Red Hat Enterprise Linux OpenStack Platform 5
Deploying OpenStack: Enterprise Environments (Red Hat Enterprise Linux OpenStack Platform Installer)

Deploying Red Hat Enterprise Linux OpenStack Platform in an enterprise environment

Scott Radvan Andrew Dahms

Red Hat Enterprise Linux OpenStack Platform 5 Deploying OpenStack: Enterprise Environments (Red Hat Enterprise Linux OpenStack Platform Installer)

Deploying Red Hat Enterprise Linux OpenStack Platform in an enterprise environment

Scott Radvan
Red Hat Customer Content Services
sradvan@redhat.com

Andrew Dahms
Red Hat Customer Content Services
adahms@redhat.com

Legal Notice

Abstract

This guide explains how to install Red Hat Enterprise Linux OpenStack Platform 5 on Red Hat Enterprise Linux in an enterprise environment using the Red Hat Enterprise Linux OpenStack Platform Installer.

Table of Contents

Chapter 1. Introduction **3**
1.1. Architecture 3
1.2. Deployment Tools and Methods 4
1.3. Supported Virtual Machine Operating Systems 4
1.4. Service Details 4

Chapter 2. Requirements **14**
2.1. User Interface Environment Requirements 14
2.2. User Interface Client Requirements 14
2.3. Host Requirements 15

Chapter 3. Installing the User Interface **19**
3.1. Subscribing to the Required Channels Using Subscription Manager 19
3.2. Preparing an Installation Medium 20
3.3. Installing the User Interface 22
3.4. Options for Installing the User Interface 26

Chapter 4. Configuring the User Interface **29**
4.1. Logging in to the User Interface 29
4.2. Changing a Password 30
4.3. Installation Media 30
4.4. Operating Systems 33
4.5. Provisioning Templates 36
4.6. Partition Tables 39
4.7. Subnets 40
4.8. Users 42
4.9. User Groups 43
4.10. Roles 45

Chapter 5. Adding Hosts **48**
5.1. Host Status Types 48
5.2. Adding Hosts 48
5.3. The Host Overview Page 50
5.4. Removing a Host 53

Chapter 6. Provisioning Red Hat Enterprise Linux OpenStack Platform **54**
6.1. Deployments 54
6.2. Assigning Hosts to Deployment Roles 66
6.3. Configuring Host Networking 69
6.4. Configuring Fencing 71
6.5. Deployment Configurations 73
6.6. Provisioning Red Hat Enterprise Linux OpenStack Platform 75

Chapter 7. Monitoring Hosts **80**
7.1. Dashboard 80
7.2. Reports 81
7.3. Facts 83
7.4. Tasks 84

Appendix A. Firewall Rules **85**

Appendix B. Configuring a Gateway **86**

Appendix C. Configuring the NTP Server **87**

Appendix D. Manual Procedures **88**
D.1. Configuring Fencing on High-Availability Nodes 88
D.2. Adding a Host Manually 89

Appendix E. Revision History **91**

Chapter 1. Introduction

Red Hat Enterprise Linux OpenStack Platform provides the foundation to build a private or public Infrastructure-as-a-Service (IaaS) cloud on top of Red Hat Enterprise Linux. It offers a massively scalable, fault-tolerant platform for the development of cloud-enabled workloads.

The current Red Hat system is based on OpenStack Juno, and packaged so that available physical hardware can be turned into a private, public, or hybrid cloud platform including:

- Fully distributed object storage
- Persistent block-level storage
- Virtual-machine provisioning engine and image storage
- Authentication and authorization mechanism
- Integrated networking
- Web browser-based GUI for both users and administration.

The Red Hat Enterprise Linux OpenStack Platform IaaS cloud is implemented by a collection of interacting services that control its computing, storage, and networking resources. The cloud is managed using a web-based interface which allows administrators to control, provision, and automate OpenStack resources. Additionally, the OpenStack infrastructure is facilitated through an extensive API, which is also available to end users of the cloud.

1.1. Architecture

The following diagram provides a high-level overview of the OpenStack architecture.

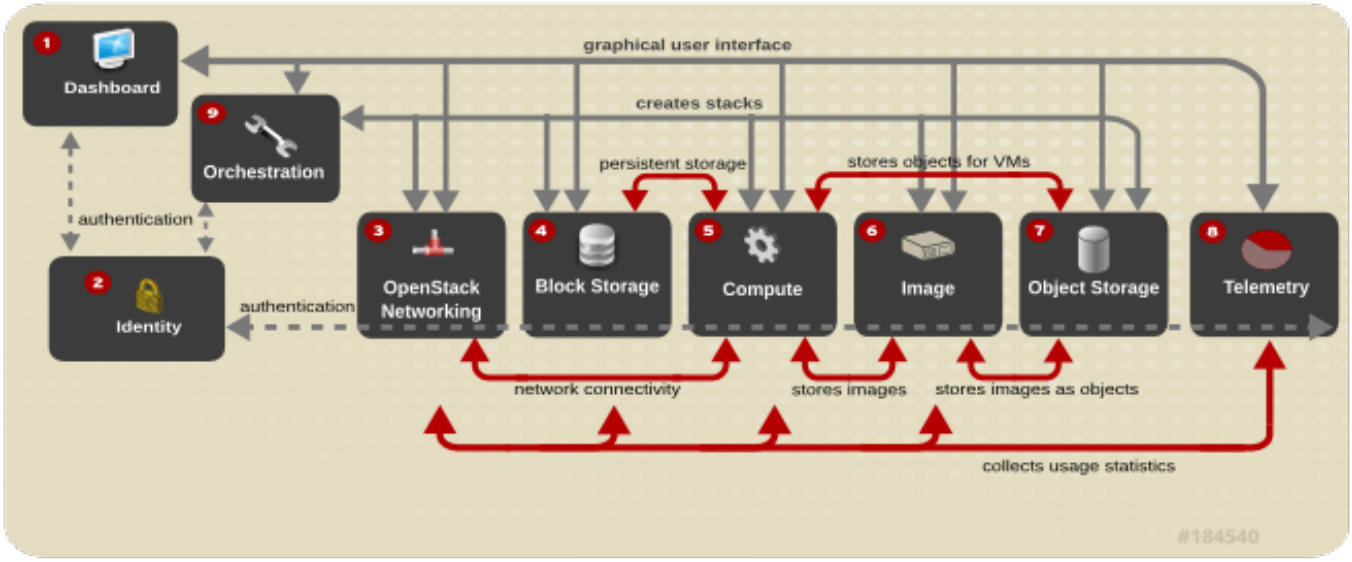

Figure 1.1. OpenStack Architecture

Each OpenStack service has a code name, which is reflected in the names of configuration files and command-line utility programs. For example, the Identity service has a configuration file called **keystone.conf**.

Table 1.1. Services

	Service	Code Name	Description
1	Dashboard	Horizon	A web-based dashboard for managing OpenStack services.
2	Identity	Keystone	A centralized Identity service that provides authentication and authorization for other services, and manages users, tenants, and roles.
3	OpenStack Networking	Neutron	A networking service that provides connectivity between the interfaces of other OpenStack services.
4	Block Storage	Cinder	A service that manages persistent block storage volumes for virtual machines.
5	Compute	Nova	A service that manages and provisions virtual machines running on hypervisor nodes.
6	Image	Glance	A registry service for storing resources such as virtual machine images and volume snapshots.
7	Object Storage	Swift	A service providing object storage which allows users to store and retrieve files (arbitrary data).
8	Telemetry	Ceilometer	A service providing measurements of cloud resources.
9	Orchestration	Heat	A service providing a template-based orchestration engine, which supports the automatic creation of resource stacks.

Each OpenStack service is comprised of a collection of Linux services, MariaDB databases, or other components, which together provide a functional group. For example, the **`glance-api`** and **`glance-registry`** Linux services, together with a MariaDB database, implement the Image service.

1.2. Deployment Tools and Methods

Deploying OpenStack: Learning Environments (Manual Setup) contains a series of chapters that describe how to manually deploy each OpenStack component. However, the settings provided by these chapters are not necessarily suitable for a production environment.

Deploying OpenStack: Proof-of-Concept Environments (PackStack) describes how to quickly deploy a demonstration version of Red Hat Enterprise Linux OpenStack Platform using PackStack. This installation tool is ideal for installing a proof-of-concept OpenStack deployment, typically for evaluation purposes.

1.3. Supported Virtual Machine Operating Systems

All guest operating systems that are certified with KVM in Red Hat Enterprise Linux 6 and Red Hat Enterprise Linux 7 are supported by RHEL OpenStack Platform 5. A detailed list of the supported guest operating systems can be found here: Certified Guest Operating Systems in Red Hat Enterprise Linux OpenStack Platform and Red Hat Enterprise Virtualization.

1.4. Service Details

1.4.1. Dashboard Service Overview

The Dashboard service provides a graphical user interface for end users and administrators, allowing operations such as creating and launching instances, managing networking, and setting access controls. Its modular design allows interfacing with other products such as billing, monitoring, and additional management tools. The service provides three basic dashboards: user, system, and settings.

The following screen capture displays a user's dashboard after OpenStack is first installed:

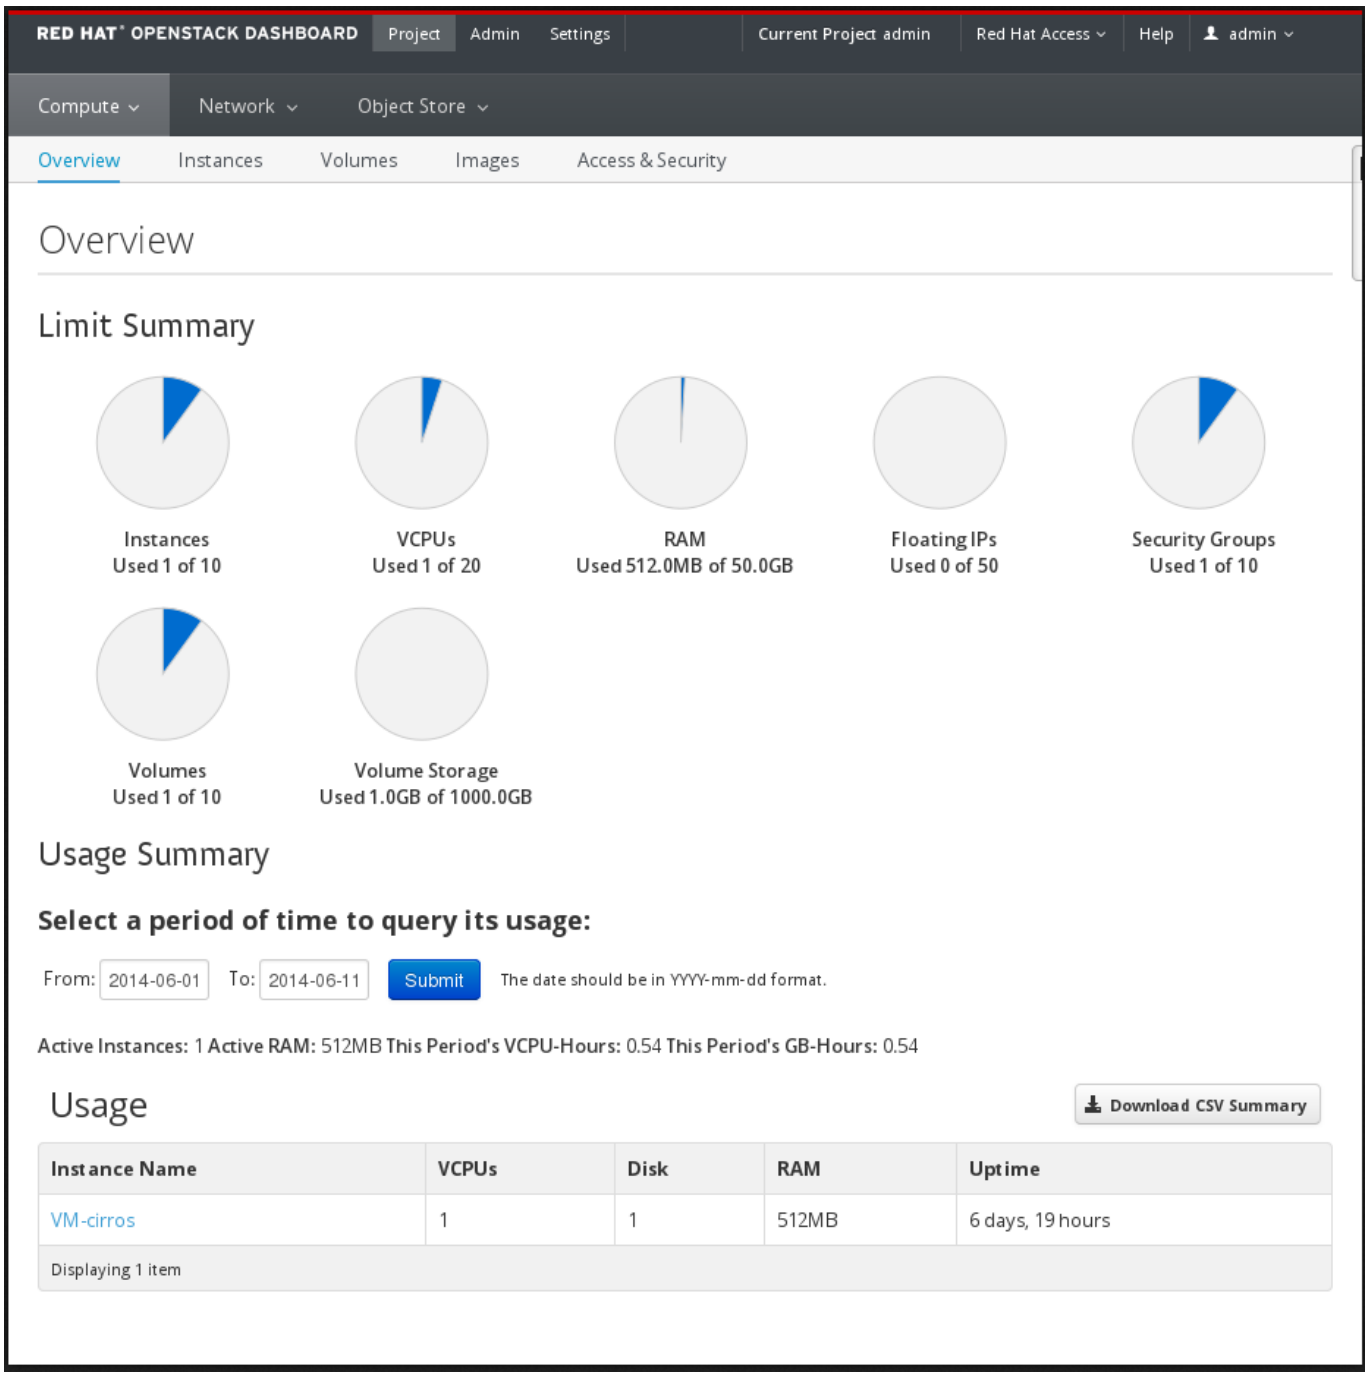

Figure 1.2. User Dashboard

The identity of the logged-in user determines the dashboards and panels that are visible in the dashboard.

Table 1.2. Dashboard service components

Component	Description
openstack-dashboard	A Django (Python) web application, provides access to the dashboard using any web browser.
An Apache HTTP server (**httpd** service)	Hosts the application.

The following diagram provides an overview of the dashboard architecture, where the dashboard service interacts with the OpenStack Identity service for authentication and authorization, the session backend for database services, the httpd service for hosting the application and all the other OpenStack services for API calls.

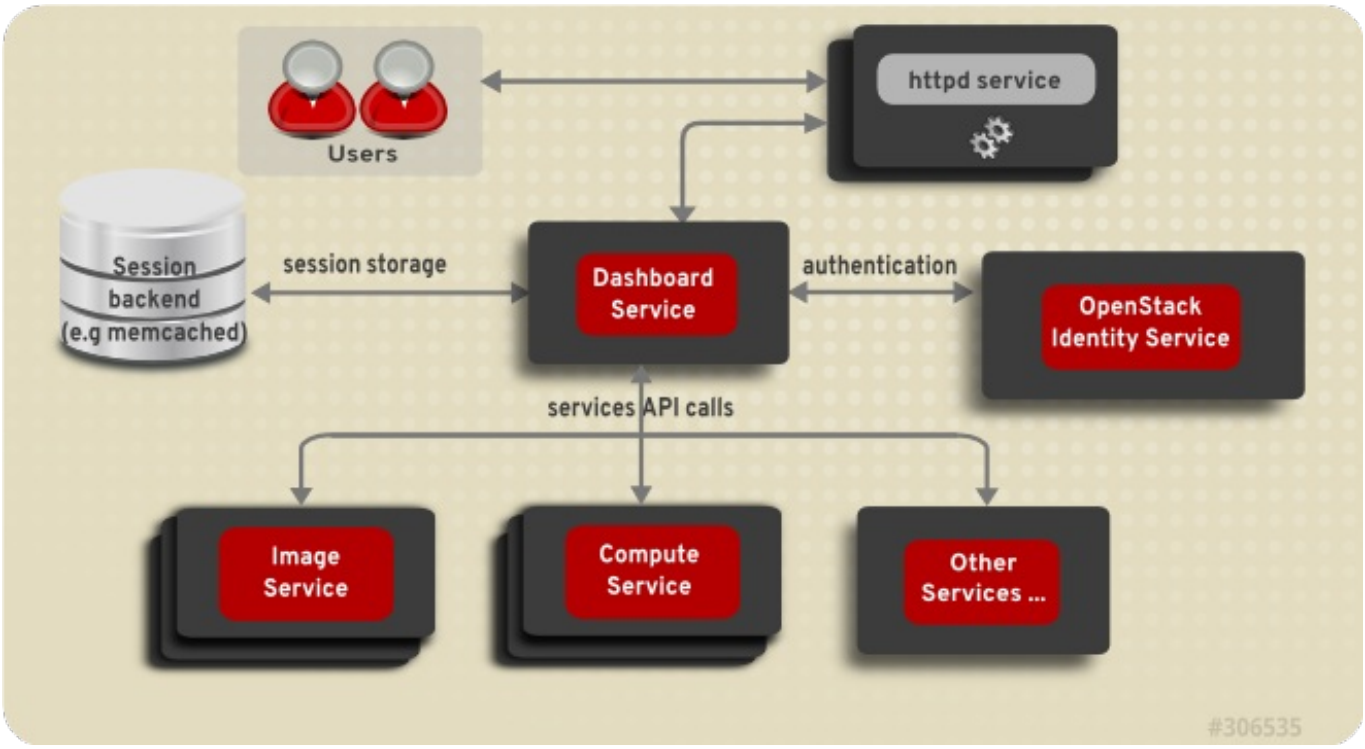

Figure 1.3. OpenStack Dashboard Architecture

1.4.2. Identity Service Overview

The Identity service authenticates and authorizes OpenStack users; the service is used by all OpenStack components. The service supports multiple forms of authentication including user name and password credentials, token-based systems, and AWS-style logins (Amazon Web Services).

The Identity service also provides a central catalog of services and endpoints running in a particular OpenStack cloud, which acts as a service directory for other OpenStack systems. OpenStack services use the following endpoints:

- **`adminURL`**, the URL for the administrative endpoint for the service. Only the Identity service might use a value here that is different from publicURL; all other services will use the same value.
- **`internalURL`**, the URL of an internal-facing endpoint for the service (typically the same as the **`publicURL`**).
- **`publicURL`**, the URL of the public-facing endpoint for the service.
- **`region`**, in which the service is located. By default, if a region is not specified, the 'RegionOne' location is used.

The Identity service uses the following concepts:

- Users, which have associated information (such as a name and password). In addition to custom users, a user must be defined for each cataloged service (for example, the 'glance' user for the Image service).
- Tenants, which are generally the user's group, project, or organization.

- Roles, which determine a user's permissions.

Table 1.3. Identity Service components

Component	Description
keystone	Provides the administrative and public APIs.
Databases	For each of the internal services.

1.4.3. OpenStack Networking Service Overview

The OpenStack Networking service handles the creation and management of a virtual networking infrastructure in the OpenStack cloud. Elements include networks, subnets, and routers; advanced services such as firewalls or virtual private networks (VPN) can also be used.

Because OpenStack Networking is software-defined, it can easily and quickly react to changing network needs (for example, creating and assigning new IP addresses). Advantages include:

- Users can create networks, control traffic, and connect servers and devices to one or more networks.
- OpenStack offers flexible networking models, so that administrators can change the networking model to adapt to their volume and tenancy.
- IPs can be dedicated or floating; floating IPs allow dynamic traffic rerouting.

Table 1.4. OpenStack Networking Service components

Component	Description
neutron-server	A Python daemon, which manages user requests (and exposes the API). It is configured with a plug-in that implements the OpenStack Networking API operations using a specific set of networking mechanisms. A wide choice of plug-ins are also available. For example, the **`openvswitch`** and **`linuxbridge`** plug-ins use native Linux networking mechanisms, while other plug-ins interface with external devices or SDN controllers.
neutron-l3-agent	An agent providing L3/NAT forwarding.
neutron-*-agent	A plug-in agent that runs on each node to perform local networking configuration for the node's virtual machines and networking services.
neutron-dhcp-agent	An agent providing DHCP services to tenant networks.
RabbitMQ server (**`rabbitmq-server`**)	Provides the AMQP message queue. This server (also used by Block Storage) handles the OpenStack transaction management, including queuing, distribution, security, management, clustering, and federation. Messaging becomes especially important when an OpenStack deployment is scaled and its services are running on multiple machines.
Database	Provides persistent storage.

1.4.4. Block Storage Service Overview

The Block Storage (or volume) service provides persistent block storage management for virtual hard drives. Block Storage allows the user to create and delete block devices, and to manage the attachment of block devices to servers. The actual attachment and detachment of devices is handled through integration with the Compute service. Both regions and zones can be used to handle

distributed block storage hosts (for details, see the Section 1.4.7, “Object Storage Service Overview”).

Block storage is appropriate for performance-sensitive scenarios such as database storage, expandable file systems, or providing a server with access to raw block-level storage. Additionally, snapshots can be taken to either restore data or to create new block storage volumes (snapshots are dependent upon driver support).

Basic operations include:

- Create, list, and delete volumes.
- Create, list, and delete snapshots.
- Attach and detach volumes to running virtual machines.

Table 1.5. Block Storage Service components

Component	Description
openstack-cinder-volume	Carves out storage for virtual machines on demand. A number of drivers are included for interaction with storage providers.
openstack-cinder-api	Responds to and handles requests, and places them in the message queue.
openstack-cinder-backup	Provides the ability to back up a Block Storage volume to an external storage repository.
openstack-cinder-scheduler	Assigns tasks to the queue and determines the provisioning volume server.
Database	Provides state information.
RabbitMQ server (`rabbitmq-server`)	Provides the AMQP message queue. This server (also used by Block Storage) handles the OpenStack transaction management, including queuing, distribution, security, management, clustering, and federation. Messaging becomes especially important when an OpenStack deployment is scaled and its services are running on multiple machines.

1.4.5. Compute Service Overview

The Compute service is the heart of the OpenStack cloud by providing virtual machines on demand. Compute schedules virtual machines to run on a set of nodes by defining drivers that interact with underlying virtualization mechanisms, and exposing the functionality to the other OpenStack components.

Compute interacts with the Identity service for authentication, Image service for images, and the Dashboard service for the user and administrative interface. Access to images is limited by project and by user; quotas are limited per project (for example, the number of instances). The Compute service is designed to scale horizontally on standard hardware, and can download images to launch instances as required.

Table 1.6. Ways to Segregate the Cloud

Concept	Description
Regions	Each service cataloged in the Identity service is identified by its region, which typically represents a geographical location, and its endpoint. In a cloud with multiple Compute deployments, regions allow for the discrete separation of services, and are a robust way to share some infrastructure between Compute installations, while allowing for a high degree of failure tolerance.

