much faster because most of the disk data does not have to be copied.

About this task

You can use any image (or template) to create a linked virtual clone. To create a linked virtual clone, follow these steps:

Procedure

1. Open IBM Cloud Manager with OpenStack and select **Images**.
2. Select the image that you want to clone and select **Deploy** > **Advanced**.
3. On the Advanced Deployment window, select the option to **Link virtual machine to image**.

What to do next

Note:

- The creation of a linked clone requires the image to contain a virtual machine snapshot. If the image used to create a linked clone does not have a virtual machine snapshot, IBM Cloud Manager with OpenStack creates a virtual machine snapshot for the image before the linked clone is created.
- If an image already has a virtual machine snapshot, IBM Cloud Manager with OpenStack does not create a new snapshot, but instead uses the current snapshot. Changes that are made to the template might not be reflected, since the clone operation is based on the snapshot and any future changes you make are outside of the snapshot If you must change the template, create a new snapshot that includes your changes.
- Storage DRS supports linked clones starting in VMware vSphere version 5.1. However, Storage DRS is not always able to make a recommendation to place linked clones on a datastore. As a result, an

attempt to deploy a linked clone to a storage DRS cluster results in the creation of a full clone, and a warning message in the log that a linked clone was not created.

- The disks of a linked clone cannot be resized. Any attempt to resize a linked disk at deployment time results in an error.
- The datastores of the image must be accessible by the deployment target or the virtual machine cannot be linked to the image. In this case, IBM Cloud Manager with OpenStack deploys a full clone.
- You cannot unlink a linked clone from the image.

Configuring image deployment properties

Image deployment customization properties that you want to apply to individual images must be configured through the IBM Cloud Manager with OpenStack web user interface. IBM Cloud Manager with OpenStack enables you to save these properties in advance so that your users do not have to know all the internal and advanced deployment details.

About this task

To set global image deployment properties, see "Configuring global image deployment" on page 135.

The values that are set in the global image deployment properties are used when deploying an image unless individual deployment properties are set for an image. The values that are set for an individual image are used unless they are explicitly overridden when deploying an image. Administrators can set which values are displayed in the basic deployment or even allow users to set advanced deployment properties. For more information about allowing users to view advanced form, see "Configuring access to advanced deployment form" on page 138.

Note: Global configurations are refreshed only when manually reset or when the deployment target changes.

To configure image default deployment customization properties to be used when deploying it from IBM Cloud Manager with OpenStack, complete the following steps:

Procedure

1. In the IBM Cloud Manager with OpenStack interface, click the **Images** tab.
2. Click the name of the image that you want to configure.

   Note: If the image that you want is not available, make sure that the correct cloud, architecture, and project are specified.
3. Click **Configure**.
4. Complete the default customization for the image properties. These properties are divided into several groups, including: Hardware, Software, Network, Storage, Image Target and Other Settings, depending on the type of image that you select.

   Note: Changes that were made in the cloud since the image was added, such as networks that were created or removed, might not display on the Configure Image panel. To ensure that the current cloud settings are available to configure the image, click **Reset to Defaults**.

   - Hardware

     Notes:

     – For OpenStack you can control the size of the virtual machine that is deployed from an image by the flavor that you select. Only flavors that meet the disk and memory size requirements for the image are listed as options.

     – For Power Systems that are running in shared mode, the minimum, desired, and maximum number of both shared virtual processors and shared processing units are paired. For each paired value, take care to ensure that values are set correctly. The number of processing units must be

less than or equal to the number of virtual processors. However, the processing units, multiplied by 10, must be greater than or equal to the number of virtual processors.

For example, if the minimum number of virtual processors is 1, then the minimum number of processing units must be less than or equal to 1 (between 0.1 and 1), and must also be greater than or equal to 1 when it is multiplied by 10.

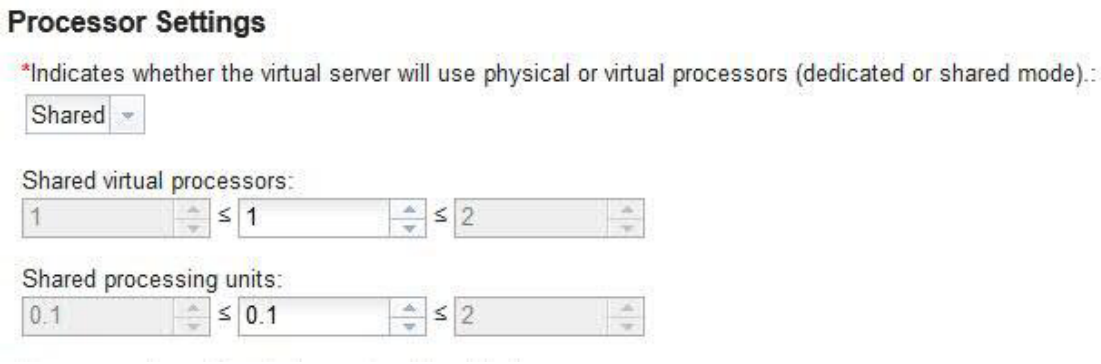

Figure 3. Processor settings

- Software
- Network

  Notes:

  – You can click **Show settings** for each network configuration setting to display configurable options. For Power Systems, VLAN configuration is available in the Adapter network configuration section. (For OpenStack images, **Show settings** is not shown or valid.)

  – The *NONE* value indicates that no network configuration is applied to the settings, in which case the values should be entered manually. (For OpenStack images, *NONE* is not shown or valid.)

  – When a network configuration (including *NONE*) is selected for use, all settings in the subsection are cleared, indicating they draw from the configuration specified.

  – When a network configuration is applied, individual settings can be specified manually by providing a value for the setting and therefore overriding the setting that is specified in the network configuration. Any settings that are blank are taken from the configuration settings. (For OpenStack networks, individual network settings cannot be specified.)

- Storage

  Notes:

  a. The `Virtual disk options for an add disk request` option is for adding a disk during the image deployment and is specific to a VMware cloud. This option is an ordered and colon-separated list of virtual disk options. The default value for an optional virtual disk option is an empty string. The options are parsed in the following order:

  Disk size in MB (Required)
  This option is parsed as a long.

  Thin provisioned (Optional)
  This option is parsed as a boolean. The default value is false which specifies thick provisioned status.

  Eagerly scrub (Optional)
  This option is parsed as a boolean. The default value is false which specifies that the underlying file system determines the scrub status.

Datastore (Optional)

This option is parsed as a string. The default value is null which specifies that you want to use the same datastore as the datastore that is used by the main disk of the virtual machine. Input value sample:

The following is an example of this option:

```
1:true:false:datastore_01
```

b. To configure the maximum number of disks that are allowed or the maximum size of the disks, see "Configuring the number and maximum size of additional storage" on page 138.

- Image target
- Other Settings

5. Depending on the image, you might be able to enter the root password for the server that is provisioned by this image, such that users that deploy the image receive the root password in an email notification.
6. Optionally, you can select specific image customization properties to display to a user on the basic deployment form.
   a. Select the **Show basic deploy panel settings** check box at the top of the configuration panel.
   b. For individual customization properties, select the associated **Show in basic deploy settings** check box. This option causes the property to be displayed when a user brings up the basic deployment form for this image. Check only those properties that you want a user to customize, for example, passwords or port numbers for a specific software product included in the image.
7. Select **Save**.

Note: You can reset an image customization to its original configuration by clicking **Reset to Defaults**.

Deploying an image

You can deploy an image with either basic configuration options or advanced configuration options. Advanced configuration options are only available if the administrator enables them for your environment.

Procedure

1. Click the name of the image you want to deploy.
2. In the Image Details properties page, click **Deploy**.

Note: The IBM Cloud Manager with OpenStack cloud administrator can configure IBM Cloud Manager with OpenStack to allow users to use the advanced deployment form when deploying an image. Click **More** > **Advanced deploy** to display the advanced deployment form.

Basic deployment

With a basic deployment, minimal configuration options, including name, description, project, flavors (if you are using OpenStack), processor information, memory, and key pairs (if you are using OpenStack and at least one key pair is configured for the user) are displayed.

Advanced deployment

Administrator Advanced deployment makes a number of different settings available when an image is deployed. For example, with advanced deployment a user can configure setting like networking, storage, and software configuration values. To enable access to these functions, you can do one of the following:

- Make the advanced deployment form available to all users.
- Choose specific values for an image by selecting the corresponding check box and exposing that on the basic deployment.

For more information about enabling advanced deployment options for users, see "Configuring image deployment properties" on page 183.

With advanced deployment, administrators can configure the options, so users can suspend and resume instances. The VMware and OpenStack Linux Kernel-based Virtual Machine (KVM) and Hyper-V environments support the suspend action by default.

If you enable the multiple instances on a single deployment operation, users can deploy multiple instances through a single deployment. The deployed instances use the deployment name as the prefix of each single instance. The new names also use **-x** as the suffix, where **x** is the index of that instance.

If the deployment approval process is enabled, you receive a single approval request. You can change the number of deployment instances while you review the request. The metering and billing functions remain for each of the single deployment instances. When deploying multiple instances on a single deployment, the instances of this deployment are not displayed immediately after you click **Deploy** or **Approve**.

You can also set fields, such as **Virtual Machine Customization**, **Virtual Machine Personality Files**, and more.

Note:

- Only the members of the selected project can see the instance that is created as a result of the image deployment.
- If approvals are enabled, deployment does not begin until the request is approved by the administrator.
- If billing is enabled, you must be a member of an account that is not delinquent for the deployment to proceed.
- The expiration period and approvals policy settings for deployment depends on the policies that are set for the cloud. If more detailed expiration and approvals are set for the project where the image is being deployed, the policies for the project are applied.
- If you are deploying multiple instances, IBM Cloud Manager with OpenStack deploys the instances one by one. If you restart IBM Cloud Manager with OpenStack before all the deployments are complete, the deployments that are not started will not be deployed. For example, if there are five instances to be deployed and three of them are complete and one is in progress when IBM Cloud Manager with OpenStack is restarted, the fourth instance will be deployed, but the fifth instance will not be deployed.
- If you are deploying an image on z/VM through OpenStack, the instance must have at least one network configuration assigned to the management network. This is because the Extreme Cloud Administration Toolkit (xCAT) is in the management network and the instance must be in the same network as the xCAT, or the deployment will fail.

  Use the Advanced Deploy option and select the first network configuration that is displayed in the IBM Cloud Manager with OpenStack user interface for the management network.

Copying image definitions

Rather than copy an entire image, you can create image copies by using just the metadata of the image.

About this task

By copying the metadata, you can make the same image available to multiple projects or provide multiple alternative configurations of the same base image. You can use the Configure Image window to modify various configuration settings for the copies. The copy image function is enabled for administrators and for project owners for images within their project.

When you copy an image definition, only the image metadata that is stored in the IBM Cloud Manager with OpenStack database is copied. As a result, any metadata that is stored with the image in the cloud is

common across the base and copied images. For example, the configuration strategy for an OpenStack image is metadata that is stored with the image in the cloud. Therefore, the same configuration strategy is used for the base and copied images. For more information about OpenStack configuration strategies, see Configuration strategies.

Note: If you delete the base image, then all copied image configurations are also deleted.

To copy an image definition, perform the following steps:

Procedure

1. On the IBM Cloud Manager with OpenStack page, click the **Images** tab.
2. On the Images page, click the base image name that you want to copy.
3. Click **Copy** to enter the image name and description that you want to assign to the copied image.

What to do next

Now you can configure the copied image and move it to different project if desired.

Viewing image properties

You can view image properties such as the image name, description, last modification date, specification version, revision comments, and logs. As an administrator, or if you have project owner authority, you can also make copies of the image, view related images (images that share the same base image), and modify the image name, description, and project.

About this task

Click the image to view or edit the details of that image. Remember that modifications that you make to an image in IBM Cloud Manager with OpenStack might not be reflected in the underlying virtualization infrastructure.

Deleting images

Using IBM Cloud Manager with OpenStack you can delete images from an OpenStack cloud and certain images from VMware clouds.

About this task

When you delete an image, it is deleted from IBM Cloud Manager with OpenStack. The image is deleted from the OpenStack cloud if it is a OpenStack base image.

The ability to delete an image varies by cloud type:

- OpenStack images can be deleted at any time. Deleting an OpenStack base image results in all of its related images, or copied images, being deleted as well.

Procedure

1. In the IBM Cloud Manager with OpenStack interface, click **Images**.
2. Select the image that you want to delete.
3. Click the delete icon.

Managing projects

You can create, manage, and request access to projects on the **Projects** page, which is available on the **Access** tab.

Projects are used to define the users that have access to a set of images and instances. Only members of a project can view images and instance within a project. In many cases, projects correspond to a department or other human organization.

To manage projects, go to the **Access** tab and click **Projects** to view the list of available projects.

IBM Cloud Manager with OpenStack self-service portal comes with a default project called the Public project, to which all users belong. All virtual images and instances created outside of the self-service portal are, by default, assigned to the Public project. You can also configure a staging project to store newly discovered images or instances. The staging project allows administrators to configure images before making them available to other users. For more information, see "Configuring a staging project" on page 137.

Overview of project membership roles

When you are added as a member of a project, one of three membership roles are assigned to you.

Owner
: A project owner has administrator authority to the project and its contents. The project owner primarily manages the contents of the project and who has authority to the project and its contents.

User
: A project user has the authority to use the project and the objects within the project. For example, a project user can deploy a virtual image to the project. A user can also view and potentially restore backup images of virtual machines that are created by other users, depending on how the project and roles were initially set up. The project user primarily handles their own deployments.

Viewer
: A project viewer has authority only to view the project and the virtual images and instances that are contained in the project.

Creating a project

If you are given authority by your administrator, you can create projects.

Before you begin

Discuss your authority level with your administrator. The com.ibm.cfs.project.creation.by.user property in the `deployment.properties` file must be set to True for you to create projects.

Procedure

1. Click **New Project**.
2. Type a project name and description in the corresponding fields.
3. Click **Create**.

Editing project properties

If you have project owner authority, you can edit the properties of an existing project, including project roles, project name, or project membership.

Procedure

1. From the list of projects, select the project you want to edit.
2. To update the project name or description, click the text field and type the new values.
3. To update project membership:
   a. Click **Project Members** to open the panel.
   b. In the Add Project Members window, select the new members and their project roles to add them to the project.
   c. Click **OK**.

d. To modify an existing member's project role, select the users you want to modify and click **Set Role to** to select the new project role.
e. To remove members from the project, select the users you want to remove and then click **Remove** to remove the users from the project.

4. To update the expiration policies:
   a. Click **Expiration Policies** to open the panel.
   b. Choose one of the following to set the expiration policy:

      Use cloud default
      The expiration of the deployment will depend on the expiration configuration of the cloud to which the image belongs.

      Customize settings
      The expiration policy you set on this panel (by setting the **Maximum expiration value** and **Maximum extension period** values) overrides the expiration policy of the cloud to which the image belongs.

5. To update the approval policies:
   a. Click **Approval Policies** to open the panel.
   b. Choose one of the following to set the approval policy:

      Use cloud default
      The project uses the approval policy of cloud groups.

      Customize settings
      The project uses the approval policy you set on this panel (by selecting checkboxes from the **Require approval for the following events** list) overrides the approval policy of the cloud groups.

6. Click **Save**.

Setting project policies

For projects that you own, you can set expiration policies and approval policies that affect the instances that are deployed in that project.

Procedure

1. Click the **Access** tab and then the **Projects** tab.
2. Click the name of the project in the table to display the project properties.
3. Click **Edit**.
4. Expand the title of the item you want to work with: **Expiration Policies** or **Approval Policies**.
5. Set your policies for your projects, or select **Use cloud default** to use the policies set by your administrator.

What to do next

For more information about expiration policies and approval policies, see the IBM SmartCloud Entry Administrators Guide.

Deleting an existing project

As a project owner, you can delete a project at any time.

About this task

When a project is deleted from IBM Cloud Manager with OpenStack, all of the virtual images and instances contained in the project are transferred to the public project.

Procedure

1. In the projects list, select the project you want to delete.

   Restriction: You cannot delete the default Public project.
2. Click the **Delete selected projects** icon.

Project management with OpenStack

Unlike other cloud types, OpenStack clouds provide native support for project management through the OpenStack keystone component. Because the projects are managed in OpenStack, the projects cannot be updated unless the OpenStack cloud is available.

Keystone is an OpenStack component that provides identity, token, catalog, and policy services to projects in the OpenStack family. Upon first connecting to an OpenStack cloud, IBM Cloud Manager with OpenStack imports all the projects that currently exist inOpenStack. The current project membership is accepted and reflected in IBM Cloud Manager with OpenStack.

After the initial OpenStack projects import, when connected to an OpenStack cloud, IBM Cloud Manager with OpenStack enters transactional mode for project management. When in transactional mode, all project management operations that are performed in IBM Cloud Manager with OpenStack are also performed in OpenStack (that is in keystone). If a project management operation (or any of the operations described in this section) fails to complete successfully in IBM Cloud Manager with OpenStack it does not occur in OpenStack. Likewise, if it fails in OpenStack, it reverts in IBM Cloud Manager with OpenStack.

IBM Cloud Manager with OpenStack enters transactional mode for project operations, while connected to OpenStack, in order to have the registries in both products synchronized. For this reason, when connected to an OpenStack cloud, IBM Cloud Manager with OpenStack cannot perform project-related operations while the OpenStack cloud is down or unavailable.

To connect to OpenStack, IBM Cloud Manager with OpenStack uses a service user account and a default service tenant. Some OpenStack installations have user accounts specific to OpenStack components (for example, nova, keystone, neutron). These and other service user accounts or service tenants in an OpenStack server that do not represent an actual user account or tenant, can be added to the list of service users and service tenants so that they are ignored by IBM Cloud Manager with OpenStack. To make this change, add the service users and tenants to the comma-separated list of users in the *com.ibm.cfs.cloud.openstack.service.users* property, or the comma-separated list of tenants in the *com.ibm.cfs.cloud.openstack.service.tenants* property, in the *openstack.properties* file.

Managing approval policies

IBM Cloud Manager with OpenStack self-service portal administrators can enable approval policy support by specifying the operations that require approval. If approval policies are enabled, the requested operation is held until the approval request is processed by the administrator.

This approval requirement ensures that self-service portal administrators control the self-service portal instance process and provides an audit trail of the requester and approver roles.

From a user standpoint, the approval lifecycle behaves similar to the following:

- Users can see only requests that they initiate.
- Users are unable to view any requests against an instance in the public project that they did not originate. Instances indicate that they are in a Pending state, but users cannot see the outstanding requests that are initiated by other users against that instance.

Setting or modifying approval policies for a cloud

Follow these steps to set or modify an approval policy for a cloud. These policies are used unless they are overridden by an approval policy for a project.

Procedure

1. In the IBM Cloud Manager with OpenStack interface, select **Configuration** > **Clouds**.
2. Click the cloud name for which you want to modify approval policies.
3. Select **Approval Policies**.
4. Set the events that require administrator approval.

Deploying an image
Approval policy that is invoked when deploying an image to create an instance in the cloud. This approval policy suspends the deployment operation until the generated request is approved or rejected.

Extending the instance expiration time frame
Approval policy that is invoked when extending the expiration date of an existing instance. This approval policy suspends the expiration operation until the generated request is approved or rejected.

Resizing an instance
Approval policy that is invoked when resizing an existing instance. This approval policy suspends the resize operation until the generated request is approved or rejected.

Capturing an instance
Approval policy that is invoked when capturing an existing instance. This approval policy suspends the capturing operation until the generated request is approved or rejected.

Deleting an instance
Approval policy that is invoked when deleting an existing instance. This approval policy suspends the delete operation until the generated request is approved or rejected.

Requesting to attach storage to a virtual machine
Approval policy that is invoked when attaching storage to a virtual machine. This approval policy suspends the attach storage operation until the generated request is approved or rejected.

Requesting to detach storage from a virtual machine
Approval policy that is invoked when detaching storage from a virtual machine. This approval policy suspends the detach storage operation until the generated request is approved or rejected.

Saving a virtual machine image
Approval policy that is invoked when saving a virtual machine image. This approval policy suspends the save image operation until the generated request is approved or rejected.

Requesting to create virtual machine snapshot
Approval policy that is invoked when creating a virtual machine snapshot. This approval policy suspends the virtual machine snapshot operation until the generated request is approved or rejected.

Restoring a virtual machine
Approval policy that is invoked when restoring a saved virtual machine image. This approval policy suspends the restore operation until the generated request is approved or rejected.

Requesting to revert virtual machine to snapshot
Approval policy that is invoked when reverting a virtual machine to snapshot version. This approval policy suspends the revert to snapshot operation until the generated request is approved or rejected.

Deleting a volume
Approval policy that is invoked when a volume is deleted. This approval policy suspends the delete volume operation until the generated request is approved or rejected.

Creating a volume
Approval policy that is invoked when a volume is created. This approval policy suspends the create volume operation until the generated request is approved or rejected.

Updating a volume
Approval policy that is invoked when a volume is updated. This approval policy suspends the update volume operation until the generated request is approved or rejected.

Setting or modifying approval policies for a project

Follow these steps to set or modify an approval policy for a project. These policies override the approval policies that are set for a cloud.

Procedure

1. In the IBM Cloud Manager with OpenStack interface, select **Access** > **Projects**.
2. Select the project for which you want to modify approval policies.
3. Select **Customize settings**.
4. Select **Approval Policies**.
5. Set the events that require administrator approval. To use setting defined for a cloud, select **Use cloud default**.

Deploying an image
Approval policy that is invoked when deploying an image to create an instance in the cloud. This approval policy suspends the deployment operation until the generated request is approved or rejected.

Extending the instance expiration time frame
Approval policy that is invoked when extending the expiration date of an existing instance. This approval policy suspends the expiration operation until the generated request is approved or rejected.

Resizing an instance
Approval policy that is invoked when resizing an existing instance. This approval policy suspends the resize operation until the generated request is approved or rejected.

Capturing an instance
Approval policy that is invoked when capturing an existing instance. This approval policy suspends the capturing operation until the generated request is approved or rejected.

Deleting an instance
Approval policy that is invoked when deleting an existing instance. This approval policy suspends the delete operation until the generated request is approved or rejected.

Requesting to attach storage to a virtual machine
Approval policy that is invoked when attaching storage to a virtual machine. This approval policy suspends the attach storage operation until the generated request is approved or rejected.

Requesting to detach storage from a virtual machine
Approval policy that is invoked when detaching storage from a virtual machine. This approval policy suspends the detach storage operation until the generated request is approved or rejected.

Saving a virtual machine image
Approval policy that is invoked when saving a virtual machine image. This approval policy suspends the save image operation until the generated request is approved or rejected.

Requesting to create virtual machine snapshot
Approval policy that is invoked when creating a virtual machine snapshot. This approval policy suspends the virtual machine snapshot operation until the generated request is approved or rejected.

Restoring a virtual machine
Approval policy that is invoked when restoring a saved virtual machine image. This approval policy suspends the restore operation until the generated request is approved or rejected.

Requesting to revert virtual machine to snapshot
Approval policy that is invoked when reverting a virtual machine to snapshot version. This approval policy suspends the revert to snapshot operation until the generated request is approved or rejected.

Deleting a volume
Approval policy that is invoked when a volume is deleted. This approval policy suspends the delete volume operation until the generated request is approved or rejected.

Creating a volume
Approval policy that is invoked when a volume is created. This approval policy suspends the create volume operation until the generated request is approved or rejected.

Updating a volume
Approval policy that is invoked when a volume is updated. This approval policy suspends the update volume operation until the generated request is approved or rejected.

Managing requests

When you deploy an image or when you initiate an action that requires approval from an administrator, a request is created and submitted to an administrator for approval. The status is set to Pending until the administrator handles the approval request.

You can set which actions require administrator approval by using the Approval policies function. For more information, see "Managing approval policies" on page 190.

Processing instance requests

When an image is deployed, initiating an instance, the deployment request may require approval by an administrator. The instance status is set to pending until the administrator handles the approval request.

About this task

You can process an instance request from the Instances tab or from the Requests tab. For more information about processing an instance request from the Instances tab, see "Processing requests from the Instances tab" on page 208

To process a pending request, follow these steps:

Procedure

1. In the IBM Cloud Manager with OpenStack interface, select **Access** > **Requests**.
2. Expand the **Request Details** section to review or update the request before approving.
3. Expand the **Comments** section to review comments or use the **Add Comment** link to provide additional comments.
   - Click **Approve** to approve the request and allow the deployment processing to start.
   - Click **Reject** to reject the request.
   - Click **Withdraw** to withdraw a request.

Clearing or archiving requests

You can clear or archive requests. Clearing requests deletes the requests while archiving requests saves them to an archive folder. By clearing requests, you can free space on your system and improve performance in the IBM Cloud Manager with OpenStack interface. Archive any requests that you may want to reference in the future.

About this task

To clear or archive a request, follow these steps:

Procedure

1. In the IBM Cloud Manager with OpenStack interface, select **Access** > **Requests**.
   - To clear requests, click **Clear**.
   - To archive requests, click **Archive**.
2. Use the Request filter to select a subset of requests to clear or archive. Filter by status or start and end date. If you filter by date, you must provide an end date.
   - To clear the selected requests, click **Clear**.
   - To archive the selected requests, click **Archive**. The filtered requests are saved in a file called `requests_<current time in milliseconds>.csv`. This file can be found in the `archives` folder, located in the IBM Cloud Manager with OpenStack configuration directory.

Managing expiration policies

Expiration policies require users to set an expiration period that specifies the maximum length of the instance lease and determine the lifecycle of expired instances.

You can set a default expiration policy for a cloud or for a project. Expiration policies set for a project override the expiration policies set for a cloud. After an expiration policy is set, you must set an expiration date whenever deploying an image from that cloud or project. However, the administrative user can set a date with no limitations.

If you are deploying an image from a cloud or project that does not have an expiration policy set, you can choose whether to set an expiration date.

The user who deployed the instance receives an email notification when the instance is about to expire. The user can extend the lease if extensions are enabled.

After the instance expires, it will be stopped. The instance can be automatically deleted after a limited time, which is specified by the grace period. If no grace period is specified, the instance is deleted immediately. This setting applies regardless of whether the instance expiration maximum is set.

Updating the default expiration policy for a cloud

You can update the default expiration policy for IBM Cloud Manager with OpenStack.

About this task

To update the default expiration policy, complete the following steps:

Procedure

1. In the IBM Cloud Manager with OpenStack interface, select **Configuration** > **Clouds**.
2. Click the name of the cloud for which you want to update the expiration policy.
3. Click **Expiration Policies** to open the form.
4. Enter information for the default expiration policy and click **Save**.

   Note: To delete expired instances immediately, set the Grace period to `0`.

Updating the default expiration policy for a project

You can update the default expiration policy for IBM Cloud Manager with OpenStack project.

Procedure

1. In the IBM Cloud Manager with OpenStack interface, select **Access** > **Projects**.
2. Select a project to open the project update page.