Concept	Description
Cells	A cloud's Compute hosts can be partitioned into groups called cells (to handle large deployments or geographically separate installations). Cells are configured in a tree. The top-level cell ('API cell') runs the **nova-api** service, but no **nova-compute** services. In contrast, each child cell runs all of the other typical **nova-*** services found in a regular installation, except for the **nova-api** service. Each cell has its own message queue and database service, and also runs **nova-cells**, which manages the communication between the API cell and its child cells. This means that: - A single API server can be used to control access to multiple Compute installations. - A second level of scheduling at the cell level is available (versus host scheduling), which provides greater flexibility over the control of where virtual machines are run.
Host Aggregates and Availability Zones	A single Compute deployment can be partitioned into logical groups (for example, into multiple groups of hosts that share common resources like storage and network, or which have a special property such as trusted computing hardware). If the user is: - An administrator, the group is presented as a Host Aggregate, which has assigned Compute hosts and associated metadata. An aggregate's metadata is commonly used to provide information for use with **nova-scheduler** (for example, limiting specific flavors or images to a subset of hosts). - A user, the group is presented as an Availability Zone. The user cannot view the group's metadata, nor which hosts make up the zone. Aggregates, or zones, can be used to: - Handle load balancing and instance distribution. - Provide some form of physical isolation and redundancy from other zones (such as by using a separate power supply or network equipment). - Identify a set of servers that have some common attribute. - Separate out different classes of hardware.

Table 1.7. Compute Service components

Component	Description
openstack-nova-api	Handles requests and provides access to the Compute services (such as booting an instance).
openstack-nova-cert	Provides the certificate manager.
openstack-nova-compute	Creates and terminates virtual instances. Interacts with the Hypervisor to bring up new instances, and ensures that the state is maintained in the Compute database.
openstack-nova-conductor	Provides database-access support for Compute nodes (thereby reducing security risks).

Component	Description
openstack-nova-consoleauth	Handles console authentication.
openstack-nova-network	Handles Compute network traffic (both private and public access). Handles such tasks as assigning an IP address to a new virtual instance, and implementing security group rules.
openstack-nova-novncproxy	Provides a VNC proxy for browsers (enabling VNC consoles to access virtual machines).
openstack-nova-scheduler	Dispatches requests for new virtual machines to the correct node.
RabbitMQ server (`rabbitmq-server`)	Provides the AMQP message queue. This server (also used by Block Storage) handles the OpenStack transaction management, including queuing, distribution, security, management, clustering, and federation. Messaging becomes especially important when an OpenStack deployment is scaled and its services are running on multiple machines.
libvirtd	The driver for the hypervisor. Enables the creation of virtual machines.
KVM Linux hypervisor	Creates virtual machines and enables their live migration from node to node.
Database	Provides build-time and run-time infrastructure state.

1.4.6. Image Service Overview

The Image service acts as a registry for virtual disk images. Users can add new images or take a snapshot (copy) of an existing server for immediate storage. Snapshots can be used as back up or as templates for new servers. Registered images can be stored in the Object Storage service, as well as in other locations (for example, in simple file systems or external web servers).

The following image formats are supported:

- raw (unstructured format)
- aki/ami/ari (Amazon kernel, ramdisk, or machine image)
- iso (archive format for optical discs; for example, CD)
- qcow2 (Qemu/KVM, supports *Copy on Write*)
- vhd (Hyper-V, common for virtual machine monitors from VMware, Xen, Microsoft, VirtualBox, and others)
- vdi (Qemu/VirtualBox)
- vmdk (VMware)

Container formats can also be used by the Image service; the format determines the type of metadata stored in the image about the actual virtual machine. The following formats are supported.

- bare (no metadata is included)
- ovf (OVF format)
- aki/ami/ari (Amazon kernel, ramdisk, or machine image)

Table 1.8. Image Service components

Component	Description
openstack-glance-api	Handles requests and image delivery (interacts with storage back-ends for retrieval and storage). Uses the registry to retrieve image information (the registry service is never, and should never be, accessed directly).
openstack-glance-registry	Manages all metadata associated with each image.
Database	Stores image metadata.
RabbitMQ server (`rabbitmq-server`)	Provides the AMQP message queue. This server (also used by Block Storage) handles the OpenStack transaction management, including queuing, distribution, security, management, clustering, and federation. Messaging becomes especially important when an OpenStack deployment is scaled and its services are running on multiple machines.

1.4.7. Object Storage Service Overview

The Object Storage service provides a storage system for large amounts of data, accessible through HTTP. Static entities such as videos, images, emails, files, or VM images can all be stored. Objects are stored as binaries on the underlying file system (using metadata stored in the file's extended attributes, xattrs). The service's distributed architecture supports horizontal scaling; redundancy as failure-proofing is provided through software-based data replication.

Because the service supports asynchronous eventual consistency replication, it is well suited to multiple data-center deployment. Object Storage uses the concept of:

- Storage replicas, which are used to maintain the state of objects in the case of outage. A minimum of three replicas is recommended.
- Storage zones, which are used to host replicas. Zones ensure that each replica of a given object can be stored separately. A zone might represent an individual disk drive or array, a server, all the servers in a rack, or even an entire data center.
- Storage regions, which are essentially a group of zones sharing a location. Regions can be, for example, servers or server farms usually located in the same geographical area. Regions have a separate API endpoint per Object Storage service installation, which allows for a discrete separation of services.

The Object Storage service relies on other OpenStack services and components. For example, the Identity Service (keystone), the rsync daemon, and a load balancer are all required.

Table 1.9. Object Storage Service components

Component	Description
openstack-swift-proxy	Exposes the public API, provides authentication, and is responsible for handling requests and routing them accordingly. Objects are streamed through the proxy server to the user (not spooled).
openstack-swift-object	Stores, retrieves, and deletes objects.
openstack-swift-account	Responsible for listings of containers, using the account database.
openstack-swift-container	Handles listings of objects (what objects are in a specific container), using the container database.

Component	Description
Ring files	Contain details of all the storage devices, and are used to deduce where a particular piece of data is stored (maps the names of stored entities to their physical location). One file is created for each object, account, and container server.
Account database	Stores account data.
Container database	Stores container data.
ext4 or XFS file system	Used for object storage.
Housekeeping processes	Replication, auditing, and updating processes.

1.4.8. Telemetry Service Overview

The Telemetry service provides user-level usage data for OpenStack-based clouds, which can be used for customer billing, system monitoring, or alerts. Data can be collected by notifications sent by existing OpenStack components (for example, usage events emitted from Compute) or by polling the infrastructure (for example, libvirt).

Telemetry includes a storage daemon that communicates with authenticated agents through a trusted messaging system, to collect and aggregate data. Additionally, the service uses a plug-in system, which makes it easy to add new monitors.

Table 1.10. Telemetry service components

Component	Description
ceilometer-agent-compute	An agent that runs on each Compute node to poll for resource utilization statistics.
ceilometer-agent-central	An agent that runs on a central management server to poll for utilization statistics about resources not tied to instances or Compute nodes.
ceilometer-collector	An agent that runs on one or more central management servers to monitor the message queues. Notification messages are processed and turned into Telemetry messages, and sent back out on to the message bus using the appropriate topic. Telemetry messages are written to the data store without modification.
ceilometer-notification	An agent that pushes metrics to the collector service from various OpenStack services.
MongoDB database	For collected usage data from collector agents. Only the collector agents and the API server have access to the database.
API Server	Runs on one or more central management servers to provide access to data in the database.
RabbitMQ server (**`rabbitmq-server`**)	Provides the AMQP message queue. This server (also used by Block Storage) handles the OpenStack transaction management, including queuing, distribution, security, management, clustering, and federation. Messaging becomes especially important when an OpenStack deployment is scaled and its services are running on multiple machines.

Telemetry service dependencies

- Each **`nova-compute`** node must have a **`ceilometer-compute`** agent deployed and running.
- All nodes running the **`ceilometer-api`** service must have firewall rules granting appropriate access.

- The **ceilometer-central-agent** cannot currently be horizontally scaled, so only a single instance of this service should be running at any given moment.
- You can choose where to locate the additional Telemetry agents, as all intra-agent communication is either based on AMQP or REST calls to the **ceilometer-api** service; as is the case for the **ceilometer-alarm-evaluator** service.

1.4.9. Orchestration Service Overview

The Orchestration service provides a template-based way to create and manage cloud resources such as storage, networking, instances, or applications.

Templates are used to create stacks, which are collections of resources (for example instances, floating IPs, volumes, security groups, or users). The service offers access to all OpenStack core services using a single modular template, with additional orchestration capabilities such as auto-scaling and basic high availability.

Features include:

- A single template provides access to all underlying service APIs.
- Templates are modular (resource oriented).
- Templates can be recursively defined, and therefore reusable (nested stacks). This means that the cloud infrastructure can be defined and reused in a modular way.
- Resource implementation is pluggable, which allows for custom resources.
- Autoscaling functionality (automatically adding or removing resources depending upon usage).
- Basic high availability functionality.

Table 1.11. Orchestration service components

Component	Description
heat	A CLI tool that communicates with the heat-api to execute AWS CloudFormation APIs.
heat-api	An OpenStack-native REST API that processes API requests by sending them to the heat-engine over RPC.
heat-api-cfn	Provides an AWS-Query API that is compatible with AWS CloudFormation and processes API requests by sending them to the heat-engine over RPC.
heat-engine	Orchestrates the launching of templates and provide events back to the API consumer.
heat-api-cloudwatch	Provides monitoring (metrics collection) for the Orchestration service.
heat-cfntools	A package of helper scripts (for example, cfn-hup, which handles updates to metadata and executes custom hooks).

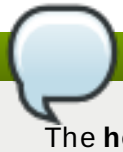

Note

The **heat-cfntools** package is only installed on images that are launched by heat into Compute servers.

Chapter 2. Requirements

This chapter outlines the main requirements for setting up an environment in which to provision a Red Hat Enterprise Linux OpenStack Platform environment using the Red Hat Enterprise Linux OpenStack Platform installer, including the requirements for setting up and accessing the installer itself, and the hardware requirements for hosts on to which the installer provisions the environment.

2.1. User Interface Environment Requirements

A typical deployment of the Red Hat Enterprise Linux OpenStack Platform installer requires:

- A private network accessible by physical machines or virtual machines on which Red Hat Enterprise Linux OpenStack Platform components can be deployed. Services such as DHCP, DNS, and PXE must be disabled on this network because these services can interfere with the installer.
- A machine running Red Hat Enterprise Linux on which to set up the installer. This guide outlines how to install the user interface on this machine. For this machine, 6 GB of memory is recommended; a minimum of 4 GB memory is required.

Important

The machine on which you set up the installer must have a fully qualified domain name that satisfies the following requirements:

- Matches the domain of the network to be provisioned.
- Does not conflict with any existing domains (to prevent resource conflicts).

Important

The installer can only run on Red Hat Enterprise Linux 6.6. To deploy Red Hat Enterprise Linux OpenStack Platform 5.0 on Red Hat Enterprise Linux 7.2 using the installer, you must first install the user interface for the installer on a Red Hat Enterprise Linux 6.6 system. From that system, you can then use the tool to deploy Red Hat Enterprise Linux OpenStack Platform 5.0 on Red Hat Enterprise Linux 7.2.

- A machine that is a member of the private network and that also has access to external networks that can act as a router or gateway. The machine on which the user interface is installed can perform this function if required; see Appendix B, *Configuring a Gateway* for information on how to configure this machine to act as a gateway.
- The details of a Customer Portal account for subscribing the machine on which to install the user interface and for subscribing all hosts in your RHEL OpenStack Platform environment, including the user name, password, the ID of pools to attach, the names of channels to enable, and the details of a HTTP proxy, if any.

2.2. User Interface Client Requirements

The user interface for the installer is accessed as a web page hosted on the same machine as the installer. You must use a supported browser to access this user interface. Browser support is divided into four levels:

- Level 1: Fully supported, preferred browsers for ideal experience.
- Level 2: Mostly supported. The user interface functions, but some design elements may not align correctly, the layout and user interface controls may not align correctly, and performance may be degraded.
- Level 3: Design elements may not align correctly.
- Level 4: Unsupported.

The following table outlines the supported browsers and their level of support:

Table 2.1. Supported Browsers

Browser	Version	Support Level
Firefox	2.6	L3
Firefox	17, 18, 19, 20	L4
Firefox	21	L2
Firefox	22, 23, 24	L1
Firefox	Latest	L1
Chrome	19, 20	L4
Chrome	21, 27	L2
Chrome	Latest	L1
Internet Explorer	7, 8	L4
Internet Explorer	9, 10, 11	L2
Safari	All	L4

2.3. Host Requirements

The following sections outline the main hardware requirements for hosts on which the installer provisions RHEL OpenStack Platform. These requirements are categorized in accordance with the three core roles that hosts perform in a RHEL OpenStack Platform environment. For more information on the these roles and the number of hosts that must be assigned to reach role to provision an environment using the installer, see Section 6.2, “Assigning Hosts to Deployment Roles”.

2.3.1. Controller Node Requirements

Controller nodes are responsible for hosting the core services in a RHEL OpenStack Platform environment, such as the Horizon dashboard, the back-end database server, and Keystone.

Processor

64-bit x86 processor with support for the Intel 64 or AMD64 CPU extensions, and the AMD-V or Intel VT hardware virtualization extensions enabled.

Memory

A minimum of 2 GB of RAM is recommended.

Disk Space

A minimum of 50 GB of available disk space is recommended.

Add additional disk space to this requirement based on the amount of space that you intend to make available to virtual machine instances. This figure varies based on both the size of each disk image you intend to create and whether you intend to share one or more disk images between multiple instances.

1 TB of disk space is recommended for a realistic environment capable of hosting multiple instances of varying sizes.

Network Interface Cards

2 x 1 Gbps Network Interface Cards.

2.3.2. Compute Node Requirements

Compute nodes are responsible for running virtual machine instances after they are launched. Compute nodes must support hardware virtualization. Compute nodes must also have enough memory and disk space to support the requirements of the virtual machine instances they host.

Processor

64-bit x86 processor with support for the Intel 64 or AMD64 CPU extensions, and the AMD-V or Intel VT hardware virtualization extensions enabled.

Memory

A minimum of 2 GB of RAM is recommended.

Add additional RAM to this requirement based on the amount of memory that you intend to make available to virtual machine instances.

Disk Space

A minimum of 50 GB of available disk space is recommended.

Add additional disk space to this requirement based on the amount of space that you intend to make available to virtual machine instances. This figure varies based on both the size of each disk image you intend to create and whether you intend to share one or more disk images between multiple instances.

1 TB of disk space is recommended for a realistic environment capable of hosting multiple instances of varying sizes.

Network Interface Cards

2 x 1 Gbps Network Interface Cards.

2.3.3. Network Node Requirements

Network nodes are responsible for hosting the services that provide networking functionality to compute instances. In particular, they host the DHCP agent, layer 3 agent, and metadata proxy services. Like all systems that handle networking in an OpenStack environment, they also host an instance of the layer 2 agent.

The hardware requirements of network nodes vary widely depending on the networking workload of the environment. The requirements listed here are intended as a guide to the minimum requirements of a network node.

Processor

No specific CPU requirements are imposed by the networking services.

Memory

A minimum of 2 GB of RAM is recommended.

Disk Space

A minimum of 10 GB of available disk space is recommended.

No additional disk space is required by the networking services other than that required to install the packages themselves. However, some disk space must be available to store log files and temporary files.

Network Interface Cards

2 x 1 Gbps Network Interface Cards.

2.3.4. Block Storage Node Requirements

Block storage nodes are nodes that host the volume service (**`openstack-cinder-volume`**) and provide volumes for use by virtual machine instances or other cloud users. The block storage API (**`openstack-cinder-api`**) and scheduling services (**`openstack-cinder-scheduler`**) can run on the same nodes as the volume service, or on separate nodes. In either case, the primary hardware requirement of the block storage nodes is that there is enough block storage available to serve the needs of the environment.

The amount of block storage required in an environment varies in accordance with the following factors:

- The number of volumes that will be created in the environment.
- The average size of the volumes that will be created in the environment.
- Whether or not the storage back end will be configured to support redundancy.
- Whether or not the storage back end will be configured to create sparse volumes by default.

Use the following formula to assist with estimating the initial block storage needs of your environment:

```
VOLUMES * SIZE * REDUNDANCY * UTILIZATION = TOTAL
```

- Replace *VOLUMES* with the number of volumes that are expected to exist in the environment at any one time.
- Replace *SIZE* with the expected average size of the volumes (in gigabytes) that will exist in the environment at any one time.
- Replace *REDUNDANCY* with the expected number of redundant copies of each volume that the back end storage will be configured to store. Use **1**, or skip this multiplication operation if no redundancy is required.
- Replace *UTILIZATION* with the expected percentage of each volume that will actually be used. Use **1**, indicating 100%, if the use of sparse volumes will not be enabled.

The resultant figure represents an **estimate** of the block storage needs of your environment in gigabytes. It is recommended that some additional space is allowed for future growth. Addition of further block storage after the environment has been deployed is facilitated by adding more block storage providers and, if necessary, additional instances of the volume service.

Chapter 3. Installing the User Interface

Setting up an enterprise Red Hat Enterprise Linux OpenStack Platform environment involves provisioning that environment using the Red Hat Enterprise Linux OpenStack Platform installer. The installer is a graphical user interface that provides functions for managing the provisioning of Red Hat Enterprise Linux OpenStack Platform components on a set of machines.

There are two main methods for setting up a machine on which to run the installer: manually installing Red Hat Enterprise Linux 6.6 on a machine and configuring the required repositories, or by using the installer live CD. The key difference between these methods is in how the base operating system and packages are installed; all steps from installing the Red Hat Enterprise Linux OpenStack Platform installer are identical for both methods. This guide outlines how to manually configure the required channels and repositories on a machine where Red Hat Enterprise Linux 6.6 is installed. For information on how to use the installer live CD, see How to Install the Red Hat Enterprise Linux OpenStack Platform Installer Using the Live CD.

After you have set up this machine, there are three key steps to the provisioning process: installing the user interface for the Red Hat Enterprise Linux OpenStack Platform installer, adding hosts to the user interface onto which to provision the environment, and then provisioning the environment. This guide outlines these steps and provides additional information on how to manually configure aspects of the installer.

Warning

Upgrading the Red Hat Enterprise Linux OpenStack Platform installer is not currently supported. To use the latest version of the installer, you must use a fresh installation.

Important

The role of the installer is to provision Red Hat Enterprise Linux OpenStack Platform environments and manage the life cycle of the hosts in that environment. After you have provisioned your environment, all additional configuration of Red Hat Enterprise Linux OpenStack Platform components on hosts you have provisioned must be performed manually on those hosts. For more information, see Section 6.6.4, “Next Steps”.

3.1. Subscribing to the Required Channels Using Subscription Manager

To install the RHEL OpenStack Platform installer user interface, you must register the system where the user interface will be installed with Red Hat Subscription Manager, and subscribe to the required channels.

Procedure 3.1. Subscribing to the Required Channels Using Subscription Manager

1. Register your system with the Customer Delivery Network, entering your Customer Portal user name and password when prompted:

```
# subscription-manager register
```

2. Find entitlement pools containing the channels required to install the Red Hat Enterprise Linux OpenStack Platform Installer.

```
# subscription-manager list --available | grep -A8 "Red Hat
Enterprise Linux Server"
# subscription-manager list --available | grep -A8 "Red Hat
Enterprise Linux OpenStack Platform"
```

3. Use the pool identifiers located in the previous step to attach the **Red Hat Enterprise Linux 6 Server** and **Red Hat Enterprise Linux OpenStack Platform** entitlements:

```
# subscription-manager attach --pool=pool_id
```

4. Enable the required channels:

```
# subscription-manager repos --enable=rhel-6-server-rpms
# subscription-manager repos --enable=rhel-6-server-openstack-
foreman-rpms
# subscription-manager repos --enable=rhel-server-rhscl-6-rpms
```

3.2. Preparing an Installation Medium

Before installing the user interface, you must prepare an installation medium that you can then specify during the installation process. An installation medium is a source of files that the installer uses to install the base operating system on a machine when you provision RHEL OpenStack Platform, and must be in the format of a Red Hat Enterprise Linux 7.2 installation tree.

3.2.1. Preparing an Installation Medium using a Web Server

Use a web server to host a Red Hat Enterprise Linux 7 installation medium. This procedure must be performed on the machine where the installation medium is to be hosted.

Procedure 3.2. Preparing an Installation Medium

1. Go to https://access.redhat.com, and log in to the Red Hat Customer Portal using your customer account details.
2. Click **Downloads** in the menu bar.
3. Click **Red Hat Enterprise Linux** to access the product download page.
4. Click **RHEL 7.2 Binary DVD**.
5. Install the Apache web server on the machine on which to host the installation medium:

```
# yum install httpd
```

6. Start the **httpd** service, and ensure it starts on boot:

   A. On Red Hat Enterprise Linux 6:

```
# service httpd start
# chkconfig httpd on
```

   B. On Red Hat Enterprise Linux 7:

```
# systemctl start httpd.service
```

7. In the root web server directory, create a directory in which to store the files for the installation medium:

```
# mkdir /var/www/html/[directory_name]
```

8. Create a temporary directory into which to mount the ISO file:

```
# mkdir /RHEL7
```

9. Mount the ISO file in the temporary directory:

```
# mount -t iso9660 -o loop rhel-server-7.2-x86_64-dvd.iso /RHEL7
```

10. Copy the contents of the temporary directory to the directory in which to store the files for the installation medium:

```
# cp -dpR /RHEL7/* /var/www/html/[directory_name]
```

11. Unmount the ISO file:

```
# umount /RHEL7
```

12. Remove the temporary directory in which you mounted the ISO file:

```
# rmdir /RHEL7
```

3.2.2. Preparing an Installation Medium using NFS

Use an NFS share to host a Red Hat Enterprise Linux 7 installation medium. This procedure must be performed on the machine where the installation medium is to be hosted.

Procedure 3.3. Preparing an Installation Medium using NFS

1. Go to https://access.redhat.com, and log in to the Red Hat Customer Portal using your customer account details.
2. Click **Downloads** in the menu bar.
3. Click **Red Hat Enterprise Linux** to access the product download page.
4. Click **RHEL 7.2 Binary DVD**.
5. Install the *nfs-utils* package:

```
# yum install nfs-utils
```

6. Create a directory to host the installation medium:

```
# mkdir /[directory_name]
```

7. Add the newly created directory to the **/etc/exports** file:

```
/[directory_name]      *(rw)
```

8. Export the directory:

```
# exportfs -r
```

9. Start the **nfs** and **rpcbind** services, and ensure they start on boot:

   A. On Red Hat Enterprise Linux 6:

```
# service nfs start
# service rpcbind start
# chkconfig nfs on
# chkconfig rpcbind on
```

   B. On Red Hat Enterprise Linux 7:

```
# systemctl start nfs.service
# systemctl start rpcbind.service
# systemctl enable nfs.service
# systemctl enable rpcbind.service
```

10. Create a temporary directory into which to mount the ISO file:

```
# mkdir /RHEL7
```

11. Mount the ISO file in the temporary directory:

```
# mount -t iso9660 -o loop rhel-server-7.2-x86_64-dvd.iso /RHEL7
```

12. Copy the contents of the temporary directory to the directory in which to store the files for the installation medium:

```
# cp -dpR /RHEL7/* /[directory_name]
```

13. Unmount the ISO file:

```
# umount /RHEL7
```

14. Remove the temporary directory in which you mounted the ISO file:

```
# rmdir /RHEL7
```

3.3. Installing the User Interface

Run the **rhel-osp-installer** command to install the RHEL OpenStack Platform user interface and configure the core parameters that the installer uses to provision RHEL OpenStack Platform 5.0. See Section 3.4, “Options for Installing the User Interface” for a full list of the options available during the installation process.

Note

The **rhel-osp-installer** command automatically configures the required SELinux permissions and adds the required firewall rules to **iptables** while preserving any existing firewall rules. See Appendix A, *Firewall Rules* for the list of firewall rules that the command configures.