   Note: To delete expired instances immediately, set the Grace period to `0`.
3. Enter information for the default expiration policy.
   - If you select **Use cloud default**, the expiration of the deployment depends on the expiration configuration of the cloud to which the image belongs.
   - If you select **Customize settings**, the expiration policy overrides the expiration policy of the cloud to which images belong.
4. Click **OK**.

Managing flavors (OpenStack only)

A flavor is the prescribed size of a provisioned virtual machine. Each flavor has a unique combination of resource configurations and sizes.

Updating the flavor for an OpenStack cloud configuration

You can update the flavor that is configured for the cloud.

Procedure

1. In the IBM Cloud Manager with OpenStack interface, select **Configuration** > **Clouds**.
2. Select the cloud for which you want to modify flavors.
3. Click **Edit**.
4. Expand the **Flavors** section. You can create a flavor based on an existing flavor or you can create a completely new flavor.
5. Click the flavor name that you want to copy or click the **Create a new flavor** icon to create a new flavor.
6. Set the following required values:
   - `Name`
   - `Virtual CPUs`
   - `Memory (MB)`
   - `Storage (GB)`
   - `Swap(MB)` This option is only supported for KVM deployments.

   Notes:
   - When updating the flavor, only integers are valid for the processor, memory, storage, and swap sizes. Any fractional data is omitted.
   - After a flavor is created or updated, it can be used to deploy OpenStack images. However, the new or updated flavor might not be immediately available when configuring an image. If this occurs, you can reset the image configuration to the default values in order to pick up the flavor changes made in the cloud. For more information, see Configuring image deployment properties.
   - Flavors with a storage size of 0 have special meaning in OpenStack and are not supported by numerous hypervisors such as Hyper-V. Hypervisors that do support such flavors use the size of the image as the storage size when provisioning a virtual machine from the image.
7. Optional: Specify extra specifications. For information, see “Extra specifications” on page 196.

Extra specifications

A flavor might include properties that are in addition to the base flavor properties. These extra specifications are key value pairs that can be used to provide advanced configuration that is in addition to the configuration provided by the base flavor properties. This configuration is specific to the hypervisor.

Advanced configuration provided with flavor extra specifications might include the following:

- CPU shares
- CPU period
- disk read and write rates per second

In addition to enabling advanced hypervisor configuration, extra specifications are used in OpenStack as a mechanism to enable advanced placement by the scheduler with specific scheduler filters such as the ComputeCapabilitiesFilter and the AggregateInstanceExtraSpecsFilter. For specific information about using these filters, see the OpenStack documentation.

KVM extra specifications

Flavor extra specifications are supported as part of the KVM support with OpenStack that is provided with IBM Cloud Manager with OpenStack.

The following extra specifications enable tuning the CPU for a virtual machine:

- `quota:cpu_shares`
- `quota:cpu_period`
- `quota:cpu_quota`

The following extra specifications enable tuning the device I/O for a virtual machine:

- `quota:disk_read_bytes_sec`
- `quota:disk_read_iops_sec`
- `quota:disk_write_bytes_sec`
- `quota:disk_write_iops_sec`
- `quota:disk_total_bytes_sec`
- `quota:disk_total_iops_sec`

The following extra specifications enable tuning the network device interfaces for a virtual machine. However, they are not supported in IBM Cloud Manager with OpenStack since they are not used with OpenVSwitch.

- `quota:vif_inbound_average`
- `quota:vif_inbound_peak`
- `quota:vif_inbound_burst`
- `quota:vif_outbound_average`
- `quota:vif_outbound_peak`
- `quota:vif_outbound_burst`

PowerVC extra specifications

Flavor extra specifications are supported as part of the PowerVC support with OpenStack that is provided with IBM Cloud Manager with OpenStack.

The following extra specifications enable tuning the CPU for a virtual machine:

- `powervm:proc_units`

- powervm:min_proc_units
- powervm:max_proc_units
- powervm:min_vcpu
- powervm:max_vcpu
- powervm:dedicated_proc
- powervm:share_idle
- powervm:uncapped
- powervm:shared_weight
- powervm:availability_priority
- powervm:processor_compatibility

The following extra specifications enable tuning the memory for a virtual machine:
- powervm:min_mem
- powervm:max_mem

The following extra specifications enable tuning the boot volumes:
- powervm:boot_volume_type

For more information, see the PowerVC Information Center.

Extra specification labels

Because it can be difficult to determine the purpose of extra specifications from their names, you can label extra specifications.

In the *openstack.properties* file any property that contains the *com.ibm.cfs.openstack.flavor.extraspec.label.* prefix is considered a localizable label for an OpenStack flavor extra spec. The suffix of the property must be the extra spec name, and the value of the property must be the list of localized extra spec labels.

For example, you can provide a label for the quota:cpu_shares extra specification that is similar to the following:

```
com.ibm.cfs.openstack.flavor.extraspecs.label.quota:cpu_shares=CPU Shares, en_us=CPU Shares,
es=Porción de Procesador
```

The first value is the label for the default locale, which is followed by a list of locale=label pairs. Each time the extra specification is displayed in the IBM Cloud Manager with OpenStack user interface, the corresponding label is displayed, based on the locale of the end user.

Managing clouds

You can add, configure, edit, and remove clouds.

In the Cloud status window, you can see the status of all of the clouds that IBM Cloud Manager with OpenStack self-service portal is connected to. To view details about a specific cloud, select **Cloud settings**.

In the **Clouds** section of the **Configuration** tab, you can add, edit, and delete cloud configurations. When you edit cloud configurations, you can also set or update expiration policies and approval policies.

You can customize each cloud to have its own approval and expiration policies. You can also configure a network for a specific cloud.

Related tasks:

"Customizing for self-service portal" on page 100
If you want to use the self-service portal to manage IBM Cloud Manager with OpenStack, you must manually include the self-service portal in a topology.

Adding a cloud configuration

You can configure a cloud in the **Clouds** section of the **Configuration** tab.

Procedure

1. Open IBM Cloud Manager with OpenStack and select **Configuration** > **Clouds**.
2. Click **Add Cloud**.
3. Enter information in each of the fields that is denoted with a red asterisk (*).
4. Click **Add**.

What to do next

After you click **Add**, you are prompted to trust the cloud server SSL certificate. If you do not accept this certificate, the cloud configuration is not added.

Note: If you do not want to accept the default Approval or expiration policies, you can edit the cloud configuration after adding it.

Configuring an OpenStack cloud

After you deploy the IBM Cloud Manager with OpenStack appliance, use the web interface to configure an OpenStack cloud.

Before you begin

Before you begin these steps, ensure that the *ceilometer* and *heat* users exist on the **com.ibm.cfs.cloud.openstack.service.users** property in the `/var/opt/ibm/.SCE41/openstack.properties` file. If the *ceilometer* and *heat* users do not exist, add them and restart the sce service before continuing. For more information, see "User management with OpenStack" on page 213.

Before you configure an OpenStack cloud, you must add an IaaS gateway. An IaaS gateway is a component that is used to manage the OpenStack cloud. In IBM Cloud Manager with OpenStack you can add only one IaaS gateway. The IaaS gateway might have multiple regions, so you can add multiple OpenStack clouds into IBM Cloud Manager with OpenStack and each OpenStack cloud belongs to one region.

To add an IaaS gateway, follow these steps:

1. Log in to IBM Cloud Manager with OpenStack web interface as an administrator.
2. Click the **Configuration** tab and select **Clouds** in the navigation pane.
3. Click **Set up the IaaS Gateway**.
4. Enter the IaaS gateway configuration settings that you want to use to manage the OpenStack cloud.

   Host name
   Select the listening IP address of the `iaasgateway` that is set in the environment file. For more information, see "Managing the IaaS gateway" on page 173.

   Port Accept the default value of **9973**.

   Administrator ID
   Enter *sceagent*, which is the user ID that is used to communicate between IBM Cloud Manager with OpenStack and OpenStack.

Password
Type the password for the sceagent user ID. For password information, see "User names and passwords" on page 223.

Tenant name
Type *Public*.

5. Click **Add** to add the IaaS gateway.

Now that you have added an IaaS gateway, you can configure an OpenStack cloud.

To configure an OpenStack cloud, follow these steps:

Procedure

1. Log in to IBM Cloud Manager with OpenStack web interface as an administrator.
2. Click the **Configuration** tab and select **Clouds** in the navigation pane.
3. Click **Create a new cloud configuration**.
4. Enter the cloud configuration settings for the OpenStack cloud.

   Name Assign a name to the OpenStack cloud that you want to create.

   Description
   Optionally, add a description for the OpenStack cloud.

   Type Select **OpenStack** for the cloud type.

   Region
   Select the region that you want to use for thisOpenStack cloud.

   QPID Settings

   Host name
   Select the host name or the IP address of the OpenStack controller node. For more information, see "Customizing Qpid messaging service attributes" on page 105.

   Port Accept the default value of **5671**.

   Secure the cloud connection using SSL
   Select if the OpenStack cloud is configured to expose an SSL connection. By default it is selected.

   User ID
   Enter *qpidclient*, the user ID that communicates with the QPID messaging server.

   Password
   Type the password for the qpidclient user ID. For password information, see "User names and passwords" on page 223.

5. Click **Add** to finish.

Results

If you receive a PKI token error when you are attempting to configure the OpenStack cloud, see "PKI error when adding OpenStack cloud with self-service portal" on page 264 for more information.

OpenStack clouds:

When you add an OpenStack cloud, IBM Cloud Manager with OpenStack enters a transactional mode for user and project operations. Also, OpenStack relies on Coordinated Universal Time (UTC) time.

When in transactional mode, all user and project operations fail if the OpenStack cloud is unavailable. These operations fail even if the user or project in question is not currently used in connection to OpenStack.

Additionally, IBM Cloud Manager with OpenStack uses OpenStack efficient polling with the **changes-since** parameter support to maintain internal caches of certain OpenStack resources. The OpenStack **changes-since** support relies on Coordinated Universal Time (UTC) time to determine if a resource has changed. As a result, the IBM Cloud Manager with OpenStack and OpenStack systems must maintain accurate and consistent UTC time to avoid caching and other problems that can occur due to incorrect system time.

Multiple cloud support

IBM Cloud Manager with OpenStack allows you to manage multiple clouds from a single instance. For example, you can have a test cloud set up to implement your latest policies before moving those policies to your production cloud.

This support also allows you to manage multiple types of clouds For example, you can have multiple VMware cloud instances available from a single IBM Cloud Manager with OpenStack user interface.

Updating an IaaS gateway

You can update an IaaS gateway in IBM Cloud Manager with OpenStack from the **Clouds** section of the **Configuration** tab.

Procedure

1. Open IBM Cloud Manager with OpenStack and select **Configuration** > **Clouds**.
2. Click **Edit**.
3. Update the IaaS gateway settings.
4. Click **Update**.

Removing an IaaS gateway

You can remove an association with anIaaS gateway from IBM Cloud Manager with OpenStack from the **Clouds** section of the **Configuration** tab.

Procedure

1. Open IBM Cloud Manager with OpenStack and select **Configuration** > **Clouds**.
2. Select the IaaS gateway that you want to remove.
3. Click **Remove**.
4. Click **Yes** to confirm.

Updating a cloud

You can update a cloud in the **Clouds** section of the **Configuration** tab.

Procedure

1. Open IBM Cloud Manager with OpenStack and select **Configuration** > **Clouds**.
2. Click the name of the cloud that you want to update.
3. Update the desired fields and click **Save**.

   Tip: From this configuration panel, you can also review or change the **Expiration Policies**, **Approval Policies**, and **Flavors (OpenStack cloud)** of the cloud. For more information, see the following topics:
   - "Managing expiration policies" on page 194
   - "Managing approval policies" on page 190
   - "Managing flavors (OpenStack only)" on page 195

Removing a cloud

You can remove an association with a cloud from IBM Cloud Manager with OpenStack from the **Clouds** section of the **Configuration** tab.

Procedure

1. Open IBM Cloud Manager with OpenStack and select **Configuration** > **Clouds**.
2. Select the cloud that you want to remove. You can either remove just the cloud, or remove the cloud and all instances that were deployed by IBM Cloud Manager with OpenStack.
3. Click **Remove**.
4. Click **Remove Cloud Only** to remove just the cloud, or click **Remove cloud configuration and all related instances** to remove the cloud and all instances that were deployed by IBM Cloud Manager with OpenStack.
5. Click **Yes** to confirm.

Results

Note: When removing a cloud configuration, all of the cloud resources created by IBM Cloud Manager with OpenStack are lost. Recreating the connection to the cloud at a later time will not recover these resources

Managing network configurations

The IBM Cloud Manager with OpenStack self-service portal provides a convenient way to manage and apply network settings by using network configurations. Network configurations are a group of network settings for a particular environment, typically a virtual network. These settings can be managed as a single entity and applied to image configurations or instance deployment settings.

For example, suppose that a cloud environment contains two virtual networks applicable to instance deployment: a public and a private virtual network. In this case, an administrator might create two network configurations, one for the public and one for the private. In the public configuration, the administrator would specify all the public network settings such as primary DNS, secondary DNS, and primary gateway. The same would be done for the private network configuration. After the configurations are created, the administrator can configure the images to use the appropriate network configuration. This action saves time by not requiring the administrator to specify each network setting in each image. It also allows an easier way to manage the network settings on a virtual network.

While the actual settings specified in a configuration are tailored to a specific environment, the network configurations themselves are a superset of all network settings regardless of image, operating system, or cloud management system. Therefore, all settings that are specified in a configuration are applicable. For example, the primary and secondary WINS settings of a network configuration are only applicable to Windows based images. So when you create a configuration for an image that is not using Windows, these values are not needed and can be left blank.

Note: With the self-service portal, you can specify the network configuration for a cloud. The self-service portal displays only the fields that are applicable for that cloud. Before you can create an OpenStack network configuration, you must select an existing OpenStack cloud.

When network configuration settings are applied to either an image configuration or during an advanced instance deployment, their individual settings can be overridden or manually specified, if wanted.

Note: You cannot override or manually specify OpenStack network configuration settings.

Managing network configurations

You can create, edit, and delete, network configurations from the IBM Cloud Manager with OpenStack web interface.

About this task

To create, edit, or delete a network configuration, follow these steps:

Procedure

1. Open IBM Cloud Manager with OpenStack and select **Configuration**.
2. Select **Network**.

   The network configurations that are defined in the property files are displayed. The Network tab provides a listing of all existing network configurations, and enables you to edit, create, or delete these network configurations.

   - The Network Configuration column shows the name of the existing network configuration.
   - The Cloud column shows the name of the cloud scope that is associated with the network configuration.
   - The Type column shows the IP address version that the network configuration supports. VMware network configurations support only IPv4 addresses, but OpenStack network configurations can support IPv4 or both IPv4 and IPv6 addresses. OpenStack networks do not support IPv6-only addresses.
   - The Available Addresses column shows the number of IP addresses available in the network.
   - The Allocated Addresses column shows the number of IP addresses that are allocated.

   You can edit, create, or delete these network configurations.

   - To view specific network configuration properties, click the network configuration name.
   - To edit specific network configuration properties, click the network configuration name and then click **Edit**.
   - To manage the IP addresses for an existing configuration, click **Manage IP Addresses** for the existing configuration.
   - To create a new network configuration, click the New icon.
   - To create a new network configuration that is based on an existing configuration, select a configuration and click the Copy icon.
   - To delete an existing configuration, select a configuration and click the Delete icon.

Adding a network configuration

You can add a network configuration from the IBM Cloud Manager with OpenStack web interface.

Before you begin

z/VM only

The `/etc/neutron/plugins/openvswitch/ovs_neutron_plugin.ini` file contains a **`network_vlan_ranges`** value similar to the following:

```
network_vlan_ranges = xcatvsw2,xcatvsw3:2:3999
```

Before you can add a network configuration for an OpenStack cloud running on z/VM, follow these steps:

1. Create two networks in the OpenStack environment as follows:
   a. First, run the following command to create the management network:

      Note: The following command must be entered on a single line, even though the example shows a line break for formatting purposes.

```
neutron net-create xcat_management --provider:network_type flat
   --provider:physical_network xcatvsw2
```

b. Next, run the following command to create the customer network:

Note: The following command must be entered on a single line, even though the example shows a line break for formatting purposes.

```
neutron net-create opnstk_compute --provider:network_type vlan
   --provider:physical_network xcatvsw3 --provider:segmentation_id 2
```

2. Define a subnet for each network as follows:
   a. Run the following command to define the subnet for the management network:

```
neutron subnet-create --allocation-pool start=10.1.0.2,end=10.1.0.199 xcat_management 10.1.0.0/16
```

Note: You can change the IP range to any value that you want. The Extreme Cloud Administration Toolkit (xCAT) MN interface can connect this range.

   b. Run the following command to change the IP range for the customer network:

```
neutron subnet-create opnstk_compute cidr
```

Note: You can use any IP address that you want.

When you use IBM Cloud Manager with OpenStack to connect to z/VM, the network configuration in OpenStack is done.

About this task

To add a network configuration, follow these steps:

Procedure

1. Open IBM Cloud Manager with OpenStack and select **Configuration**.
2. Select **Network**.
3. To create a network configuration, click **New**.
4. Specify a cloud scope.

   When you specify a cloud scope, the network configuration that you are adding is only available when you deploy an image to that cloud. If you specify VMware for the cloud scope, the configuration is available to all VMware images. When you specify a cloud scope, this page displays only the fields that are applicable to the selected cloud scope.
5. Enter a name for the configuration.
6. Follow the steps for the cloud scope that you selected.
   - **VMware**
     a. Optionally, enter a description.
     b. Select a unique Network ID.
     c. Select one of the following IP Settings:
        – **Use DHCP**
        – **Use IP address pool**

          If you select **Use IP address pool**, follow these steps:
          1) Specify a Subnet mask and Gateway address. You can also provide an alternative gateway address.
          2) Specify a number of IP addresses and the starting address for allocation. The ending address is calculated based on the number of addresses and starting address.

Note: If you specify a number of IP addresses, the number must be at least 2. To create a single IP address, you must first create then network configuration, and then add the single IP address.

3) Specify DNS Settings.

d. Specify System-wide settings, including Linux and AIX Network Settings and Windows Network Settings.

e. Choose to be a member of a domain or a workgroup:

– **Domain**

If you select Domain, specify the domain name, user, and password.

– **Workgroup**

If you select Workgroup, specify the workgroup name.

f. If you selected **Use IP address pool**, you can also select **Obtain host name and domain name from DNS server**. If you select this option, the DNS used by the system must correlate with the DNS used by this application. If it does not, the names that are obtained might be different from the name that is resolved by the system DNS. The DNS server must be configured correctly for the operating system of the IBM Cloud Manager with OpenStack server. If the names cannot be resolved, the host name prefix and domain name that are provided in this configuration are used.

Note: Only the host name and domain name are configured when you select **Obtain host name and domain name from DNS server**. For other setup, such as the DNS IP address, you must configure those settings manually when an image is deployed.

g. Click **Save**.

- **OpenStack**

a. Select one of the following IP address versions:

– IPv4 only

– IPv4 and IPv6

If you select IPv4 and IPv6, you can enter separate IP address settings for IPv4 and IPv6 addresses. However, the number of IPv6 addresses to allocate must be the same as the number of IPv4 addresses.

b. Specify a subnet mask (for IPv4) or prefix length (for IPv6) and gateway address.

c. Specify a number of IP addresses and the starting address for allocation. The ending address is calculated from the number of addresses and the starting address.

Notes:

1) There must be at least two IP addresses.

d. Specify DNS settings.

e. Specify provider network settings as follows:

– Specify one of the following network types:

- **None selected**

A network is created based on the **tenant_network_type** property in the `/etc/neutron/plugin.ini` file. This value is set to vlan in the SCE image. If this option is used, the physical network name and vlan ID are automatically selected based on the "network_vlan_ranges" property in `/etc/neutron/plugin.ini` file. This property is set to **default:1:4094** in the SCE image.

- **Flat**

A virtual network that is realized as packets on a specific physical network that contains no IEEE 802.1Q header. Each physical network can realize at most one flat network.

- **Local**

A virtual network that allows communication within a host, but not across the network. Local networks are intended mainly for single-node test scenarios.

- **VLAN**

  A virtual network that is realized as packets on a specific physical network that contains IEEE 802.1Q headers with a specific VLAN id. VLAN networks that share a physical network are isolated from each other. Each distinct physical network that supports VLAN networks is treated as a separate VLAN trunk, with a distinct space of VLAN id values.

– If you select **Flat** or **VLAN** for the network type, enter the physical network name.

This physical network name must match the name that is specified in the **network_vlan_ranges** property of the `/etc/neutron/plugin.ini` file.

Note: You can create only one Flat network on each physical network.

– If you select **VLAN**, enter the VLAN ID.

Valid VLAN ID values are 1 through 4094.

f. Click **Save**.

Editing network configurations

You can edit a network configuration from the IBM Cloud Manager with OpenStack self-service portal.

About this task

To edit a network configuration, follow these steps:

Procedure

1. Open IBM Cloud Manager with OpenStack and select **Configuration**.
2. Click **Network**. The network configurations that are defined in the property files are displayed.
3. Select a network configuration that you want to edit from the list of available configurations. The current properties are displayed for the selected configuration. The properties that are displayed depend on the cloud management system for which the network configuration was created.
4. Click **Edit**. You can edit only certain network configuration properties.
5. Change the properties of the configuration. If you want to edit the IP addresses for this configuration, click **Manage IP Addresses**. For more information about setting up an IP address pool, see "Managing IP address pools."
6. Click **Save** to save your changes, or **Cancel** to exit the screen without saving your changes.

Managing IP address pools

IBM Cloud Manager with OpenStack can automatically select the IP address (or IP addresses) to be used when provisioning a virtual machine from a list of predetermined IP addresses known as an IP address pool. IP addresses are managed and assigned automatically to an instance so that the user requesting the deployment does not need to specify them.

About this task

An IP address is marked as "In Use" when IBM Cloud Manager with OpenStack selects that IP addresses from the network configuration and uses it for the deployment of an instance. When the instance is deleted by IBM Cloud Manager with OpenStack, the IP address "In Use" indication is cleared so that the IP address can be reused by another instance deployment. If IBM Cloud Manager with OpenStack detects that the instance has failed and no longer exists in the cloud, the IP address is unlocked immediately and the "In Use" flag cleared.

The administrator can also mark an IP address or a range of IP addresses as "Locked". "Locked" IP addresses are not selected by IBM Cloud Manager with OpenStack for instance deployment. The purpose of "Locked" IP addresses is to allow the administrator to mark certain IP addresses in the network as

reserved or "In Use" by other applications. If the administrator later wants to enable the IP address so that it can be used by IBM Cloud Manager with OpenStack for instance deployment, the "Unlock" option can be used to remove the "Locked" indicator.

The main difference between "In Use" and "Locked" is conceptual; addresses that are "In Use" are being used by the IBM Cloud Manager with OpenStack application, while addresses that are "Locked" are being used by an external application or are not available as specified by an administrator.

Each network configuration contains its own IP address pool, which allows IP addresses to be managed on a per network configuration basis. If a configuration is applied to the deployment settings of an instance (and the configuration is not set to use DHCP), the IBM Cloud Manager with OpenStack automatically uses the pool that is associated with the configuration.

Notes:

1. Network configurations typically represent a VLAN or a virtual network. While a network configuration cannot contain the same IP address more than once, different network configurations can contain the same IP addresses. This behavior was added to allow multiple VLANs to use the same IP address ranges. If the same IP address ranges are specified in multiple network configurations, care must be taken to ensure that these network configurations are used on different networks or VLANs.
2. OpenStack network configurations cannot contain the same IP addresses. Each of the IP subnets that are defined in the OpenStack network configurations must be unique and must not overlap.
3. The IP addresses for an OpenStack network configuration are specified when the OpenStack network configuration is first created. IP addresses cannot be added to or removed from an OpenStack network configuration. Lock and unlock of IP addresses is supported.

The following steps describe how an administrator can manage the IP address pools that are used by the IBM Cloud Manager with OpenStack application.

Procedure

1. Open IBM Cloud Manager with OpenStack and select **Configuration** > **Network**.
2. From the Network page, select **Edit under IP Addresses**.
3. The **IP Addresses** view is displayed. Use this view to add, remove, lock, or unlock IP addresses.
4. To add IP addresses, select **Add**.
   a. Add an individual or range of IP addresses to the pool.
   b. Select **OK** to add the IP address or range, or select **Cancel** to cancel the operation.

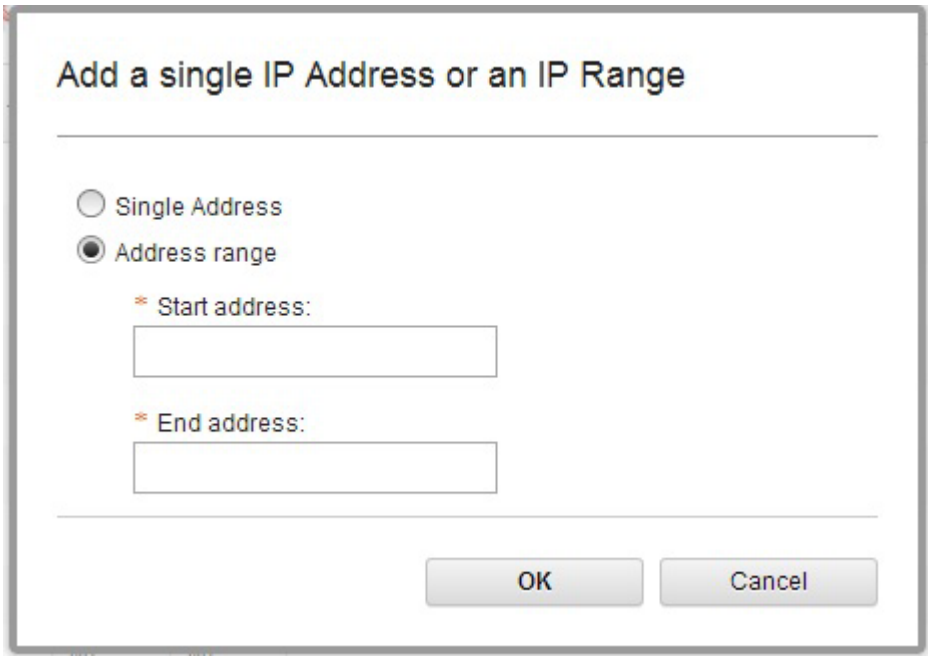

5. To remove, lock, or unlock specific IP addresses, select the IP addresses to which to apply the operation, then select **Remove**, **Lock** or **Unlock** from the Manage IP addresses page to apply the operation.

   Note: The IP addresses page allows for smart selection of IP addresses to which to apply the Remove, Lock, and Unlock operations. When Remove, Lock, or Unlock is selected, smart selection determines whether any addresses are selected on the page. If addresses are selected, the operation is applied to the selected addresses without an extra dialog. If no addresses are selected, a dialog is displayed which allows either individual or range based operations for remove, lock, or unlock.

Managing instances

Use the Instances tab in the IBM Cloud Manager with OpenStack self-service portal to manage instances after they are created.

VMware: An instance in the self-service portal is equivalent to a VMware virtual machine. All of the VMware virtual machines are displayed on the self-service portal Instances tab.