Procedure 3.4. Installing the Red Hat Enterprise Linux OpenStack Platform Installer User Interface

1. Install the *rhel-osp-installer* package:

```
# yum install rhel-osp-installer
```

2. Start the installation:

```
# rhel-osp-installer
```

3. Enter the number for the network interface that the installer will use to provision RHEL OpenStack Platform, and press **Enter**:

```
Please select NIC on which you want Foreman provisioning enabled:
1. eth1
2. eth0
?
```

4. Configure networking options:

   a. Enter the number for the configuration option to change, and press **Enter**.

   b. Enter a new value, and press **Enter**.

   c. When you have specified the preferred value for each configuration option, enter **1**, and press **Enter**.

```
Networking setup:
       Network interface: 'eth1'
              IP address: 'XX.XX.XX.XX'
            Network mask: 'XX.XX.XX.XX'
         Network address: 'XX.XX.XX.XX'
            Host Gateway: 'XX.XX.XX.XX'
        DHCP range start: 'XX.XX.XX.XX'
          DHCP range end: 'XX.XX.XX.XX'
            DHCP Gateway: 'XX.XX.XX.XX'
           DNS forwarder: 'XX.XX.XX.XX'
                  Domain: 'mydomain.example.com'
             Foreman URL: 'https://host.mydomain.example.com'
           NTP sync host: '0.rhel.pool.ntp.org'
                Timezone: 'America/New_York'
Configure networking on this machine: ✓
Configure firewall on this machine: ✓

The installer can configure the networking and firewall rules on
```

```
this machine with the above configuration. Default values are
populated from the this machine's existing networking
configuration.

If you DO NOT want to configure networking please set 'Configure
networking on this machine' to No before proceeding. Do this by
selecting option 'Do not configure networking' from the list below.

How would you like to proceed?:
1. Proceed with the above values
2. Change Network interface
3. Change IP address
4. Change Network mask
5. Change Network address
6. Change Host Gateway
7. Change DHCP range start
8. Change DHCP range end
9. Change DHCP Gateway
10. Change DNS forwarder
11. Change Domain
12. Change Foreman URL
13. Change NTP sync host
14. Change Timezone
15. Do not configure networking
16. Do not configure firewall
17. Cancel Installation
```

Important

The name of the domain must match that of the fully qualified domain name of the machine on which the user interface is being installed.

Important

By default, the address of the machine on which the user interface is being installed is specified as the DHCP gateway, which is the gateway the installer configures on hosts it provisions. If you have not configured the machine on which the user interface is being installed to act as a gateway, you must edit this value and specify the address of a machine that can perform this function.

5. Configure client authentication:

    a. Enter the number for the configuration option to change, and press **Enter**.

    b. Enter a new value, and press **Enter**.

    c. When you have specified either a SSH public key or a root password, enter **1**, and press **Enter**.

```
Configure client authentication
          SSH public key: ''
           Root password: '**********'
```

```
Please set a default root password for newly provisioned machines.
If you choose not to set a password, it will be generated randomly.
The password must be a minimum of 8 characters. You can also set a
public ssh key which will be deployed to newly provisioned
machines.

How would you like to proceed?:
1. Proceed with the above values
2. Change SSH public key
3. Change Root password
4. Toggle Root password visibility
```

6. Specify the details of an installation medium:

    a. Enter **1**, and press **Enter**.

    b. Enter the address of a Red Hat Enterprise Linux Server 7.2 installation tree that the machine on which the user interface is being installed can access via a web server, and press **Enter**.

    c. Enter **2**, and press **Enter**.

```
Now you should configure installation media which will be used for
provisioning.
Note that if you don't configure it properly, host provisioning
won't work until you configure installation media manually.

Enter RHEL repo path:
1. Set RHEL repo path (http or https URL): http://
2. Proceed with configuration
3. Skip this step (provisioning won't work)
```

Important

If you choose not to configure installation media in this step, you will not be able to provision RHEL OpenStack Platform unless you manually configure an installation media entry in the user interface.

7. Specify the details of a Subscription Manager account:

    a. Enter the number for the configuration option to change, and press **Enter**.

    b. Enter a new value, and press **Enter**.

    c. When you have specified the preferred value for each configuration option, enter **9**, and press **Enter**.

```
Enter your subscription manager credentials:
1. Subscription manager username:
2. Subscription manager password:
3. Comma separated repositories:        rhel-7-server-openstack-
5.0-rpms
4. Subscription manager pool (recommended):
```

```
5. Subscription manager proxy hostname:
6. Subscription manager proxy port:
7. Subscription manager proxy username:
8. Subscription manager proxy password:
9. Proceed with configuration
10. Skip this step (provisioning won't subscribe your machines)
```

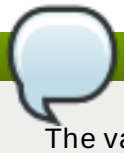

Note

The value for the Subscription Manager pool must be in the format of a Subscription Manager entitlement pool ID. Moreover, you can only specify a single entitlement pool ID. If you leave the value for this configuration item blank, the installer attempts to auto-attach the required entitlements on systems it provisions.

Important

If you choose not to specify your Subscription Manager account details in this step, you will not be able to provision RHEL OpenStack Platform unless you manually configure an operating system entry with the required details in the user interface.

The user interface is installed, and the entries in the user interface required to provision RHEL OpenStack Platform are automatically generated based on the details you provided. After the installation is complete, the user name of and a randomly generated password for the default administrative user account are displayed. The address for accessing the user interface is also displayed.

3.4. Options for Installing the User Interface

The following tables outline the options available when installing the user interface for the installer.

3.4.1. Networking Setup Options

The following table outlines the options available in the networking setup step of installing the user interface.

Table 3.1. Networking Setup Options

Option	Description
Network interface	The network interface on which the provisioning network is to be configured. The default values of other networking options, such as the **IP address** and **Network address**, are automatically specified based on the network interface you select.
IP Address	The IP address of the machine where the user interface is being installed.
Network mask	The network mask of the machine where the user interface is being installed. This value is also used to populate the network mask of the default subnet that the installer configures to act as the provisioning network.

Option	Description
Network address	The network address of the default subnet that the installer configures to act as the provisioning network. This address must be in CIDR format. For example, **XX.XX.XX.0/24**.
Host gateway	The address of a machine that can act as a gateway for the machine where the user interface is being installed.
DHCP range start	The first address in the range of IP addresses that the installer can assign to machines on the provisioning network. This value is also used to populate the DHCP range start of the default subnet that the installer configures to act as the provisioning network.
DHCP range end	The last address in the range of IP addresses that the installer can assign to machines on the provisioning network. This value is also used to populate the DHCP range end of the default subnet that the installer configures to act as the provisioning network.
DHCP Gateway	The address of a machine that can act as a gateway for machines that the installer provisions. By default, the address of the machine where the user interface is being installed is specified as the DHCP gateway. If you have not configured the machine on which the user interface is being installed to act as a gateway, you must edit this value and specify the address of a machine that can perform this function.
DNS forwarder	The address of a machine that can resolve IP addresses and host names. This value is also used to populate the primary DNS of the default subnet that the installer configures to act as the provisioning network.
Domain	The name of the domain that the installer provides. The name of the domain must match that of the fully qualified domain name of the machine where the user interface is being installed.
Foreman URL	The address for accessing the user interface. By default, this value is set to the fully qualified domain name of the machine where the user interface is being installed.
NTP sync host	The address of an NTP server that the installer can use to synchronize the time on machines that it provisions.
Timezone	The timezone that the installer applies to machines that it provisions. Timezones must be in the format of an IANA time zone identifier such as **America/New_York** or **Asia/Tokyo**.

3.4.2. Client Authentication Options

The following table outlines the options available in the client authentication step of installing the user interface.

Table 3.2. Client Authentication Options

Option	Description
SSH public key	A public SSH key that is copied to machines that the installer provisions.
Root password	The password for the **root** user on machines the installer provisions. The password must be at least eight characters in length. If you do not manually specify a password, a random password is generated.

3.4.3. Installation Media Options

The following table outlines the options available in the installation media step of installing the user interface.

Table 3.3. Installation Media Options

Option	Description
Set RHEL repo path (http or https URL)	The address of a Red Hat Enterprise Linux 7.2 installation tree. The path must end in the directory that contains directories such as **isolinux**, **LiveOS**, and **repodata**. This option is used to populate the default installation media entry that the installer uses to install the base operating system on machines that it provisions.

3.4.4. Subscription Manager Options

The following table outlines the options available in the Subscription Manager step of installing the user interface.

Table 3.4. Subscription Manager Options

Option	Description
Subscription manager username	The username of a Customer Portal account that can be used to register machines that the installer provisions with the Content Delivery Network.
Subscription manager password	A password that can be used to authenticate the above account.
Comma separated repositories	A comma-separated list of repositories to enable on machines that the installer provisions. By default, the **rhel-7-server-openstack-5.0-rpms** channel, which provides the packages required to install Red Hat Enterprise Linux OpenStack Platform components, is provided.
Subscription manager pool (recommended)	The ID of an entitlement pool to attach to machines that the installer provisions. Only a single entitlement pool ID can be specified, and if no entitlement pool ID is specified, the installer attempts to auto-attach the required entitlement pool.
Subscription manager proxy hostname	The host name of a machine that can be used as a Subscription Manager proxy.
Subscription manager proxy port	The port by which to connect with the Subscription Manager proxy.
Subscription manager proxy username	A username by which to connect with the Subscription Manager proxy.
Subscription manager proxy password	The password by which the authenticate the above username.

Chapter 4. Configuring the User Interface

The configuration required to provision Red Hat Enterprise Linux OpenStack Platform is performed automatically when you install the user interface, but you can also manually configure aspects of the user interface such as the location of an installation media entry, or the credentials of a Customer Portal account used to register machines with the Customer Delivery Network. The information in this chapter is provided as a reference for updating the existing values specified during the installation process or for performing additional configuration.

4.1. Logging in to the User Interface

Log in to the user interface for the installer.

Note

The first time you log in to the user interface, you must use the default administrative user account with the name **admin** and the password randomly generated during the installation process. If you do not have a copy of this password, you can retrieve it by running the following command:

```
# grep admin_password /etc/foreman/rhel-osp-installer.answers.yaml
```

Procedure 4.1. Logging in to the User Interface

1. In a web browser, navigate to the URL provided on completing installation of the user interface. By default, this URL comprises the fully qualified domain name of the machine on which the user interface is installed. For example, **https://myhost.example.com/**.
2. Enter your username in the **Username** text field, and your password in the **Password** text field.

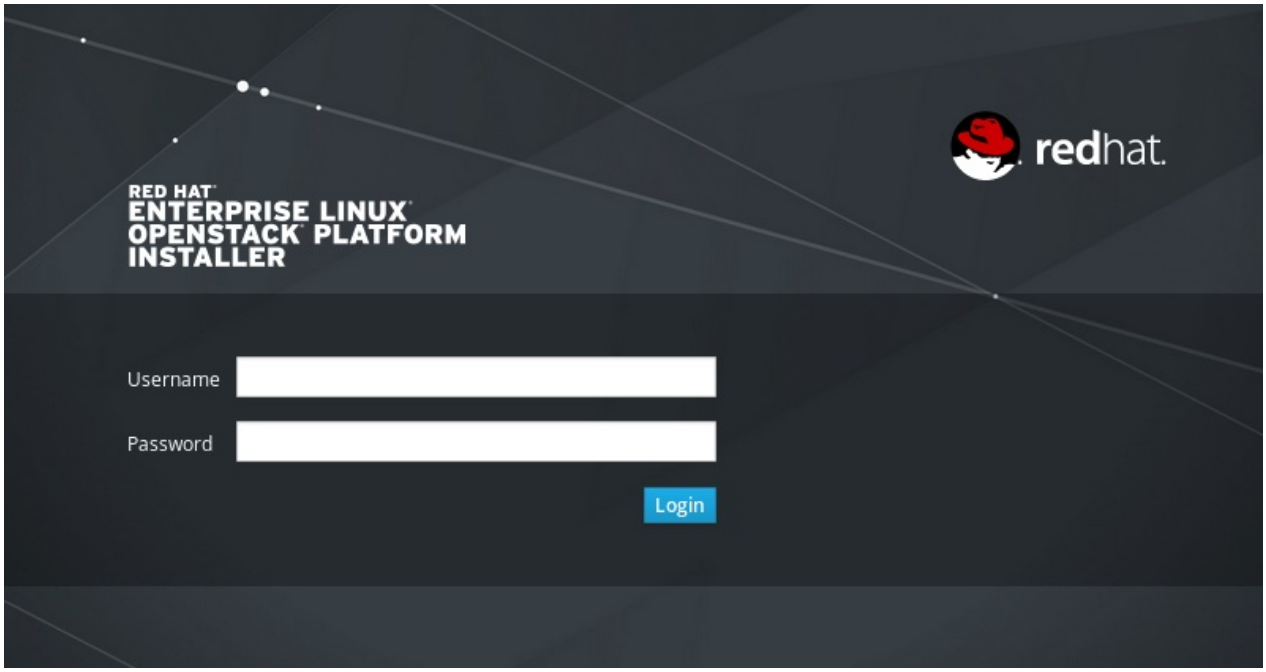

Figure 4.1. The Red Hat Enterprise Linux OpenStack Platform Installer Login Screen

3. Click the **Login** button.

You have logged in to the user interface, and the **Dashboard** page is displayed. If no reports on hosts that the installer manages are available, a welcome page is displayed that includes instructions on editing general settings. No action needs to be taken.

4.2. Changing a Password

You can change the password you use to log in to the user interface.

Procedure 4.2. Changing a Password

1. Click ***User Name*** → **My account** to open the **Edit User** window.

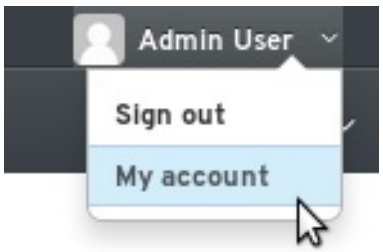

Figure 4.2. Accessing Account Settings

2. Enter a new password for in the **Password** text field and again in the **Verify** text field.
3. Click **Submit**.

4.3. Installation Media

Installation media are sources of files that the installer can use to install the base operating system on a machine when you provision RHEL OpenStack Platform. Installation media must be in the format of an operating system installation tree, and must be accessible to the machine on which the user interface is installed via a URL or an NFS share. The installer accesses installation media via installation media entries in the user interface that define details regarding the installation media such as their location and the operating system family they represent.

Note

For information on how to use a CD or ISO file to prepare an installation tree, see Section 3.2, “Preparing an Installation Medium”.

Note

By default, the **`rhel-osp-installer`** command prompts you to specify the location of an installation medium during the installation process, and creates a default installation media entry based on this location. If you skipped this step when you installed the user interface, you must manually edit the **`RedHat mirror`** installation media entry and specify the location of a Red Hat Enterprise Linux 7 installation medium before you can provision RHEL OpenStack Platform.

4.3.1. Creating an Installation Media Entry

Create an installation media entry in the user interface to represent an installation medium for installing a base operating system on a machine when you provision RHEL OpenStack Platform.

Procedure 4.3. Creating an Installation Media Entry

1. From the title bar in the main screen of the user interface, click **Hosts** → **Installation Media**.

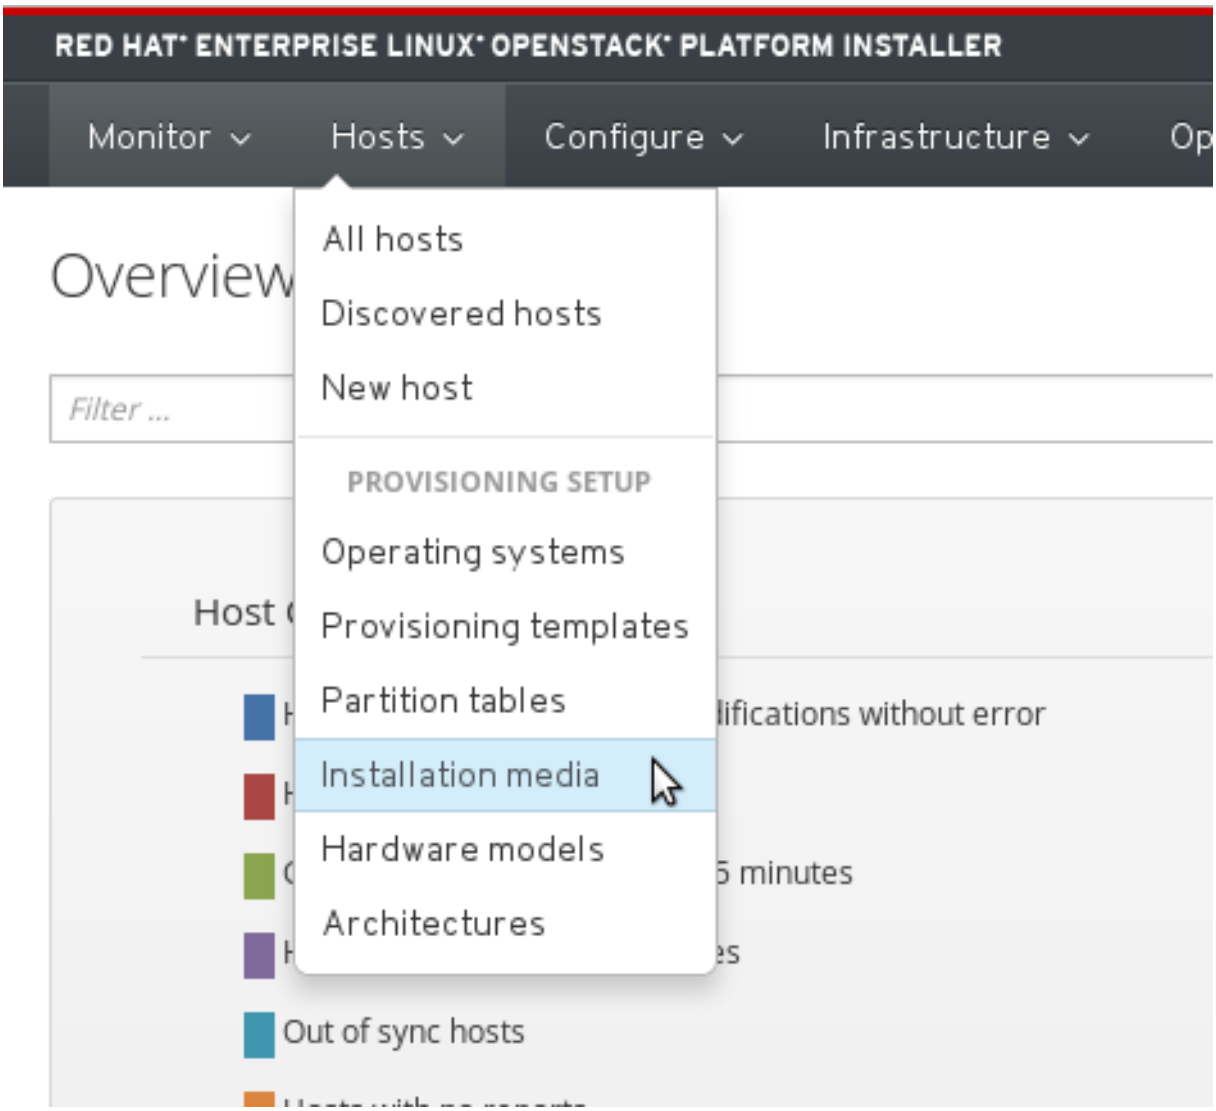

Figure 4.3. The Installation Media Menu Entry

2. Click **New Medium**.
3. In the **Name** text field, enter a name to represent the installation media entry in the user interface.
4. In the **Path** text field, enter the path of a URL or an NFS share containing the installation tree. The following variables can be used in the path to represent multiple different system architectures and versions:

   $arch

   The system architecture, for example **x86_64**.

   $version

   The operating system version, for example **7.2**.

   $major

   The operating system major version, for example **7**.

   $minor

   The operating system minor version, for example **0**.

   Example 4.1. HTTP Path

   http://download.example.com/rhel/$version/Server/$arch/os/

   Example 4.2. NFS Path

   nfs://download.example.com:/rhel/$version/Server/$arch/os/

5. Select **Red Hat** from the **Operating system family** list.
6. Click **Submit**.

4.3.2. Specifying the Installation Media Entry for Provisioning

If you did not configure an installation media entry when you installed the user interface, you must manually specify an installation media entry for Red Hat Enterprise Linux 7 that the installer can use for provisioning. The following procedure assumes you have already created an installation media entry that you can specify.

Procedure 4.4. Specifying the Installation Media Entry for Provisioning

1. From the title bar in the main screen of the user interface, click **Configure → Host groups**.
2. Click the **base_RedHat_7** host group.
3. Click the **Operating System** tab.
4. Select an installation media entry from the **Media** list.
5. Click **Submit**.

You have specified the installation media entry that the installer uses to install the base operating system on hosts when provisioning a RHEL OpenStack Platform environment.

4.3.3. Removing an Installation Media Entry

Remove an existing installation media entry from the user interface if you no longer require that entry or the installation medium it represents is no longer available.

Procedure 4.5. Removing an Installation Media Entry

1. From the title bar in the main screen of the user interface, click **Hosts** → **Installation Media** to open the **Installation Media** page.
2. Click the **Delete** button that corresponds to the installation media entry to delete.
3. Click **OK** when prompted.

Note

Removing an installation media entry only removes the association with the installation medium from the user interface, and does not delete the installation medium itself.

4.4. Operating Systems

An operating system entry is a collection of options that defines the way in which the base operating system is installed on a host when you provision Red Hat Enterprise Linux OpenStack Platform. Operating system entries combine other entries in the user interface, such as installation media for installing the base operating system on a host and partition tables that define the partitioning scheme for the storage available to that host, and allow you to specify parameters specific to the operating system entry, such as the details of a Subscription Manager account.

Note

By default, the **rhel-osp-installer** command creates operating system entries for Red Hat Enterprise Linux 6.6 and Red Hat Enterprise Linux 7.2 that contain the details of the subscription manager account you entered during the installation process. These operating system entries can be edited as necessary.

4.4.1. Creating an Operating System Entry

You can create new operating system entries.

Procedure 4.6. Creating an Operating System Entry

1. Click **Hosts** → **Operating systems**.
2. Click **New Operating system**.

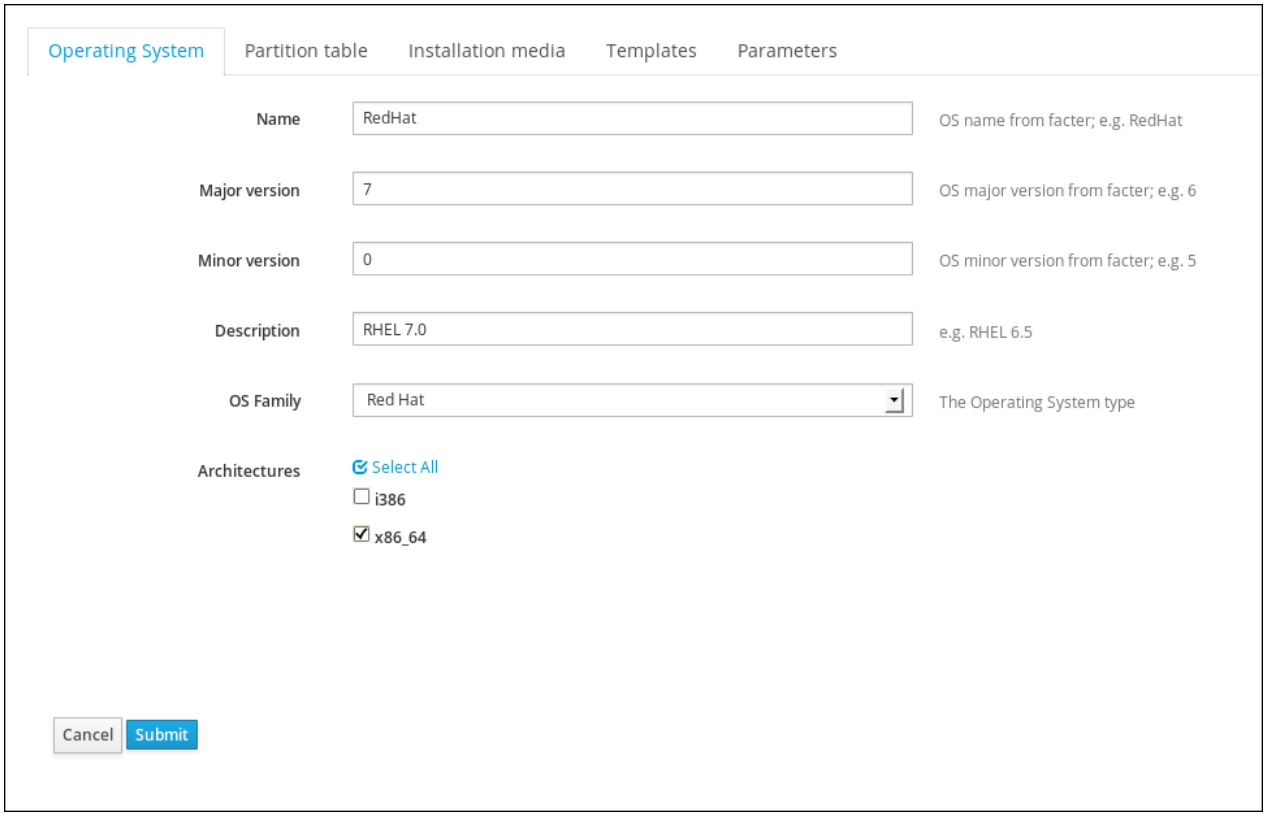

Figure 4.4. The New Operating System Page

3. Configure general operating system details:

    a. In the **Name** text field, enter a name to represent the operating system entry in the Red Hat Enterprise Linux OpenStack Platform Installer.

    b. In the **Major version**, enter the number corresponding to the major version of the operating system.

    c. In the **Minor version**, enter the number corresponding to the minor version of the operating system.

    d. In the **Description** text field, enter a description to represent the operating system entry in the Red Hat Enterprise Linux OpenStack Platform Installer.

    e. From the **OS Family** drop-down menu, select the operating system family to which the operating system belongs.

    f. In the **Architectures** section, select the check boxes for the architectures that the operating system entry can provision.

4. Configure the partition tables that the operating system entry can access:

    a. Click the **Partition table** tab.

    b. In the **All items** list, click the name of a partition table entry to move that partition table entry to the **Selected items** list.

5. Configure the installation media that the operating system entry can access:

    a. Click the **Installation media** tab.

    b. In the **All items** list, click the name of an installation media entry to move that installation media entry to the **Selected items** list.

6. Click **Submit**.

You have created a new operating system entry that can be used to provision a host.

4.4.2. Updating the Values of Operating System Entry Parameters

By default, the installer creates an operating system entry for Red Hat Enterprise Linux 7.2 and populates the values for this operating system entry automatically based on the information you provide during the installation process. You do not need to manually configure this operating system entry to provision Red Hat Enterprise Linux OpenStack Platform 5.0 on Red Hat Enterprise Linux 7.2, but you can edit this operating system entry to update the values of the parameters such as Subscription Manager credentials and the repositories to enable on machines that the installer provisions.

Note

See Section 4.4.3, “Common Operating System Entry Parameters” for a list of common parameters.

Procedure 4.7. Updating the Values of Operating System Entry Parameters

1. Click **Hosts** → **Operating systems**.
2. Click **RedHat 7.2** to open the details page for the Red Hat Enterprise Linux 7.2 operating system entry.
3. Click the **Parameters** tab.
4. Update the values of the parameters using the text fields for each parameter.
5. Click **Submit**.

You have updated the values of the parameters. These parameters are applied to any machines newly provisioned using the Red Hat Enterprise Linux 7.2 operating system entry.

4.4.3. Common Operating System Entry Parameters

The following is a list of common parameters you can specify in an operating system entry. Of these, to provision Red Hat Enterprise Linux OpenStack Platform based on an operating system entry, you must specify the parameters for a Subscription Manager account that can be used to install packages on the hosts being provisioned. You can also specify parameters for other options such as the location of a NTP server.

Table 4.1. Common Operating System Entry Parameters

Parameter	Description
ntp-server	The address of an NTP server with which to synchronize the system clock.
subscription_manager	Whether or not to use Subscription Manager. Accepted values are **true** or **false**.
subscription_manager_repos	A comma-separated list of repositories to enable. The repository required to provision RHEL OpenStack Platform on Red Hat Enterprise Linux 7.2 is **rhel-7-server-openstack-5.0-rpms**.

Parameter	Description
subscription_manager_pool	The ID of an entitlement pool to attach. You can only specify the ID of a single entitlement pool. This value is optional, but recommended; if you do not specify an entitlement pool ID, the installer attempts to auto-attach the required entitlements to hosts it provisions.
subscription_manager_username	The username of a Customer Portal account by which to register the host with the Customer Delivery Network.
subscription_manager_password	The password of a Customer Portal account by which to authenticate the above username.
http-proxy	The IP address or fully qualified domain name of a HTTP proxy for Subscription Manager.
http-proxy-port	The port by which to connect to the HTTP proxy for Subscription Manager.
http-proxy-user	The name of the user by which to authenticate with the HTTP proxy for Subscription Manager.
http-proxy-password	The password by which to authenticate with the HTTP proxy for Subscription Manager.
time-zone	The time zone to apply to the host. Time zones must be in the format of an IANA time zone identifier such as **America/New_York** or **Asia/Tokyo**.

4.4.4. Removing an Operating System Entry

You can remove existing operating system entries from the user interface.

Procedure 4.8. Removing an Operating System Entry

1. Click **Hosts** → **Operating systems**.
2. Click the **Delete** button for the operating system entry to delete.
3. Click **OK** when prompted.

You have removed an operating system entry, and that operating system entry is no longer accessible.

4.5. Provisioning Templates

A provisioning template is a collection of settings that defines the way in which a base operating system is installed on a host, and corresponds to a kickstart file for Red Hat Enterprise Linux systems. In the user interface, you can create and work with standalone provisioning templates in which all required settings are contained in a single template, and blocks of code that can be used across multiple provisioning templates, such as snippets, post-installation scripts, and generic installation scripts.

Note

The default provisioning template used to provision hosts in a RHEL OpenStack Platform environment can be edited as necessary to configure custom options. However, options that are configured using dedicated user interface elements, such as the root password, time zone, or custom repos, must be configured using the relevant operating system parameter or deployment option.

4.5.1. Creating a Provisioning Template

Create a new provisioning template entry.

Procedure 4.9. Creating a Provisioning Template

1. From the title bar in the main screen of the user interface, click **Hosts → Provisioning templates**.
2. Click **New Template**.
3. Configure general provisioning template details:
    a. Enter a name for the provisioning template in the **Name** field.
    b. Enter the body of the provisioning template in the **Template editor** text area.
    c. Enter a comment describing the creation of the provisioning template in the **Audit Comment** field.
4. Configure the provisioning template type:
    a. Click the **Type** tab.
    b. Select **Snippet** to designate the provisioning template as a snippet. A snippet is not a standalone provisioning template, but a part of a provisioning template that can be inserted into other provisioning templates.
    c. If you did not select the **Snippet** check box, select the type of the provisioning template from the **Type** list:
        - **PXELinux**: A PXELinux template.
        - **PXEGrub**: A PXEGrub template.
        - **iPXE**: An iPXE template.
        - **Provision**: The main provisioning template, such as a kickstart file for Red Hat Enterprise Linux systems.
        - **finish**: A post-installation script such as that specified in a kickstart file for Red Hat Enterprise Linux systems.
        - **script**: A generic script that can be run during the provisioning process, such as that specified in a kickstart file for Red Hat Enterprise Linux systems.
        - **user_data**: A **user_data** block.
        - **ZTP**: A zero touch provisioning template.

5. Configure the provisioning template associations:

    a. Click the **Associations** tab.

    b. From the **All items** list in the **Applicable Operating Systems** section, click the name of an operating system entry to move that operating system entry to the **Selected items** list and make the provisioning template available to that operating system entry.

    c. Optionally, click **Add combination** and select a host group from the **Host Group** list or an environment from the **Environment** list to make the provisioning template available to the specified combination of host groups and environments.

6. Click **Submit**.

4.5.2. Locking a Provisioning Template

Lock a provisioning template to prevent users from editing its properties.

Procedure 4.10. Locking a Provisioning Template

1. From the title bar in the main screen of the user interface, click **Hosts → Provisioning templates**.
2. Click the disclosure arrow next to the **Clone** button for the provisioning template.
3. Click **Lock**.

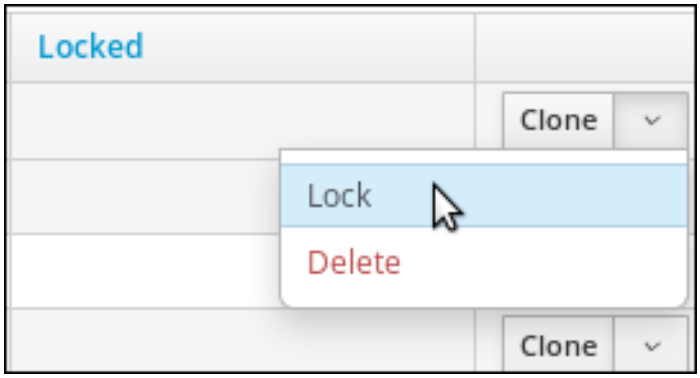

Figure 4.5. The Lock button

4. Click **OK** when prompted.

You have locked the provisioning template, and users cannot edit the properties of that provisioning template other than its associations. You must unlock the provisioning template to edit any other properties.

4.5.3. Removing a Provisioning Template

Remove a provisioning template that is no longer required.

Procedure 4.11. Removing a Provisioning Template

1. From the title bar in the main screen of the user interface, click **Hosts → Provisioning templates**.

2. If the provisioning template is locked, you must first unlock it:
    a. Click the disclosure arrow next to the **Clone** button for the provisioning template.
    b. Click **Unlock** and click **OK** when prompted.

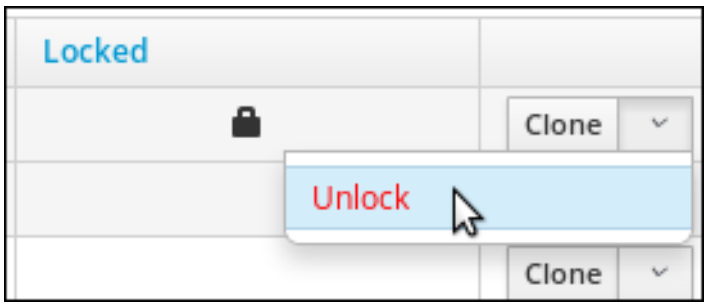

Figure 4.6. The Unlock button

3. Click the disclosure arrow next to the **Clone** button.
4. Click **Delete** and click **OK** when prompted.

4.6. Partition Tables

A partition table is a set of directives that defines the way in which the Red Hat Enterprise Linux OpenStack Platform Installer configures the disks available to machines it provisions. Partition tables are applied to operating system entries.

Note

The default partition table entry applied to Red Hat Enterprise Linux systems is **Kickstart default**. This partition table entry uses only the first disk available to the machine being provisioned; advanced configuration involving multiple disks must be manually configured. You can edit this partition table entry to configure the preferred partitioning scheme, or create a new partition table entry and specify that entry in the **RedHat 7.2** operating system entry.

4.6.1. Creating a Partition Table Entry

You can create new partition table entries that define custom partitioning schemes.

Procedure 4.12. Creating a Partition Table Entry

1. Click **Hosts** → **Partition tables**.
2. Click **New Partition Table**.
3. In the **Name** text field, enter a name to represent the partition table in the user interface.
4. In the **Layout** text area, enter the layout for the disk partition:

> **Example 4.3. Red Hat Enterprise Linux Partition Table**

```
zerombr
clearpart --all --initlabel
autopart
```

Note

The format of the layout must match that for the intended operating system to which the partition table is applied. For Red Hat Enterprise Linux 6.6 and Red Hat Enterprise Linux 7.2, the layout must match that of a kickstart file.

5. From the **Os family** drop-down menu, select the operating system family to which the partition table applies.
6. Click **Submit**

You have created a new partition table that can be applied to an operating system entry.

4.6.2. Removing a Partition Table Entry

You can remove existing partition table entries from the user interface.

Procedure 4.13. Removing a Partition Table Entry

1. Click **Hosts → Partition tables**.
2. Click the **Delete** button for the partition table entry to delete.
3. Click **OK** when prompted.

You have removed a partition table entry, and that partition table entry is no longer accessible.

4.7. Subnets

A subnet is the basic unit by which network settings are applied to hosts in a deployment. While a default subnet is created based on the values you specify when you install the user interface, you can also manually create new subnets to support complex networking setups. Later, when you create deployments, you can assign certain types of network traffic to specific subnets and configure the network interfaces on hosts to use specific subnets.

Note

The settings for subnets created in the user interface are provided for reference only, and do not update the DHCP or DNS settings on the machine where the user interface is installed.

4.7.1. Creating a Subnet

Create a new subnet to carry network traffic.

Procedure 4.14. Creating a Subnet

1. From the title bar in the main screen of the user interface, click **Infrastructure → Subnets**.
2. Click **New Subnet**.
3. Configure general subnet details:
    a. Enter a name for the subnet in the **Name** field.
    b. Enter the network address of the subnet in the **Network address** field.
    c. Enter the network mask for the subnet in the **Network mask** field.
    d. Enter the address of a gateway in the **Gateway** address field.
    e. Enter the address of the primary DNS server in the **Primary DNS server** field.
    f. Enter the address of the secondary DNS server, if any, in the **Secondary DNS server** field.
    g. Select the IP address management source from the **IPAM** list. If the subnet contains a DHCP server, select **DHCP**. If no DHCP capability exists on the subnet, select **Internal DB** and the Installer manages the IP address assignment and records IP addresses in its internal database. Select **None** for no IPAM management.
    h. Optionally, enter the start of the IP address range that the installer can auto-suggest for hosts in the **Start of IP range** field.
    i. Optionally, enter the end of the IP address range that the installer can auto-suggest for hosts in the **End of IP range** field.
    j. Optionally, enter the ID of a VLAN for the subnet.
    k. Select the default boot mode for interfaces assigned to the subnet from the **Boot mode** list. This sets the boot protocol for client network interfaces. Select **DHCP** to use the DHCP protocol or **Static** to manually assign IP address either manually through the host's network interface settings or using the Installer's IPAM Internal Database option.
4. Configure domain membership:
    a. Click the **Domains** tab.
    b. Select the check box for all domains of which the subnet is a member.
5. Click the **Capsules** tab to view the capsule settings. A capsule acts as a proxy for certain Installer services, such as DHCP, TFTP, and DNS.

    The Capsule settings are only required for the **default** subnet. This is because the Installer uses this subnet for host provisioning and deployment of Red Hat OpenStack Platform services. Any new subnets added to the Installer do not require Capsule settings.
6. Click **Submit**.

4.7.2. Removing a Subnet

Remove an existing subnet from the user interface if you no longer require that subnet.

Procedure 4.15. Removing a Subnet

1. From the title bar in the main screen of the user interface, click **Infrastructure → Subnets**.

2. Click the **Delete** button for the subnet to delete.
3. Click **OK** when prompted.

4.8. Users

A user is an account that can be used to log in to the user interface and administer resources. You can create new users in the internal directory that the installer manages or in an external directory service, and can apply fine-grained permissions to a user that define the actions that user can perform.

4.8.1. Adding a User

You can add new users to the user interface from the internal directory or an externally configured directory.

Procedure 4.16. Adding a User

1. Click **Administer** → **Users**.
2. Click the **New User** button to open the new user window.

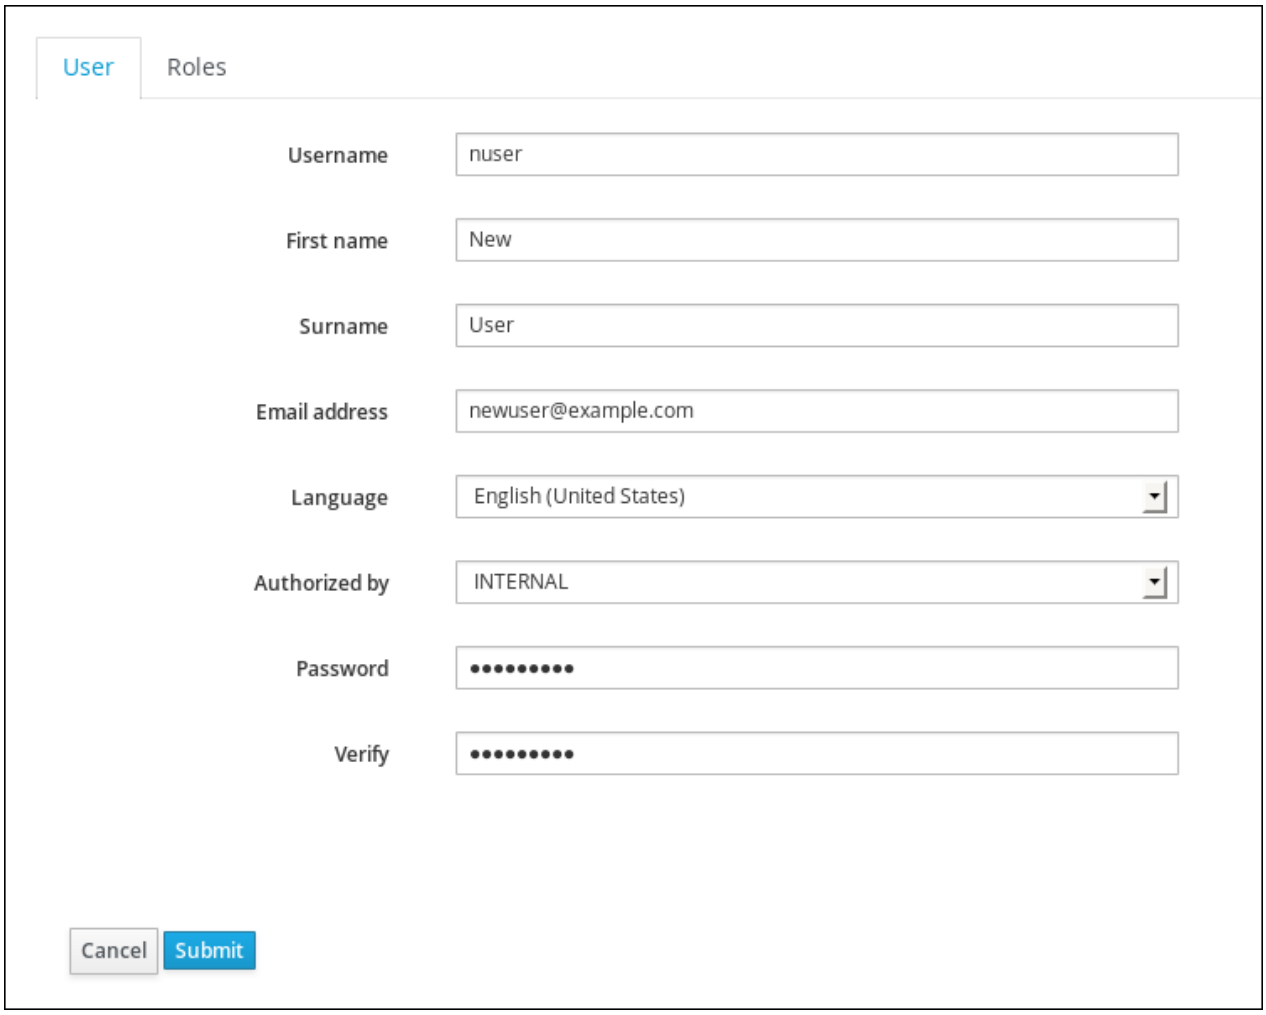

Figure 4.7. The New User Window

3. In the **Username** text field, enter the user name by which the user logs in to the user interface.

4. In the **First name** and **Surname** text fields, enter the real first name and surname of the user.
5. In the **Email address** text field, enter an email address by which the user can be contacted.
6. From the **Language** drop-down menu, select the language by which the user interface is displayed to the user.
7. Set a password for the user:
    a. From the **Authorized by** drop-down menu, select the directory by which the user will be authenticated.
    b. Enter an initial password for the user in the **Password** text field and again in the **Verify** text field.
8. Click **Submit**.

You have added a user to the user interface, and that user can log in and administer the user interface using the newly configured credentials.

4.8.2. Removing a User

Remove an existing user from the user interface.

Procedure 4.17. Removing a User

1. Click **Administer** → **Users**.
2. Click the **Delete** button for the user to delete.
3. Click **OK** when prompted.

You have removed an existing user from the user interface.

4.9. User Groups

User groups are collections of users inside the user interface. User groups allow you to apply a consistent set of permissions to a group of users at the same time, with any changes in permissions to the user group being automatically applied to all users that belong to that user group.

4.9.1. Adding a User Group

You can create user groups to apply a set of roles to a subset of the users in the user interface.

Procedure 4.18. Adding a User Group

1. Click **Administer** → **User Groups**.
2. Click the **New User Group** button.

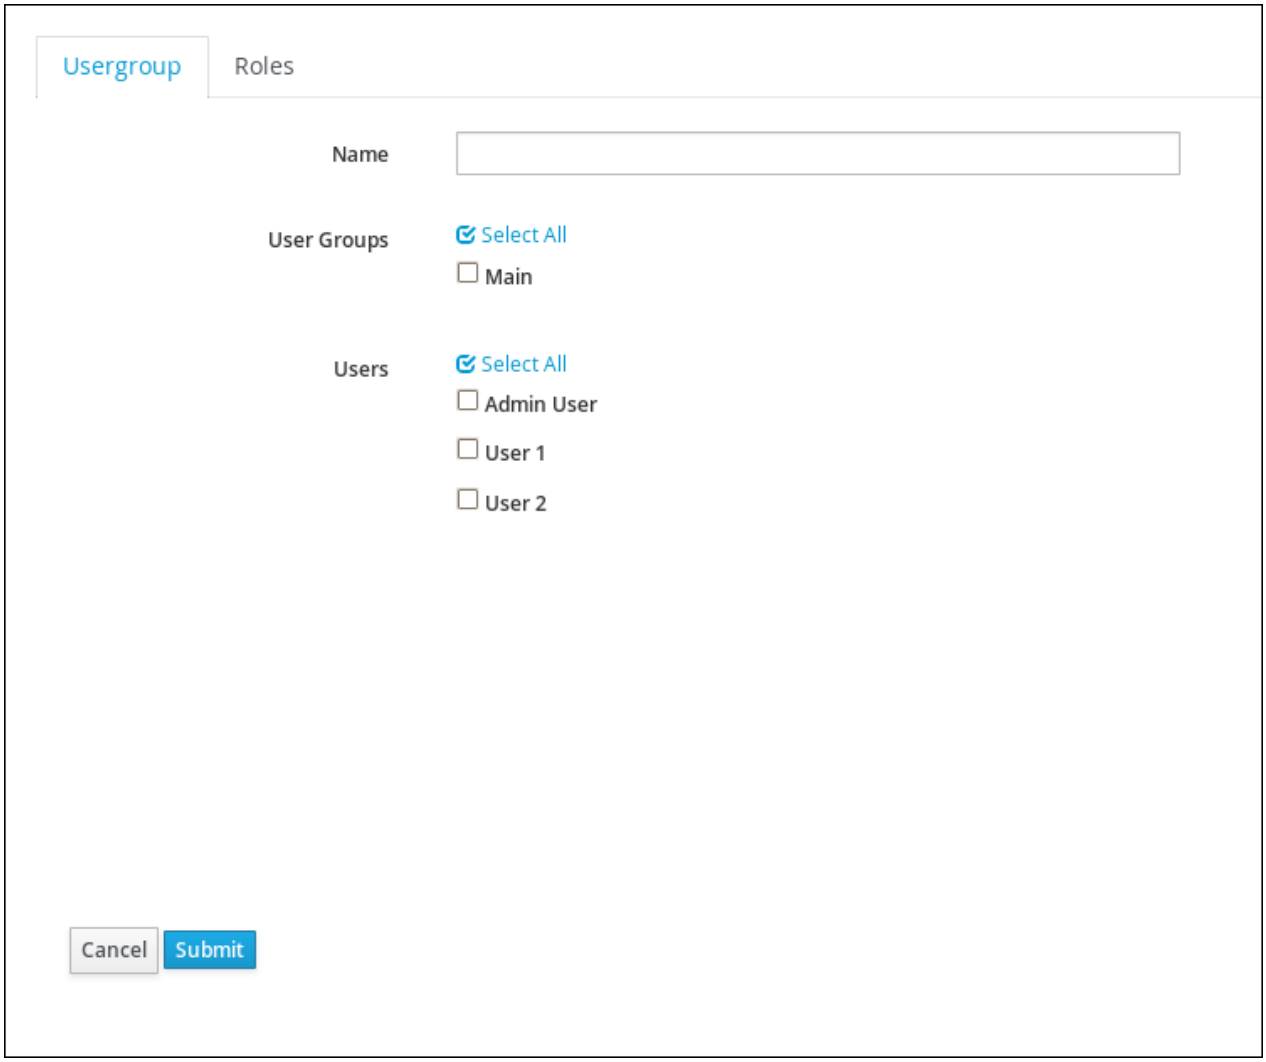

Figure 4.8. The New User Group Window

3. In the **Name** text field, enter a name by which to identify the user group in the user interface.
4. From the list of check boxes in the **User Groups** section, select the check boxes to include the corresponding user groups in the new user group.
5. From the list of check boxes in the **Users** section, select the check boxes to include the corresponding users in the new user group.
6. Click the **Roles** tab.
7. Select the **Admin** check box to grant administrator privileges to the users in the user group.
8. From the list of roles in the **All items** list, click the name of a role to apply that role to the user group.
9. Click **Submit**.