You can filter the list of instances by Cloud, Projects, or Architectures.

As an administrator, you can act on pending instance requests and hide specific instances from appearing to other users.

Note: When you start or restart the self-service portal on a high scale cloud, the synchronization between the self-service portal and the cloud might take longer than expected. This resynchronization might cause operations such as deploying, deleting, or resizing an instance to be delayed or even fail. Wait for the synchronization to complete before you attempt these actions.

Capturing an instance

You can capture an instance or workload to create an image.

When you capture an OpenStack PowerVC instance, a message is displayed to indicate that capturing this instance requires some manual preparation. For information about preparing an OpenStackPowerVC instance for capture, see Capture requirements in the PowerVC information center at the following web

page: http://www.ibm.com/support/knowledgecenter/SSXK2N_1.2.1/
com.ibm.powervc.standard.help.doc/powervc_standard.html

Note: For OpenStack PowerVM, you can only capture instances that are stopped.

For information about capturing an OpenStack instance, see "Considerations for capturing an OpenStack instance" on page 153

Pinning an instance

In a deployed instance, you can pin a virtual machine to a specific physical host to prevent the server from being relocated. However, an instance or workload that is set to be highly available cannot have pinned virtual machines.

Procedure

1. In the IBM Cloud Manager with OpenStack interface, select **Instances**.
2. Select an instance to open the properties.
   - To pin an instance, select **Pin**.
   - To unpin an instance, select **Unpin**.
3. Click **Close**.

Migrating an instance (OpenStack)

In a deployed instance, you can migrate a virtual machine to a specific physical host. You can also migrate multiple instances to a specific physical host.

Before you begin

Make sure that the instances that you want to migrate have a status of OK. If you are migrating multiple instances, all the instances must be running on a common OpenStack cloud.

Note: If you are using a PowerVC virtualization environment, overcommitting a disk and migrating non-shared storage are not supported for PowerVM instances. For more information about PowerVC live migration settings and capabilities, see the PowerVC Information Center.

Procedure

1. In the IBM Cloud Manager with OpenStack interface, select **Instances**.
2. Select the instances that you want to migrate.
3. Click **More** > **Migrate to new host** to open the live migration page.
4. Select a destination host as follows:
   - To specify that the system select a destination host, select **Allow system**.

     If the associated instance is a PowerVM instance and its **Use PowerVC placement policy** virtual machine property is *true*, then the *Allow system* property indicates that the PowerVC scheduler selects a host.
   - To manually select a destination host, select **Manually select** and then select a destination host.
5. Click **Migrate**. The instance status changes to Migrating.
6. Click the refresh arrow to update the status. When the status changes from Migrating to OK, the migration is complete and the instance is available

Processing requests from the Instances tab

When an image is deployed, initiating an instance, the deployment request may require approval by an administrator. In this case, the instance status is set to pending until the administrator handles the approval request.

About this task

You can process an instance request from the Instances tab or from the Requests tab. For more information about processing an instance request from the Requests tab, see “Processing instance requests” on page 193

To process a pending request, follow these steps:

Procedure

1. In the IBM Cloud Manager with OpenStack interface, select **Instances**.
2. Select an instance name to view the instance details. Find the request list in the instance details and select a request to display. The Request properties page appears.
3. Expand the **Request Details** section to review or update the request before approving.
4. Expand the **Comments** section to review comments or use the **Add Comment** link to provide additional comments.
   - Click **Approve** to approve the request and allow the deployment processing to start.
   - Click **Reject** to reject the request.
   - Click **Withdraw** to withdraw a request.

Hiding or showing an instance

Follow these steps to show or hide an instance.

Procedure

1. In the IBM Cloud Manager with OpenStack interface, select **Instances**.
2. Select an instance and click **Hide/Show** to hide or show the instance in the instance list for all non-administrative users.
3. After an instance is hidden, a non-administrative user does not see the instance in the instance list, but administrative users can choose to display the hidden instance. To display hidden instances in the instance list, select **Include hidden instances**.

Resizing an instance (VMware)

You can modify the amount of resources used by the virtual machines provisioned by your instance running on VMware. Depending on how your VMware virtual machines are configured, you can add memory and virtual processors while your virtual machine is running.

About this task

Increasing the size of the virtual machine disks makes more space available on the disk, but does not change the size of the partitions and the file systems. There are commands that must be run on the guest operating system to increase the size of the file system. For more information about how to change the size of the file system after storage is added, see your operating system documentation.

For more information about how a running virtual machine handles changes in memory and processor, see the VMware documentation and your operating system documentation.

Procedure

1. Click the name of the instance that you want to resize.
2. Click **More** > **Resize**.
3. Update the number of processors and memory resources to be allocated to the virtual machine in your instance.

   The settings that can be resized when a virtual machine is in the started state depend on how the virtual machine is configured on VMware:

Notes:

a. If the instance is started and the virtual machine is not configured to allow memory or processor changes, those fields are not displayed. To change those values, you must first stop the instance.
b. For memory, the virtual machine must have the Memory Hot Add option enabled. Memory is only allowed to be increased, and the maximum amount that is allowed, and the valid values, are determined by VMware.
c. For processors, the virtual machine must have the processor Hot Plug option enabled. To remove processors, the virtual machine must have the processor Hot Add and Remove option enabled. The maximum number of processors that are allowed is determined by the number of logical processors on the vSphere machine that is running the virtual machine.
d. If you are changing the storage size, you can update only to a larger disk size.

4. Increase the disk size.
5. Click **Resize**.

Note:

- If approvals are enabled, then the approval must be completed before the instance is resized.
- Linked clone disks or disks that are using an IDE controller cannot be resized.

Resizing an instance (OpenStack)

You can modify the amount of resources that are used by the virtual machines.

About this task

Stop the instance before you continue the procedure.

Procedure

1. Click the name of the instance that you want to resize.
2. Click **More** > **Resize...** to open the Resizing instance page.
3. Under the **Hardware** section, update the **OpenStack Flavor** to be allocated to the virtual machine in your instance.

Notes:

- The flavor details change depending on the size flavor that you select.
- When you update the flavor, the processor, memory, and storage size fields accept integers only. Any fractional data is omitted.
- (PowerVM and Hyper-V) If you are changing the storage size, you can update to a larger disk size only.
- (KVM and PowerKVM only) If you choose a flavor with a smaller storage size, KVM skips to storage resize if it cannot be completed. The other resources are resized accordingly.
- (KVM and PowerKVM only) Resizing an instance only supports local storage.
- (z/VM only) You can change only the CPU and memory resources assigned to the virtual machine.

4. Click **Resize**.

Notes:

- If approvals are enabled, then the approval must be completed before the instance is resized. To verify that the instance was resized, check the virtual machine **flavor ID** property of the instance.
- If the instance is running, OpenStack stops, resizes, and restarts the instance after the hypervisors are resized (except PowerVM instance).
- Active resize is supported for PowerVC instances. The instances stay active during the resizing progress. The active resize action is limited by instance resource ranges (maximum and minimum values of processor, memory, and processor units), and unsupported flavors are filtered. In addition,

the active resize action is not allowed unless the health status of the instance is *OK*. For more information about resource ranges, see “Extra specifications” on page 196.

If you are resizing an instance on the Hyper-V hypervisor, the **IBMComputeNodeService** service that is deployed with the Hyper-V agent installer must run with domain credentials and configure Kerberos constrained delegation. You can set the service credentials by using the following command:

```
C:\sc config "IBM SmartCloud Entry Compute Service" obj="DOMAIN\username" password="password"
```

.

To configure the Kerberos constrained delegation setting, see step 1 in the following guide: Configure constrained delegation.

Managing storage volumes

If you have an OpenStack cloud, you can use the **Volumes** tab in IBM Cloud Manager with OpenStack self-service portal to create and manage storage volumes for your virtual machines. You can manage the volumes by sorting on the cloud and project to which the volumes and associated virtual machines belong.

About this task

From the **Volumes** tab you can create, delete, edit, and attach and detach storage volumes for virtual machines in your OpenStack environment.

When you add a volume, you can choose the source for the new volume. You can create an empty volume, or you can create a volume that is sourced from an existing image or snapshot of a volume. To create an image from a snapshot, you must capture an existing volume and use the captured volume as the snapshot source.

When you delete a volume, all data is lost from the volume. Alternatively, you can detach a volume from a virtual machine to retain the content of the volume and use it in the future by attaching it to a different instance.

Procedure

1. On the **Volumes** tab, use the **Cloud** and **Project** controls to sort the volumes that you want to manage.
2. Use the **Add** and **Delete** icons to add and delete volumes.
3. Use the **More** menu to complete tasks such as**Attach**, **Detach**, and **Capture** on a selected volume.
4. To edit an existing volume, click the volume and select **Edit**.

What to do next

Note: You also can manage storage volumes for a specific virtual machine from the **Instances** tab. Click the virtual machine for which you want to manage storage, and expand **Storage Volumes**.

Managing users

Administrator The **Users** tab in the IBM Cloud Manager with OpenStack self-service portal is enabled for administrative users and is used for creating, viewing, and managing users.

Creating a user

Complete the following steps to create a user.

Procedure

1. In the IBM Cloud Manager with OpenStack self-service portal, select **Access**.
2. Select **Users**.
3. Click **New User**.
4. Enter information for the new user.
5. Click **Create**.

Notes:

- You can only create valid user accounts when using local authentication. When using LDAP authentication, user accounts are created and managed directly through the LDAP server.
- A user can also request a new user account. To approve or reject these requests, go to **Access** > **Requests**.

Viewing or updating a user

About this task

To view or update information about a user, follow these steps:

Procedure

1. In the IBM Cloud Manager with OpenStack interface, select **Access**.
2. Select **Users**.
3. To view or update information about a user, select the user you want to view.

Locking a user

If you want to prevent a user from accessing the IBM Cloud Manager with OpenStack self-service portal, you can lock the user account.

About this task

To lock a user, follow these steps:

Procedure

1. Open the IBM Cloud Manager with OpenStack self-service portal and select **Access**.
2. Select **Users**.
3. Select the user to unlock and click **More** > **Lock**.

Unlocking a user

If a user has three invalid login attempts in a 24 hour period, the user account becomes locked and requires an administrator to unlock it.

About this task

To unlock a user, follow these steps:

Procedure

1. Open the IBM Cloud Manager with OpenStack self-service portal and select **Access**.
2. Select **Users**.
3. Select the user to unlock and click **More** > **Unlock**.

What to do next

Note: If the default administrator account becomes locked, it is unlocked when the server is restarted.

Deleting a user

Complete the following steps to delete a user.

Procedure

1. In the IBM Cloud Manager with OpenStack interface, select **Access**.
2. Select **Users**.
3. Select the user you want to delete from the list of users and click **Delete**.
4. To confirm the user deletion, select **Yes**. To cancel the user deletion, select **No**.

   Note: A user can also request to delete his user account. To approve or reject these requests, go to **Access** > **Requests**.

User management with OpenStack

Unlike other cloud types, OpenStack clouds provide native support for user management through the OpenStack keystone component.

When you first connect to an OpenStack cloud, IBM Cloud Manager with OpenStack imports all the user accounts that currently exist in OpenStack. All user roles and project membership are accepted and reflected in IBM Cloud Manager with OpenStack.

After IBM Cloud Manager with OpenStack imports the initial OpenStack users and connects to an OpenStack cloud, IBM Cloud Manager with OpenStack enters transactional mode for user management. When in transactional mode, all operations that are performed in IBM Cloud Manager with OpenStack are also performed in OpenStack (for example, keystone). If a user management operation (such as any of the operations that are described in this section) fails to complete successfully in IBM Cloud Manager with OpenStack, it does not occur in OpenStack. Likewise, if it fails in OpenStack it reverts in IBM Cloud Manager with OpenStack.

IBM Cloud Manager with OpenStack enters transactional mode for user operations while connected to OpenStack so that the user registries in both products are always synchronized. For this reason, when connected to an OpenStack cloud, it is not possible to perform user-related operations while the OpenStack cloud is down or unavailable.

To connect to OpenStack, IBM Cloud Manager with OpenStack uses a service user account and a default service tenant. Some installations of OpenStack have user accounts specific to OpenStack components (for example, nova, keystone, neutron). These and other service user accounts or service tenants in an OpenStack server that do not represent an actual user account or tenant, can be added to the list of service users and service tenants so that they are ignored by IBM Cloud Manager with OpenStack. To make this change, add the service users and tenants to the comma-separated list of users in the *com.ibm.cfs.cloud.openstack.service.users* property, or the comma-separated list of tenants in the *com.ibm.cfs.cloud.openstack.service.tenants* property, in the *openstack.properties* file.

Managing accounts

You can view information for those accounts of which you are either an owner or a member.

Accounts are required when IBM Cloud Manager with OpenStack self-service portal billing is enabled. See the following guidelines for self-service portal billing:

- Only self-service portal administrators can create accounts, but you can be made an account owner.
- You can deploy instances only if you are an account member and the account has a positive balance with which to pay for server use.
- Only account owners and self-service portal administrators can manage accounts.
- Accounts have a balance, an owner, an account balance threshold, account members, and invoices.

- The *balance* is a monetary balance of the account. The cost of each request and running deployment is subtracted from the balance over time.
- The account *owner* is the self-service portal user profile that is accountable for crediting and paying the account.
- The *account balance threshold* is a value that represents the amount at which the account balance becomes a *low balance*. If the balance drops to zero, the account is delinquent.
- The *account members* are self-service portal users that belong to the account. When account members deploy instances in the self-service portal, the instances are billed to their account.
- Each instance has an *invoice*. An account can have many invoices, which are viewable from the Account properties window.

Creating an account

You can create an account at any time.

Procedure

1. Click **New Account**.
2. Enter information for the new account. Both the **Account name** field and the **Account owner** field are required.
3. Click **Create.**

Add members to an account

You can add members to your account at any time, however, users can only be members of one account at a time.

Procedure

1. In the account table, select the account to which you want to add members.
2. To open the account member management window, click **Edit list**.
3. To add a member, select the member to be added from the **Available users** list and click **Add**.

Viewing or managing an account

You can view the properties of any account, or manage the accounts that you own.

About this task

To view account properties or manage accounts that you own, select the **Access** tab and click **Accounts**. Then, you can select the account that you want to work with in the account table.

Deleting an account

You can delete an account only if you are the owner of the account, and only when the account is not associated with any active instances.

Procedure

1. In the account table, select the account you want to delete.
2. Click the **Delete** icon and confirm the deletion.

Clearing or archiving events

From the Events tab, you can see events such as instance completion, instance failure, new account requests, and new accounts created. You can also clear or archive events. Clearing an event deletes it. Archiving an event saves it to an archive folder. By clearing events, you can free space on your system and improve performance in the IBM Cloud Manager with OpenStack self-service portal. Archive any events that you might want to reference in the future.

About this task

To clear or archive an event, follow these steps:

Procedure

1. In the self-service portal, select **Reports** > **Events**.
   - To clear an event, click **Clear**.
   - To archive an event, click **Archive**.
2. Use the Events filter to select a subset of events to clear or archive. Filter by severity or start and end date. If you filter by date, you must provide an end date.
   - To clear the selected events, click **Clear**.
   - To archive the selected events, click **Archive**. The archived events are saved to a file called `events_<current time in milliseconds>.csv`. This file is in the `archives` folder, which is in the self-service portal configuration directory.

Viewing capacity statistics

Using the Capacity view, you can identify the current capacity of the resources in your virtualization environment. Understanding the capacity of resources within the cloud helps you gauge the health of the entire cloud. It also helps you determine suitable targets to which you might deploy instances.

The **Capacity** tab shows the total, allocated available resources of a host or resource pool, including the number of virtual processors, memory, and storage size. The usage rate shows the real-time metrics of hosts and virtual machines, like processor, memory, storage usage, available storage size, disk I/O requests, network I/O packets, and more.

To access the Capacity view, click the **Reports** tab and then select **Capacity** from the left navigation.

Used This field shows a summary of all allocated resources, regardless of the state of the guest operating system.

Physical
This field shows the physical capacity.

Draft comment
Do we need to mention/explain the overcommit rate here?

The color of the capacity indicator can be green or yellow. Green indicates that the used resources are less than the physical resources. Yellow indicates that the used resource is overcommitted on the available physical resources, but you can still deploy.

You can also access the individual instance to see the target deployment grid that displays live metrics for the resources you want to deploy or migrate.

Notes:

1. The allocated processors and memory that is displayed for a IBM PowerVC host might not match the actual allocation in the PowerVC environment. This is because the IBM Cloud Manager with OpenStack appliance environment with OpenStack does not support fractional processor units and manages only resources that are owned by the Storage Connectivity Group that is defined in the PowerVC driver configuration.
2. The physical capacity for processors and memory that is displayed for a PowerVC host does not include resources that are reserved in the PowerVC environment.
3. The Capacity view does not display storage data for PowerVC clouds.

Managing with OpenStack dashboard

Use this information if you are managing your cloud using the OpenStack dashboard, which is based on OpenStack Horizon. The IBM Cloud Manager with OpenStack dashboard is intended for administrator use only.

If you choose, you can use the OpenStack dashboard to manage your cloud environment. After you install and deploy IBM Cloud Manager with OpenStack, access the dashboard. From a supported browser, type `https://controller.fqdn.com/`, where **`controller.fqdn.com`** is the fully qualified domain name of the controller node in your topology.

Use the default user name and password for initial login, depending upon which topology you deployed:

Minimal topology

- User name: admin
- Password: admin

Controller +*n* compute topology

- User name: admin
- Password: openstack1

Tip: The user name and password might be different if you specified alternative values during the IBM Cloud Manager with OpenStack deployment process. Consider changing these values after initial login for security reasons.

You can configure some of the properties in the dashboard as well. For specific configuration options, see "Configuring IBM Cloud Manager with OpenStack dashboard properties" on page 166.

The OpenStack dashboard is based on OpenStack Horizon. For information about the OpenStack dashboard, see Horizon: The OpenStack Dashboard Project.

Backing up and restoring IBM Cloud Manager with OpenStack

To protect your IBM Cloud Manager with OpenStack data, you must back up critical files in case the server enters an undesired state. Before you back up your data, determine the circumstances in which you intend to restore your data.

Backing up and restoring the deployment server

While there are many techniques for backing up and restoring servers, the following procedure backs up the minimal set of data needed for a successful recovery.

How you choose to back up the deployment server data varies depending on how you deployed the deployment server and the backup and recovery strategy and infrastructure of your organization. Regardless of the technique that is chosen, it is important to test both your backup and restore processes.

Backing up the deployment server

1. Make a directory to copy the backup data. This directory can be anything that you choose, but you must ensure that it has enough disk space to hold the backup data. The commands in this document use the following directory:

   ```
   mkdir /tmp/backup
   ```

2. If you changed the Chef server configuration files, you must back up your changes. IBM Cloud Manager with OpenStack sets these values during installation; therefore, no changes are typically required. If you used a response file for installation, you must backup the response file and keep it with your backups so it can be used to reinstall IBM Cloud Manager with OpenStack.

```
cp ~/cmwo_install.rsp /tmp/backup
```

3. Back up the certificates of the webserver.

```
mkdir /tmp/backup/nginx
cp -a /var/opt/chef-server/nginx/ca /tmp/backup/nginx/ca
```

4. Back up the bookshelf data.

```
mkdir /tmp/backup/bookshelf
cp -a /var/opt/chef-server/bookshelf/data /tmp/backup/bookshelf/data
```

5. Back up the postgres database. The **pg_dump** command produces a consistent backup even if the server processes are running.

```
mkdir /tmp/backup/database

su - opscode-pgsql
pg_dump -c opscode_chef > /tmp/chefserver.sql
exit

cp /tmp/chefserver.sql /tmp/backup/database
```

6. Back up your topology, node attribute, and encrypted data bag, secret key files. By default, these files are in your `/root/your-deployment-name` directory.

```
cp -a /root/your-deployment-name /tmp/backup/
```

7. Back up the `/root/.chef/knife.rb` file.

```
cp /root/.chef/knife.rb /tmp/backup
```

8. Back up the server SSL certificate.

```
cp /root/.chef/os-management-server.pem /tmp/backup
```

9. If you are using identity files rather than passwords, back up identity files of the node. The location depends on your configuration. When you deployed the node, you needed to specify either the node's password or identity file in your topology files. Check your topology files to see whether you used identity files and find their locations.
10. If you customized the `file-repo` or the `yum-repo`, back up those changes.
11. Compress the `/tmp/backup` directory and store in a safe location.

```
zip -r backup /tmp/backup
```

Restoring the deployment server

Start with a clean system and do the following:

1. Retrieve your backup archive file that you want to restore and extract it to `/tmp/backup`.

```
unzip backup.zip
```

2. Install IBM Cloud Manager with OpenStack. Use the response file if you used one during the initial installation.
3. Apply any IBM Cloud Manager with OpenStack fix packs.
4. Apply any `file-repo` or `yum-repo` changes that were made or restore those changes from the backup archive file.
5. Restore the postgres database.

```
/opt/chef-server/embedded/bin/psql -U opscode-pgsql opscode_chef
< /tmp/backup/database/chefserver.sql
```

6. Stop the server processes.

```
chef-server-ctl stop
```

7. Restore the bookshelf data.

```
cp -a /tmp/backup/bookshelf/data/* /var/opt/chef-server/bookshelf/data/
```

8. Restore chef configuration data.

```
cp -a /tmp/backup/nginx/ca /var/opt/chef-server/nginx
```

9. Restore the `knife.rb` file.

```
cp -a /tmp/backup/knife.rb /root/.chef/
```

10. Restore the certificate of the webserver.
```
cp -a /tmp/backup/os-management-server.pem /root/.chef
```
11. Restore your topology, node attribute, and data bag encryption, secret key files.
```
cp -a /tmp/backup/your-deployment-name /root
```
12. If you are using identity files rather than passwords, restore the identity files of the node. The location depends on your configuration.
13. Restart the chef server.
```
chef-server-ctl start
```
14. Reindex the chef server.
```
chef-server-ctl reindex
```

Your deployment server restore is now complete.

Backing up and restoring the OpenStack databases

The backup procedure is different depending on the database that is being used.

Backing up database data for recovery

All of the database data that is related to IBM Cloud Manager with OpenStack users, such as projects, networks, instances, images, are stored in the database. The backup procedure is different depending on the database that is being used.

1. Stop the IBM Cloud Manager with OpenStack services to ensure that the backup data is complete.
2. Follow the instructions that pertain to your specific database.

MySQL

For more information, see Backup and Recovery in the OpenStack community documentation.

DB2 OpenStack database

DB2 can also be used to host the OpenStack databases on the cloud controller. The following shell script backs up a list of DB2 OpenStack databases. This shell script must be modified to match the environment. The variable *DB* is where the DB2 program is located. The variable *DBLIST* is a list of OpenStack databases to back up. The variable *DBBACKUPDIR* is a directory to save the backups in. Since this shell script quiesces each database, while this shell script runs, calls to these databases fail. The backup directory contains the backup files that appear similar to the following example: `NOVA.0.db2inst1.DBPART000.20140523134214.001` The following shell script is an example DB2 OpenStack database backup script:

```
#!/bin/bash
#########################################################
# DB2 OpenStack DB backup Linux shell script
# MODIFY these variables to fit your environment
#    DB2 - home directory of DB2
#    DBLIST - list of Openstack DB to be backed up
#    DBBACKUPDIR - directory holding backed up copies
#########################################################
DB2="/opt/ibm/db2/V10.5/bin/db2"
DBLIST="nova glance"
DBBACKUPDIR="/tmp/DB2backupOpenstack"

mkdir -p $DBBACKUPDIR
for DB in $DBLIST
do
su - db2inst1 <<EOSU
echo "------DB $DB is being backed up ------"
$DB2 connect to $DB
$DB2 quiesce database immediate force connections
$DB2 connect reset
```

```
echo $DB2 backup database $DB to $DBBACKUPDIR without prompting
$DB2 backup database $DB to $DBBACKUPDIR without prompting
$DB2 connect to $DB
$DB2 unquiesce database
$DB2 connect reset
EOSU
done
```

For more information, see Backup and Recovery in the OpenStack community documentation.

Restoring the DB2 OpenStack databases

DB2 can also be used to host the OpenStack databases on the cloud controller. The following shell script can restore a DB2 OpenStack database. This shell script must be modified to match the environment. The variable *DB2* is where the DB2 program is located. The variable *DBBACKUPDIR* is a directory where the backup files are located. The Variable *DB* is the OpenStack database to restore. The variable *DBBACKUPTIME* is the time stamp part of the backup file name, found in the backup directory. Since this shell script quiesces each database, while this shell script runs, calls to these databases fail. Here is an example DB2 OpenStack database restore shell script:

```
#!/bin/bash
#########################################################
# DB2 OpenStack DB restore Linux shell script
# MODIFY these variables to fit your environment
#    DB2HOME - home directory of DB2
#    DB - database name to be restored
#    DBBACKUPDIR - directory holding backed up copies
#    DBBACKUPTIME - timestamp part of backupname
#        example : 20140522162140
#########################################################
DB2="/opt/ibm/db2/V10.5/bin/db2"
DBBACKUPDIR="/tmp/DB2backupOpenstack"
DB="nova"
DBBACKUPTIME="20140523135648"

su - db2inst1 <<EOSU
$DB2 connect to $DB
$DB2 quiesce database immediate force connections
$DB2 connect reset
echo $DB2 restore database $DB from $DBBACKUPDIR taken at $DBBACKUPTIME without prompting
$DB2 restore database $DB from $DBBACKUPDIR taken at $DBBACKUPTIME without prompting
$DB2 connect to $DB
$DB2 unquiesce database
$DB2 connect reset
EOSU
```

For more information, see Backup and Recovery in the OpenStack community documentation.

Backing up and restoring the self-service portal

To protect your self-service portal data, you must back up critical files in case the server enters an undesired state. Before you back up your data, determine the circumstances in which you intend to restore your data.

Backing up self-service portal data for recovery

There are two kinds of data to back up. The first set of data is for server configuration and the second set of data is used by the database. When you consider what data to back up, review both sets of data.

Note: This procedure backs up the self-service portal only. It does not back up the underlying virtualization managers, such as VMware vCenter or storage devices.

1. Stop the IBM Cloud Manager with OpenStack server to ensure that the backup data is complete.

2. Back up the following configuration files.
   - The `.SCE41` folder.
   - In the `installation` folder: `skc.ini`

     Note: If any values are changed in this file, you must back up the updated file after you change default values.

A copy of all these files is required to ensure a complete backup.

Backing up database data for recovery

All of the database data that is related to IBM Cloud Manager with OpenStack users, such as projects, networks, instances, images, are stored in the database. The backup procedure is different depending on the database that is being used.

1. Stop the IBM Cloud Manager with OpenStack server to ensure that the backup data is complete.
2. Follow the instructions that pertain to your specific database.

   Derby database
   If you are using the Derby database, backup the `.SCE41/database folder` that stores all the database data.

   DB2 database
   If DB2 is configured, backup the database in the DB2 server. For more information about how to back up the DB2 server, see the DB2 product information.