You have added a user group to the user interface, and all users in that user group and its child user groups, if any, are granted the roles that the user group defines.

4.9.2. Removing a User Group

Remove an existing user group from the user interface.

Procedure 4.19. Removing a User Group

1. Click **Administer** → **Users**.
2. Click the **Delete** button for the user group to delete.
3. Click **OK** when prompted.

You have removed an existing user group from the user interface, and the roles defined by that user group defined are no longer applied to the users that belonged to that user group and its child user groups.

4.10. Roles

A role is a set of permissions that defines the actions a user can perform in the user interface. Separate lists of permissions are available for each type of resource in the user interface, such as users, reports and operating systems, and can be further filtered to restrict the scope of the permissions to resources that match certain criteria. Roles are applied to users and user groups.

4.10.1. Creating a Role

You can create roles for granting a custom set of permissions to users in the user interface.

Procedure 4.20. Creating a Role

1. Click **Administer** → **Roles**.
2. Click **New Role**.

Figure 4.9. The New Role Window

3. In the **Name** text field, enter a name by which to identify the role in the user interface.
4. Click **Submit**.

You have created a role for granting a custom set of permissions to users in the user interface and must configure the permissions that role grants.

4.10.2. Adding Permissions to a Role

You can add new permissions to a role in the user interface to specify additional actions that users to which the role is assigned can perform.

Procedure 4.21. Adding Permissions to a Role

1. Click **Administer** → **Roles**.

2. Click **New Role**.
3. From the **Filters and permissions** list for the role to edit, click **Add permission** to open the **New Filter** page.

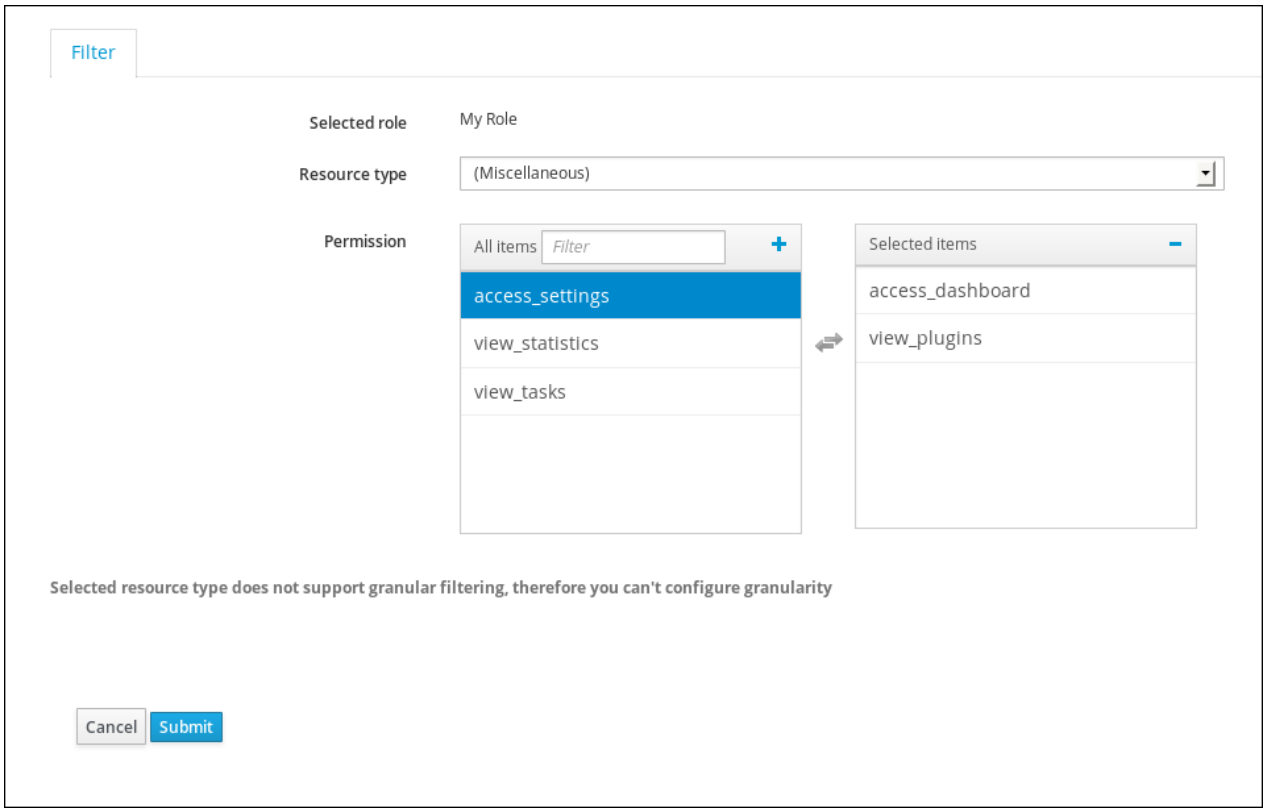

Figure 4.10. The New Filter Page

4. From the **Resource type** list, select the resource type to edit.
5. From the list of permissions in the **All items** list, select the permissions to add to the role.
6. Optionally, clear the **Unlimited?** check box and use the **Search** text field to specify a filter using search syntax.
7. Click **Submit**.

You have added permissions to a role, and users to which that role is assigned can perform the actions that role defines.

4.10.3. Removing Permissions from a Role

You can remove existing permissions from a role.

Procedure 4.22. Removing Permissions from a Role

1. Click **Administer** → **Roles**.
2. Click the **Filters and permissions** list for the role to edit to open the **Filters** page.
3. From the **Edit** list for the permission to remove, click **Delete**.
4. Click **OK** when prompted.

You have removed permissions from a role, and users to which that role is assigned can no longer perform the actions defined by the permissions that were removed.

4.10.4. Removing a Role

You can remove existing roles from the user interface.

Procedure 4.23. Removing a Role

1. Click **Administer** → **Roles**.
2. From the **Filters and permissions** list for the role to remove, click **Delete**.
3. Click **OK** when prompted.

You have removed a role, and the role is removed from all users to which it was applied.

Chapter 5. Adding Hosts

To provision a Red Hat Enterprise Linux OpenStack Platform environment, you must add hosts to the installer to act as nodes in that environment. Adding a host to the installer allows you to manage the host via the user interface, from collecting details on the hardware, operating system, and current status of the host to assigning that host to a role in a Red Hat Enterprise Linux OpenStack Platform environment and installing the base operating system and services required for the host to act as a node in that environment.

5.1. Host Status Types

Each host in the user interface is assigned a status type in accordance with the most recent action performed on or upcoming changes to be applied to that host. The following table outlines the status types to which hosts can be assigned. You can view the status of a host at any time by clicking **Hosts → All hosts** in the main title bar of the user interface.

Table 5.1. Host Status Types

Icon	Status	Description
D	Disabled	Alerts have been disabled for the host. The installer continues to collect reports from the host, but does not report the current status of the host.
B	Build	The host has been added to the user interface, but has not been provisioned.
S	Sync	The installer was unable to communicate with the host over the last report interval.
N	No Reports	No reports are available for the host.
E	Error	The installer has detected an error when attempting to collect reports from the host.
A	Active	Changes were applied to the host over the last report interval.
P	Pending	There are changes that are scheduled to be applied to the host.
O	OK	The host encountered no errors over the last report interval, no changes were applied to the host during that interval, and there are no pending changes.

5.2. Adding Hosts

The recommended method for adding hosts to the user interface is via discovery. With this function, machines in the private network that the installer defines are booted over the network using a discovery image, and are automatically registered in the user interface. During the registration process, details regarding the host such as the names of network interfaces attached to the host are also registered in the user interface, and are later used when configuring the host to act as a node in a RHEL OpenStack Platform environment.

It is also possible to add a host by manually entering the details of the host such as the MAC address and other networking details. However, details regarding the host are not automatically registered in the user interface as per discovery, and this method is not recommended in conjunction with the installer. For reference, see Section D.2, “Adding a Host Manually”.

Note

The discovery image is an ISO file that is installed alongside the *rhel-osp-installer* package. This ISO file is located in the **/usr/share/foreman-discovery-image/** directory on the machine on which the user interface is installed. Access to the discovery image over the network is automatically configured when you install the user interface; no user input is required.

5.2.1. Adding a Host via Discovery

Add a host to the user interface via the discovery function. This procedure assumes that the host is connected to the private network on which the installer provides the DHCP service.

Procedure 5.1. Adding a Host via Discovery

1. Start the host, and elect to start over the network from the boot options menu to start the host using the PXE service that the installer provides.
2. Select **Foreman Discovery** from the PXE boot options menu. The host starts into the **Foreman Discovery** screen and is automatically registered in the user interface.

Figure 5.1. The Foreman Discovery Screen

3. Log in to the user interface and confirm that the host has been registered:

    a. Click **Hosts → Discovered hosts**.

    b. Click the name of the newly registered host to open the details page for the host, and review the details.

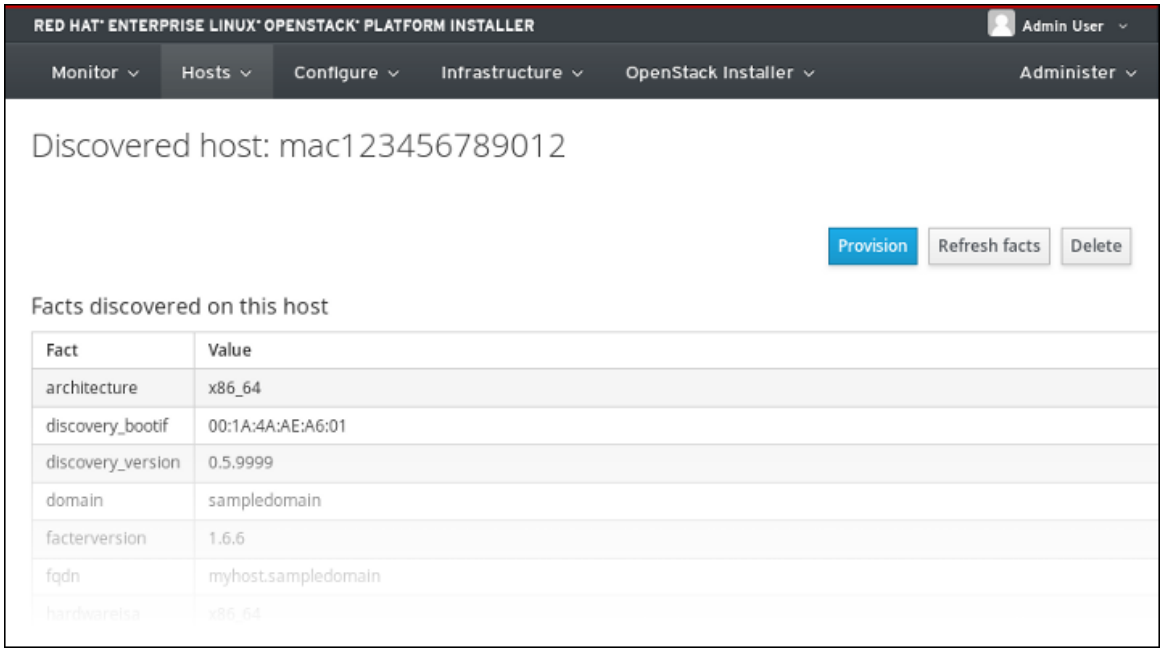

Figure 5.2. The Discovered Host Details Page

The host is added to the user interface, and is automatically included in the list of free hosts when you assign hosts to a deployment. Moreover, the host is automatically added to the **`discovery`** environment, which is the default Puppet environment for discovered hosts in which no Puppet classes are defined. This environment acts as a holding area that identifies hosts that have not yet been provisioned.

5.3. The Host Overview Page

The host overview page displays information about a given host and the connection between the host and the installer. You can view the host overview page for a host by clicking **Hosts → All hosts**, and clicking the name of a host.

Note

The information displayed in the host overview page is based on facts and reports collected from the host. For more information on facts and reports, see Chapter 7, *Monitoring Hosts*.

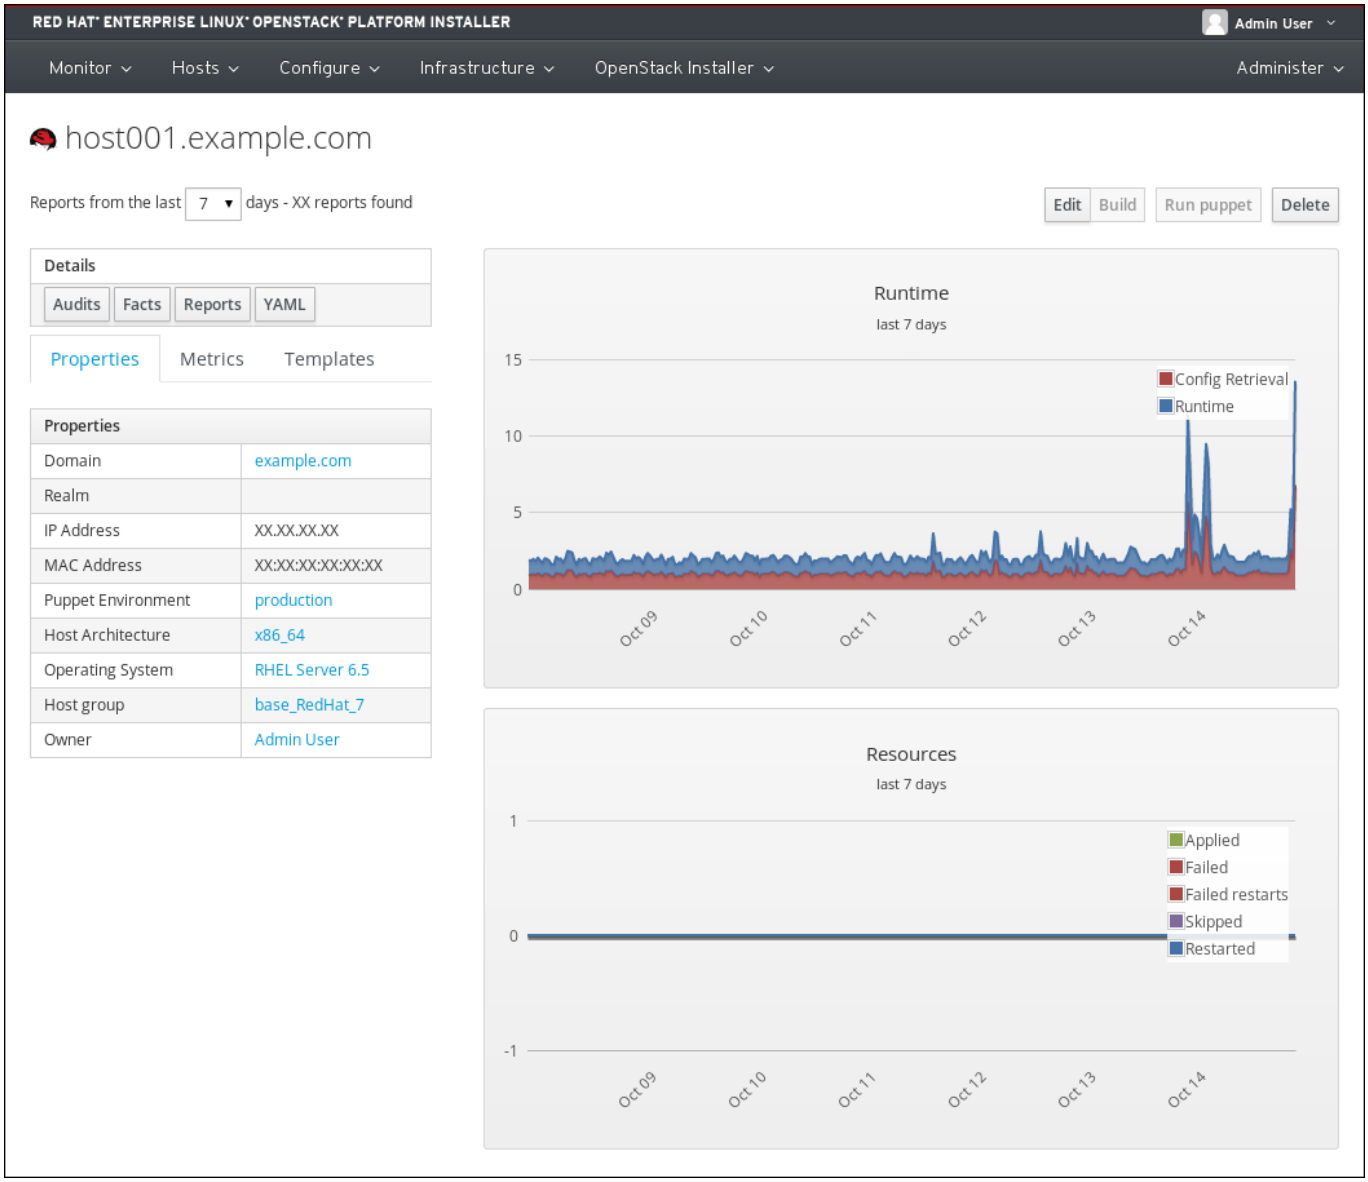

Figure 5.3. The host overview page

The Details Bar

The details bar contains a row of buttons that provide shortcuts to more information about the host, and tabs that display summaries of important details and events.

Audits

Displays a page containing audit entries for the current host, if any.

Facts

Displays the **Facts** page, and applies a filter so that only the facts for the current host are displayed. This button is only available after the installer has collected facts from the host.

Reports

Displays the **Reports** page, and applies a filter so that only the reports for the current host are displayed. This button is only available after the installer has collected reports from the host.

YAML

Displays a page containing details about the host in YAML format, such as its IP address, MAC address, and the name and values of parameters that have been applied to the host.

Properties

A list of general details about the host, such as its IP address, MAC address, and the operating system entry that has been applied to the host.

Metrics

A table displaying a summary of all events reported for the host.

Templates

A list of all provisioning templates currently accessible by the host. The provisioning templates include in this list are automatically configured in accordance with the operating system entry applied to the host.

Host Actions

In addition to displaying information about the host, the host overview page also provides a number of buttons for performing common actions on the host.

Edit

Opens the host details page for the host and allows you to configure settings for the host. However, in principle, the installer configures all the details required for provisioning the host, and no manual configuration is required.

Build

Flags the host to be provisioned the next time it is restarted. After you click this button, the label changes to **Cancel build**, which allows you to cancel the provisioning operation. The button reverts to **Build** after provisioning is complete. However, in principle, the installer manages all aspects of the provisioning process, and you do not need to manually provision hosts.

Run puppet

Executes a Puppet run on the host. This button is disabled by default, and must be manually enabled by setting the value of the **puppetrun** option to **true** in the **Puppet** tab of the **Settings** page accessed by clicking **Administer → Settings**.

Delete

Deletes the host from the user interface.

Host Graphs

The host overview page contains two graphs that display the status of recent Puppet runs executed on the host.

Runtime

This graph tracks two data points: **Config Retrieval**, and **Runtime**. The **Config Retrieval** data point represents the amount of time taken to collect information about the host during a given Puppet run, and the **Runtime** data point represents the amount of time required to execute the Puppet run. Both data points are measured in seconds.

Resources

This graph tracks the number of actions performed on the host during a Puppet run. The categories displayed in this graph are identical to those displayed in the **Reports** page, and are measured using the number of actions in each category.

5.4. Removing a Host

You can remove hosts from the user interface.

Procedure 5.2. Removing a Host

1. Click **Hosts** → **All hosts**.
2. From the **Edit** list for the host to remove, click **Delete**.
3. Click **OK** when prompted.

You have removed a host, and that host is no longer managed by the installer.

Chapter 6. Provisioning Red Hat Enterprise Linux OpenStack Platform

The installer provisions Red Hat Enterprise Linux OpenStack Platform using deployments. A deployment is a collection of settings that defines the hosts on which services are to be provisioned, options such as whether to use Neutron networking or Nova networking as the networking back end, and key parameters for several of the services to be provisioned.

Two of the key considerations when you create a new deployment are the networking back end, and the volume driver for the Block Storage (Cinder) service:

Networking Back Ends

RHEL OpenStack Platform offers two networking back ends: Nova networking, and Neutron networking. Nova networking has been deprecated in the OpenStack technology roadmap, but still remains available. For information on each of these back ends, including how to decide which back end is suitable for your environment and how to configure the options available to each back end, see "Choose a Network Back End" in Deploying OpenStack: Learning Environments (Manual Setup).

Volume Drivers

A volume driver is the storage back end that the Block Storage service uses to provide block storage to Compute nodes in a RHEL OpenStack Platform environment. The default implementation is to use LVM, but you can also choose from NFS storage, Ceph storage, or EqualLogic storage. For more information on each of the volume drivers and how to configure the options available to each volume driver, see "Volume Drivers" in the Configuration Reference Guide.

6.1. Deployments

A deployment maps OpenStack services to hosts in your environment and installs these services to their assigned hosts. This includes various options for high availability controllers, Neutron or Nova networking, and a choice of messaging service. A deployment also provides specific parameters to customize each service. Deployments provide mechanisms to define all aspects of an OpenStack environment and help scale the environment in the future.

The following diagram outlines the topology of a sample deployment that uses high availability:

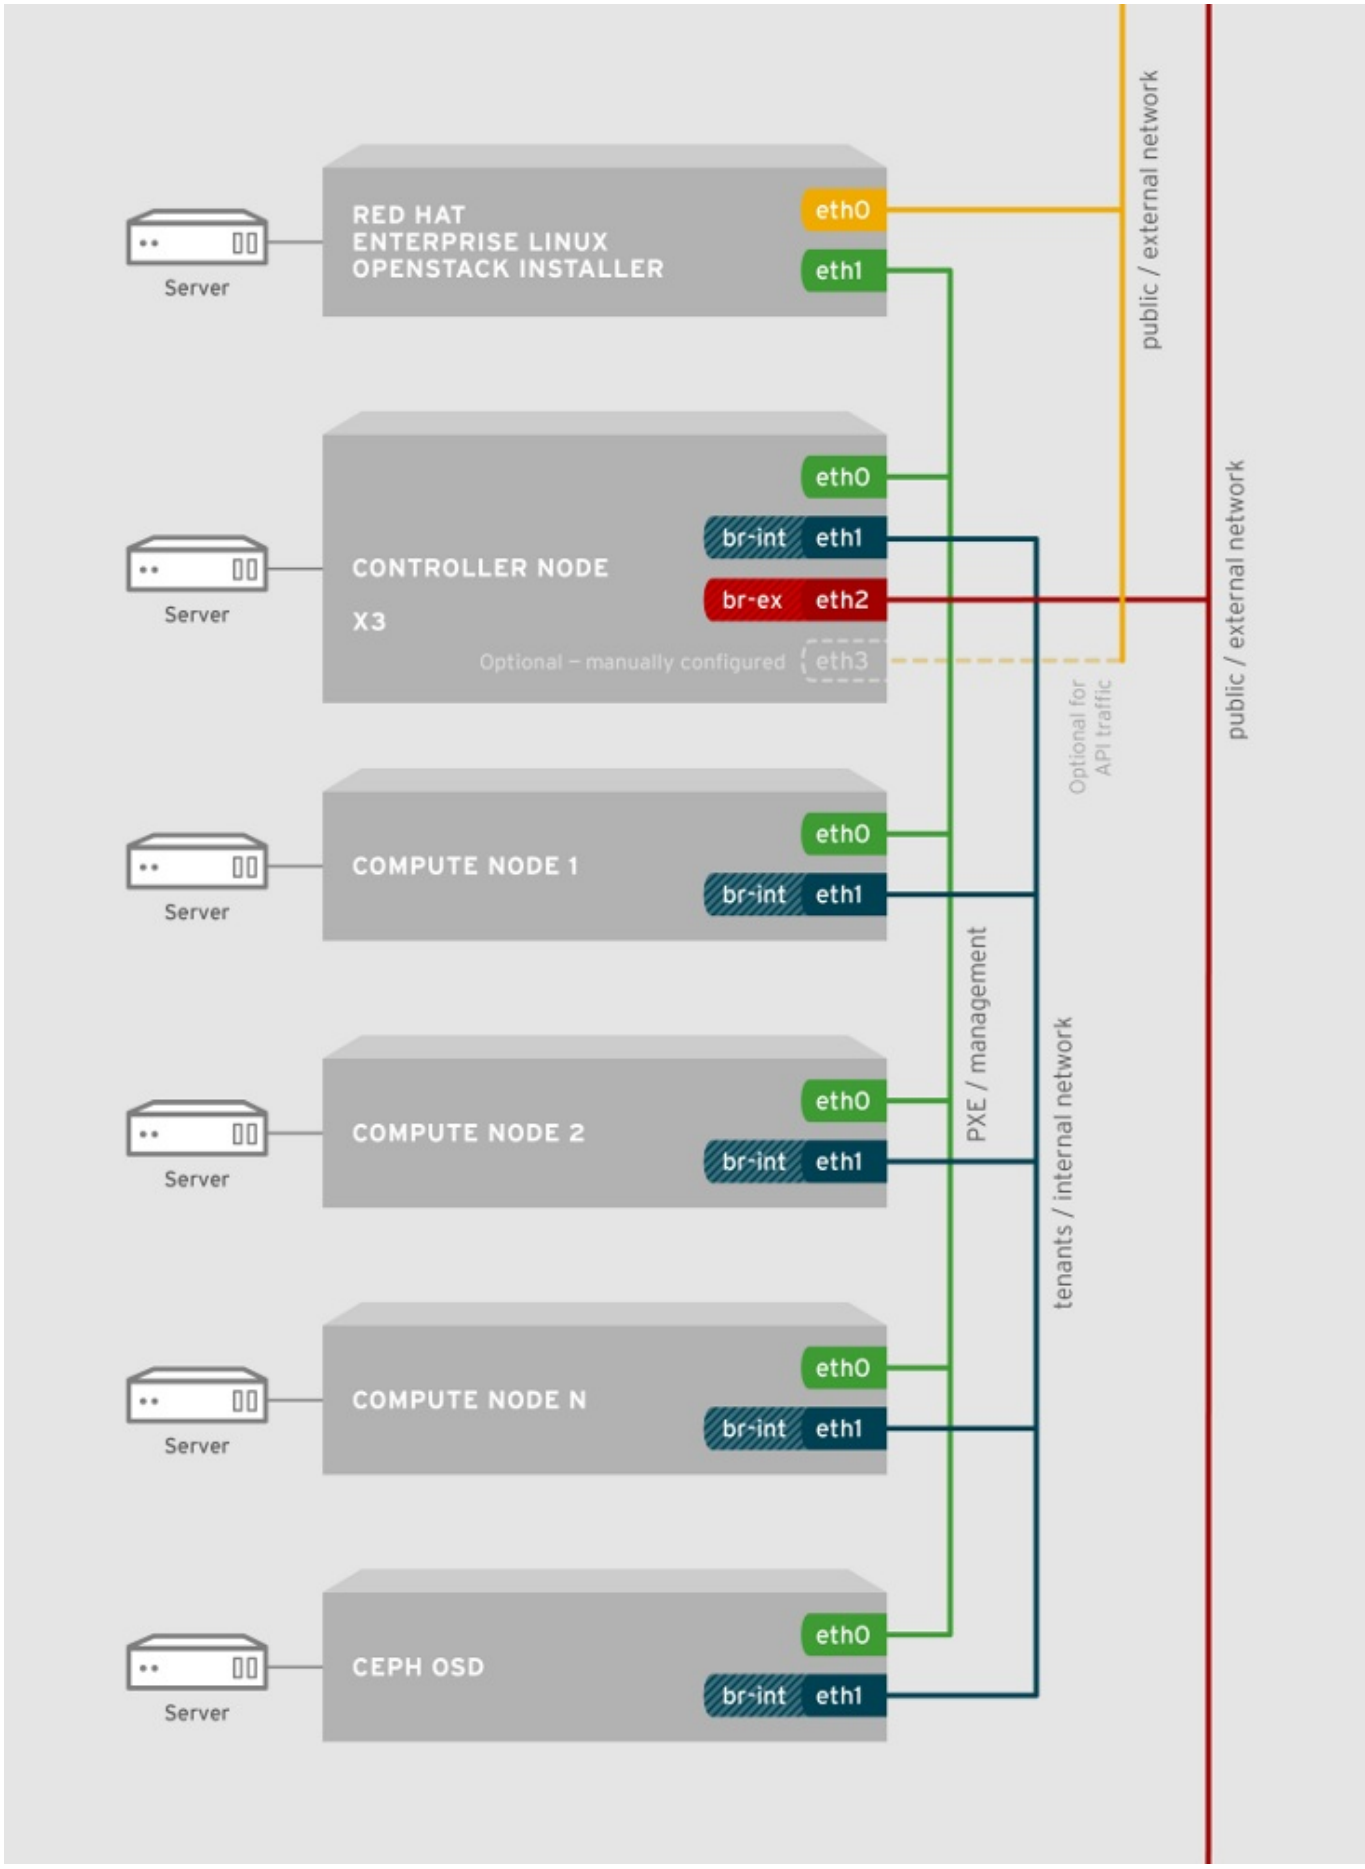

Figure 6.1. The topology of a sample deployment with high availability

There are seven hosts in this deployment. The server icon at the top of the diagram represents the

installer. The server icon beneath the installer represents controller nodes. Because this deployment uses high availability, there are three hosts assigned to this role so that the environment can continue to function in the event of a failure on any of the three nodes. The three server icons directly beneath the controller nodes represent compute nodes. The first two compute nodes were added when the RHEL OpenStack Platform environment was provisioned. The third compute node, **`Compute N`**, represents additional compute nodes that can be added to the environment after it has been provisioned to enable horizontal scaling. The final server icon in the diagram represents a Ceph OSD node.

The nodes in this deployment communicate with each other and with external networks over the four following networks:

- The yellow network represents a public network that only the installer uses. This network allows the installer to communicate with external network resources such as the Content Delivery Network. It is configured manually on the host where the user interface is installed. Moreover, access to the dashboard from external networks can also be enabled by placing public API traffic on the same subnet as the provisioning network described below, or by optionally creating a dedicated subnet for this network traffic type and configuring external routing.

- The green network represents the provisioning network that the installer defines when you install the user interface. This network carries traffic between the installer and the hosts in the deployment. This is the network that the installer uses to boot hosts using PXE, and to set up and configure hosts when you provision a RHEL OpenStack Platform environment. While the installer defines this network when you install the user interface, you must also create a subnet to carry this network traffic when you create a deployment. Moreover, while this network traffic type is carried over a dedicated subnet in this deployment, this network traffic type can be shared with any of the network traffic types carried by the blue network outlined below. The subnet that carries this network traffic type must be assigned to every host in the deployment.

- The blue network represents a private network that carries network traffic between the nodes in a RHEL OpenStack Platform environment. This network carries management, cluster management, administrative API, storage, and storage clustering network traffic. The subnet that carries these network traffic types must be assigned to every host in the deployment.

- The red network represents a public network that carries network traffic between instances in the RHEL OpenStack Platform environment and external networks. This network provides external network connectivity to virtual machines. A dedicated subnet must be created to carry this type of network traffic. The subnet that carries this network traffic type is only assigned to controller nodes.

6.1.1. Creating a Deployment

Create a deployment for provisioning a RHEL OpenStack Platform environment. For information on each of the options available when configuring networking and services, see Section 6.1.2, “New Deployment Settings”.

Procedure 6.1. Creating a Deployment

1. Click **OpenStack Installer → New deployment**.

2. Configure deployment settings:

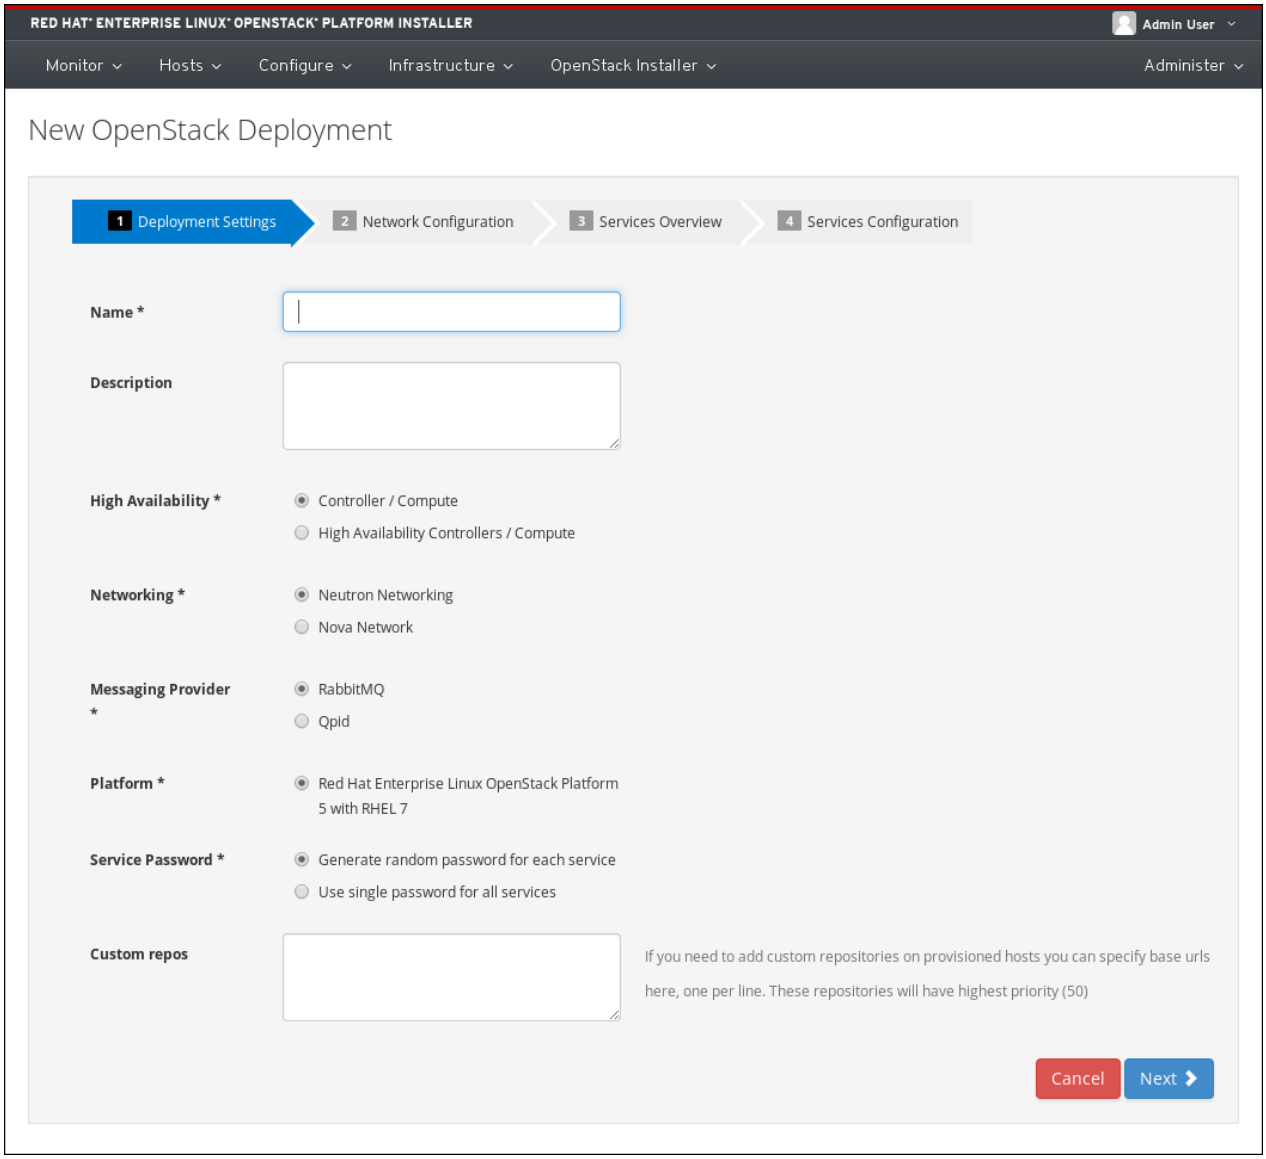

Figure 6.2. The Deployment Settings step

a. Enter a name for the deployment in the **Name** field.

b. Enter a description in the **Description** field.

c. Select either **Controller / Compute** or **High Availability Controllers / Compute** in the **High Availability** section.

d. Select either **Neutron Networking** or **Nova Network** in the **Networking** section.

e. Select either **RabbitMQ** or **Qpid** as the message broker in the **Messaging Provider** section.

Note

As of Red Hat Enterprise Linux OpenStack Platform 5, RabbitMQ replaces QPid as the default (and recommended) message broker.

f. Ensure **Red Hat Enterprise Linux OpenStack Platform 5 with RHEL 7** is selected in the **Platform** section.

g. Select **Generate random password for each service** in the **Service Password** section to generate a random password for each service, or select **Use single password for all services** and enter a password in the **Password** and **Verify** fields to set the same password for all services.

h. Optionally, enter a list of custom repositories in the **Custom repos** text area. Multiple custom repositories must be added on separate lines, and each custom repository must take the form of a base URL.

i. Click **Next**.

3. Configure network traffic:

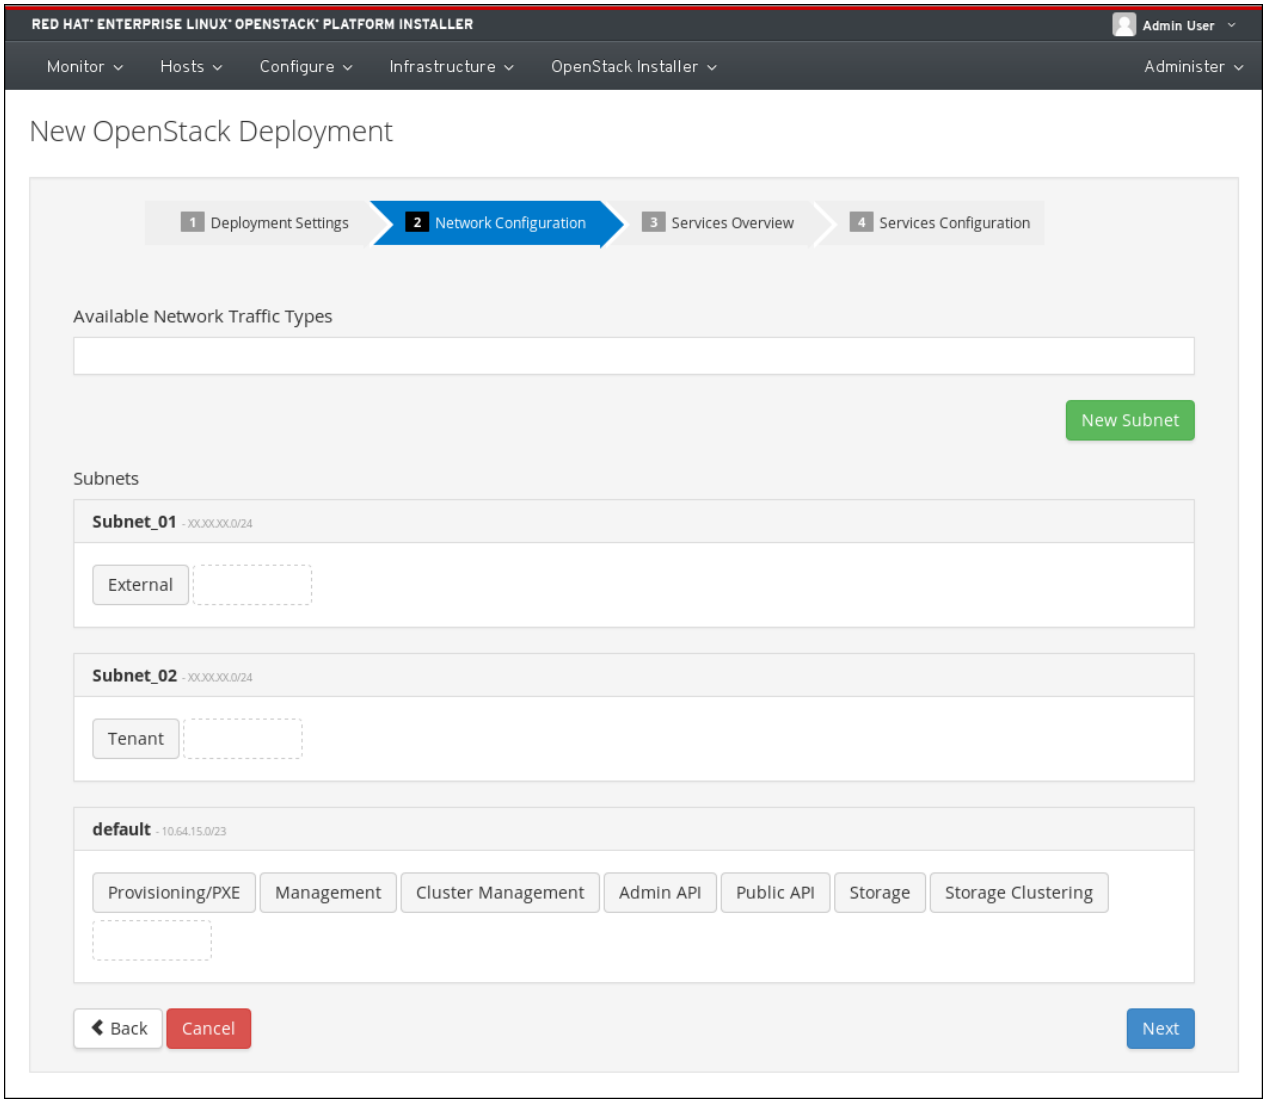

Figure 6.3. The Network Configuration step

a. Drag and drop the available network traffic types into the section for a subnet.

b. Create new subnets if you require a new subnet to which to assign a network traffic type:

   a. Click **New Subnet**.

   b. Enter a name to represent the subnet in the user interface in the **Name** field.

   c. Select **External DHCP** or **No existing DHCP** from the **DHCP server** list.

Setting **DHCP server** to **External DHCP** means an external DHCP server exists on the subnet and the Installer sets the boot mode of clients on this subnet to use DHCP.

Setting **DHCP server** to **No existing DHCP** means no DHCP server exists on the subnet and the Installer assigns static IP addresses to clients on the subnet. These static IP addresses are recorded in the Installer's database, which gives the Installer a DHCP-like capability. If you select **No Existing DHCP** and the subnet is to carry the **Public API** network traffic type, you must enter the address of a machine that can act as a gateway in the **Gateway** field. If you select **No existing DHCP**, you can also specify the **IP Range Start** and the **IP Range End**.

d. Enter the network address in the **Network Address** field. This address must be in CIDR notation. For example, **XX.XX.XX.0/24**.

e. Optionally, enter a VLAN ID for the subnet in the **VLAN** field.

f. Click **Create Subnet**.

c. Click **Next**.

Note

By default, all network traffic types except **External** is assigned to the default subnet. Because the **External** network traffic type is required and cannot be assigned to the same subnet as other network traffic types, you must create a dedicated subnet for this network traffic type. You can then drag and drop network traffic types to subnets as required, or to the **Available Network Traffic Types** section to disable optional network traffic types.