   Note: Ensure that the information referenced matches the version of DB2 that you are using. If not, reference the appropriate support documentation for similar information.

Restoring the server

To restore a backup of the server and the Derby database, copy all the saved files back to the original server. After the copy is complete, start the IBM Cloud Manager with OpenStack server.

If you are using the DB2 database, there are some extra steps.

1. Ensure the path that is specified in the `database.properties` configuration file, by the property *database.db2.path*, is correct.

   ```
   # If db2 then the path to the DB2 dtatbase needs to be provided. This will be ignored for derby.
   #database.db2.path=//localhost:50000/cfs:.
   ```

Essentially, creating a backup of the entire `home` folder and `skc.ini` file ensures a complete backup of the IBM Cloud Manager with OpenStack server. Copying the files back to their original location restores the data.

Important: There is a limitation to be aware of when you use this best practice. Any incremental data that occurs after backing up the IBM Cloud Manager with OpenStack server is lost after you restore the server. Therefore, some functions might not work as expected. For example, consider the circumstance where you create an instance after you complete a capture, and then you restore the server. The IP address that was assigned to the instance (after the backing up) is still available in the IP address pool. It might be assigned to another instance.

Chapter 8. Security

IBM Cloud Manager with OpenStack offers security options such as secure sockets layer (SSL), Lightweight Directory Access Protocol (LDAP), and user administration. This section provides information on managing passwords that are associated with the security options.

Passwords

IBM Cloud Manager with OpenStack uses keys to encrypt and decrypt passwords and other sensitive information. If using the self-service portal, there is a protected file named `cfs.keystore` that stores the randomly-generated Data Encryption Standard (DES) key that IBM Cloud Manager with OpenStack uses.

The following list provides links to various sections in this document that describe default passwords and places where passwords are entered and stored in IBM Cloud Manager with OpenStack.

- "Customizing passwords and secrets" on page 102
- Chapter 3, "Installing and uninstalling IBM Cloud Manager with OpenStack," on page 19
- "Configuring secure shell (SSH) communication" on page 125
- "Configuring REST API authentication" on page 131
- "Configuring database" on page 132
- "Configuring global image deployment" on page 135
- Set secure access during deployment (Linux on VMware)
- "Creating a configuration strategy" on page 148
- "Configuring the default administrator user account" on page 175
- "Configuring LDAP authentication using the web interface" on page 175
- "Adding a network configuration" on page 202

Port usage

The number of ports that IBM Cloud Manager with OpenStack uses depends upon the various components and which network interfaces allow access to them.

The following tables list the default ports that might apply to your IBM Cloud Manager with OpenStack environment, depending upon your configuration. You can change these values to customize them for your environment as well.

Table 30. Port usage for IBM Cloud Manager with OpenStack single controller

Port	Service	Notes
22	sshd	SSH access that uses the customer network can be enabled.
67	DHCP server	
68	DHCP server	
80	openstack-dashboard-server	Provides access to the Horizon dashboard.
443	openstack-dashboard-server	Provides https access to the Horizon dashboard.
5000	openstack-identity-api	
35357	openstack-identity-admin	

© Copyright IBM Corp. 2012, 2014

Table 30. Port usage for IBM Cloud Manager with OpenStack single controller (continued)

Port	Service	Notes
5671	openstack-messaging-server	
6080	openstack-compute-novnc	
8000	openstack-orchestration-api-cfn	
8003	openstack-orchestration-api-cloudwatch	
8004	openstack-orchestration-api	
8774	openstack-compute-api	
8776	openstack-block-storage-api	
8777	openstack-telemetry-api	
9191	openstack-image-registry	
9292	openstack-image-api	
9696	openstack-network-api	
9973	ibm-openstack-iaas-gateway	
27017	openstack-database-nosql	DB2 NoSQL wire protocol listener for access to ceilodb2 database
50000	openstack-database-server	

Table 31. Port Usage - IBM Cloud Manager with OpenStack self-service portal

Port	Service	Notes
7777	sce	OSGi console, access from localhost only
18080	sce	HTTP port

Table 32. Port usage - Compute nodes

Port	Service	Notes
22	ssh	SSH port. This port must be accessible from the Chef server.
5900 - 5999	vnc-server	Only applicable to KVM/QEMU or PowerKVM compute nodes.

Table 33. Port usage - Chef server

Port	Notes
1443	The secure (HTTPS) port for accessing the Chef server. The port number is configurable.
1480	The non-secure (HTTP) port for accessing the Chef server. The port number is configurable.

User names and passwords

The following default user names and passwords are included with IBM Cloud Manager with OpenStack.

Default user names and passwords

Table 34. OpenStack credentials

User name	Password	Default data bag item (that contains the password)	Description
ceilodb2	ceilodb2 (minimal topology) and ceilometer (controller +*n* compute or distributed database topology)	db_passwords/ ceilometer	Used for internal communication to Ceilometer database.
cinder	cinder	db_passwords/cinder	Used for internal communication to the Cinder database.
glance	glance	db_passwords/glance	Used for internal communications to the Glance database.
heat	heat	db_passwords/heat	Used for internal communications to the Heat database.
dash	horizon	db_passwords/horizon	Used for internal communications to the Horizon database.
keystone	keystone	db_passwords/keystone	Used for internal communications to the Keystone database.
neutron	neutron	db_passwords/neutron	Used for internal communications to the Neutron database.
nova	nova	db_passwords/nova	Used for internal communications to the Nova database.
powervc	powervc	db_passwords/powervc	Used for internal communications to the PowerVC driver database.
admin	admin (minimal topology) openstack1 (controller +*n* compute or distributed database topology)	user_passwords/admin	OpenStack and (if installed) the self-service portal administrator user.
ceilometer	openstack-ceilometer	service_passwords/ openstack-ceilometer	OpenStack service user for Ceilometer.
cinder	openstack-block-storage	service_passwords/ openstack-block-storage	OpenStack service user for Cinder.
glance	openstack-image	service_passwords/ openstack-image	OpenStack service user for Glance.
heat	openstack-orchestration	service_passwords/ openstack-orchestration	OpenStack service user for Heat.
neutron	openstack-network	service_passwords/ openstack-network	OpenStack service user for Neutron.
nova	openstack-compute	service_passwords/ openstack-compute	OpenStack service user for Nova.

Table 34. OpenStack credentials (continued)

User name	Password	Default data bag item (that contains the password)	Description
powervc	openstack-powervc-driver	service_passwords/ openstack-powervc-driver	OpenStack service user for PowerVC driver.
gwagent	openstack-iaas-gateway	service_passwords/ openstack-iaas-gateway	OpenStack service user for IaaS gateway.
db2das1	db2das1 (minimal topology) openstack1 (controller +*n* compute or distributed database topology)	user_passwords/db2das1	DB2 administrator user.
db2fenc1	db2fenc1 (minimal topology) openstack1 (controller +*n* compute or distributed database topology)	user_passwords/ db2fenc1	DB2 fenced user.
db2inst1	db2inst1 (minimal topology) openstack1 (controller +*n* compute or distributed database topology)	user_passwords/db2inst1	DB2 instance user.
qpidadmin	qpidadmin (minimal topology) openstack1 (controller +*n* compute or distributed database topology)	user_passwords/ qpidadmin	Qpid administrator user. Used for communication with the Qpid messaging server.
qpidclient	qpidclient (minimal topology) openstack1 (controller +*n* compute or distributed database topology)	user_passwords/ qpidclient	Qpid client user. Used for communication with the Qpid messaging server.
qpidssl	qpidssl (minimal topology) openstack1 (controller +*n* compute or distributed database topology)	user_passwords/qpidssl	Qpid SSL user. Used for communication with the Qpid messaging server.
sceagent	sceagent (minimal topology) openstack1 (controller +*n* compute or distributed database topology)	user_passwords/sceagent	IBM Cloud Manager with OpenStack self-service portal user. Used to set up the IaaS Gateway and add an OpenStack cloud.

Note: You can customize all passwords when you deploy a topology.

Chapter 9. Reference

The following topics provide reference information for IBM Cloud Manager with OpenStack.

Commands

IBM Cloud Manager with OpenStack offers a command-line interface.

Refer to the following commands and their descriptions. For instructions to use the command, open a command line and enter **--help**. For example, **knife os manage deploy topology --help**. To assist in analyzing problems with the commands, enter **-VV** to obtain debug level information. For example, **knife os manage deploy topology simple_topology.json -VV**. You can also review the sample topology file, keys, and descriptions provided by IBM Cloud Manager with OpenStack.

Table 35. Commands and descriptions

Command	Description
knife os manage deploy evaluation [NODE_FQDN] (options) where: • **NODE_FQDN** is the host fully qualified domain name or IP address for the node. The fully qualified domain name (FQDN) of the node is optional. However, if it is not specified, you are prompted for it and it defaults to the FQDN of the host the command is running on. The command also prompts for the SSH root user password of the node being deployed to.	Used to deploy an evaluation node.
knife os manage deploy node NODE_FQDN RUN_LIST_ENTRY[,RUN_LIST_ENTRY] (options) where: • **NODE_FQDN** is the host fully qualified domain name or IP address for the node. • **RUN_LIST_ENTRY** is a comma-separated list of run list items such as roles, and recipes. If the run list entries are specified, they are set on the node. Otherwise, they are not required.	Used to deploy a single node.
knife os manage deploy topology TOPOLOGY_FILE_NAME (options) where: • **TOPOLOGY_FILE_NAME** is a JSON file that describes the topology. If the topology file is in the current directory, only the topology file name is required, otherwise, the fully qualified topology file name is accepted as well.	Used to deploy a topology.
knife os manage update node NODE_FQDN [RUN_LIST_ENTRY[,RUN_LIST_ENTRY]] (options) where: • **NODE_FQDN** is the host fully qualified domain name or IP address for the node. • **RUN_LIST_ENTRY** is a comma-separated list of run list items such as roles, and recipes. If the run list entries are specified, they are set on the node. Otherwise, they are not required.	Used to update or refresh a single node.

© Copyright IBM Corp. 2012, 2014

Table 35. Commands and descriptions (continued)

Command	Description
knife os manage update topology TOPOLOGY_FILE_NAME (options) where: • **TOPOLOGY_FILE_NAME** is a JSON file that describes the topology. If the topology file is in the current directory, only the topology file name is required, otherwise, the fully qualified topology file name is accepted as well.	Used to update or refresh a topology.
knife os manage set environment [ENVIRONMENT_NAME] (options) where: • **ENVIRONMENT_NAME** is the name of a new environment to create, or and existing environment to modify. New environments are created based on one of the existing example environments that are stored in the `chef-repo`. The command prompts for several environment values, and when done, saves the environment that is created or modified on the chef server. It can then be used to deploy a single node or a topology. The **ENVIRONMENT_NAME** field is optional. If it is not provided, the command prompts for one.	Used to create a new, or modify an existing Chef environment that is used during a topology or node deployment.
knife os manage set topology [TOPOLOGY_FILE_NAME] (options) where: • **TOPOLOGY_FILE_NAME** is the name of a new topology JSON file to create. The command prompts for all of the information that is required to deploy a topology of nodes. The resulting JSON file is saved to the current directory, and can then be used to deploy a topology using the knife **os manage deploy topology** command.	Used to create a new topology JSON file that can be used during the deployment of a topology. One of the required topology items is the environment name. If you enter a new environment name in this command while entering the topology information results in the **knife os manage set environment** command being called automatically to create the new environment and upload it to the Chef server when it is complete. Thus, this command can be used to create a new environment, and upload to the chef server, as well as, create a new topology JSON file in one step.
knife os manage appliance migrate node NODE_FQDN (options) where: • **NODE_FQDN** is the host fully qualified domain name of the controller node.	Used to migrate the data and configuration from a IBM SmartCloud Entry 3.2 appliance to a IBM Cloud Manager with OpenStack single control node.

Sample topology JSON file

A topology JSON looks similar to the following layout:

```
{
    "name":"Topology1",
    "description":"This is a topology",
    "environment":"os-single-controller-n-compute",
    "secret_file":"data_bags/example_data_bag_secret",
    "run_sequentially":false,
    "concurrency":10,
    "nodes": [
        {
            "name": "controller",
            "description": "This is the controller node",
            "fqdn":"controllername.company.com",
            "password":"passw0rd",
```

```
            "run_order_number":1,
            "quit_on_error":true,
            "chef_client_options": [
                "-i 3600",
                "-s 600"
            ],
            "runlist": [
                "role[ibm-os-single-controller-node]"
            ]
        },
        {
            "name": "KVM qemu compute",
            "description": "This is a KVM qemu compute node",
            "fqdn":"computename.company.com",
            "user":"admin",
            "identity_file":"/root/identity.pem",
            "run_order_number":2,
            "allow_update": false,
            "runlist": [
                "role[ibm-os-compute-node-kvm]"
            ],
            "attributes":"{\"openstack\":{\"compute\":{\"libvirt\":{\"virt_type\":\"qemu\"}}}}"
        }
    ]
}
```

Table 36. Topology Keys

Topology Keys	Description
name	A user-defined name of the topology.
description	A user-defined description of the topology.
environment	The environment to use with all nodes in the topology. The environment must be on the chef server.
secret_file	The secret file that contains the symmetric key for encrypted data bags. This file is only required if encrypted data bags are used to store the OpenStack passwords.
run_sequentially	An optional flag that is used to control how nodes in a node **run_order_number** run. Set to *true* to run each node sequentially. Set to *false* to run all nodes in a **run_order_number** in parallel. The default value is *false*.
concurrency	An optional integer value that is used to specify the maximum number of nodes that can run in parallel if **run_sequentially** is set to *false*. If not set, all nodes in with a **run_order_number** run in parallel, without limit. The default is there is no limit on the number of nodes that can run in parallel.
nodes	A list of nodes in the topology. At least one node is required

Table 37. Node Keys

Node Keys	Description
name	A user-defined name of the node.
description	A user-defined description of the node.
fqdn	The host fully qualified domain name or IP address of the node.
chef_node_name	The optional setting for the node's name on the Chef server. This setting defaults to the fully qualified domain name value of the node if not explicitly set.
user	The ssh user of the node. This setting defaults to *root*.
password	The ssh user password of the node. This might be used in place of **identity_file**. Either the password, or the **identity_file** are required for ssh access. The password value can be either plain text, or encrypted. Encrypted passwords are written by the **os manage set topology** command when it is used to build the topology JSON file.

Table 37. Node Keys (continued)

Node Keys	Description
identity_file	The identity file that contains the ssh public key. This file might be used in place of password for ssh access. Either the password, or the **identity_file** are required for ssh access to the node.
run_order_number	Each node has an optional **run_order_number** setting that is an integer value greater than or equal to zero. The default **run_order_number** value, if not specified, is zero. The **run_order_number** is used to group one or more nodes to be deployed together as a collection of nodes. If the topology setting of **run_sequentially** is not specified or is set to false, all nodes with a given **run_order_number** are run in parallel as a group. If the topology setting of **run_sequentially** is set to *true*, all nodes with **run_order_number** are run sequentially as a group. The numbers are processed by the smallest one first, and the largest one last. When a **run_order_number** is processed, all nodes with that **run_order_number** are deployed in parallel or sequentially, depending on the topology **run_sequentially** setting. Within a given **run_order_number**, nodes with different environment values must be bootstrapped sequentially, but nodes with the same environment value can be bootstrapped in parallel. All nodes within a given **run_order_number** deploy (**run chef_client**) in parallel after all nodes in the **run_order_number** are bootstrapped. The default value is *zero*.
quit_on_error	Each node has an optional **quit_on_error** boolean value that defaults to *false*. The **quit_on_error** setting applies to a given **run_order_number**. If set to *true*, any node with that **run_order_number** causes the run to stop on failure at the end of the run. Flow would not continue to the next **run_order_number** set of nodes. If any of the nodes within a given **run_order_number** has its **quit_on_error** value set to *true*, then the process stops at the end of the run, and does not proceed to the set of nodes in the next **run_order_number**. The default value is *false*.
chef_client_options	An optional list of run options to be passed to the chef-client program run on the node.
runlist	The run list for the node. The run list is a list of roles or cookbook recipes to be applied to the node. The roles and cookbook recipes must be on the Chef server.
attributes	The optional JSON attribute string to pass to the node during the bootstrap process. The attributes that are specified in the JSON string are applied the first-time chef-client is run on the node.
attribute_file	The optional name of an attribute JSON file to use in place of the attributes setting. The attributes in the JSON file are applied the first-time chef-client is run on the node. Either **attribute_file** or attributes can be specified, but not both.
allow_update	An optional Boolean value that is used during updates only that indicates whether a node should be updated. Set to *true* to allow the node to be updated. Set to *false* if the node should not be updated. The default is *true*.

The following settings are required topology settings.

- The following topology keys are required during a topology deploy operation. The other keys are optional.

  ```
  environment
  ```

- The following node keys are required for each node during a topology deploy operation. The other keys are optional.

  ```
  fdqn
  runlist
  either the password or the identity_file.
  ```

- None of the topology keys are required during a topology update operation.
- The following node keys are required for each node during a topology update operation. The other keys are optional.

  ```
  fdqn
  either the password or the identity_file.
  ```

Data bags

The following example data bags are provided in support of the IBM Cloud Manager with OpenStack topologies.

These data bags contain passwords and secrets for your cloud operating environments running IBM Cloud Manager with OpenStack. By default, all data bags are encrypted using the example secret key that is provided with your IBM Cloud Manager with OpenStack installation (`/opt/ibm/scen/chef-repo/data_bags/example_data_bag_secret`).

- db_passwords
- secrets
- service_passwords
- user_passwords
- appliance32_passwords

db_passwords

Contains IBM Cloud Manager with OpenStack database passwords for your deployment. By default, all passwords match the data bag item name.

- **ceilometer**: Database password for the telemetry component.
- **cinder**: Database password for the block-storage component.
- **glance**: Database password for the image component.
- **heat**: Database password for the orchestration component.
- **horizon**: Database password for the dashboard component.
- **keystone**: Database password for the identity component.
- **neutron**: Database password for the network component.
- **nova**: Database password for the compute component.
- **powervc**: Database password for the PowerVC driver component.

secrets

Contains IBM Cloud Manager with OpenStack secrets. By default, all secrets match the data bag item name except, **openstack_simple_token**.

- **neutron_metadata_secret**: Metadata secret for the network component.
- **openstack_identity_bootstrap_token**: Bootstrap token for the identity component.
- **openstack_metering_secret**: Metering secret for the telemetry component.
- **openstack_simple_token**: Simple token for the IBM Cloud Manager with OpenStack identity component. The secret must be a `base64` encoded value. To generate a secret, run the following command:

  ```
  dd if=/dev/urandom bs=16 count=1 2>/dev/null | base64
  ```

- **openstack_vmware_secret_name**: Secret for the VMware driver.

service_passwords

Contains IBM Cloud Manager with OpenStack service user passwords. By default, all passwords match the data bag item name.

- **openstack-block-storage**: Cinder service user password for the block-storage component.
- **openstack-ceilometer**: Ceilometer service user password for the telemetry component.
- **openstack-compute**: Nova service user password for the compute component.
- **openstack-image**: Glance service user password for the image component.
- **openstack-network**: Neutron service user password for the network component.
- **openstack-orchestration**: Heat service user password for the orchestration component.
- **openstack-iaas-gateway**: IaaS gateway service user password for the iaas-gateway component.

- **openstack-powervc-driver**: PowerVC driver service user password for the PowerVC driver component.

user_passwords

Contains IBM Cloud Manager with OpenStack user passwords. This includes passwords for the IBM Cloud Manager with OpenStack administrator along with the Qpid and DB2 users. By default, all passwords are *openstack1*.

- **admin**: Password for the OpenStack *admin* user and the IBM Cloud Manager with OpenStack self-service portal *admin* user.
- **db2inst1**: Password for the DB2 *db2inst1* user.
- **db2das1**: Password for the DB2 *db2das1* user.
- **db2fenc1**: Password for the DB2 *db2fenc1* user.
- **qpidadmin**: Password for the Qpid *qpidadmin* user.
- **qpidclient**: Password for the Qpid *qpidclient* user.
- **qpidssl**: Password for the Qpid *qpidssl* certificate user
- **pvcadmin**: Password for the PowerVC *admin* user
- **pvcqpid**: Password for the PowerVC Qpid *qpidclient* user
- **sceagent**: Password for the IBM Cloud Manager with OpenStack self-service portal *sceagent* user.
- **xcat**: Password for the z/VM xcat *admin* user
- **zlinuxroot**: Password for the instances that are created by z/VM *root* user

appliance32_passwords

Contains the IBM SmartCloud Entry 3.2 related passwords. These passwords are used when you need to migrate the existing data in IBM SmartCloud Entry 3.2 environment to the new deployed IBM Cloud Manager with OpenStack nodes.

- **db2inst1**: Password for the DB2 *db2inst1* user in IBM SmartCloud Entry 3.2.
- **openstack**: Password for the user which is identified in the `/home/sysadmin/.keystonerc` file in OpenStack in IBM SmartCloud Entry 3.2.
- **sysadmin**: Password for the *sysadmin* user to log into IBM SmartCloud Entry 3.2.

Related tasks:

"Customizing passwords and secrets" on page 102
You can customize the passwords and secrets that are used during the deployment process.

Mapping attributes to services

The following environment and node attributes are mapped to specific options for IBM Cloud Manager with OpenStack services.

Qpid messaging service attributes

Review mappings for the Qpid messaging service attributes

Table 38. Qpid messaging service attributes

Attribute	File	Option
qpid.broker.max-connections	/etc/qpid/qpidd.conf	max-connections
qpid.broker.connection-backlog	/etc/qpid/qpidd.conf	connection-backlog
qpid.broker.worker-threads	/etc/qpid/qpidd.conf	worker-threads
qpid.broker.interface	/etc/qpid/qpidd.conf	interface

OpenStack identity service attributes

Review mappings for the OpenStack identity service attributes.

Table 39. OpenStack identity service attributes

Attribute	File	Section	Option
openstack.identity.verbose	/etc/keystone/keystone.conf	DEFAULT	verbose
openstack.identity.debug	/etc/keystone/keystone.conf	DEFAULT	debug

OpenStack image service attributes

Review mappings for the OpenStack image service attributes

Table 40. OpenStack image service attributes

Attribute	File	Section	Option
openstack.image.verbose	• /etc/glance/glance-api.conf • /etc/glance/glance-cache.conf • /etc/glance/glance-registry.conf • /etc/glance/glance-scrubber.conf	DEFAULT	verbose
openstack.image.debug	• /etc/glance/glance-api.conf • /etc/glance/glance-cache.conf • /etc/glance/glance-registry.conf • /etc/glance/glance-scrubber.conf	DEFAULT	debug
openstack.image.filesystem_store_datadir	/etc/glance/glance-api.conf	DEFAULT	filesystem_store_datadir

OpenStack block storage service attributes

Review mappings for the OpenStack block storage service attributes.

Table 41. OpenStack block storage service attributes

Attribute	File	Section	Option
openstack.block-storage.verbose	/etc/cinder/cinder.conf	DEFAULT	verbose
openstack.block-storage.debug	/etc/cinder/cinder.conf	DEFAULT	debug
openstack.block-storage.iscsi_ip_address	/etc/cinder/cinder.conf	DEFAULT	iscsi_ip_address
openstack.block-storage.san.san_ip	/etc/cinder/cinder.conf	DEFAULT	san_ip

Table 41. OpenStack block storage service attributes (continued)

Attribute	File	Section	Option
openstack.block-storage.san.san_login	/etc/cinder/cinder.conf	DEFAULT	san_login
openstack.block-storage.san.san_private_key	/etc/cinder/cinder.conf	DEFAULT	san_private_key
openstack.block-storage.storwize. storwize_svc_volpool_name	/etc/cinder/cinder.conf	DEFAULT	storwize_svc_volpool_name
openstack.block-storage.storwize. storwize_svc_vol_rsize	/etc/cinder/cinder.conf	DEFAULT	storwize_svc_vol_rsize
openstack.block-storage.storwize. storwize_svc_vol_warning	/etc/cinder/cinder.conf	DEFAULT	storwize_svc_vol_warning
openstack.block-storage.storwize. storwize_svc_vol_autoexpand	/etc/cinder/cinder.conf	DEFAULT	storwize_svc_vol_autoexpand
openstack.block-storage.storwize. storwize_svc_vol_grainsize	/etc/cinder/cinder.conf	DEFAULT	storwize_svc_vol_grainsize
openstack.block-storage.storwize. storwize_svc_vol_compression	/etc/cinder/cinder.conf	DEFAULT	storwize_svc_vol_compression
openstack.block-storage.storwize. storwize_svc_vol_easytier	/etc/cinder/cinder.conf	DEFAULT	storwize_svc_vol_easytier
openstack.block-storage.storwize. storwize_svc_flashcopy_timeout	/etc/cinder/cinder.conf	DEFAULT	storwize_svc_flashcopy_timeout
openstack.block-storage.storwize. storwize_svc_vol_iogrp	/etc/cinder/cinder.conf	DEFAULT	storwize_svc_vol_iogrp
openstack.block-storage.storwize. storwize_svc_connection_protocol	/etc/cinder/cinder.conf	DEFAULT	storwize_svc_connection_protocol
openstack.block-storage.storwize. storwize_svc_iscsi_chap_enabled	/etc/cinder/cinder.conf	DEFAULT	storwize_svc_iscsi_chap_enabled
openstack.block-storage.storwize. storwize_svc_multipath_enabled	/etc/cinder/cinder.conf	DEFAULT	storwize_svc_multipath_enabled
openstack.block-storage.storwize. storwize_svc_multihostmap_enabled	/etc/cinder/cinder.conf	DEFAULT	storwize_svc_multihostmap_ enabled
openstack.block-storage.ibmnas.nas_ip	/etc/cinder/cinder.conf	DEFAULT	nas_ip
openstack.block-storage.ibmnas.nas_login	/etc/cinder/cinder.conf	DEFAULT	nas_login
openstack.block-storage.ibmnas.nas_ssh_port	/etc/cinder/cinder.conf	DEFAULT	nas_ssh_port
openstack.block-storage.ibmnas.shares_config	/etc/cinder/cinder.conf	DEFAULT	nfs_shares_config
openstack.block-storage. ibmnas.mount_point_base	/etc/cinder/cinder.conf	DEFAULT	nfs_mount_point_base
openstack.block-storage. ibmnas.nfs_sparsed_volumes	/etc/cinder/cinder.conf	DEFAULT	nfs_sparsed_volumes
openstack.block-storage. gpfs.gpfs_mount_point_base	/etc/cinder/cinder.conf	DEFAULT	gpfs_mount_point_base
openstack.block-storage. gpfs.gpfs_images_dir	/etc/cinder/cinder.conf	DEFAULT	gpfs_images_dir
openstack.block-storage. gpfs.gpfs_images_share_mode	/etc/cinder/cinder.conf	DEFAULT	gpfs_images_share_mode
openstack.block-storage. gpfs.gpfs_sparse_volumes	/etc/cinder/cinder.conf	DEFAULT	gpfs_sparse_volumes
openstack.block-storage. gpfs.gpfs_max_clone_depth	/etc/cinder/cinder.conf	DEFAULT	gpfs_max_clone_depth
openstack.block-storage. gpfs.gpfs_storage_pool	/etc/cinder/cinder.conf	DEFAULT	gpfs_storage_pool

OpenStack orchestration service attributes

Review mappings for the OpenStack orchestration service attributes.