4. In the **Services Overview** step, review the list of services to be provisioned, and click **Next**.

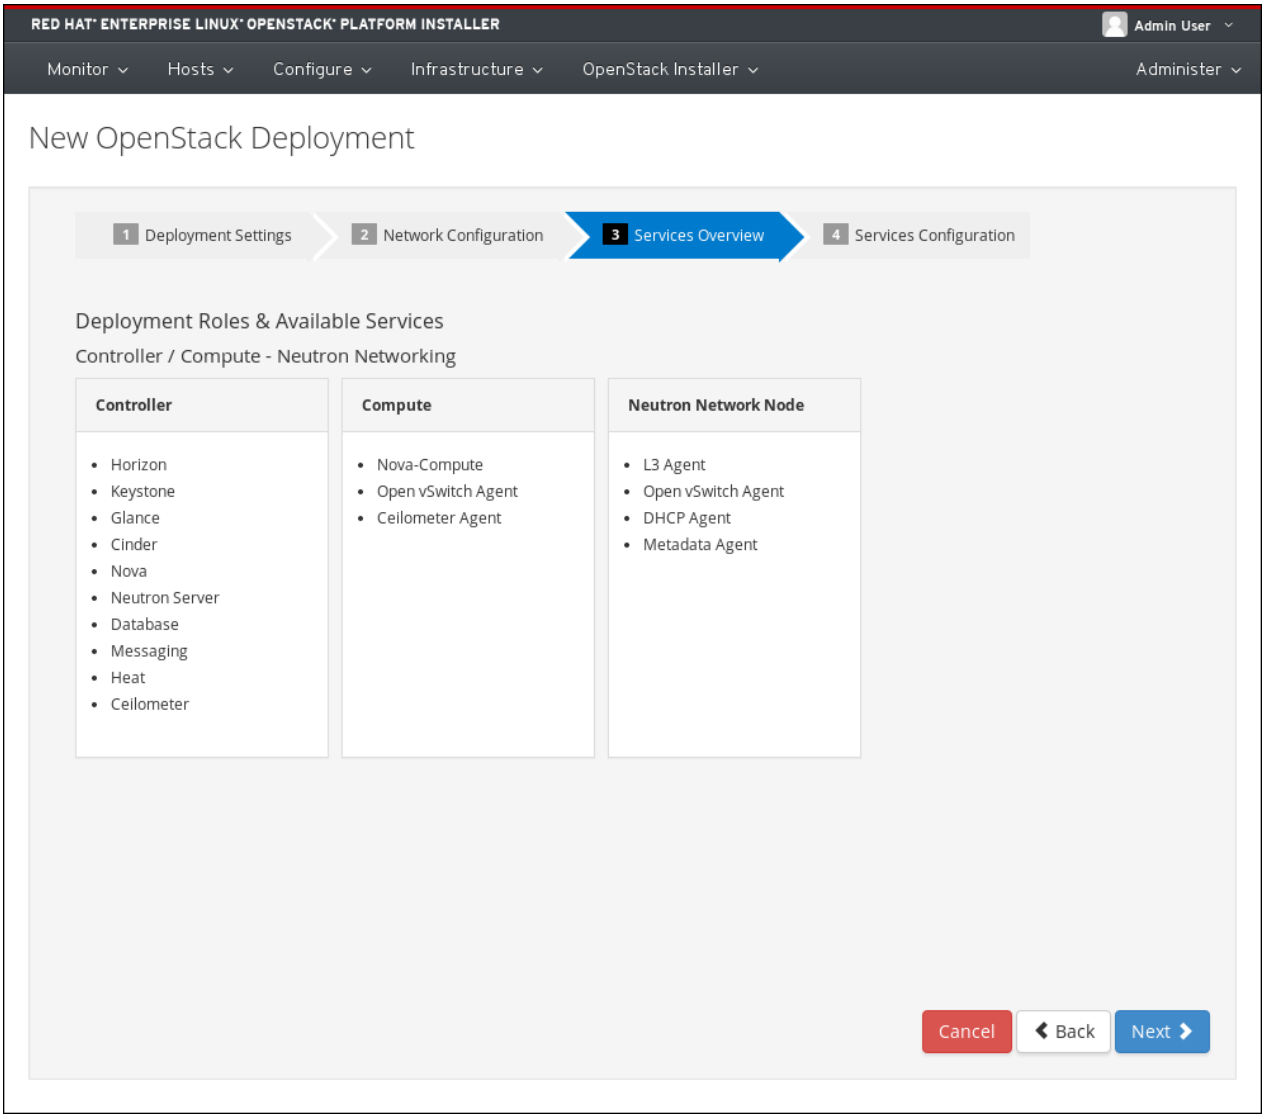

Figure 6.4. The Services Overview step

5. Configure service options:

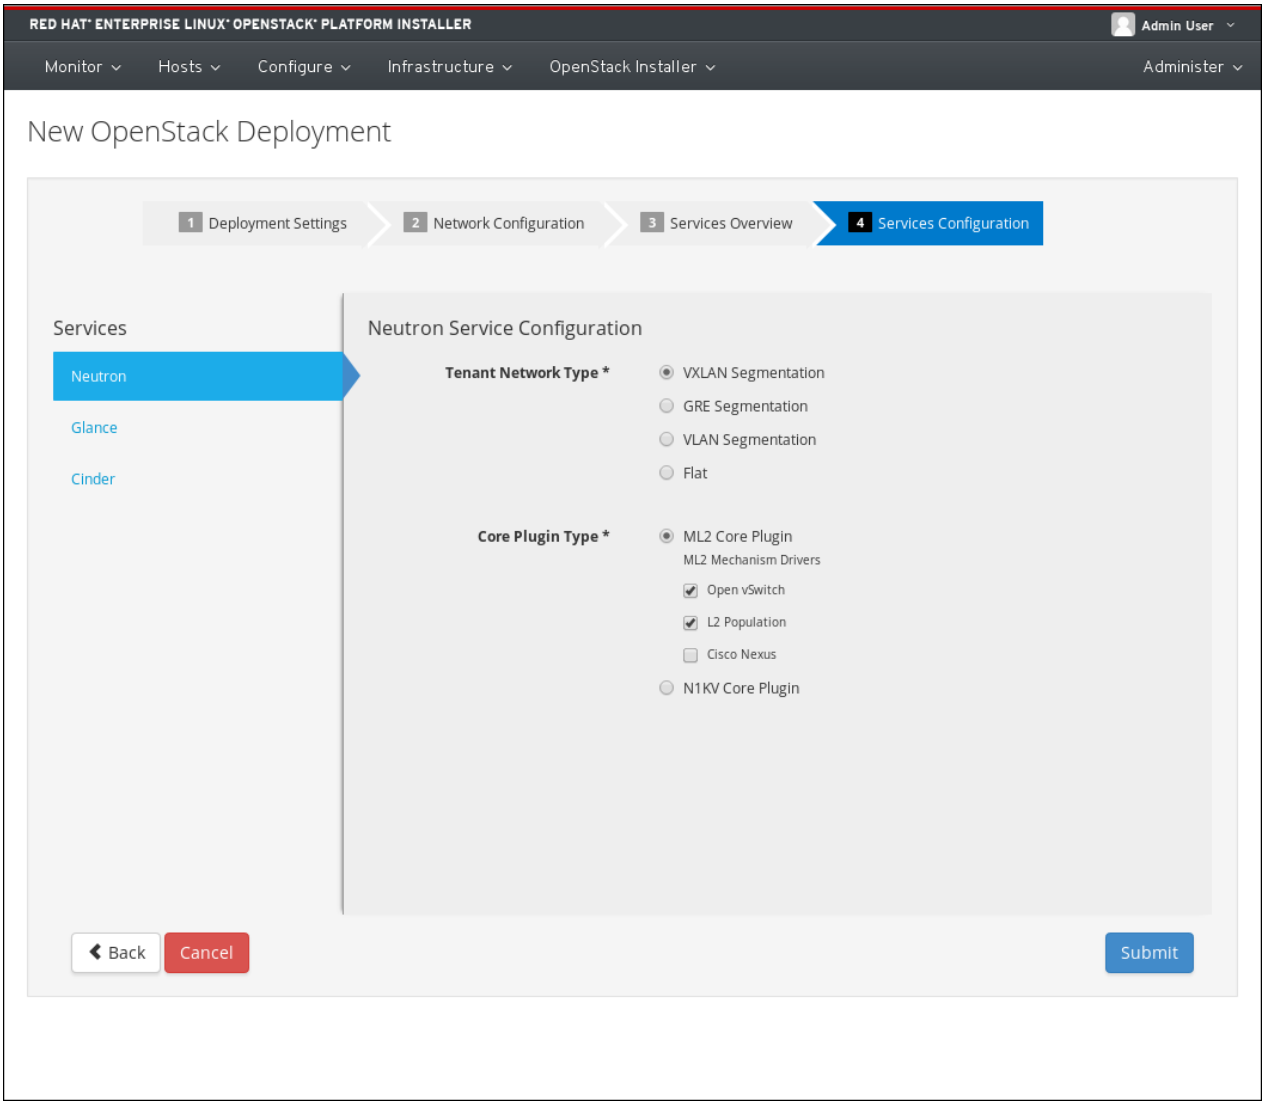

Figure 6.5. The Services Configuration step

A. **Neutron**

Options for this service are only available if you selected **Neutron Networking** in the **Deployment Settings** step.

a. Select a **Tenant Network Type**.

b. Select a **Core Plugin Type**.

B. **Nova**

Options for this service are only available if you selected **Nova Network** in the **Deployment Settings** step.

a. Select a **Tenant Network Type**.

b. Enter a **Floating IP range for external network**.

c. Enter a **Fixed IP range for tenant networks**.

C. **Glance**

a. Select a driver backend from the **Choose Driver Backend** radio boxes.

D. **Cinder**

a. Select one or more driver backends from the **Choose Driver Backend** check boxes.

6. Click **Submit**.

6.1.2. New Deployment Settings

This section outlines the options available when you create a new deployment.

6.1.2.1. Network Configuration Settings

The following table outlines the different network traffic types available in the **Network Configuration** step of creating a new deployment.

Table 6.1. Network Configuration Settings

Network Traffic Type	Description
Provisioning/PXE	Communication with the installer. For example, to provision hosts and to allow hosts to boot over the network using the PXE service that the installer provides.
Management	Private communication amongst services.
External	Communication between instances in the RHEL OpenStack Platform environment and external network resources. This network traffic type cannot be placed on the same subnet as other network traffic types.
Cluster Management	Pacemaker cluster management communication.
Admin API	Administrative communication amongst services.
Public API	Access to the REST API and the Horizon dashboard.
Tenant	Tenant network communication for instances in the RHEL OpenStack Platform environment.
Storage	Storage-related communication. This network traffic type is optional.
Storage Clustering	Storage-related communication for replication and data checks in Ceph clusters. This network traffic type is optional.

6.1.2.2. Services Configuration Settings

The following table outlines the options available in the **Services Configuration** step of creating a new deployment.

Table 6.2. Services Configuration Settings

Service	Option	Description

Service	Option	Description
Neutron	**Tenant Network Type**	The type of network used for tenant networks. You can choose from the following options: - **VXLAN Segmentation**: Virtual extensible LAN. A network layout that extends the functionality of a VLAN by providing a larger network segment ID range and other features. - **GRE Segmentation**: Generic routing encapsulation segmentation. A network layout in which tunnels are used to segregate and carry network traffic over individual tenant networks. - **VLAN Segmentation**: A network layout in which multiple isolated networks are defined. If you choose this option, you must also specify the tenant VLAN range. - **Flat**: A simple network layout in which all instances are members of the same network.

Service	Option	Description
	Core Plugin Type	The core network plugin. You can choose from the following options: **ML2 Core Plugin** The modular layer 2 plugin. If you select this plugin, you can select the following ML2 mechanism drivers: • **Open vSwitch**: A software-defined networking virtual switch designed to supersede the heritage Linux software bridge. Open vSwitch provides switching services to virtualized networks with support for industry-standard NetFlow, OpenFlow, and sFlow. • **L2 Population**: Enables horizontal scaling of broadcast, multicast, and unicast traffic on large overlay networks. L2 population implements a partial mesh for ARP resolution and MAC learning traffic whereby this traffic is sent only to the necessary agent by encapsulating it as a targeted unicast. • **Cisco Nexus**: Enables configuration of Cisco Nexus switches. If you select this option, you must enter the **Switch Hostname**, **Switch IP Address**, **Switch Login**, **Switch Password**, **SSH Port**, and **Port Mappings**. You can add additional switches by clicking **Add Another Switch**. **N1KV Core Plugin** The Cisco 1000v plugin. If you select this plugin, you must also specify the details of a virtual supervisor module: • **VSM IP**: The IP address of the virtual supervisor module. • **VSM Password**: A password with which to authenticate with the virtual supervisor module.
Nova	**Tenant Network Type**	• **Flat with DHCP**: A simple network layout in which all instances are members of the same network, and the DHCP service is also provided. • **VLAN**: A network layout in which multiple isolated networks are defined. If you choose this option, you must also specify the tenant VLAN range.

Service	Option	Description
	Floating IP range for external network	A range of IP addresses to be used for external networks. Ranges must be entered using CIDR notation. For example, **XX.XX.XX.0/24**.
	Fixed IP range for tenant networks	A range of IP addresses to be used for tenant networks. Ranges must be entered using CIDR notation. For example, **XX.XX.XX.0/24**.
Glance	**Choose Driver Backend**	The driver to use for back-end storage of images. You can choose from the following options: - **Local File**: The local file system. This option is only available in deployments without high availability. - **Ceph**: Ceph storage. - **NFS**: Network file system storage. If you select this option, you must also specify the NFS address in the **network path** text field.
Cinder	**Choose Driver Backend**	The volume driver to use for back-end storage. You can choose from the following options: - **LVM**: Logical volume manager storage. This option is only available for deployments without high availability. - **NFS**: Network file system storage. If you select this option, you must also specify the NFS address in the **NFS URI** text field. - **Ceph**: Ceph storage. - **EqualLogic** for EqualLogic storage attached network systems. If you select this option, you must also specify the details of the storage attached network, including the **SAN IP Addr**, **SAN Login**, **SAN Password**, **Pool**, and **Group**.

6.1.3. Editing a Deployment

Edit the properties of an existing deployment. The new properties are applied to Red Hat Enterprise Linux OpenStack Platform environments newly provisioned using that deployment.

Important

Editing the properties of a deployment and repeating the provisioning of a Red Hat Enterprise Linux OpenStack Platform environment to apply those change is not supported.

Procedure 6.2. Editing a Deployment

1. Click **OpenStack Installer** → **Deployments** to open the **OpenStack Deployments** page.
2. Click the name of a deployment to open the details page for that deployment.

3. Click the name of the deployment to edit the name, and click **Save** when complete to make the change persistent.
4. Click the description of the deployment to edit the description, and click **Save** when complete to make the change persistent.
5. Optionally, click the **Revisit Setup Wizard** button to open the setup wizard and configure further options such as the messaging broker and the password with which to authenticate services.

6.1.4. Removing a Deployment

Remove an existing deployment from the user interface. Removing a deployment does not affect any existing environments that have been provisioned based on that deployment.

Procedure 6.3. Removing a Deployment

1. Click **OpenStack Installer → Deployments**.
2. Click the **Delete** button for the deployment to remove.
3. Click **OK** when prompted.

6.2. Assigning Hosts to Deployment Roles

To provision RHEL OpenStack Platform using a deployment, you must assign hosts you have added to the user interface to roles in that deployment. Assigning a host to a role automatically populates configuration options for that host, such as the operating system entry, Subscription Manager details, and puppet modules to install on the host.

The following list outlines the number of hosts that must be assigned to each deployment role in accordance with the key options available when you create a deployment.

Without High Availability (Neutron)

- One controller node.
- One or more compute nodes.
- One Neutron networking node.

Without High Availability (Nova)

- One controller node.
- One or more compute nodes.

With High Availability (Nova and Neutron)

- Three controller nodes.
- One or more compute nodes.

6.2.1. Assigning a Host to a Deployment Role

Assign a host you have added to the user interface to one of the roles in a deployment.

Procedure 6.4. Assigning Hosts to Deployment Roles

1. Click **OpenStack Installer → Deployments**.
2. Click the name of the deployment to which to assign hosts to open the details page for that deployment.

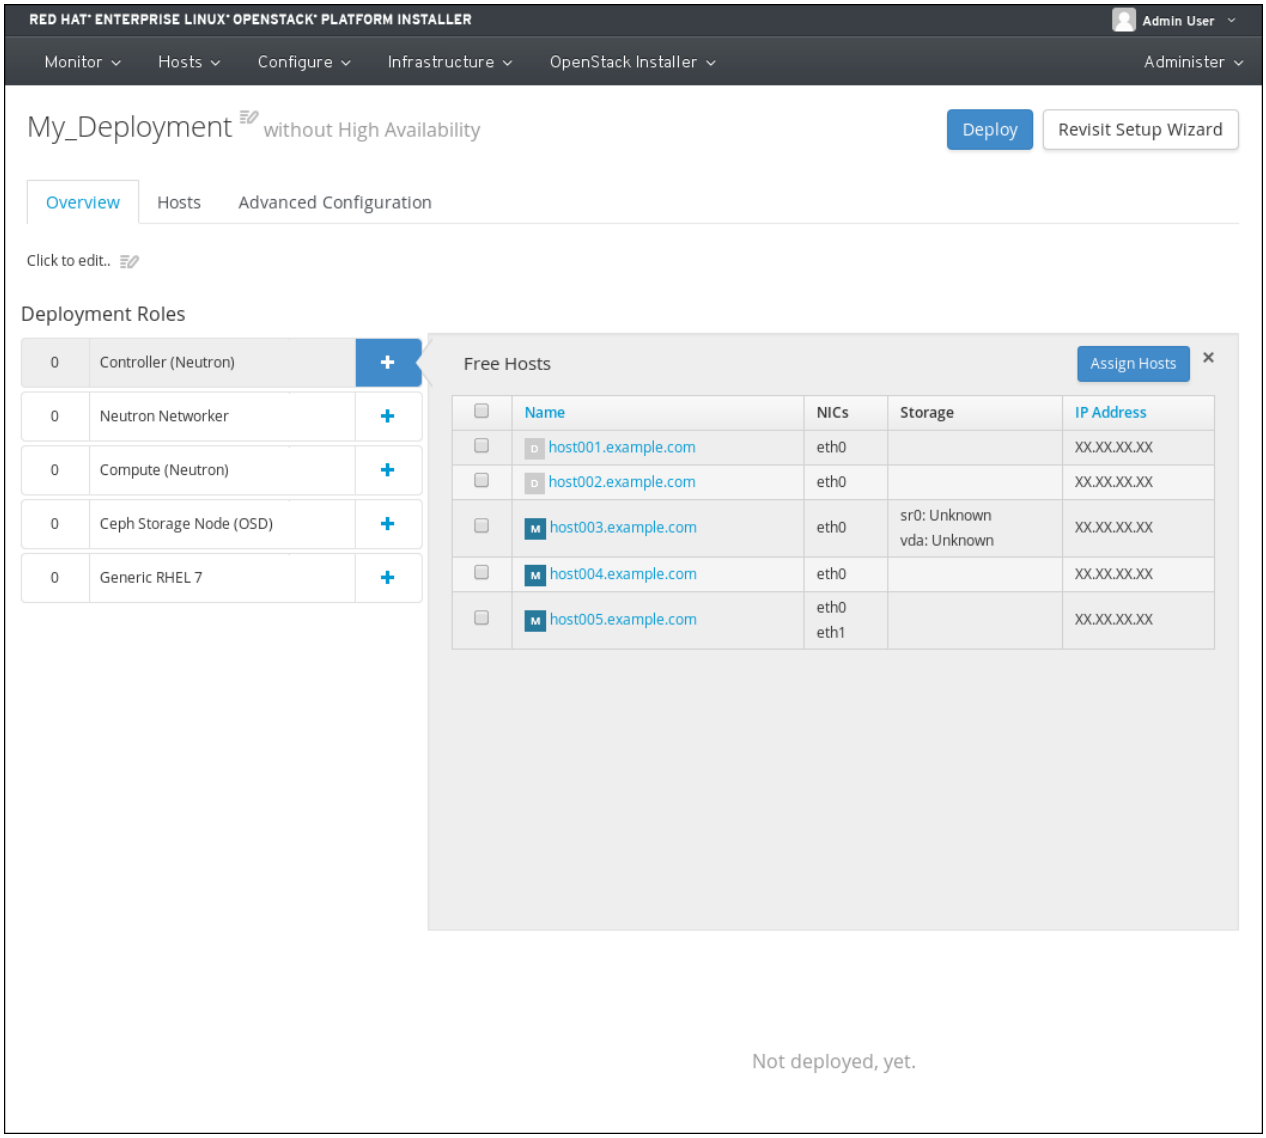

Figure 6.6. The deployment details page

3. In the **`Deployment Roles`** section, click **+** for a deployment role to display the **`Free Hosts`** section for that deployment role:

   Controller

   Provides key services such as the MySQL database for storing data about your environment, Horizon, Keystone, and Glance.

   Compute

   A host that acts as a hypervisor, providing the processing capabilities required for running virtual machines in the environment. You can add more Compute nodes to your environment at any time by assigning additional hosts to this deployment role and repeating the provisioning process; the installer ignores all hosts that have already been provisioned and provisions only the new host.

Neutron Networking

Provides Neutron networking services. This deployment role is only available when you select Neutron networking as the networking back end for your deployment and have selected not to use high availability.

Ceph Storage Node (OSD)

A generic Red Hat Enterprise Linux 7.2 host that can be manually configured after deployment to act as a Ceph storage server node. This deployment role is optional.

Generic RHEL 7

A generic Red Hat Enterprise Linux 7.2 host that can be manually configured after deployment to provide services not defined by any of the pre-existing deployment roles. This deployment role is optional.

4. Select the check box for a host in the **Free Hosts** section.
5. Click **Assign Hosts** to assign the host to the selected deployment role.

6.2.2. Unassigning a Host from a Deployment

Unassign a host you have added to the user interface from one of the roles in a deployment.

Procedure 6.5. Unassigning Hosts from Deployment Roles

1. Click **OpenStack Installer → Deployments**.
2. Click the name of the deployment from which to unassign hosts to open the details page for that deployment.

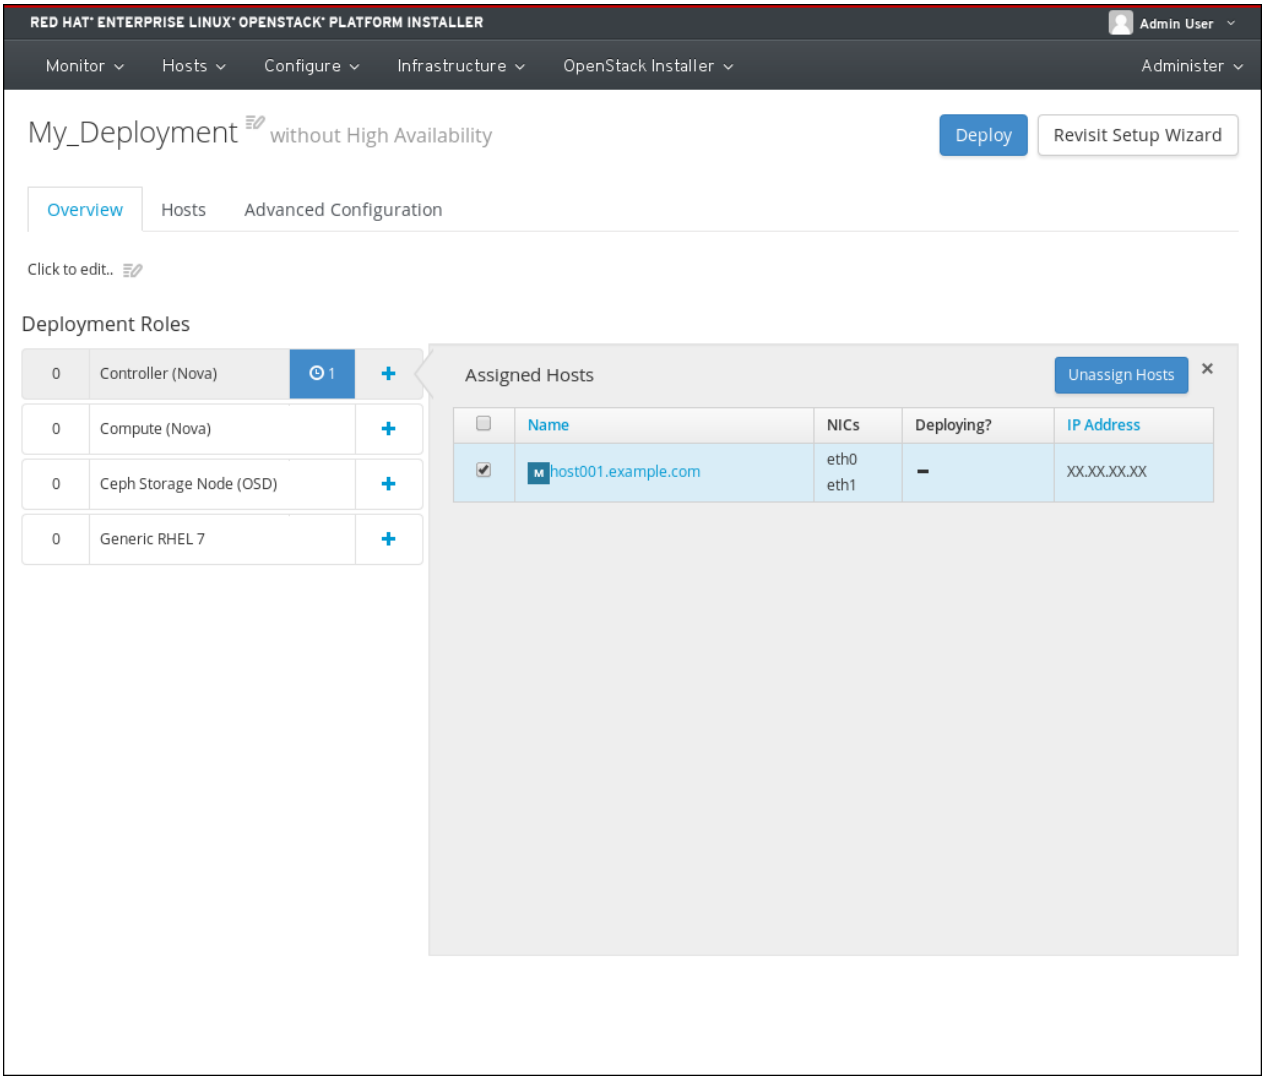

Figure 6.7. The Deployment Details Page

3. In the **Deployment Roles** section, click the assigned hosts ( ) button for a deployment role to display the **Assigned Hosts** section for that deployment role.
4. Select the check box for a host in the **Assigned Hosts** section.
5. Click **Unassign Hosts** to unassign the host from the selected deployment role.

6.3. Configuring Host Networking

After you assign a host to a deployment role in a deployment, you can configure the network traffic that each network interface on the host carries. Network traffic types are divided amongst subnets in accordance with the options you select when you create a deployment.