Table 42. OpenStack orchestration service attributes

Attribute	File	Section	Option
openstack.orchestration.verbose	/etc/heat/heat.conf	DEFAULT	verbose
openstack.orchestration.debug	/etc/heat/heat.conf	DEFAULT	debug

OpenStack telemetry service attributes

Review mappings for the OpenStack telemetry service attributes.

Table 43. OpenStack telemetry service attributes

Attribute	File	Section	Option
openstack.telemetry.verbose	/etc/ceilometer/ ceilometer.conf	DEFAULT	verbose
openstack.telemetry.debug	/etc/ceilometer/ ceilometer.conf	DEFAULT	debug

OpenStack compute service attributes

Review mappings for the OpenStack compute service attributes.

Table 44. OpenStack compute service attributes

Attribute	File	Section	Option
openstack.compute.verbose	/etc/nova/nova.conf	DEFAULT	verbose
openstack.compute.debug	/etc/nova/nova.conf	DEFAULT	debug
openstack.compute.rpc _thread_pool_size	/etc/nova/nova.conf	DEFAULT	rpc_thread_pool_size
openstack.compute.rpc _conn_pool_size	/etc/nova/nova.conf	DEFAULT	rpc_conn_pool_size
openstack.compute.rpc _response_timeout	/etc/nova/nova.conf	DEFAULT	rpc_response_timeout
openstack.compute .state_path	/etc/nova/nova.conf	DEFAULT	state_path
openstack.compute .lock_path	/etc/nova/nova.conf	DEFAULT	lock_path
openstack.compute .instances_path	/etc/nova/nova.conf	DEFAULT	instances_path
openstack.compute.log_dir	/etc/nova/nova.conf	DEFAULT	log_dir
openstack.compute.config .quota_driver	/etc/nova/nova.conf	DEFAULT	quota_driver
openstack.compute.config .quota_cores	/etc/nova/nova.conf	DEFAULT	quota_cores
openstack.compute.config .quota_instances	/etc/nova/nova.conf	DEFAULT	quota_instances
openstack.compute.config .quota_ram	/etc/nova/nova.conf	DEFAULT	quota_ram
openstack.compute.config .quota_floating_ips	/etc/nova/nova.conf	DEFAULT	quota_floating_ips
openstack.compute.config .quota_fixed_ips	/etc/nova/nova.conf	DEFAULT	quota_fixed_ips
openstack.compute .config.quota_security_groups	/etc/nova/nova.conf	DEFAULT	quota_security_groups
openstack.compute.config .quota_security_groups_rules	/etc/nova/nova.conf	DEFAULT	quota_security_groups_rules

Table 44. OpenStack compute service attributes (continued)

Attribute	File	Section	Option
openstack.compute.config .quota_metadata_items	/etc/nova/nova.conf	DEFAULT	quota_metadata_items
openstack.compute.config .quota_injected_files	/etc/nova/nova.conf	DEFAULT	quota_injected_files
openstack.compute. config.quota_injected_file _path_bytes	/etc/nova/nova.conf	DEFAULT	quota_injected_file_path_bytes
openstack.compute .config.quota_injected _file_content_bytes	/etc/nova/nova.conf	DEFAULT	quota_injected_file_content_bytes
openstack.compute.config .quota_key_pairs	/etc/nova/nova.conf	DEFAULT	quota_key_pairs
openstack.compute.libvirt .virt_type	/etc/nova/nova.conf	libvirt	virt_type
openstack.compute.osapi _compute_workers	/etc/nova/nova.conf	DEFAULT	osapi_compute_workers
openstack.compute .metadata_workers	/etc/nova/nova.conf	DEFAULT	metadata_workers
openstack.compute .conductor.workers	/etc/nova/nova.conf	conductor	workers
openstack.region	/etc/nova/nova.conf	DEFAULT	os_region_name

OpenStack network service attributes

Review mappings for the OpenStack network service attributes.

Table 45. OpenStack network service attributes

Attribute	File	Section	Option
openstack.network.verbose	/etc/neutron/ neutron.conf	DEFAULT	verbose
openstack.network.debug	• /etc/neutron/ neutron.conf • /etc/neutron/ metadata_agent.ini • /etc/neutron/ dhcp_agent.ini • /etc/neutron/ l3_agent.ini • /etc/neutron/ lbaas_agent.ini	DEFAULT	debug
openstack.network.rpc _thread_pool_size	/etc/neutron/ neutron.conf	DEFAULT	rpc_thread_pool_size
openstack.network.rpc_conn _pool_size	/etc/neutron/ neutron.conf	DEFAULT	rpc_conn_pool_size
openstack.network.rpc _response_timeout	/etc/neutron/ neutron.conf	DEFAULT	rpc_response_timeout
openstack.network.core _plugin	/etc/neutron/ neutron.conf	DEFAULT	core_plugin

Table 45. OpenStack network service attributes (continued)

Attribute	File	Section	Option
openstack.network.service _plugins	/etc/neutron/ neutron.conf	DEFAULT	service_plugins
openstack.network.ml2 .mechanism_drivers	/etc/neutron/plugins/ ml2/ml2_conf.ini	ml2	mechanism_drivers
openstack.network.ml2 .type_drivers	/etc/neutron/plugins/ ml2/ml2_conf.ini	ml2	type_drivers
openstack.network.ml2 .tenant_network_types	/etc/neutron/plugins/ ml2/ml2_conf.ini	ml2	tenant_network_types
openstack.network.ml2 .flat_networks	/etc/neutron/plugins/ ml2/ml2_conf.ini	ml2_type_flat	flat_networks
openstack.network.ml2 .network_vlan_ranges	/etc/neutron/plugins/ ml2/ml2_conf.ini	ml2_type_vlan	network_vlan_ranges
openstack.network.ml2 .tunnel_id_ranges	/etc/neutron/plugins/ ml2/ml2_conf.ini	ml2_type_gre	tunnel_id_ranges
openstack.network .openvswitch.network_vlan_ranges	/etc/neutron/plugins/ openvswitch/ ovs_neutron_plugin.ini	OVS	network_vlan_ranges
openstack.network .openvswitch.enable_tunneling	/etc/neutron/plugins/ openvswitch/ ovs_neutron_plugin.ini	OVS	enable_tunneling
openstack.network .openvswitch.tunnel_type	/etc/neutron/plugins/ openvswitch/ ovs_neutron_plugin.ini	OVS	tunnel_type
openstack.network .openvswitch.tunnel_id_ranges	/etc/neutron/plugins/ openvswitch/ ovs_neutron_plugin.ini	OVS	tunnel_id_ranges
openstack.network .openvswitch.bridge_mappings	/etc/neutron/plugins/ openvswitch/ ovs_neutron_plugin.ini	OVS	bridge_mappings
openstack.network.quota .driver	/etc/neutron/ neutron.conf	QUOTAS	quota_driver
openstack.network.quota .items	/etc/neutron/ neutron.conf	QUOTAS	quota_items
openstack.network.quota .default	/etc/neutron/ neutron.conf	QUOTAS	default_quota
openstack.network.quota .network	/etc/neutron/ neutron.conf	QUOTAS	quota_network
openstack.network.quota .subnet	/etc/neutron/ neutron.conf	QUOTAS	quota_subnet
openstack.network.quota .port	/etc/neutron/ neutron.conf	QUOTAS	quota_port
openstack.network.quota .security_group	/etc/neutron/ neutron.conf	QUOTAS	quota_security_group
openstack.network.quota .security_group_rule	/etc/neutron/ neutron.conf	QUOTAS	quota_security_group _rule

IBM OpenStack IaaS gateway attributes

Review mappings for the IBM OpenStack IaaS gateway attributes.

Table 46. IBM OpenStack IaaS gateway attributes

Attributes	File	Section	Option
ibm-openstack.iaas-gateway.logging.verbose	/etc/iaasgateway/iaasgateway.conf	DEFAULT	verbose
ibm-openstack.iaas-gateway.logging.debug	/etc/iaasgateway/iaasgateway.conf	DEFAULT	debug

IBM OpenStack PowerVC driver service attributes

Review mappings for the IBM OpenStack PowerVC driver service attributes.

Table 47. IBM OpenStack PowerVC driver service attributes

Attributes	File	Section	Option
ibm-openstack.powervc-driver.powervc.admin_user	/etc/powervc/powervc.conf	powervc	admin_user
ibm-openstack.powervc-driver.powervc.auth_url	/etc/powervc/powervc.conf	powervc	auth_url
ibm-openstack.powervc-driver.powervc.qpid.host	/etc/powervc/powervc.conf	powervc	qpid_hostname
	/etc/powervc/powervc-neutron.conf	DEFAULT	qpid_hostname

IBM OpenStack z/VM driver service attributes

Review mappings for the IBM OpenStack z/VM driver service attributes.

Table 48. IBM OpenStack z/VM driver service attributes

Attributes	File	Section	Option
ibm-openstack.zvm-driver.xcat.server	/etc/nova/nova.conf	DEFAULT	zvm_xcat_server
	/etc/neutron/plugins/zvm/neutron_zvm_plugin.ini	AGENT	zvm_xcat_server
ibm-openstack.zvm-driver.xcat.username	/etc/nova/nova.conf	DEFAULT	zvm_xcat_username
	/etc/neutron/plugins/zvm/neutron_zvm_plugin.ini	AGENT	zvm_xcat_username
ibm-openstack.zvm-driver.xcat.zhcp_nodename	/etc/nova/nova.conf	DEFAULT	zhcp_nodename
	/etc/neutron/plugins/zvm/neutron_zvm_plugin.ini	AGENT	xcat_zhcp_nodename
ibm-openstack.zvm-driver.xcat.master	/etc/nova/nova.conf	DEFAULT	zvm_xcat_master
ibm-openstack.zvm-driver.xcat.mgt_ip	/etc/neutron/plugins/zvm/neutron_zvm_plugin.ini	AGENT	xcat_mgt_ip
ibm-openstack.zvm-driver.xcat.mgt_mask	/etc/neutron/plugins/zvm/neutron_zvm_plugin.ini	AGENT	xcat_ mgt_mask

Table 48. IBM OpenStack z/VM driver service attributes (continued)

Attributes	File	Section	Option
ibm-openstack.zvm-driver.diskpool	/etc/nova/nova.conf	DEFAULT	zvm_diskpool
ibm-openstack.zvm-driver.diskpool_type	/etc/nova/nova.conf	DEFAULT	zvm_diskpool_type
ibm-openstack.zvm-driver.zvm_host	/etc/nova/nova.conf	DEFAULT	zvm_host
	/etc/neutron/plugins/zvm/neutron_zvm_plugin.ini	AGENT	zvm_host
ibm-openstack.zvm-driver.host	/etc/nova/nova.conf	DEFAULT	host
ibm-openstack.zvm-driver.user_profile	/etc/nova/nova.conf	DEFAULT	zvm_user_profile
ibm-openstack.zvm-driver.config_drive.inject_password	/etc/nova/nova.conf	DEFAULT	zvm_config_drive_inject_password
ibm-openstack.zvm-driver.scsi_pool	/etc/nova/nova.conf	DEFAULT	zvm_scsi_pool
ibm-openstack.zvm-driver.fcp_list	/etc/nova/nova.conf	DEFAULT	zvm_fcp_list
ibm-openstack.zvm-driver.zhcp_fcp_list	/etc/nova/nova.conf	DEFAULT	zvm_zhcp_fcp_list

Roles

The following roles are provided in support of the reference topologies.

You add one or more of these roles to the run list of a node system. After the roles are run on a node system, the node system provides the described IBM Cloud Manager with OpenStack services.

- `ibm-os-allinone-kvm`
- `ibm-os-single-controller-node`
- `ibm-os-single-controller-powervc-driver`
- `ibm-os-single-controller-distributed-database-node`
- `ibm-os-compute-node-kvm`
- `ibm-os-compute-node-powerkvm`
- `ibm-os-zvm-driver-node`
- `ibm-os-powervc-driver-node`
- `ibm-os-client-node`
- `ibm-os-database-server-node`
- `ibm-sce-node`

ibm-os-allinone-kvm
Installs and configures all default IBM Cloud Manager with OpenStack services on a single node system, including the compute node services for the KVM or QEMU hypervisor types. This role combines the `ibm-os-single-controller-node` and `ibm-os-compute-node-kvm` roles to support the minimal topology.

ibm-os-single-controller-node
Installs and configures all default IBM Cloud Manager with OpenStack services on a single controller node system except for the compute node services. This role supports both the minimal and controller *+n* compute topologies.

ibm-os-single-controller-powervc-driver
Installs and configures all default IBM Cloud Manager with OpenStack services on a single controller node system with IBM Cloud Manager with OpenStack PowerVC driver services to provide manage-to PowerVC capabilities. This role combines the `ibm-os-single-controller-node` and `ibm-os-powervc-driver-node` roles to support the controller +*n* compute and distributed database topology.

ibm-os-single-controller-distributed-database-node
This role is similar to the `ibm-os-single-controller-node` role. The only difference is that it does not run the database server service. This role supports the distributed database topology.

ibm-os-compute-node-kvm
Installs and configures the compute node services for the KVM or QEMU hypervisor types.

ibm-os-compute-node-powerkvm
Installs and configures the compute node services for the PowerKVM hypervisor type.

ibm-os-zvm-driver-node
Installs and configures the compute driver services for the z/VM hypervisor type.

ibm-os-powervc-driver-node
Installs and configures the PowerVC driver services to manage-to the PowerVC virtualization manager.

ibm-os-client-node
Installs and configures the command-line client interface. This role is included in the `ibm-os-single-controller-node` role and needs to be used only if you would like another node system to provide a command-line client interface to IBM Cloud Manager with OpenStack.

ibm-os-database-server-node
Installs and configures the database server service. This includes creating the required databases for IBM Cloud Manager with OpenStack. This role supports the distributed database topology.

ibm-sce-node
Installs and configures the IBM Cloud Manager with OpenStack self-service portal. This role supports all topologies. You can manually add this role to the run list for a node.

Related tasks:

"Customizing Platform Resource Scheduler" on page 102
If you want to use IBM Platform Resource Scheduler (PRS) instead of the default OpenStack compute scheduler filters, then you need to complete more customization tasks.

"IBM Cloud Manager with OpenStack services" on page 169
The IBM Cloud Manager with OpenStack services provided throughout your topology deployment are determined by the roles that are defined for each node.

"Customizing for self-service portal" on page 100
If you want to use the self-service portal to manage IBM Cloud Manager with OpenStack, you must manually include the self-service portal in a topology.

"Deploying multiple hypervisor types" on page 98
Use this information to deploy the components that are necessary to create a cloud environment with multiple hypervisor types.

Best practices when using the self-service portal

This section contains some tips and techniques when using the IBM Cloud Manager with OpenStack self-service portal.

IBM Cloud Manager with OpenStack FAQ

The frequently asked questions (FAQ) topic is a list of questions and answers about IBM Cloud Manager with OpenStack.

Q: **How do I find my home directory?**

A:
1. In Windows, enter `% HOMEPATH%` in the address bar of a Windows Explorer window
2. For Linux, type `echo $HOME` in a command window.

Q: **I created a trusted certificate. Why am I still getting an exception that says the connection is untrusted?**

A: When the CA is not trusted by clients automatically and you are attempting to access IBM Cloud Manager with OpenStack with the https protocol, an exception is encountered that says the connection is untrusted. You must confirm that the risks are understood and must add an exception to continue. Even with a trusted certificate, when you are using Internet Explorer, a similar exception is likely to occur.

Q: **I want to have spaces in my directory name, but it keeps failing when I try to create it. How can I have spaces?**

A: If you have spaces in a directory name, then you must have double quotation marks around it as shown in the following example: `sce240_windows_installer.exe -i silent -f "c:\My Directory\ installer.properties"`

Q: **My user ID is locked! How do I unlock it?**

A: If you have three invalid attempts to log in to IBM Cloud Manager with OpenStack in a 24 hour period, your user ID is locked and must be unlocked by an administrator. If your administrator ID becomes locked, you can either wait 24 hours without logging in or restart IBM Cloud Manager with OpenStack and then try logging in again.

Q: **How does IBM Cloud Manager with OpenStack store the passwords for local users or clouds?**

A: The passwords are encrypted and stored in either property files or a database.

Q: **IBM Cloud Manager with OpenStack GUI looks distorted. How can I fix that?**

A: See the information in “Display issue with Internet Explorer (self-service portal)” on page 262.

Q: **I upgraded/installed IBM Cloud Manager with OpenStack, but I’m still seeing the previous version in my browser. How can I fix that?**

A: Clear the cache in your browser and try again. You might have to close your browser after you clear the cache and then reopen your browser and try connecting to IBM Cloud Manager with OpenStack again.

Q: **My image is not visible in the window. Where is it?**

A: Make sure that your image is deployed and that the correct project is specified. If it still is not visible, contact the administrator to ensure that you have access.

Q: **The product charges that I set are incorrect or are not updating. What do I do?**

A: First of all, verify that the currencies for all configurable products are the same. You cannot mix currencies. To change your currency for a product, see the “Configuring billing” on page 158. Make sure that you are restarting IBM Cloud Manager with OpenStack after saving.

Q: **The instances for a user were moved to a different project. Now when the user logs on, he cannot see his instances. How can the user access his instances?**

A: The project where the instances were moved might need to be edited to grant the user access to the project. When you have ensured that the user has access to the new project, have the user check again to see whether the instances display.

Q: **IBM Cloud Manager with OpenStack will not start for me. I am running Windows. Why am I having problems?**

A: You must be the Windows Administrator to run IBM Cloud Manager with OpenStack. However, if you are in the administrator group, you can right-click the IBM Cloud Manager with OpenStack icon and select **Run as Administrator**.

Q: When updating IBM Cloud Manager with OpenStack to a new release, can I migrate data and configurations from two releases previous to the current release? For example, can I migrate data in IBM Cloud Manager with OpenStack from version 2.2 to version 2.4?

A: No, you must migrate sequentially. For example, migrate from IBM Cloud Manager with OpenStack version 2.2 to version 2.3. Then you can migrate from IBM Cloud Manager with OpenStack version 2.3 to version 2.4.

Q: Does the IBM Cloud Manager with OpenStack infocollect command support collecting a database log such as DB2?

A: No, you must check with the administrator of the database and collect the log manually.

Q: Why does my login fail with the session timing out?

A: If your user login fails because the session times out, there might be a problem with the timezone setting. Verify that the IBM Cloud Manager with OpenStack server and client time and timezone match. For example, on the server, if the timezone is Coordinated Universal Time +08:00, the time is 11:27. For the client, the timezone is Coordinated Universal Time +07:00, and the time should be 10:27.

Q: Why can't I access the IBM Cloud Manager with OpenStack GUI page after I start it?

A: Verify that a firewall is not blocking the http or https port that you accessed. To check whether it worked in the IBM Cloud Manager with OpenStack host, access for example, `http://localhost:18080/cloud/web/login.html` or use the UNIX command **`wget http://localhost:18080/cloud/web/login.html`**.

Using the `screen` command

The **`screen`** command can be used to start or shut down the IBM Cloud Manager with OpenStack server or to access the OSGI console when the server is up and running when running Linux.

For example, enter **`screen`** and then run the command to start the server. After the server is started, type `ctrl+a`, then `d` to disconnect and leave the IBM Cloud Manager with OpenStack server running,

To get back to the IBM Cloud Manager with OpenStack OSGI prompt to perform other actions, such as enabling additional logging, enter `screen -r`.

Using the `nohup` command

On AIX or Linux, if you start a process from the command line and then log off, the processes you started are generally terminated, even if you spawn them as background processes, because each process is sent a hang up signal (`SIGHUP`). The **`nohup`** command allows you to start a process with the hang up signal disabled so that the process is not ended when you log off.

The **`nohup`** command is used frequently for starting services, such as ssh daemon or DB2 instances.

For example, to start IBM Cloud Manager with OpenStack as a background service, run the following command:

```
nohup /opt/ibm/SCE24/skc -nosplash < /dev/null > /dev/null &
```

The options in this command include the following:

-nosplash
Prevents the process from displaying a splash screen.

< /dev/null
Disconnects the process from terminal input. This option can prevent the process from entering a *stopped* state as can sometimes happen when started from the command line on AIX. This option is not needed when starting the command from a shell script.

> /dev/null
Redirects the OSGI console output. For example, you might want to redirect the output to a log file.

& Runs the command as a background process.

Chapter 10. Troubleshooting and support for IBM Cloud Manager with OpenStack

To isolate and resolve problems with your IBM products, you can use the troubleshooting and support information. This information contains instructions for using the problem-determination resources that are provided with your IBM products, including IBM Cloud Manager with OpenStack.

Techniques for troubleshooting problems

Troubleshooting is a systematic approach to solving a problem. The goal of troubleshooting is to determine why something does not work as expected and how to resolve the problem. Certain common techniques can help with the task of troubleshooting.

The first step in the troubleshooting process is to describe the problem completely. Problem descriptions help you and the IBM technical-support representative know where to start to find the cause of the problem. This step includes asking yourself basic questions:

- What are the symptoms of the problem?
- Where does the problem occur?
- When does the problem occur?
- Under which conditions does the problem occur?
- Can the problem be reproduced?

The answers to these questions typically lead to a good description of the problem, which can then lead you to a problem resolution.

What are the symptoms of the problem?

When starting to describe a problem, the most obvious question is "What is the problem?" This question might seem straightforward; however, you can break it down into several more-focused questions that create a more descriptive picture of the problem. These questions can include:

- Who, or what, is reporting the problem?
- What are the error codes and messages?
- How does the system fail? For example, is it a loop, hang, crash, performance degradation, or incorrect result?

Where does the problem occur?

Determining where the problem originates is not always easy, but it is one of the most important steps in resolving a problem. Many layers of technology can exist between the reporting and failing components. Networks, disks, and drivers are only a few of the components to consider when you are investigating problems.

The following questions help you to focus on where the problem occurs to isolate the problem layer:

- Is the problem specific to one platform or operating system, or is it common across multiple platforms or operating systems?
- Is the current environment and configuration supported?
- Do all users have the problem?
- (For multi-site installations.) Do all sites have the problem?

© Copyright IBM Corp. 2012, 2014

If one layer reports the problem, the problem does not necessarily originate in that layer. Part of identifying where a problem originates is understanding the environment in which it exists. Take some time to completely describe the problem environment, including the operating system and version, all corresponding software and versions, and hardware information. Confirm that you are running within an environment that is a supported configuration; many problems can be traced back to incompatible levels of software that are not intended to run together or have not been fully tested together.

When does the problem occur?

Develop a detailed timeline of events leading up to a failure, especially for those cases that are one-time occurrences. You can most easily develop a timeline by working backward: Start at the time an error was reported (as precisely as possible, even down to the millisecond), and work backward through the available logs and information. Typically, you need to look only as far as the first suspicious event that you find in a diagnostic log.

To develop a detailed timeline of events, answer these questions:

- Does the problem happen only at a certain time of day or night?
- How often does the problem happen?
- What sequence of events leads up to the time that the problem is reported?
- Does the problem happen after an environment change, such as upgrading or installing software or hardware?

Responding to these types of questions can give you a frame of reference in which to investigate the problem.

Under which conditions does the problem occur?

Knowing which systems and applications are running at the time that a problem occurs is an important part of troubleshooting. These questions about your environment can help you to identify the root cause of the problem:

- Does the problem always occur when the same task is being performed?
- Does a certain sequence of events need to happen for the problem to occur?
- Do any other applications fail at the same time?

Answering these types of questions can help you explain the environment in which the problem occurs and correlate any dependencies. Remember that just because multiple problems might have occurred around the same time, the problems are not necessarily related.

Can the problem be reproduced?

From a troubleshooting standpoint, the ideal problem is one that can be reproduced. Typically, when a problem can be reproduced you have a larger set of tools or procedures at your disposal to help you investigate. Consequently, problems that you can reproduce are often easier to debug and solve.

However, problems that you can reproduce can have a disadvantage: If the problem is of significant business impact, you do not want it to recur. If possible, re-create the problem in a test or development environment, which typically offers you more flexibility and control during your investigation.

- Can the problem be re-created on a test system?
- Are multiple users or applications encountering the same type of problem?
- Can the problem be re-created by running a single command, a set of commands, or a particular application?

Related tasks:

"Searching knowledge bases" on page 248
You can often find solutions to problems by searching IBM knowledge bases. You can optimize your results by using available resources, support tools, and search methods.

Logging tasks

The IBM Cloud Manager with OpenStack log files are a source of information for additional details about IBM Cloud Manager with OpenStack errors.

By default, IBM Cloud Manager with OpenStack creates a log file in the `<home directory>/logs` directory and saves 9 log files of 50 MB each. The latest log file is called `skc-0.log`.

Change logging levels from the OSGi command prompt

The logging levels can be changed dynamically while the server is running by using the **log** command from the IBM Cloud Manager with OpenStack (OSGi) command prompt.

About this task

The logging levels can be changed dynamically while the server is running by using the **log** command from the IBM Cloud Manager with OpenStack (OSGi) command prompt. Changes made using the **log** command are not saved and are only in effect while the server is running. If the server is restarted, the logging levels are reset to their initial values as specified in the `logging.properties` file. For more information about changing these values in the `logging.properties` file, see "Configuring logging" on page 156.

To run the **log** command, follow these steps:

Procedure

1. Access the IBM Cloud Manager with OpenStack OSGi console.
2. At the OSGi command prompt enter `log <action> <action parameters>`, where the following actions are supported:

   help Displays the online help.

   list Lists the available loggers and their current levels.

   setlevel <logger name>=<logger level>
   Sets the specified logger to the specified logging levels. To set more than one logger, separate the logger name=logger level pair with a space.

Results

See the following examples for typical log commands:

```
log help
log list
log setlevel com.ibm.cfs.cloud=finest
log setlevel com.ibm.cfs.cloud=info default=finest
```

The most common log message level that an IBM Cloud Manager with OpenStack administrator might want to change is `com.ibm.cfs.rest.client.internal=FINE`. Changing the message level causes the output of HTTP requests and responses to be sent to and from the OpenStack REST API.

In a production environment, keep a backup of the log files for at least two weeks to help resolve problems that are not reported immediately or that go undetected.

Note: The property file values are not case-sensitive so a property such as `com.ibm.cfs.rest.client.internal=FINE` is the same as `com.ibm.cfs.rest.client.internal=fine`.

Retrieve log and system files

IBM Cloud Manager with OpenStack provides a command-line utility that enables you to gather logs and system information. If IBM Cloud Manager with OpenStack self-service portal and OpenStack are installed in the same node, then both self-service portal and standard OpenStack logs can be collected, otherwise you can collect the former. When you use standard OpenStack logs you do not have to use **log_config** to customize the configuration. This tool runs independently of IBM Cloud Manager with OpenStack and is available even when IBM Cloud Manager with OpenStack is not running.

To use the command-line utility, run one of the following commands:

- **infocollect.bat**: on a Windows system.
- **infocollect.sh**: on Linux or AIX systems

Note: These scripts can be found in the SELF_SERVICE_PORTAL_INSTALLATION_DIR/program/bin directory.

The command accepts the following options:

-c Specifies the configuration directory, for example: SELF_SERVICE_PORTAL_HOME, where all the IBM Cloud Manager with OpenStack configuration and log files are saved. This argument is mandatory. You must provide an existing directory path. The command exits with an error if the specified directory cannot be found.

-d Specifies the destination directory for the result files. If this argument is used, provide an existing directory path. This command exists with an error if the specified. If this argument is not provided, the HOME directory of the caller is used. If the HOME directory is not found, for example, the corresponding environment variable is not set correctly, the system TEMP directory is used as the default output directory. For example, in Linux /tmp is the system TEMP directory and in Windows 7, the %USER_HOME%\AppData\Local\Temp is the TEMP directory.

-h Prints usage information.

When this utility is started, the following files are created:

openStackLog.zip(Optional)

Contains all of the OpenStack log files:

*.log files

*.gz files

sceHome.zip

Contains all of the IBM Cloud Manager with OpenStack configurations:

*.properties files

*.log files

*.udr files

Billing configurations: .xml files under products/

All .xml and .txt files under SELF_SERVICE_PORTAL_HOME

basicSysInfo.txt

Contains basic OS information:

CPU

Memory

OS name

Note: This information is retrieved by calling OS shell commands (for example, **DOS** commands for Windows), so the results vary depending on the concrete OS.

Example 1 (Windows)

Collect the configurations, logs, and system information, and save the result to the `F:\documents\sce\diagnostic` directory. The `SELF_SERVICE_PORTAL_HOME` is `C:\Users\Administrator\.SCE41`.

```
infocollect.bat -c C:\Users\Administrator\.SCE41 -d F:\documents\sce\diagnostic
```

Example 2 (Linux)

Collect the configurations in Linux or AIX, logs, and system information, and save the result to the directory of `/tmp/diagnostic`. The `SELF_SERVICE_PORTAL_HOME` is `/var/opt/ibm/.SCE41`.

```
./infocollect.sh -c /var/opt/ibm/.SCE41 -d /tmp/diagnostic
```

Troubleshooting using the OSGi console

Use the Open Services Gateway initiative (OSGi) console to review information about IBM Cloud Manager with OpenStack.

By default, IBM Cloud Manager with OpenStack starts an OSGi command-line console when the IBM Cloud Manager with OpenStack executable is run. You can access the console directly in the command window started by the executable.

You can also run the console in the background and assign a specific port for telnet connections. To assign a port, modify the `skc.ini` file and add an unused port number on a new line after the `-console` option.

```
-console
<port number>
```

For example, to assign port *7777* for telnet connections, change the option to the following:

```
-console
7777
```

To connect to the OSGi console, type the following:

```
telnet localhost 7777
```

Logging chef-client output

You can obtain chef-client logs from the nodes when you use the IBM Cloud Manager with OpenStack commands.

The IBM Cloud Manager with OpenStack commands log chef-client console output that is captured from each node on the Chef server in a central location. The log file name includes the node name, as well as, an integer time stamp. Up to 9 log files are kept for each node in the `/var/log/chef-server/nodes` folder on the Chef server. Each IBM Cloud Manager with OpenStack deployment or update command typically writes 3 log files. One is for the initial **ping** to the node. The second is for the bootstrap, and the final is for the actual chef-client run.

The log file generation can be controlled by two settings in the chef server `/root/.chef/knife.rb` file:

```
knife[:node_log_max_files] = 9
knife[:node_log_path] = '/var/log/chef-server/nodes'
```

These settings default in the IBM Cloud Manager with OpenStack command-line interface code to the preceding values shown above. The preceding values can be set in `knife.rb`, if the defaults are not adequate.

To turn of node logging, set **`knife[:node_log_max_files] = 0`**

Log files are created for each node to which the IBM Cloud Manager with OpenStack command connects. The log file name is `<node_name>_<integer_timestamp>.log`. For example:

```
host.my.company.com_1401283587.log
```

By default, up to 9 log files are saved for each node.

Logging chef-client output to a log file on the node itself can also be accomplished by specifying chef-client options when you issue one of the IBM Cloud Manager with OpenStack **os manage deploy** or **os manage update** commands. When chef-client writes to a log file, it does not affect its console output.

To enter the chef-client options on a IBM Cloud Manager with OpenStack CLI 'node' type command, use the `-O` option followed by a string that contains the chef-client options. For example:

```
knife os manage update node host.my.company.com -P passw0rd -O '-l info -L /var/log/chef-client.log
```

The `-l` chef-client option is the log level (*info* is specified in the preceding example). The `-L` chef-client option is the log file location.

Note: The log file path must exist or an error occurs when the log file is written.

To enter the chef-client options on a IBM Cloud Manager with OpenStack CLI 'topology' type command, use the **chef-client-options** JSON key in the topology file that contains a list of chef-client options. For example, the JSON key and value for a node in the topology file might look like the following example:

```
 "chef_client_options": ["-l info", "-L /var/log/chef-client.log"],
```

The preceding options are passed to chef-client to cause it to perform *info* level logging and store the log information in `/var/log/chef-client.log`.

When forcing the chef-client to write to a log file, the output is not the same as that seen on the console. That console output is logged on the chef server as previously mentioned.

Another option to forcing chef-client to write to a log file is to set the following values in the `client.rb` file on the node:

```
log_location "/var/log/chef-client.log"
verbose :info
log_level :info
```

That would cause *info* level logging to occur for all chef-client calls to the log file `/var/log/chef-client.log`.

Note: The chef-client log is not cleared. It is always appended to by chef-client.

Searching knowledge bases

You can often find solutions to problems by searching IBM knowledge bases. You can optimize your results by using available resources, support tools, and search methods.

About this task

You can find useful information by searching the knowledge center for IBM Cloud Manager with OpenStack. However, sometimes you need to look beyond the knowledge center to answer your questions or resolve problems.

Procedure

- Find the content that you need by using the IBM Support Portal.

The IBM Support Portal is a unified, centralized view of all technical support tools and information for all IBM systems, software, and services. The IBM Support Portal lets you access the IBM electronic support portfolio from one place. You can tailor the pages to focus on the information and resources that you need for problem prevention and faster problem resolution. Familiarize yourself with the IBM Support Portal by viewing the demo videos (https://www.ibm.com/blogs/SPNA/entry/the_ibm_support_portal_videos) about this tool. These videos introduce you to the IBM Support Portal, explore troubleshooting and other resources, and demonstrate how you can tailor the page by moving, adding, and deleting portlets.

- The Techdocs Web site at www.ibm.com/support/techdocs/ provides white papers, technotes, tips, and other documents related to IBM Cloud Manager with OpenStack.

  Tip: To locate the information that you need, either select the categories that you want to search or select *All of the Techdocs library* to search all categories. Then enter *IBM SmartCloud Entry* in the **for:** field and click **Search**.
- Search for content on the forum: Forum
- Search for content by using the IBM masthead search. You can use the IBM masthead search by typing your search string into the Search field at the top of any ibm.com® page.
- Search for content by using any external search engine, such as Google, Yahoo, or Bing. If you use an external search engine, your results are more likely to include information that is outside the ibm.com domain. However, sometimes you can find useful problem-solving information about IBM products in newsgroups, forums, and blogs that are not on ibm.com.

  Tip: Include "IBM" and the name of the product in your search if you are looking for information about an IBM product.

Related concepts:

"Techniques for troubleshooting problems" on page 243
Troubleshooting is a systematic approach to solving a problem. The goal of troubleshooting is to determine why something does not work as expected and how to resolve the problem. Certain common techniques can help with the task of troubleshooting.

Related reference:

"Known problems and solutions for an Administrator" on page 259
If you are using IBM Cloud Manager with OpenStack with the *Administrator* role, review theses known problems and solutions that you might encounter.

"Known problems and solutions for a User" on page 274
If you are using IBM Cloud Manager with OpenStack with the *User* role, review theses known problems and solutions that you might encounter.

Getting fixes from Fix Central

You can use Fix Central to find the fixes that are recommended by IBM Support for a variety of products, including IBM Cloud Manager with OpenStack. With Fix Central, you can search, select, order, and download fixes for your system with a choice of delivery options. A IBM Cloud Manager with OpenStack product fix might be available to resolve your problem.

Procedure

To find and install fixes:

1. Obtain the tools that are required to get the fix. If it is not installed, obtain your product update installer. You can download the installer from Fix Central. This site provides download, installation, and configuration instructions for the update installer.
2. Click **Find Product**.
3. Begin typing IBM Cloud Manager with OpenStack in the **Product Selector** field.
4. Select IBM Cloud Manager with OpenStack from the list. For Installed version, select All.

5. Identify and select the fix that is required.
6. Download the fix.
   a. Open the download document and follow the link in the "Download Package" section.
   b. When downloading the file, ensure that the name of the maintenance file is not changed. This change might be intentional, or it might be an inadvertent change that is caused by certain web browsers or download utilities.
7. Apply the fix using the instructions in the readme.txt file that comes with the fix package.

Related concepts:

"Applying fixes and updates" on page 38
You can apply fixes and updates for the IBM Cloud Manager with OpenStack product.

Contacting IBM Support

IBM Support provides assistance with product defects, answers FAQs, and helps users resolve problems with the product.

Before you begin

After trying to find your answer or solution by using other self-help options such as technotes, you can contact IBM Support. Before contacting IBM Support, your company or organization must have an active IBM software maintenance agreement (SWMA), and you must be authorized to submit problems to IBM. For information about the types of available support, see the Support portfolio topic in the "*Software Support Handbook*": http://www14.software.ibm.com/webapp/set2/sas/f/handbook/offerings.html.

Procedure

To contact IBM Support about a problem:

1. Define the problem, gather background information, and determine the severity of the problem. For more information, see the Getting IBM support topic in the *Software Support Handbook* http://www14.software.ibm.com/webapp/set2/sas/f/handbook/getsupport.html.
2. Gather diagnostic information.
3. Submit the problem to IBM Support in one of the following ways:
   - Online through the IBM Support Portal (http://www.ibm.com/software/support/): You can open, update, and view all of your service requests from the Service Request portlet on the Service Request page.
   - By phone: For the phone number to call in your region, see the Directory of worldwide contacts (http://www.ibm.com/planetwide/) web page.

Results

If the problem that you submit is for a software defect or for missing or inaccurate documentation, IBM Support creates an Authorized Program Analysis Report (APAR). The APAR describes the problem in detail. Whenever possible, IBM Support provides a workaround that you can implement until the APAR is resolved and a fix is delivered. IBM publishes resolved APARs on the IBM Support website daily, so that other users who experience the same problem can benefit from the same resolution.

What to do next

Be prepared to work with the IBM technical-support representative by using IBM Assist On-site, which is a remote-assistance plug-in that you can download to your computer. The IBM technical-support representative can use IBM Assist On-site to view your desktop and share control of your mouse and keyboard. This tool can shorten the time that it takes to identify the problem, collect the necessary data,

and solve the problem. For more information, see IBM Assist On-site: http://www.ibm.com/support/assistonsite/.

Exchanging information with IBM

To diagnose or identify a problem, you might need to provide IBM Support with data and information from your system. In other cases, IBM Support might provide you with tools or utilities to use for problem determination.

Sending information to IBM Support

To reduce the time that is required to resolve your problem, you can send trace and diagnostic information to IBM Support.

Procedure

To submit diagnostic information to IBM Support:

1. Open a problem management record (PMR). You can use the Service Request Tool to create a PMR: http://www.ibm.com/support/servicerequest
2. Collect the diagnostic data that you need. Diagnostic data helps reduce the time that it takes to resolve your PMR. Manually collect and compress the required diagnostic information by following the directions that you receive from IBM Support.

   Note: See "Retrieve log and system files" on page 246 for directions on collecting logs and system information for IBM Cloud Manager with OpenStack.
3. Transfer the files to IBM. Use the information on the IBM Enhanced Customer Data Repository (ECuRep) website: http://www.ibm.com/software/support/exchangeinfo.html. While ECuRep is the primary method for uploading data, you can also use one of the following methods to transfer the files to IBM. All of these data exchange methods are described on the ECuRep website as well.
   - The Service Request tool
   - Standard data uploads methods: FTP, HTTP
   - Secure data upload methods: FTPS, SFTP, HTTPS
   - Email

Receiving information from IBM Support

ECuRep is the primary method for exchanging data. However, on occasion, an IBM technical-support representative might ask you to download diagnostic tools or other files. You can use FTP to download these files.

Before you begin

Ensure that your IBM technical-support representative provided you with the preferred server to use for downloading the files and the exact directory and file names to access.

Procedure

To download files from IBM Support:

1. Use FTP to connect to the site that your IBM technical-support representative provided and log in as `anonymous`. Use your email address as the password.
2. Change to the appropriate directory:
   a. Change to the `/fromibm` directory.
      ```
      cd fromibm
      ```
   b. Change to the directory that your IBM technical-support representative provided.

```
cd nameofdirectory
```

3. Enable binary mode for your session.
```
binary
```
4. Use the **get** command to download the file that your IBM technical-support representative specified.
```
get filename.extension
```
5. End your FTP session.
```
quit
```

Subscribing to Support updates

To stay informed of important information about the IBM products that you use, you can subscribe to updates.

About this task

By subscribing to receive updates about IBM Cloud Manager with OpenStack, you can receive important technical information and updates for specific IBM Support tools and resources. You can subscribe to updates by using My Notifications:

My Notifications
With My Notifications, you can subscribe to Support updates for any IBM product. (My Notifications replaces My Support, which is a similar tool that you might have used in the past.) With My Notifications, you can specify that you want to receive daily or weekly email announcements. You can specify what type of information you want to receive (such as publications, hints and tips, product flashes (also known as alerts), downloads, and drivers). My Notifications enables you to customize and categorize the products about which you want to be informed and the delivery methods that best suit your needs.

Procedure

To subscribe to My Notifications, complete the following steps:

1. Go to the IBM Support Portal (http://www.ibm.com/software/support/) and sign in by using your IBM ID and password.
2. In the **Notifications** portlet, click **Manage all my subscriptions**, then click the **Subscribe** tab.
3. Click **Other Software**.
4. Select IBM Cloud Manager with OpenStack, then click **Continue**.
   a. Select your preferences for how to receive updates, whether by email, online in a designated folder, or as an RSS or Atom feed.
   b. Select the types of documentation updates that you want to receive, for example, new information about product downloads and discussion group comments.
   c. Click **Submit**.

Results

Until you modify your My Notifications preferences, you receive notifications of updates that you have requested. You can modify your preferences when needed.

Related reference:

Subscribe to My Notifications support content updates

My Notifications for IBM technical support

My Notifications for IBM technical support overview

Troubleshooting errors when deploying topologies

If a topology deployment fails, review the log output from the deployment command for more information. Identify errors, starting at the beginning of the log output. The first error found is typically the most useful. Errors found at the end of the log output might be the result of earlier errors. The following topics might provide additional information to solve the problem.

Related tasks:

"Deploying multiple hypervisor types" on page 98
Use this information to deploy the components that are necessary to create a cloud environment with multiple hypervisor types.

"Deploying with KVM or QEMU compute nodes" on page 82
Deploy the components that are necessary to create a cloud environment with KVM or QEMU compute nodes.

"Deploying with PowerKVM compute nodes" on page 86
Deploy the components that are necessary to create a cloud environment with IBM PowerKVM compute nodes.

"Deploying with z/VM compute nodes" on page 90
Deploy the components that are necessary to create a cloud environment with z/VM compute nodes.

"Deploying with Hyper-V compute nodes" on page 80
Deploy the components that are necessary to create a cloud environment with Hyper-V compute nodes.

"Deploying with KVM or QEMU compute nodes" on page 82
Deploy the components that are necessary to create a cloud environment with KVM or QEMU compute nodes.

Data bag item not found

You might see an error in the log that a data bag item is missing.

Symptoms

```
### Data bag item not found ###

    ================================================================================
    Recipe Compile Error in /var/chef/cache/cookbooks/openstack-identity/recipes/server.rb
    ================================================================================


    Net::HTTPServerException
    ------------------------
    404 "Object Not Found"


    Cookbook Trace:
    ---------------
      /var/chef/cache/cookbooks/openstack-common/libraries/passwords.rb:46:in `secret'
      /var/chef/cache/cookbooks/openstack-common/libraries/passwords.rb:52:in `get_secret'
      /var/chef/cache/cookbooks/openstack-identity/recipes/server.rb:113:in `from_file'


    Relevant File Content:
    ----------------------
    /var/chef/cache/cookbooks/openstack-common/libraries/passwords.rb:
```

Causes

The data bag attribute in your environment file might be incorrect or missing on the Chef server.

Resolving the problem

Verify that the data bag attributes in your environment file are correct and that the data bags and data bag items for your topology deployment exist on the Chef server. You can use the **knife data bag list** and **knife data bag show my-data-bag-name** commands to determine the data bags and data bag items available on the Chef server.

Related tasks

"Customizing passwords and secrets" on page 102
You can customize the passwords and secrets that are used during the deployment process.

Node verification failed

You might see an error in the log that the topology node verification failed.

Symptoms

```
ERROR: RuntimeError: Topology node verification failed.
```

Causes

An error in the topology file might cause this error.

Resolving the problem

Verify that the following information is correct in your topology file:

- Fully qualified domain name for each node
- Login information for each node

Correct any errors and try the task again.

Deploying a client node fails

You might see an error in the log that the topology node failed to deploy.

Symptoms

```
ERROR 14: Peer cert cannot be verified or peer cert invalid.
```

Causes

The deployment can fail if an SSL certificate exists on the client node.

Resolving the problem

Ensure that the client node does not have an existing SSL certificate.

1. On the Chef server, run the following commands:

   ```
   knife node delete -y  chef-client-fqdn
   knife client delete -y chef-client-fqdn
   ```

2. On the client node, run the following commands:
   a. Run the **rm -rf /etc/chef** command.
   b. Run the **rm -rf /etc/yum.repos.d/ibmos*** command.
   c. List the yum repos and validate that you do not have OpenStack yum repos configured:
      1) Run **yum repolist** to list the repos.
      2) Delete OpenStack yum repos. To delete all yum repos, run **yum clean all** or to delete an individual repo, run **yum erase <repo>**.
   d. Delete the SSL certificates.
      1) Export **NSS_DEFAULT_DB_TYPE="sql"**.
      2) To list the named certificates in the NSS db, run the following command:

         ```
         certutil -L -d /etc/pki/nssdb
         ```

3) To extract the Chef certificate and view who issued it, run the following command:

```
certutil -L -d /etc/pki/nssdb -a -n 'chef-server-cert' > chef.crt
openssl x509 -noout -text -in chef.crt | grep Issuer
```

4) To delete the Chef certificate, run the following command:

```
certutil -D -d /etc/pki/nssdb -n 'chef-server-cert'
```

Cookbook not found

You might see an error in the log file about missing cookbooks.

Symptoms

You might see an error similar to the following example:

```
================================================================================
Error Resolving Cookbooks for Run List:
================================================================================


Missing Cookbooks:
------------------
The following cookbooks are required by the client but don't exist on the server:
* openstack-common
```

Resolving the problem

Verify that the cookbook version constraints in the environment file are correct. The cookbook might exist on the Chef server, however, it might be ignored because of a version constraint. For example, a cookbook version constraint of *~> 9.2.0* would support versions *>= 9.2.0* and *< 9.3.0*. You can use **knife cookbook show cookbook-name** to determine the versions of a cookbook that are available.

Network interface not found

You might encounter an error in the log if the network interface cannot be found.

Symptoms

You might see information similar to the following example if a network interface is not defined correctly.

```
================================================================================
Recipe Compile Error in /var/chef/cache/cookbooks/openstack-common/recipes/set_endpoints_by_interface.rb
================================================================================


NoMethodError
-------------
undefined method `[]' for nil:NilClass


Cookbook Trace:
---------------
  /var/chef/cache/cookbooks/openstack-common/libraries/network.rb:29:in `address_for'
  /var/chef/cache/cookbooks/openstack-common/libraries/endpoints.rb:96:in `address'
```

Resolving the problem

Verify that the network interface attributes in the environment file and node specific attributes files are correct. For example, a node may not have an eth0 network interface which is the default value for some attributes.

Operating system repository not found

You might see an error in the log file that the yum repository is not available.

Symptoms

The error might be similar to the following example:

```
================================================================================
Error executing action `run` on resource 'ruby_block[operating system yum repository not available]'
================================================================================


#<Class:0x000000044e86b8&gt;::NoOperatingSystemRepositoryException
----------------------------------------------------------------
There are no yum repositories available that contain operating system packages that the openstack
installation depends on. You must either configure your nodes to have access to a yum repository
or you can create this repository on the deployment server. You can create your own repository by
copying the Packages and repodata directory from this node system's installation media to the
following location: https://mychefserver.com:14443/yum-repo/operatingsystem/rhel6.5/x86_64.
Once the yum repository is created on the server or you have updated the node systems yum repository
list to include a repository which contains operating system packages, re-run the deployment process.
The package used to test the repository is libvirt. For more information on creating your own
operating system yum repository, see the product documentation.
```

Resolving the problem

Follow the instructions that are outlined in the error to set up an operating system yum repository for your topology deployment.

Related tasks

"Configuring operating system yum repositories on the deployment server" on page 36
After you install IBM Cloud Manager with OpenStack on the deployment server, you must determine how your Linux nodes can access yum repositories that contain operating system packages that are required when OpenStack is deployed.

Package failed to install or upgrade

You might see an error in the log file about being unable to upgrade a resource.

Symptoms

The error might be similar to the following example:

```
================================================================================
Error executing action `upgrade` on resource 'package[python-openstackclient]'
================================================================================
```

Resolving the problem

The node might be unable to connect to the Chef server or the operating system repository. Log into the node and attempt to install or upgrade the package manually. For example, **yum install python-openstackclient**). This command should provide more information about why the error occurred and how to resolve it.

PowerVC driver deploys successfully; later is dead

It appears that you successfully deployed a controller node that includes the PowerVC driver, but later the PowerVC driver service is dead.

Symptoms

After successfully deploying a controller node you might subsequently notice that PowerVC services are dead and the system is unusable.

Causes

A configuration error such as a mistake for the IP address of the PowerVC host system can cause the PowerVC driver services to eventually fail.

Diagnosing the problem

You can use the following command to check the status of the PowerVC driver service.

```
service openstack-nova-powervc status
```

Resolving the problem

If PowerVC services are dead, ensure that the settings in the environment file are correct. Then try the deployment task again.

Compute node is deployed or updated successfully; later Open vSwitch agent is down

You successfully deploy, update, or apply a fix pack to a compute node, but later the Open vSwitch agent on the compute node is down.

Symptoms

After you successfully deploy, update, or apply a fix pack to a compute node, you are unable to start a new virtual machine.

You check the status of the Open vSwitch agent by running the following command on your controller node:

```
neutron agent-list
```

The Open vSwitch agent is down.

Causes

This problem can occur for various reasons, one of which is that there is an error in the `/etc/neutron/plugins/openvswitch/ovs_neutron_plugin.ini` file on the compute node.

Resolving the problem

To resolve the problem, try to restart the Open vSwitch agent by running the following command on the compute node:

```
service neutron-openvswitch-agent restart
```

Take one of the following actions:

- If the agent status on the controller is active, try to start the virtual machine.
- If the agent status is not active, see the `/var/log/neutron/openvswitch-agent.log` file on the compute node.

If you applied a fix pack to multiple PowerKVM compute nodes, it is likely that the agent is down on all of them. You can resolve the problem similarly for each compute node.

Fingerprint error during deploy

When redeploying to a node where a topology has previously been deployed, you receive an SSH fingerprint error.

Symptoms

When updating the environment during the deployment process, you receive an error similar to the following:

```
ERROR: fingerprint 1d:51:e2:2d:79:85:9c:c0:7c:9d:82:c7:3f:94:ad:b4 does not
match for "your_system.domain.com,x.x.x.x"
ERROR: /opt/chef/embedded/lib/ruby/gems/1.9.1/gems/net-ssh-2.9.0/lib/net/ssh
/verifiers/secure.rb:50:in `process_cache_miss'
/opt/chef/embedded/lib/ruby/gems/1.9.1/gems/net-ssh-2.9.0/lib/net/ssh/
verifiers/secure.rb:35:in `verify'
/opt/chef/embedded/lib/ruby/gems/1.9.1/gems/net-ssh-2.9.0/lib/net/ssh/
verifiers/strict.rb:16:in `verify'
/opt/chef/embedded/lib/ruby/gems/1.9.1/gems/net-ssh-2.9.0/lib/net/ssh/
verifiers/lenient.rb:15:in `verify'
/opt/chef/embedded/lib/ruby/gems/1.9.1/gems/net-ssh-2.9.0/lib/net/ssh
```

```
/transport/kex/diffie_hellman_group1_sha1.rb:173:in `verify_server_key'
/opt/chef/embedded/lib/ruby/gems/1.9.1/gems/net-ssh-2.9.0/lib/net/ssh/
transport/kex/diffie_hellman_group1_sha1.rb:68:in `exchange_keys'
/opt/chef/embedded/lib/ruby/gems/1.9.1/gems/net-ssh-2.9.0/lib/net/ssh/
transport/algorithms.rb:364:in `exchange_keys'
/opt/chef/embedded/lib/ruby/gems/1.9.1/gems/net-ssh-2.9.0/lib/net/ssh/
transport/algorithms.rb:205:in `proceed!'
/opt/chef/embedded/lib/ruby/gems/1.9.1/gems/net-ssh-2.9.0/lib/net/ssh/
transport/algorithms.rb:196:in `send_kexinit'
/opt/chef/embedded/lib/ruby/gems/1.9.1/gems/net-ssh-2.9.0/lib/net/ssh/
transport/algorithms.rb:151:in `accept_kexinit'
/opt/chef/embedded/lib/ruby/gems/1.9.1/gems/net-ssh-2.9.0/lib/net/ssh/
transport/session.rb:195:in `block in poll_message'
/opt/chef/embedded/lib/ruby/gems/1.9.1/gems/net-ssh-2.9.0/lib/net/ssh/
transport/session.rb:173:in `loop'
/opt/chef/embedded/lib/ruby/gems/1.9.1/gems/net-ssh-2.9.0/lib/net/ssh/
transport/session.rb:173:in `poll_message'
/opt/chef/embedded/lib/ruby/gems/1.9.1/gems/net-ssh-2.9.0/lib/net/ssh/
transport/session.rb:210:in `block in wait'
/opt/chef/embedded/lib/ruby/gems/1.9.1/gems/net-ssh-2.9.0/lib/net/ssh/
transport/session.rb:208:in `loop'
/opt/chef/embedded/lib/ruby/gems/1.9.1/gems/net-ssh-2.9.0/lib/net/ssh/
transport/session.rb:208:in `wait'
/opt/chef/embedded/lib/ruby/gems/1.9.1/gems/net-ssh-2.9.0/lib/net/ssh/
transport/session.rb:87:in `initialize'
/opt/chef/embedded/lib/ruby/gems/1.9.1/gems/net-ssh-2.9.0/lib/net/
ssh.rb:202:in `new'
/opt/chef/embedded/lib/ruby/gems/1.9.1/gems/net-ssh-2.9.0/lib/net/
ssh.rb:202:in `start'
/opt/chef/embedded/lib/ruby/gems/1.9.1/gems/net-ssh-multi-1.2.0/lib/net/
ssh/multi/server.rb:185:in `new_session'
/opt/chef/embedded/lib/ruby/gems/1.9.1/gems/chef-11.12.4/lib/chef/
monkey_patches/net-ssh-multi.rb:79:in `next_session'
/opt/chef/embedded/lib/ruby/gems/1.9.1/gems/net-ssh-multi-1.2.0/lib/
net/ssh/multi/server.rb:138:in `session'
/opt/chef/embedded/lib/ruby/gems/1.9.1/gems/net-ssh-multi-1.2.0/lib/
net/ssh/multi/session_actions.rb:36:in `block (2 levels) in sessions'
ERROR: Node 'powervc-rhel64' at your_system.domain.com verification error:
Could not connect to your_system.domain.com
as user 'root' with the specified password : Ssh execution of command
'2>/dev/null 1>&2 echo hello' failed:
fingerprint 1d:51:e2:2d:79:85:9c:c0:7c:9d:82:c7:3f:94:ad:b4 does not match
for "your_system.domain.com,x.x.x.x"
ERROR: RuntimeError: Topology node verification failed.
```

Resolving the problem

You might need to delete the SSH fingerprint on your deployment server. The deployment knife commands use SSH to communicate with the node system to deploy OpenStack. Use one of the following methods to delete SSH fingerprints from the `.ssh/known_hosts` file on the deployment server:

- Run the following command:

  ```
  ssh-keygen -R node fqdn
  ```

- Edit the `.ssh/known_hosts` file in a text editor and delete the line that contains your node system.

Error restarting neutron when deploying a topology

When you attempt to deploy a topology, you receive an error about restarting the neutron-plugin-openvswitch-agent service.

Symptoms

You see something similar to the following error:

```
ESC[31mError executing action `restart` on resource 'service[neutron-plugin-openvswitch-agent]'ESC
```

Causes

You might not have the right version of a package, such as OpenSSL.

Resolving the problem

Ensure that the yum repositories are configured correctly so that package updates are accessible. For more information, see “Configuring operating system yum repositories on the deployment server” on page 36.

Then, try the deployment task again.

Cinder error occurs when you deploy a controller node

You see a cinder error in the log after you try to deploy a controller node.

Symptoms

You try to deploy a controller node and it fails. You check the log and see the following error message:

```
Error executing action `run` on resource 'execute[cinder-manage db sync]'
IOError: [Errno 13] Permission denied: '/var/log/cinder/cinder.log'
```

Causes

This error can occur if Service Pack 1 is not installed before you deploy a controller node with storage support.

Resolving the problem

To resolve the problem, manually delete the `/var/log/cinder/cinder.log` file and reconverge the controller node.

Known problems and solutions for an Administrator

If you are using IBM Cloud Manager with OpenStack with the *Administrator* role, review theses known problems and solutions that you might encounter.

If your particular problem is not represented, see “Searching knowledge bases” on page 248 for a list of other references, particularly the Techdocs Web site, where more recent solutions or workarounds might instead reside.

Related tasks:

“Searching knowledge bases” on page 248
You can often find solutions to problems by searching IBM knowledge bases. You can optimize your results by using available resources, support tools, and search methods.

Known issues

There are several known issues with the current release of IBM Cloud Manager with OpenStack.

If your particular problem is not represented, see “Searching knowledge bases” on page 248 for a list of other references, particularly the Techdocs Web site, where more recent solutions or workarounds might instead reside.

Installation fails after uninstall

When you attempt to install IBM Cloud Manager with OpenStack on the deployment server, after performing an uninstall, the installation fails with a `NONFATAL_ERROR` during console and graphical user interface mode installs or with a nonzero return code during silent mode installs.

Symptoms

The installation fails with a `NONFATAL_ERROR` during console and graphical user interface mode installs or with a nonzero return code during silent mode installs after performing an uninstall.

Causes
Occasionally, the IBM Cloud Manager with OpenStack uninstall process does not remove the `openstack-utils` RPM. This symptom appears when the installation process attempts to install this package again. To verify this is the case, review the `/tmp/cmwo-installer.log` to determine whether the `openstack-utils` was already installed. If a line similar to the following is found in the `/tmp/cmwo-installer.log` then this is the cause of the symptom:

```
package openstack-utils-2014.1.1-201405072052.ibm.el6.15.noarch is already installed
```

If the symptom was caused for other reasons, contact your support representative.

Resolving the problem
To resolve the problem, complete the following steps.

- Uninstall IBM Cloud Manager with OpenStack by running, **cmwo_uninstall**.
- Uninstall the `openstack-utils` RPM by running, **yum erase openstack-utils -y**.
- Install IBM Cloud Manager with OpenStack.

Cannot resize z/VM instance (self-service portal)
You cannot resize a z/VM instance through IBM Cloud Manager with OpenStack.

Symptoms
While you attempt to resize a z/VM instance through IBM Cloud Manager with OpenStack, it might take a long time and appear to fail.

Causes
The virtual machine status in the OpenStack command line displays the following message:

```
VERIFY_RESIZE
```

However, you can access the virtual machine, and the virtual machine resize is successful. The resize operation can take a long time (for example, 2 hours or more), depending on the size of the virtual machine. IBM Cloud Manager with OpenStack times out while waiting for the resize operation to finish.

Resolving the problem
You can extend the timeout value for this operation in the `openstack.properties` file. Modify the **com.ibm.cfs.cloud.openstack.client.wait.for.server.resize.timeout.in.seconds** value.

Cannot resize PowerVC instance
You cannot resize a PowerVC instance through IBM Cloud Manager with OpenStack.

Symptoms
While you attempt to resize a PowerVC instance through IBM Cloud Manager with OpenStack, you might receive an error similar to the following message.

```
Exception during message handling:
Get error: PVCExpendvdiskFCMapException:
Flashcopy is in progress for boot volume, volume size didn't change.
Please try again later. (HTTP 409) (Request-ID: req-xxxx)
```

Resolving the problem
After you create the virtual machine in PowerVC from IBM Cloud Manager with OpenStack, the create process completes, but the FlashCopy might not complete in the backend of PowerVC. You cannot complete the resize operation during this time. You must wait to resize the instance until the FlashCopy completes as well.

Cannot complete active resize of PowerVC instance

You cannot complete an active resize of a PowerVC instance through IBM Cloud Manager with OpenStack.

Symptoms

While you attempt to actively resize a PowerVC instance through IBM Cloud Manager with OpenStack, you might receive an error.

Resolving the problem

If the virtual machine (running PowerVC from IBM Cloud Manager with OpenStack), does not meet the requirements for an active restart, you must stop the virtual machine before you resize the instance. For more information, see "Resizing an instance (OpenStack)" on page 210.

Instance status on PowerVC driver hangs and the migration fails

Instance `migrating` status hangs on PowerVC driver and the migration fails.

Symptoms

When you are using IBM Cloud Manager with OpenStack with the PowerVC driver to migrate an instance on PowerVC, you might encounter a problem where live migration fails. The instance status might hang and stay in `migrating` status for a long time. When this occurs, you cannot use the instance.

Resolving the problem

To resolve this problem, use **ssh** to access the IBM Cloud Manager with OpenStack controller node and check the log file `nova-powervc.log` in the directory `/var/log/powervc`. Fix the problem that is described in the log file. After the problem is fixed, use the command **nova reset-state** to reset the instance status. After the status of the instance is changed to `active` by the PowerVC driver code, rerun the live migration operation.

Instance resizing status on PowerVC driver hangs (self-service portal)

Instance resizing status hangs on PowerVC driver.

Symptoms

When you are using IBM Cloud Manager with OpenStack with the PowerVC driver to resize an instance, the status in the self-service portal hangs and does not update. The confirmation step does not occur.

Resolving the problem

To resolve this issue, you must update an attribute that automatically completes the confirmation resize step. Update the `/etc/nova/nova.conf` file, in the [DEFAULT] section with *resize_confirm_window=1*. For example:

```
openstack.compute.misc_nova: ['api_rate_limit=false',
'flat_injected=true', 'use_ipv6=true', "resize_confirm_window=1"]
```

Finally, run **service openstack-nova-powervc restart** to resolve the issue.

Cannot attach disk to virtual machine without restarting the virtual machine

After you attach a disk to a virtual machine that is running PowerKVM, the new volume is not visible on the associated instance unless you restart the virtual machine.

Symptoms

You attached a disk to a PowerKVM virtual machine by using IBM Cloud Manager with OpenStack, but when you check the instance, the volume does not display. If you restart the virtual machine, the attached disk does display.

Causes

The packages that support hotplug on PowerKVM are out of date.

Resolving the problem

The packages on the guest that you are using for the virtual machine must be updated. To support hotplug, you must update the following packages to the latest versions:

- `powerpc-util`
- `librtas`
- `ppc64-diag`

To download the latest versions of the PowerKVM packages for your Linux distribution, see Index of /software/server/POWER/Linux/yum/IBM.

Two users within the same browser

Logging in as two different users within the same browser shows only the most recent user.

Details

Different tabs or windows of the same browser instance share the same session and browser cookies so the user does not really have two independent sessions. If a user logs in with two different user IDs at the same time, the browser will use information based on that most recent login, and there is no clear indication that one just superseded the other. For example, if a user logs in as UserA in one browser window and UserB in another browser window, both windows are UserB, and all content and settings displayed belong to UserB.

Solution

To log in as a different user with a browser, log out and close all browser instances before logging in as the alternate user. To log in as two different users at the same time, two different browsers are needed, for example, Internet Explorer and Mozilla Firefox.

Display issue with Internet Explorer (self-service portal)

IBM Cloud Manager with OpenStack self-service portal layout and format sometimes appear to be out of place and hard to navigate in Internet Explorer 9 and Internet Explorer 10.

Details

When you are using IBM Cloud Manager with OpenStack self-service portal in Internet Explorer 9 and Internet Explorer 10, you might experience that the layout and format of the screen is difficult to navigate.

Solution

The display issue occurs because Internet Explorer 9 and Internet Explorer 10, by default, display the self-service portal in Internet Explorer Compatibility View mode. To resolve this issue, you must switch from Internet Explorer Compatibility View mode to the standard mode.

1. To switch from Internet Explorer Compatibility View mode to the standard Internet Explorer mode, click **Compatibility View**.

   The **Compatibility View** icon is found to the right of the address bar.

2. If the **Compatibility View** icon is not visible, press **F12**.
3. Depending on which version of Internet Explorer you are using, continue with one of the following steps:

- If you are using Internet Explorer 9, click **Browser Mode: IE9 > Internet Explorer 9** to select the standard mode.

  Note: The only check mark in the menu is in front of Internet Explorer 9.

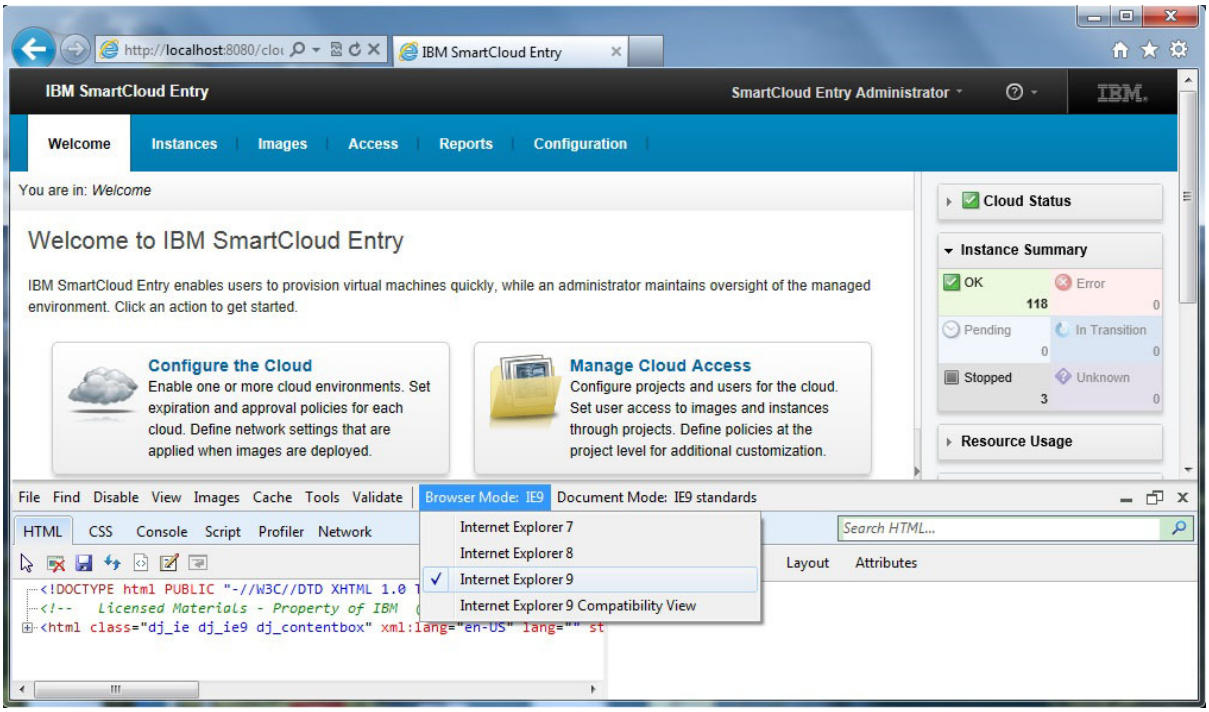

- If you are using Internet Explorer 10, click **Browser Mode: IE10 > Internet Explorer 10** to select the standard mode.

  Note: The only check mark in the menu is in front of Internet Explorer 10.

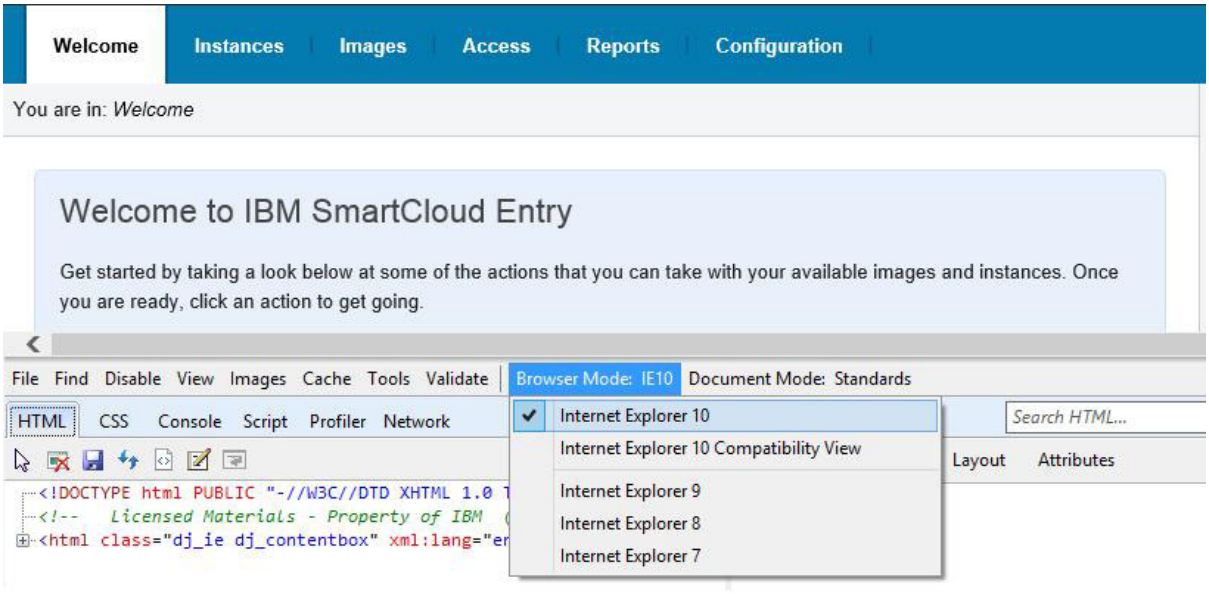

Tip: If IBM Cloud Manager with OpenStack switches from standard mode to compatibility view mode automatically, clear the option in **Tools** > **Internet options** > **Advanced** > **Automatically recover from page layout errors with Compatibility View**.

Delete of an instance while a storage flashcopy is running against the instance will cause the delete to fail

An IBM Cloud Manager with OpenStack instance cannot be deleted immediately after the same instance has been deployed or captured.

Details

A recently deployed or captured IBM Cloud Manager with OpenStack instance has a status of "OK" and the cloud manager instance status also shows "OK", but the storage flashcopy may still be in progress, and an attempt to delete the instance while the flashcopy is running will cause the delete to fail.

Solution

To prevent an error from happening on a delete soon after a deploy or capture, you must either monitor the flashcopy until it ends, or wait for some conservative period of time, for example, 20 minutes, before attempting the delete.

The best way to tell if the flashcopy has finished is by accessing the storage subsystem user interface (UI). Here is how to use the Storwize storage subsystem UI to determine when the flashcopy is finished:

1. Access Storwize UI.
2. From IBM Cloud Manager with OpenStack, perform the deploy or capture.
3. From the Storwize UI left navigation pane, select: Copy Services, then Flashcopy Mappings.

   A flash copy should be in progress. If a flash copy is not in progress, wait for it to be displayed; in our testing the flashcopy usually appeared about 30-60 seconds after the IBM Cloud Manager with OpenStack deploy or capture was started.
4. Wait until the flash copy completes. Once the flash copy has completed, the IBM Cloud Manager with OpenStack instance can be deleted.

PKI error when adding OpenStack cloud with self-service portal

You encounter an error message when you attempt to add an OpenStack cloud to IBM Cloud Manager with OpenStack using the self-service portal.

Details

You receive the following message when you are attempting to add an OpenStack cloud to IBM Cloud Manager with OpenStack. `Error: Unreachable Cloud :CYX6154E: An error occurred while making the OpenStack identity service token request for user 'sceagent'. The identity service responded with the following status: 500 - Error occurred when dealing with PKI token. The internal reason is '__init__() got an unexpected keyword argument 'output'' Verify that the identity service is running, and that the user name, password and tenant name are correct. Contact your system administrator for assistance.`

Solution

Due to time change on the deployment server, or because a Network Time Protocol (NTP) server is not being used, the self-signed certificates that are used for PKI tokens can become invalid. To fix the issue, ensure that the deployment server operating system has the correct date and time. Then, use the following steps to regenerate the tokens:

1. Stop the OpenStack keystone service. For more information, see “Restarting IBM Cloud Manager with OpenStack services” on page 172.

   Note: Stop only the keystone service.
2. Clear the keystone certificates by using the following command:

   ```
   # rm -rf keystone base signing directory
   ```

Where *keystone base signing directory* is found in the `/etc/keystone/keystone.conf` file of the controller, under the [signing] section. By default, the base signing directory is `/etc/keystone/ssl`.

3. Generate new PKI keys. Alternatively, you can import external certificates.
   a. Find the **keystone user** and **keystone group** by running the following command on the OpenStack controller:

      ```
      ls -la /etc/keystone/ssl
      ```

      The **keystone user** is the owner of that directory. The **keystone group** is the group that is identified for that directory.
   b. Use the following command to generate new PKI keys:

      ```
      # keystone-manage pki_setup --keystone-user keystone user --keystone-group keystone group
      ```
4. Verify the new ticket by using the following command.

   ```
   # openssl x509 -in  /etc/keystone/ssl/certs/signing_cert.pem -text -noout
   ```

   The output lists the current date and time in the output for the **Not Before** value.
5. Clear the PKI signing caches on the OpenStack controller by using the following commands:

   ```
   rm /var/cache/ceilometer/api/*
   rm /var/cache/cinder/api/*
   rm /var/cache/glance/api/*
   rm /var/cache/glance/registry/*
   rm /var/cache/nova/api/*
   rm /var/lib/neutron/keystone-signing/*
   ```

   Note: If the `signing_dir` directory is not present, temporary caches are used. If temporary caches are used, restarting the service creates a new cache.
6. Restart OpenStack services and the self-service portal. For more information, see "Restarting IBM Cloud Manager with OpenStack services" on page 172.

Error opening sockets to server when using DB2

When you are running IBM Cloud Manager with OpenStack with DB2 as the database, you intermittently encounter an unexpected condition that the server cannot fulfill the request.

Details

The following specific error is displayed:

```
The server encountered an unexpected condition which prevented it from fulfilling the request
```

If you check the IBM Cloud Manager with OpenStack log you might see an exception similar to the following:

```
[05/16/13 04:38:17:009] 33904 SEVERE: CYX1846E: Internal database error.
<openjpa-2.1.0-r422266:1071316 fatal general error> org.apache.openjpa.persistence.
PersistenceException: [jcc][t4][2043][11550][4.8.87] Exception java.net.NoRouteToHostException:
Error opening socket to server localhost/127.0.0.1 on port 50,000 with message: Cannot assign
requested address. ERRORCODE=-4499, SQLSTATE=08001. Stack Trace: <openjpa-2.1.0-r422266:1071316
fatal general error> org.apache.openjpa.persistence.PersistenceException:
[jcc][t4][2043][11550][4.8.87] Exception java.net.NoRouteToHostException: Error opening socket to
server localhost/127.0.0.1 on port 50,000 with message: Cannot assign requested address.
ERRORCODE=-4499, SQLSTATE=08001
```

DB2 socket connections in TIMED_WAIT state are exhausting the available ports for DB2 connections. Normally, the system releases sockets in TIMED_WAIT state after 2 minutes. However, in some cases, this wait is too long and there are no more available sockets to use. You can reduce the amount of wait time by adjusting the TIMED_WAIT reuse and recycle values on the IBM Cloud Manager with OpenStack server.

Solution

To adjust the TIMED_WAIT reuse and recycle values on the controller node for the deployed topology, add the following to the `/etc/sysctl.conf` file:

```
net.ipv4.tcp_tw_reuse = 1
net.ipv4.tcp_tw_recycle = 1
```

Make the TIMED_WAIT reuse and recycle values effective by using the following command:

```
/sbin/sysctl -p
```

Error occurs when installing qpid

During the installation of IBM Cloud Manager with OpenStack or when deploying a cloud topology, you might encounter an error that the Qpid package did not install correctly.

Symptoms

During installation or when deploying a topology you might see an error similar to the following:

```
[2014-02-12T10:43:31-06:00] DEBUG: Chef::Exceptions::Exec: package[qpid-cpp-server]
(/var/chef/cache/cookbooks/qpid/providers/setup.rb line 15) had an error:
Chef::Exceptions::Exec:  returned 1, expected 0
```

Causes

A conflict might exist with Red Hat Enterprise Linux GNOME Desktop Environment (gnote), KDE Desktop Environment (Akonadi - the PIM Storage Service), or another external package that depends on version 1.41.0 of Boost C++ libraries.

Resolving the problem

To remove the conflict, identify and remove the conflicting package. Run the following command to identify any conflicting packages:

```
rpm -qa | grep -e "boost-.*-1\.41\.0.*" | xargs rpm --test -e
```

Then run the following command to remove any conflicting packages:

```
yum remove "package_name"
```

where *package_name* is the name of the conflicting package.

After any conflicting packages are removed, try the installation or deployment task again.

Authentication required error with Red Hat Enterprise Linux 6.4

After installing IBM Cloud Manager with OpenStack, you see authentication errors when running OpenStack commands.

Symptoms

When you run commands such as **`glance image-list`** or **`neutron agent-list`**, you encounter authentication required errors.

Causes

The default version of OpenSSL for Red Hat Enterprise Linux version 6.4 is `openssl 1.0.0-27`. To resolve the authentication error, the OpenSSL version must be updated.

Resolving the problem

Update the OpenSSL version to eliminate the authentication errors. You can find an updated version of OpenSSL in the Red Hat Enterprise Linux version 6.5 ISO. See the following example for updating the OpenSSL version:

```
# rpm -qa|grep openssl
openssl-devel-1.0.1e-15.el6.ppc64
openssl-1.0.1e-15.el6.ppc
openssl-1.0.1e-15.el6.ppc64
```

Installation fails when SSH is interrupted

When attempting to install IBM Cloud Manager with OpenStack on the deployment server, the installation fails and you are unable to uninstall the product.

Symptoms

You might see an error similar to the following message: `This Application has Unexpectedly Quit`.

Causes

InstallAnywhere cannot tolerate a Secure Shell (SSH) interruption. The results of the attempted installation are unpredictable. The installation is incomplete and the uninstall utility is unusable.

Resolving the problem

You must manually clean up the system.

1. Uninstall the Chef and Chef server RPMs if they were installed.
2. Delete the `/etc/chef` and `/etc/chef-server` directories if they exist.
3. Delete the `/root/.chef` directory if it exists.
4. Delete the `/opt/ibm/cmwo` directory.
5. Delete the `/var/..com.zerog.registry.xml` file if it exists. This file is the InstallAnywhere registry.
6. Trying running the installation again

SSL connection error occurs when certificate expires

An SSL connection error occurs due to an expired certificate.

Symptoms

You receive an SSL connection error.

Causes

The error can occur when your SSL certificate is expired.

Resolving the problem

Determine whether the certificate is expired. If it expired, you must delete the certificate and create a new certificate with a specific expiration date. To resolve the problem, complete the following steps:

1. Determine the name of your certificate by going to the following directory:

   ```
   /etc/pki/
   ```

   Run the following command:

   ```
   # certutil -L -d qpid/
   ```

   Check the output for the name of the certificate, for example, `qpidssl` or `openssl`.
2. Determine whether the certificate is expired by running the following command:

   ```
   # certutil -V -n certificate_name -u C -d qpid/
   ```

   where *certificate_name* is the name of your certificate.
3. Determine the expiration date of the certificate by running the following command:

   ```
   # certutil -L -d qpid/ -n certificate_name
   ```
4. Delete the expired certificate by running the following command:

   ```
   # certutil -D -d qpid/  qpid/cert.password -n certificate_name
   ```
5. Create a certificate by running the following command:

```
# certutil -S -d ./qpid/ -n certificate_name -s 'CN=certificate_name' -t 'CT,,' -x -v 10
-f ./qpid/cert.password -z /usr/bin/certutil
```

IBM Knowledge Center seems to cut off bottom of page

While you are browsing content about IBM Cloud Manager with OpenStack in IBM Knowledge Center, you encounter a page that seems to cut off part of the information at the bottom of the page.

Symptoms

You might see extra white space at the bottom of an IBM Knowledge Center topic that contains a table. The link to a parent topic might be missing.

Causes

The framework has a known issue.

Resolving the problem

Try displaying the content in a different web browser.

Unstable HTTPS access to Gateway

If you deploy many virtual machines, you might see intermittent operation failures that are related to a known Python bug.

Symptoms

You see the following error when you use the IBM Cloud Manager with OpenStack self-service portal, or in a log file.

```
Message: CYX1417E - Could not connect to cloud.
Explanation: Can not connect to cloud or response timeout, this may caused by the
incorrect cloud settings or cloud server is too busy to response. Check the logs
```

However, the network connection between IBM Cloud Manager with OpenStack and the Gateway appears to be normal.

Causes

The intermittent task failures that you see might be related to the following known Python bug: http://bugs.python.org/issue8240

Resolving the problem

The bug is not fixed in Python 2.6, but with the following actions, you can avoid the instability or see an improvement in performance:

- The better that your network performance is, the less you experience this issue. If you improve network performance, the chance for intermittent operation failures decreases.
- If HTTPS is not required in your environment, you can disable HTTPS in the Gateway to avoid this issue.
- The fix in the Python bug link is verified. You can build and apply the fix if you have the required skills.

Using IPv6 and Python version 2.6

If you want to use the IPv6 management network and you are also using Python, version 2.6, you might need to patch Python with an update to version 2.6.

Symptoms

For example, you attempt to configure glance_api_servers in `nova.conf` to [2001:db8::22]:9696. You receive en error similar to the following output:

```
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
  File "urlparse.py", line 105, in port
    return int(port, 10)
ValueError: invalid literal for int() with base 10: 'db8::22]:9696'
```

Causes

The root cause of this problem is that the urlparse lib in Python 2.6 does not parse IPv6 URLs correctly.

Resolving the problem

To resolve this issue, you must patch the Python 2.6 code to successfully parse the IPv6 management network URL. Use the following instructions to patch Python.

1. For more information about the IPv6 URL parsing problem, see http://bugs.python.org/issue2987.
2. For more information about the fix patch for the bug, see http://bugs.python.org/file10451/python-urlparse-rfc2732-fix.patch.
3. Find the following code in `/usr/lib64/python2.6/urlparse.py`:

```
@property
  def hostname(self):
      netloc = self.netloc
      if "@" in netloc:
          netloc = netloc.rsplit("@", 1)[1]
      if ":" in netloc:
          netloc = netloc.split(":", 1)[0]
      return netloc.lower() or None

  @property
  def port(self):
      netloc = self.netloc
      if "@" in netloc:
          netloc = netloc.rsplit("@", 1)[1]
      if ":" in netloc:
          port = netloc.split(":", 1)[1]
          return int(port, 10)
      return None
```

4. Change the bold lines in step 3 into:

```
@property
   def hostname(self):
       netloc = self.netloc
       if "@" in netloc:
           netloc = netloc.rsplit("@", 1)[1]
     if netloc and netloc[-1] != ']' and ":" in netloc:
        netloc = netloc.rsplit(":", 1)[0]
       return netloc.lower() or None

   @property
   def port(self):
       netloc = self.netloc
       if "@" in netloc:
           netloc = netloc.rsplit("@", 1)[1]
       if netloc and netloc[-1] != ']' and ":" in netloc:
        port = netloc.rsplit(":", 1)[1]
           return int(port, 10)
       return None
```

5. If you need to apply the fix to more than one system, you can generate a patch by completing the following steps:
   a. `cp /usr/lib64/python2.6/urlparse.py /usr/lib64/python2.6/urlparse.py.ori`
   b. Update `/usr/lib64/python2.6/urlparse.py` according to the previous steps.
   c. Generate the patch by using the following command: `diff /usr/lib64/python2.6/urlparse.py.ori /usr/lib64/python2.6/urlparse.py > urlparse.patch`

d. Apply the patch on other system by using the following command: `patch /usr/lib64/python2.6/urlparse.py urlparse.patch`

6. After the code is changed or the patch is applied, clean up the pyc and pyo files for `urlparse.py`:

```
rm /usr/lib64/python2.6/urlparse.pyc
rm /usr/lib64/python2.6/urlparse.pyo
```

7. Restart the OpenStack services that are configured with IPv6 management network URLs.

Error accessing OpenStack resources (self-service portal)

After deploying a Minimal topology with the IBM Cloud Manager with OpenStack self-service portal enabled, you see errors accessing OpenStack resources (for example, volumes, images, flavors, etc.) from the self-service portal.

Symptoms

When you access OpenStack resources from the self-service portal, you encounter OpenStack identity service token request errors.

Causes

The OpenStack simple token is not valid.

Resolving the problem

Generate a new simple token using the following command:

```
$ dd if=/dev/urandom bs=16 count=1 2>/dev/null | base64
```

Then set the **simple_token_secret** option in the `/etc/keystone/keystone.conf` file to the new simple token. Do the same for the **com.ibm.cfs.token.key** option in the `/var/opt/ibm/.SCE41/authentication.properties` file. Finally, restart the openstack-keystone and sce services.

```
$ service openstack-keystone restart
$ service sce restart
```

If the Chef client is run on the node in the future, by running the knife os manage update command on the deployment server, you must reapply these changes.

Error viewing OpenStack resources with multiple regions (self-service portal)

After you deploy a topology that uses multiple regions with the IBM Cloud Manager with OpenStack self-service portal enabled, you see errors when you attempt to list OpenStack resources.

Symptoms

When you view OpenStack resources from the self-service portal, you encounter errors such as the following:

```
Failed to get the endpoint url of the volumes.
```

Causes

PKI tokens are being used for the authentication strategy and the tokens have become too large.

Resolving the problem

Switch the authentication strategy to uuid by setting this value in your chef environment file:

```
['openstack']['auth']['strategy'] = 'uuid'
```

Then redeploy the topology.

Limitations

There are several known limitations with the current release of IBM Cloud Manager with OpenStack.

If your particular problem is not represented, see “Searching knowledge bases” on page 248 for a list of other references, particularly the Techdocs Web site, where more recent solutions or workarounds might instead reside.

Unable to configure Keystone SSL/TLS from the deployment server

You cannot configure Keystone SSL/TLS from the deployment server.

Symptoms

Although OpenStack Keystone supports SSL/TLS encrypted communications between the Keystone service and the LDAP service, you cannot update the configuration from the deployment server.

Resolving the problem

To enable SSL/TLS, you must use the "**misc_keystone**" support to add arbitrary configuration values to `keystone.conf`.

Update the environment file of the topology to include the **misc_keystone** property contain the following attributes:

```
"identity": {
       "verbose": false,
       "debug": false,
       "misc_keystone": [
         "[ldap]",
         "use_tls=true",
         "tls_cacertfile=/etc/keystone/ssl/certs/cacert.pem",
         "tls_cacertdir=/etc/keystone/ssl/certs/",
         "tls_req_cert=demand"
       ]
     },
```

Note:

1. Duplicate `[ldap]` sections are created in the `keystone.conf` file, but the Keystone service merges all properties together.
2. If the environment values change, on subsequent updates the updated values in the environment file completely overwrites the old values.

Billing limitation with OpenStack Cinder service

In the IBM Cloud Manager with OpenStack self-service portal, there is a limitation when you use the Cinder service with the billing function.

Symptoms

Billing does not work with the volume configurable product if you use the Cinder service API, version 1.

Resolving the problem

If you want to use the billing function for volumes, you must use the Cinder service API, version 2.

IP namespace limitation

Linux IP namespaces are not supported by IBM Cloud Manager with OpenStack

By default, the Red Hat Enterprise Linux 6.5 kernel enables namespace support; however, to allow its use by OpenStack, a new version of the **ip** command is required. The new version must support the `netns` option. This support is provided in OpenStack on Red Hat Enterprise Linux (RDO), but running RDO with IBM Cloud Manager with OpenStack is not supported.

Starting IBM Cloud Manager with OpenStack on a high scale cloud

When starting or restarting IBM Cloud Manager with OpenStack on a high scale cloud, the synchronization between IBM Cloud Manager with OpenStack and the cloud may take longer than

expected. This resynchronization may cause operations such as deploying, deleting, or resizing an instance to be delayed or even fail. Wait for the synchronization to complete before attempting these actions.

Limitations when using VMware within the self-service portal

- The vCenter linked mode is not supported due to limited testing. You are responsible for any errors while you are using linked vCenters mode.
- The size of a virtual disk can be increased either at deployment time or through the IBM Cloud Manager with OpenStack self-service portal resize function. This increases the size of the disk, however, the guest file system is not changed and does not automatically use the increased size. To increase the guest file system to use the larger disk size, see the VMware documentation and guest operating system documentation.
- If your virtual system disks contain a Logical Volume Manager, then you must install VMware Tools inside the guest so that vCenter can customize the image during the deploy operation. For more information about LVM support, see the VMware documentation.
- All snapshots must be integrated before you create a template.
- Preferably, use shared storage for the hosts that are part of a cluster deployment target. If this is not possible, then remove hosts from the cluster while they are in maintenance mode. This prevents the default storage selection algorithm from selecting a data store that is only available from the host in maintenance mode. Selecting such a data store would cause a deployment to fail. For more information about the default storage selection algorithm, see "VMware datastore assignment during deployment" on page 141.

VMware capture instance:

The capture instance function is implemented using the VMware clone to template feature. The IBM Cloud Manager with OpenStack allows you to configure the optional properties in `vmware.properties` for controlling where the new template is created when a capture request is initiated:

`com.ibm.cfs.cloud.vmware.capture.image.datastore.names`
Datastore(s) used when capturing the image of an instance, for example when creating a template. This list is a series of datastore names separated by commas.

`com.ibm.cfs.cloud.vmware.capture.image.destination.name`
The destination host or cluster for where the new template will be placed.

`com.ibm.cfs.cloud.vmware.capture.image.destination.type`
The type of the destination for the new template, either HOST or CLUSTER.

`com.ibm.cfs.cloud.vmware.capture.image.folder`
The folder path for where to place the new template, for example, `. /DatacenterName/vm/FolderName`.

If these properties are not specified, the default behavior is to create the new template in the same location as the existing virtual machine.

VMware storage:

- When attaching a new storage volume in a cloud using VMware vCenter Server cloud, the space in assigned storage name will be ignored in the attached storage volume.

Include only ASCII characters in configuration files

When editing self-service portal configuration files, only use ASCII characters. If non-ASCII characters are used, the original characters are not preserved and display as garbled text in the IBM Cloud Manager with OpenStack self-service portal user interface. The configuration files include all the *.properties* files in the home directory.

Disk resize support

In a Shared Storage Pool environment, changing the disk size when deploying is not supported. In addition, when deploying an IBM i workload, disk resize is not available. In these cases, the disk size property is not displayed in the output of `POST /cloud/api/workload` and the disk resize field is not displayed on the Advanced deployment window.

The install path cannot contain non-English characters

When installing IBM Cloud Manager with OpenStack, the path to the installer cannot contain non-English characters.

Unable to restore a virtual server with fixed VHD type

When a virtual server is deployed by using an image with a fixed VHD type on Hyper-V, that virtual server cannot be successfully restored after a backup.

The limitation is due to the following error condition:

```
CYX6161E: An error occurred at '2013-05-06T21:05:04Z' for virtual machine with identifier
'fd319c6e-91e9-4350-b896-9fdcd3236282'. The error code is '500'. The detailed error is:
ImageTooLarge - Image is larger than instance type allows
File "C:\Program Files (x86)\IBM\SmartCloud Entry\Hyper-V Agent\Python27\lib\site-packages\
nova\compute\manager.py", line 230, in decorated_function
return function(self, context, *args, **kwargs)
File "C:\Program Files (x86)\IBM\SmartCloud Entry\Hyper-V Agent\Python27\lib\site-packages\
nova\compute\manager.py", line 1230, in run_instance
do_run_instance()
File "C:\Program Files (x86)\IBM\SmartCloud Entry\Hyper-V Agent\Python27\lib\site-packages\
nova\openstack\common\lockutils.py", line 242, in inner
retval = f(*args, **kwargs)
File "C:\Program Files (x86)\IBM\SmartCloud Entry\Hyper-V Agent\Python27\lib\site-packages\
nova\compute\manager.py", line 1229, in do_run_instance
admin_password, is_first_time, node, instance)
File "C:\Program Files (x86)\IBM\SmartCloud Entry\Hyper-V Agent\Python27\lib\site-packages\
nova\compute\manager.py", line 877, in _run_instanceself._set_instance_error_state(context,
instance['uuid'])
File "C:\Program Files (x86)\IBM\SmartCloud Entry\Hyper-V Agent\Python27\lib\contextlib.py",
line 24, in __exit__
self.gen.next()
File "C:\Program Files (x86)\IBM\SmartCloud Entry\Hyper-V Agent\Python27\lib\site-packages\
nova\compute\manager.py",
line 798, in _run_instance
image_meta = self._check_image_size(context, instance)
File "C:\Program Files (x86)\IBM\SmartCloud Entry\Hyper-V Agent\Python27\lib\site-packages\
nova\compute\manager.py", line 1040, in _check_image_size
raise exception.ImageTooLarge()
Instance could no longer be found in the Cloud.
It could have been purposely deleted from the Cloud.
```

This issue is tracked at https://bugs.launchpad.net/nova/+bug/1177927.

Validation errors when updating a configuration strategy

When you update the configuration strategy for an image, all fields on the page are validated after you select a local file for the template, user metadata, or mapping. If any of the fields are not valid, an error message displays. Error messages display even if you have not yet provided a value for a required field, such as the mapping field. Proceed by specifying the required fields.

Problem choosing correct datastore during concurrent deployments (VMware vCenter)

The IBM Cloud Manager with OpenStack self-service portal is unable to choose the correct target datastore if VMware vCenter does not provide real-time datastore usage information.

VMware vCenter does not refresh the data of datastore usage before cloning is completed. When you run concurrent deployments, if VMware vCenter cannot provide real-time datastore usage, IBM Cloud

Manager with OpenStack self-service portal cannot decide the correct target datastore under concurrent deployment scenario. As a result, the datastore is not enough when the concurrent deployments need a total datastore resource that is bigger than the datastore nodes left. With insufficient resources, deployments fail as the resources do not meet the minimum amounts. The self-service portal tries to suggest other datastore nodes. It attempts three times, by default. If you want to configure a value greater than three, open the `vmware.properties` file and add the following line: **`"com.ibm.vmware.client.clone.retry.attempts=5"`**, for example.

The failed virtual machine sends a log to the self-service portal interface, as follows:

```
The following instance logs were found in the Cloud. CYX0886E: Unable to determine if the deployed
virtual machine '*' is started and ready for use in the allotted time of 2,700 seconds.
The deployed virtual machine may still be starting or may have had customization problems.
Check the log and the virtual machine for more information.
If the virtual machine was started and is ready for use,
increase the wait time in the vmware.properties file."
```

XFS dependency packages required for Swift

If you want to use Swift on an IBM Power system or System z system that is running IBM Cloud Manager with OpenStack, you must first manually install several XFS file system packages.

Symptoms

Before you can use Swift with IBM Cloud Manager with OpenStack on an IBM Power system or System z system, you must locate and install several XFS file system packages.

Resolving the problem

The XFS packages that are required to run Swift on an IBM Power system or System z system with IBM Cloud Manager with OpenStack are included in the following list:

- xfsdump
- xfsprogs-devel
- xfsprogs

Locate and install these XFS file system packages before attempting to run Swift on an IBM Power system or System z system with IBM Cloud Manager with OpenStack.

Migration information not available

You cannot find instructions for migrating from IBM SmartCloud Entry, version 3.2 to IBM Cloud Manager with OpenStack, version 4.1.

Symptoms

You are not able to locate instructions for migrating from IBM SmartCloud Entry, version 3.2 to IBM Cloud Manager with OpenStack, version 4.1.

Causes

Instructions for migrating from IBM SmartCloud Entry, version 3.2 to IBM Cloud Manager with OpenStack, version 4.1 are not available at this time.

Known problems and solutions for a User

If you are using IBM Cloud Manager with OpenStack with the *User* role, review theses known problems and solutions that you might encounter.

Related tasks:

"Searching knowledge bases" on page 248
You can often find solutions to problems by searching IBM knowledge bases. You can optimize your results by using available resources, support tools, and search methods.

Cached display

When you install or upgrade to IBM Cloud Manager with OpenStack, your browser might not update the images from a previous version of the software.

Be sure to clear your browser's cache after you upgrade to, or install, IBM Cloud Manager with OpenStack.

Saving image exception

If the virtual machine is not created by the current IBM Cloud Manager with OpenStack, you receive an error message when you save an image backup of the virtual machine.

IBM Cloud Manager with OpenStack requires the details of the virtual machine to create an image backup. If this error occurs, you receive the following message:

```
CYX4755E: Instance "instance name" does not have a customization and its virtual machine cannot be saved.
```

Note: *Instance name* is the name of the instance that contains the virtual machine.

Error message language

Some error messages appear in a language other than the language that you set for the IBM Cloud Manager with OpenStack user interface.

These error messages appear in the language set in your operating system. If the error message does not appear in the language that you want, verify your settings in both the IBM Cloud Manager with OpenStack user interface and the operating system.

Internet Explorer display

When you are using IBM Cloud Manager with OpenStack in Internet Explorer 9, 10, or Internet Explorer 11, you might see that layout and formatting makes the screen difficult to navigate.

When you are accessing IBM Cloud Manager with OpenStack with Internet Explorer 9, 10, or Internet Explorer 11, the browser might render the web user interface in Compatibility View mode by default.

Switch from Internet Explorer Compatibility View mode to the standard mode with the following steps:

1. To switch from Compatibility View mode to the standard mode, click the Compatibility View button, which is located on the right side of the address bar (highlighted in green in the following image).

Figure 4. Compatibility button

2. If the Compatibility View button is not visible, press F12.
3. Depending on which version of Internet Explorer you are using, continue with one set of the following steps:
   - If you are using Internet Explorer 9, click **Browser Mode: IE9** > **Internet Explorer 9** to select the standard mode of viewing. Notice that the only check mark in the menu is in front of **Internet Explorer 9**.
   - If you are using Internet Explorer 10, click **Browser Mode: IE10** > **Internet Explorer 10** to select the standard mode of viewing. Notice that the only check mark in the menu is in front of **Internet Explorer 10**.

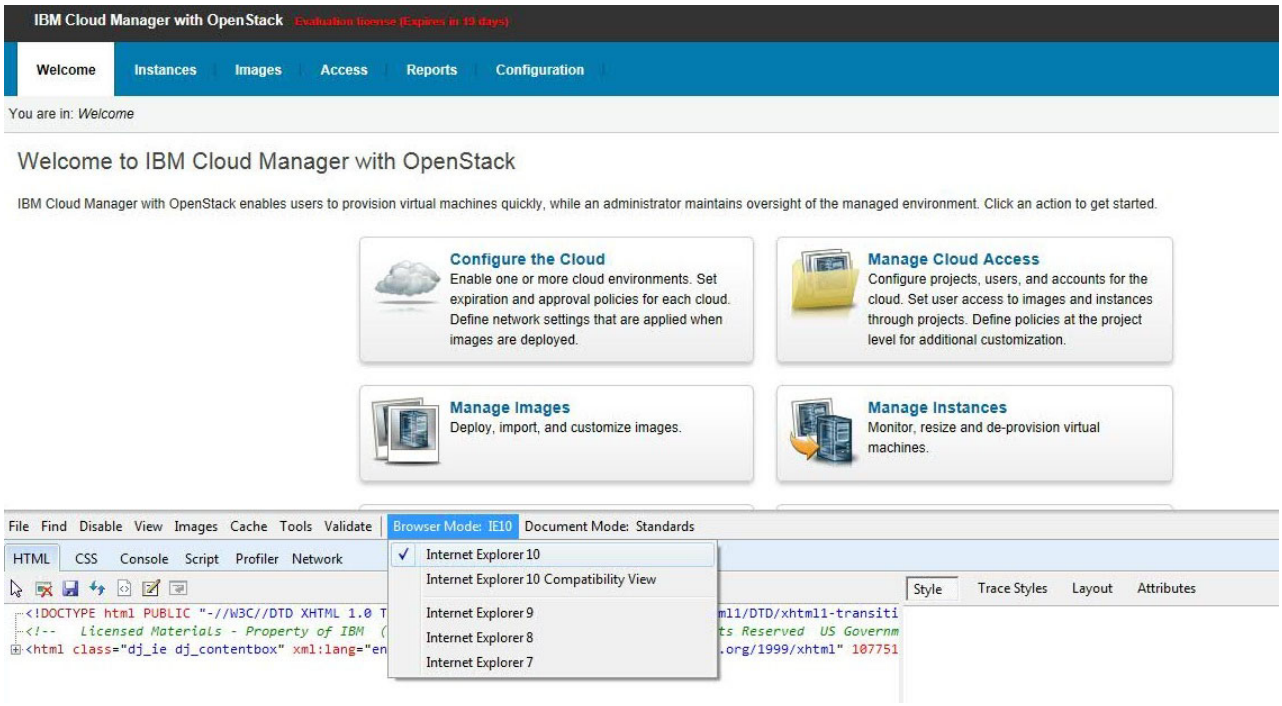

Figure 5. Browser Mode: IE10 menu

- If you are using Internet Explorer 11, click **Document Mode** to select the standard mode of viewing.

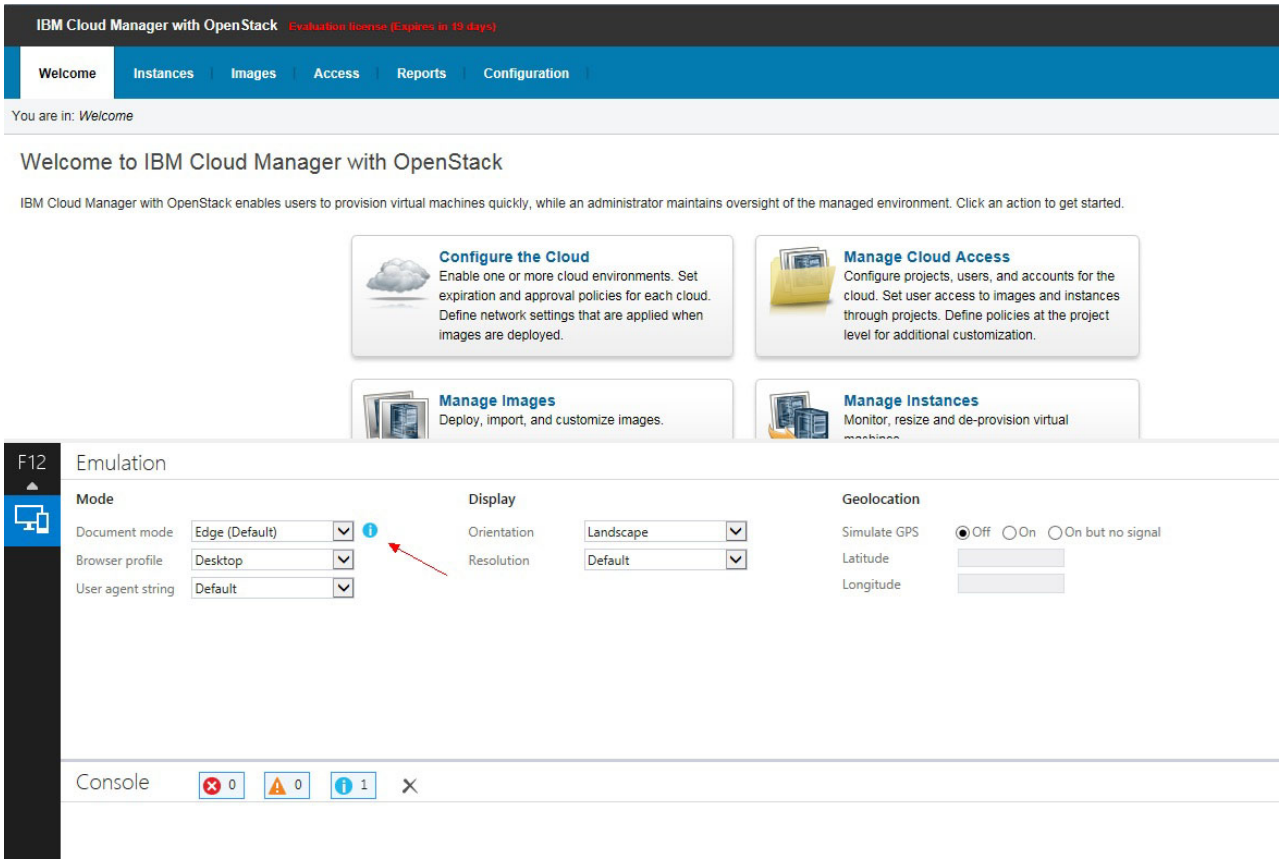

Figure 6. Browser Mode: IE11 menu

Tip: If Internet Explorer, which is accessing IBM Cloud Manager with OpenStack, switches from standard mode to Compatibility View mode automatically, clear the option in **Tools** > **Internet options** > **Advanced** > **Automatically recover from page layout errors with Compatibility View**.

Login fails

If your user login fails because the session times out, there might be a problem with the timezone setting.

Check the following possible resolutions.

- Verify that the IBM Cloud Manager with OpenStack server and client time and timezone match. For example, on the server, if the timezone is Coordinated Universal Time +08:00, the time is 11:27. For the client, the timezone is Coordinated Universal Time +07:00, and the time should be 10:27.

  Note: The client is the clock for the system where the IBM Cloud Manager with OpenStack user interface is being run (such as a personal computer).
- Verify that the user is not locked. The administrator can lock and unlock user accounts from the self-service portal. If a user is locked out, that person cannot log in. If the default administrator account is locked, it unlocks when the server is restarted.

Accessibility

IBM Cloud Manager with OpenStack does not interfere with the accessibility features for supported browsers. For a comprehensive list of accessibility features please visit the accessibility support page for the supported browser that you are using. For a list of supported browsers, see Supported web browsers.

No hardcopy publications are shipped with this program. The IBM Knowledge Center is a worldwide central repository of IBM technical publications hosted in a single application, located at a single URL:

http://www-01.ibm.com/support/knowledgecenter/

Find and navigate technical content more efficiently and easily with improved search, filtering and user experience. Create your own collections of IBM documents with PDF output on-demand.

Note: You can find the IBM Cloud OpenStack Platform product collection here:

http://www-01.ibm.com/support/knowledgecenter/SSUTA8/welcome

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© Copyright IBM Corp. 2012, 2014

Notices

This information was developed for products and services offered in the U.S.A.

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
Armonk, NY 10504-1785
U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: 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 states 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.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Dept. LRAS/Bldg. 903
11501 Burnet Road
Austin, TX 78758-3400
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

© Copyright IBM Corp. 2012, 2014

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

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

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.

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

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 the names and addresses used by an actual business enterprise is entirely coincidental.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® and ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at Copyright and trademark information at www.ibm.com/legal/copytrade.shtml

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

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

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

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

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

Depending upon the configurations deployed, this Software Offering may use session and persistent cookies that collect each user’s user name and password for purposes of session management, authentication, and enhanced user usability. These cookies cannot be disabled.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, See IBM’s Privacy Policy at http://www.ibm.com/privacy and IBM’s Online Privacy Statement at http://www.ibm.com/privacy/details the section entitled “Cookies, Web Beacons and Other Technologies” and the “IBM Software Products and Software-as-a-Service Privacy Statement” at http://www.ibm.com/software/info/product-privacy.

Printed in USA