6.3.1. Configuring Networking on a Host

Specify the network traffic carried by the network interfaces on a host.

Procedure 6.6. Configuring Networking on a Host

1. Click **OpenStack Installer → Deployments**.
2. Click the name of a deployment to open the details page for that deployment.

3. Click the **Hosts** tab.
4. Click the **Assigned** sub-tab.
5. Select the check box for a host and click **Configure Networks**.
6. Drag and drop subnets between the sections for network interfaces to change the network traffic carried by those interfaces.

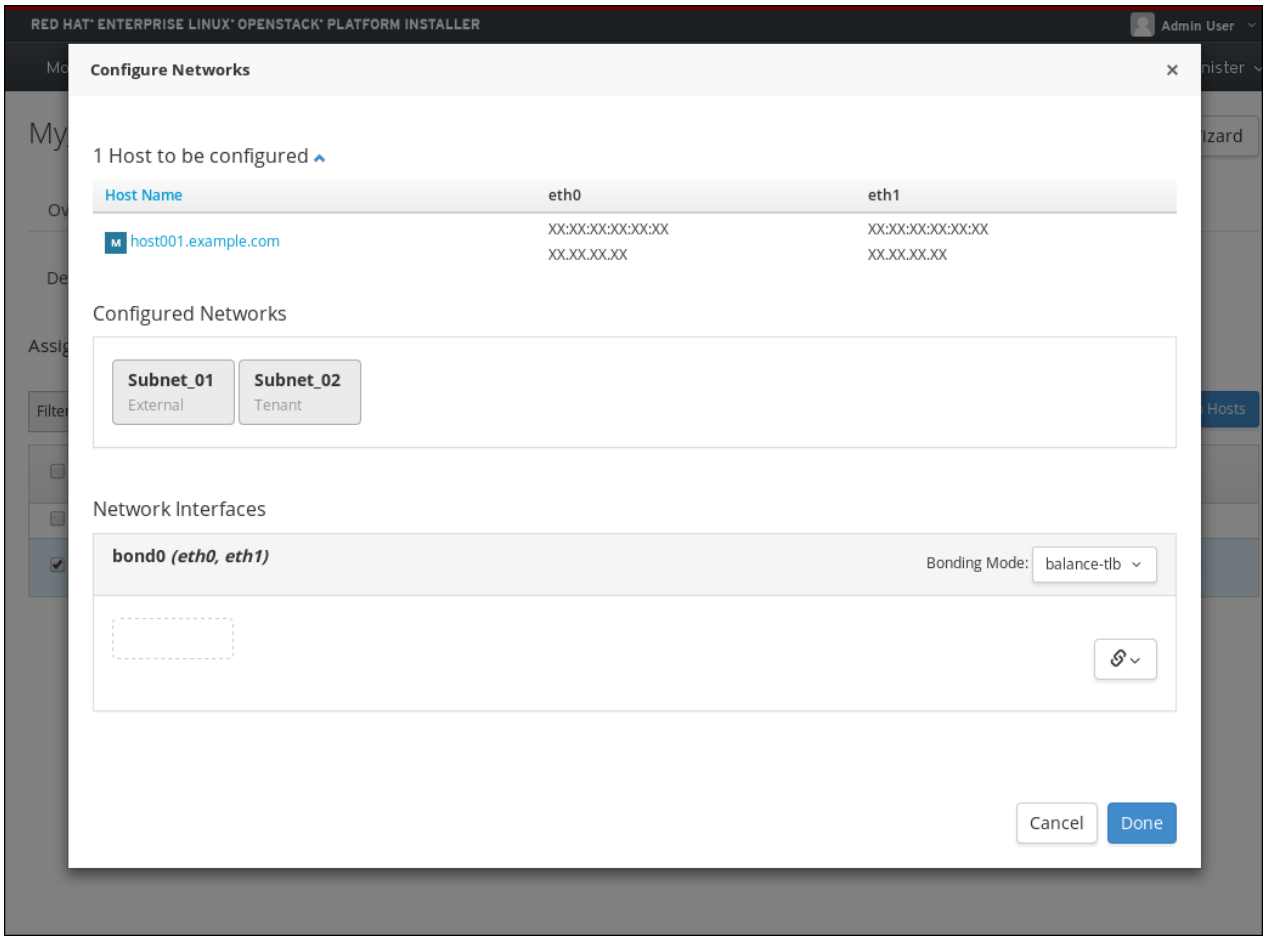

Figure 6.8. The Configure Networks Screen

7. Optionally, configure a bond between two network interfaces:
    a. Click the bonding button for a network interface and click the name of a network interface to bond the two network interfaces.
    b. Select the bonding mode from the **Bonding Mode** list.
8. Click **Done**.

6.3.2. Bonding Modes

The following table outlines the bonding modes that you can assign to bonded network interfaces.

Table 6.3. Bonding Mode

Bonding Mode	Description

Bonding Mode	Description
balance-rr	The packets in all inbound and outbound network traffic handled by the bond is transmitted sequentially through each of the network interfaces in the bond. For example, if there are two network interfaces in the bond, the first packet is handled by the first network interface, the second packet is handled by the second network interface, and the third packet is handled by the first network interface.
active-backup	One of the network interfaces in the bond is designated as the active network interface, and the other network interface is designated as a backup. Inbound and outbound network traffic are both handled only by the active network interface. Network traffic is only handled by the backup network interface when the active network becomes unavailable and the backup network interface becomes the active network interface.
balance-xor	Network traffic is handled by the network interfaces in the bond in accordance with an XOR formula. With this bonding mode, network traffic from a given source MAC address is always handled by the same destination MAC address in the network interface.
broadcast	All inbound and outbound network traffic is passed through all network interfaces in the bond.
802.3ad	All inbound and outbound network traffic is handled in accordance with aggregation groups that each have the same speed and duplex settings.
balanced-tlb	Inbound network traffic is handled by the current network interface. Outbound network traffic is distributed amongst the active network interfaces in the bond in accordance with the load on each network interface.
balance-alb	In addition to distributing outbound network traffic in a similar fashion to **balance-tlb**, inbound IPv4 network traffic is also distributed amongst the active network interfaces in the bond.

6.4. Configuring Fencing

If you selected **High Availability Controllers / Compute** in the **Deployment Settings** step when creating a deployment, you must manually configure fencing on each of the hosts you assign to the controller deployment role. The following procedure must be performed after you have assigned hosts to deployment roles and before you provision the RHEL OpenStack Platform environment.

Note

For more information about configuring high availability in Red Hat Enterprise Linux 7.2, see High Availability Add-On Reference.

Important

To configure fencing, you must have a fencing device in your RHEL OpenStack Platform environment that you can use to fence the nodes in that environment.

6.4.1. Configuring Fencing on High-Availability Nodes

Configure fencing on hosts that have been assigned to the controller deployment role in deployments where you selected **High Availability Controllers / Compute** in the **Deployment Settings** step when creating that deployment.

Procedure 6.7. Configuring Fencing on High-Availability Nodes

1. From the title bar in the main screen of the user interface, click **Hosts** → **All hosts**.
2. Click the name of the host to configure.
3. Click **Edit**.
4. Click the **Fencing** tab.

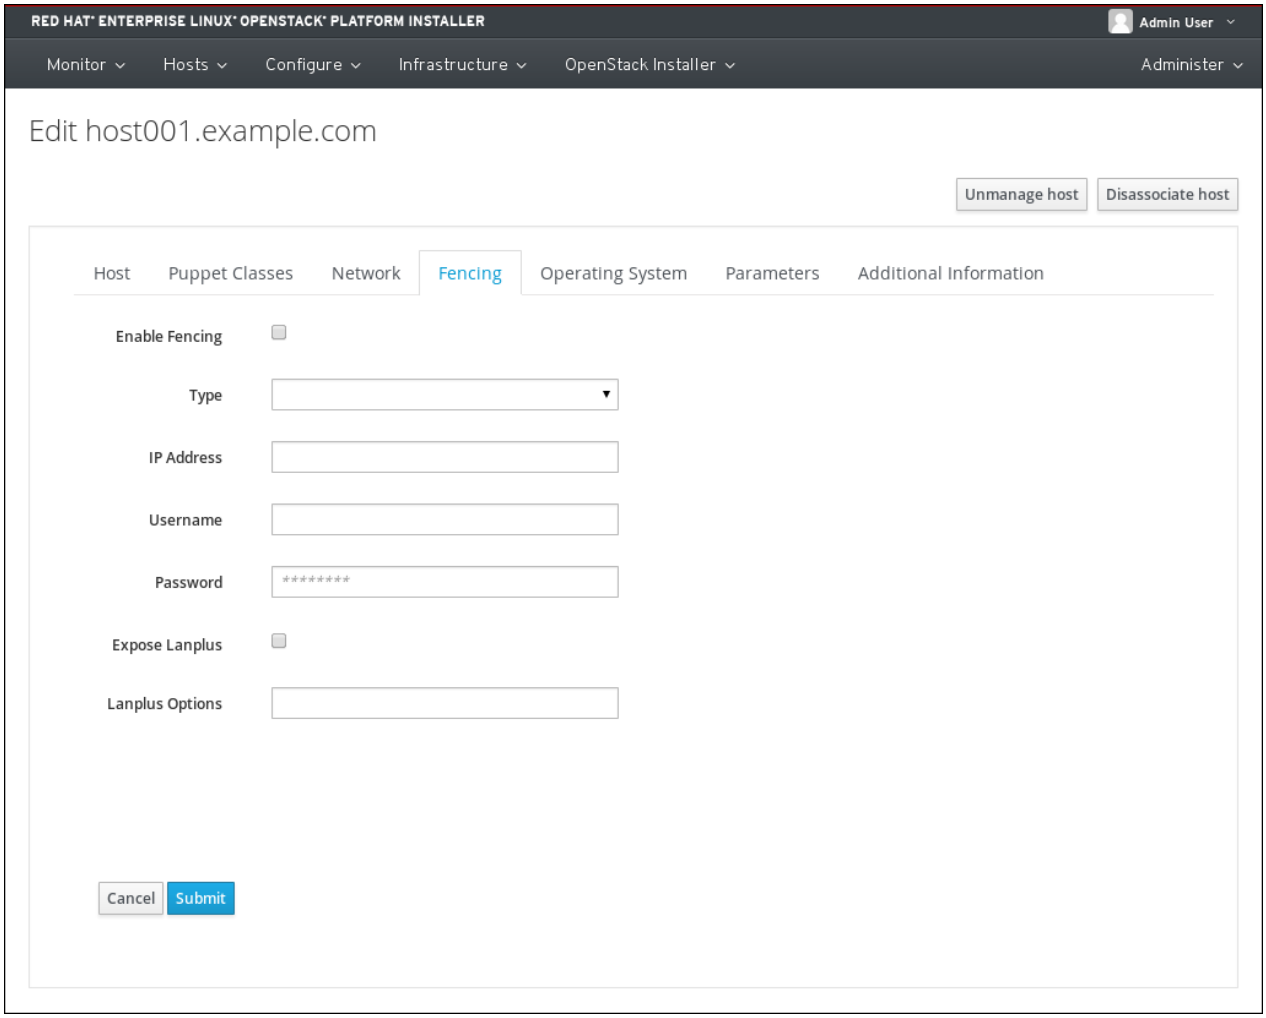

Figure 6.9. The Fencing tab

5. If the message **Fencing is disabled. To enable, first you need to add a BMC interface to the Host** is displayed, you must manually configure the details of a baseboard management controller network interface controller attached to the host:
    a. Click the **Network** tab.
    b. Click **Add Interface**.
    c. Select **BMC** from the **Type** list.
    d. Enter the MAC address for the network interface controller in the **MAC address** field.

e. Click the **Fencing** tab.

6. Select **Enable Fencing**.
7. Select **IPMI** from the **Type** list.
8. Enter the IP address of the fencing device in the **IP Address** field.
9. Enter the name of the IPMI user on the fencing device in the **Username** field.
10. Enter the password by which to authenticate the above user in the **Password** field.
11. Select **Expose Lanplus** to use the **lanplus** interface.
12. If you selected the **Expose Lanplus** check box, enter any options in the **Lanplus Options** field.
13. Click **Submit**.

You have configured fencing for the host, and the settings you specified are applied to the host when you provision RHEL OpenStack Platform.

6.5. Deployment Configurations

You can use the user interface to back up the configuration of an existing deployment or overwrite the values of an existing deployment based on a backed up configuration file.

Important

While it is possible to edit the values of the parameters for each of the services in a deployment, you are asked to input the values of key parameters for each of the services during the process of creating a deployment. All other parameters do not need to be manually configured.

6.5.1. Viewing the Configuration of a Deployment

You can view a summary of the values assigned to each parameter for the services in a deployment.

Procedure 6.8. Viewing the Configuration of a Deployment

1. Click **OpenStack Installer → Deployments**.
2. Click the name of a deployment.
3. Click **Advanced Configuration**.
4. Click the name of a service to view the parameters for that service.
5. Click **Back to Deployment** to return to the deployment details page when finished.

6.5.2. Editing Service Parameters

While the parameters for the services in a deployment are automatically generated when you create that deployment, you can manually edit those parameters to override the existing values.

1. Click **OpenStack Installer → Deployments**.
2. Click the name of a deployment.
3. Click **Advanced Configuration**.
4. Click **Edit**.
5. Click the name of a service to view the parameters for that service.
6. Enter a new value for the parameter in the text field for that parameter.
7. Click **Apply**

Your changes are saved, and the new values are applied to any RHEL OpenStack Platform environments you provision using the deployment.

6.5.3. Importing a Deployment Configuration

Import a deployment configuration file from your local system to specify the values of parameters for the services in a deployment.

Procedure 6.9. Importing a Deployment Configuration

1. Click **OpenStack Installer → Deployments**.
2. Click the name of a deployment to open the details page for that deployment.
3. Click **Advanced Configuration**.
4. Click **Import** to open the **Import Config** pop-up window.
5. Click **Browse** and search your local system for the configuration file to upload.
6. Click **Import Config**.

The selected configuration file is imported and the values of the parameters defined by the configuration file overwrite those defined for the deployment.

Important

Importing a configuration file only overwrites the values of parameters that exist both in the current list of parameters for the deployment and in the imported configuration file. All other parameters retain their originally defined value and must be manually updated.

6.5.4. Exporting a Deployment Configuration

Export the values of parameters for the services in a deployment to a YAML file on your local system.

Procedure 6.10. Exporting a Deployment Configuration

1. Click **OpenStack Installer → Deployments**.
2. Click the name of a deployment.
3. Click **Advanced Configuration**.

4. Click **Export** from the **Configuration** list.

A YAML file containing the values of parameters for each service in the deployment is downloaded to your local system.

6.6. Provisioning Red Hat Enterprise Linux OpenStack Platform

6.6.1. Provisioning Red Hat Enterprise Linux OpenStack Platform

Use a deployment to provision RHEL OpenStack Platform on one or more hosts.

Procedure 6.11. Provisioning Red Hat Enterprise Linux OpenStack Platform

1. Click **OpenStack Installer → Deployments**.
2. Click the name of the deployment to provision.

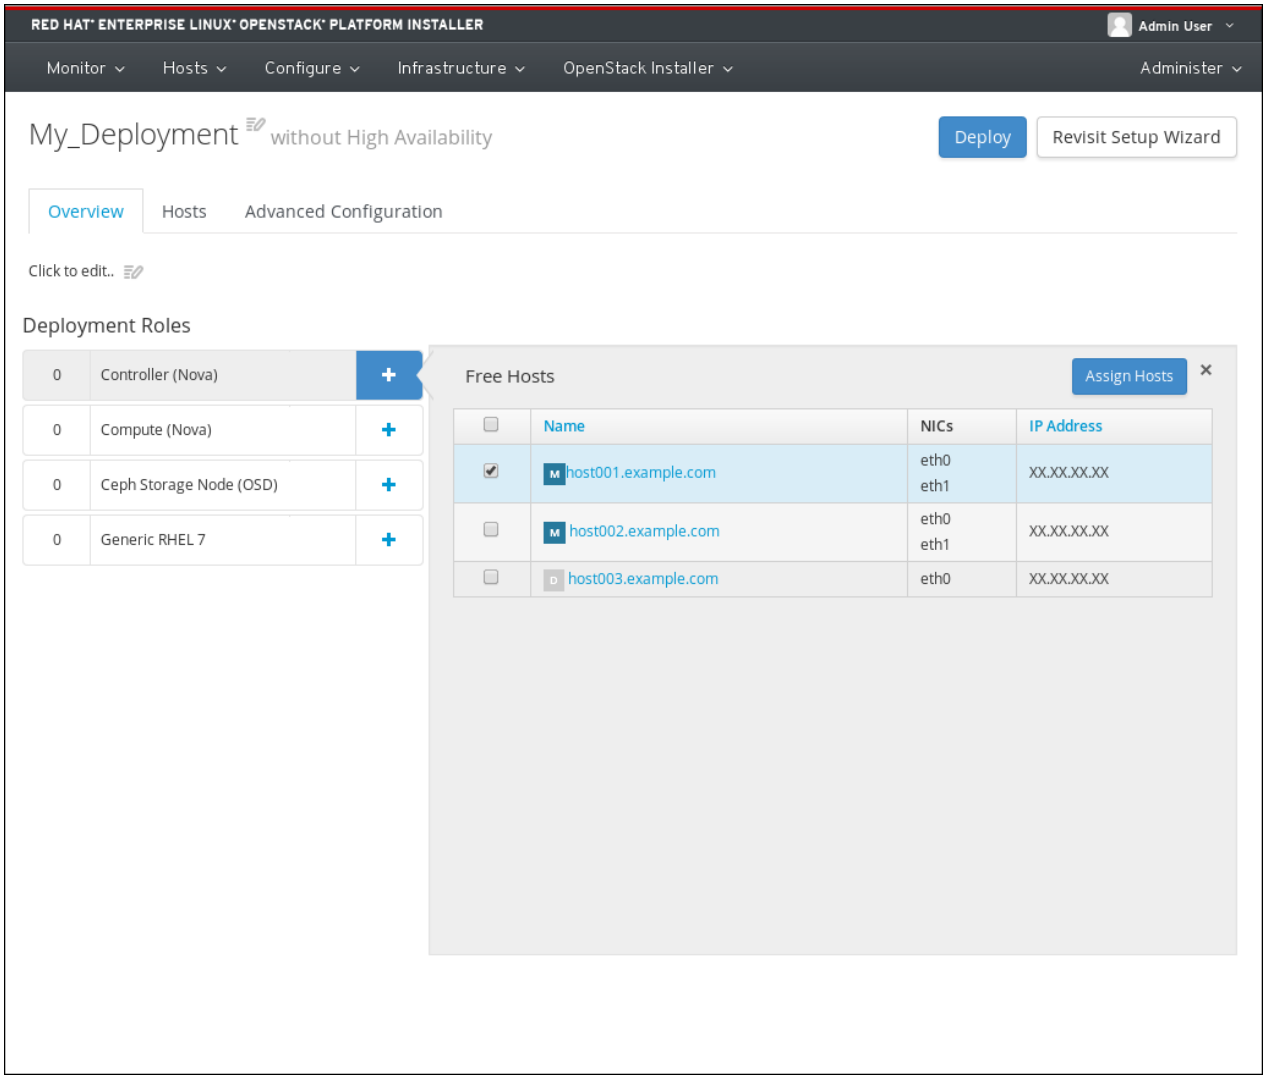

Figure 6.10. The deployment details page

3. Click **Deploy** to open the deployment confirmation screen.

Figure 6.11. The deployment confirmation screen

4. Click **Deploy**.

Hosts assigned to roles in the deployment are moved to the **production** environment, which contains the Puppet classes required to provision RHEL OpenStack Platform. The hosts are provisioned in accordance with the role to which they are assigned in the deployment, and the progress of provisioning is displayed.

6.6.2. Retrieving Service Details

After you have provisioned your RHEL OpenStack Platform environment, the installer provides a list of details for that environment such as service passwords and the location of the Horizon dashboard.

Procedure 6.12. Retrieving Service Details

1. Click **OpenStack Installer → Deployments**.
2. Click the name of a deployment you have provisioned.
3. Review the details of the deployment.

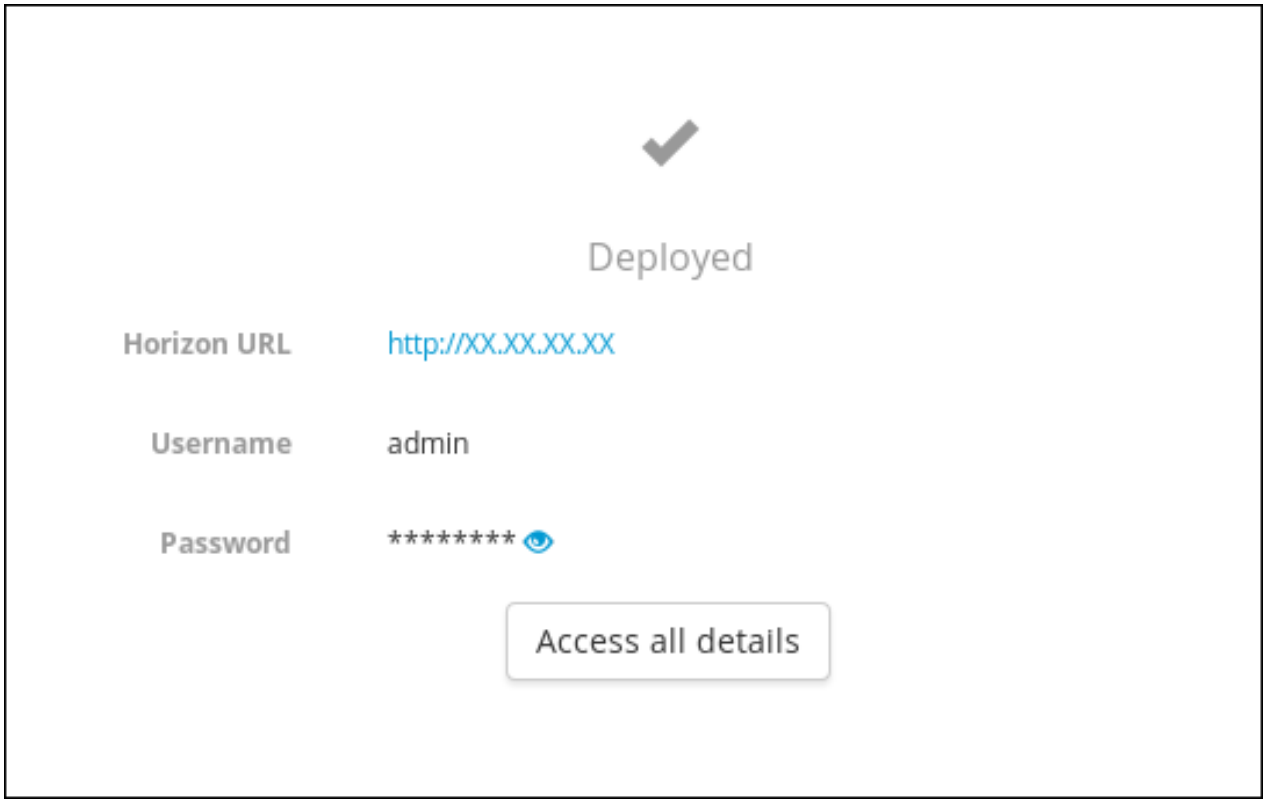

Figure 6.12. The service details page

4. Optionally, click the **Access all details** button to display the **All details** window, which displays a list of the passwords for all services in your environment.

6.6.3. Logging In

Once at least one controller node and one compute node are successfully installed and configured, you can access the user interface with a web browser. Replace *HOSTNAME* with the host name or IP address of the server acting as the controller:

- HTTPS

```
https://HOSTNAME/dashboard/
```

- HTTP

```
http://HOSTNAME/dashboard/
```

When prompted, log in using the credentials of the **admin** user. To begin using the OpenStack deployment, see the *Administration User Guide* and *Cloud Administrator Guide*.

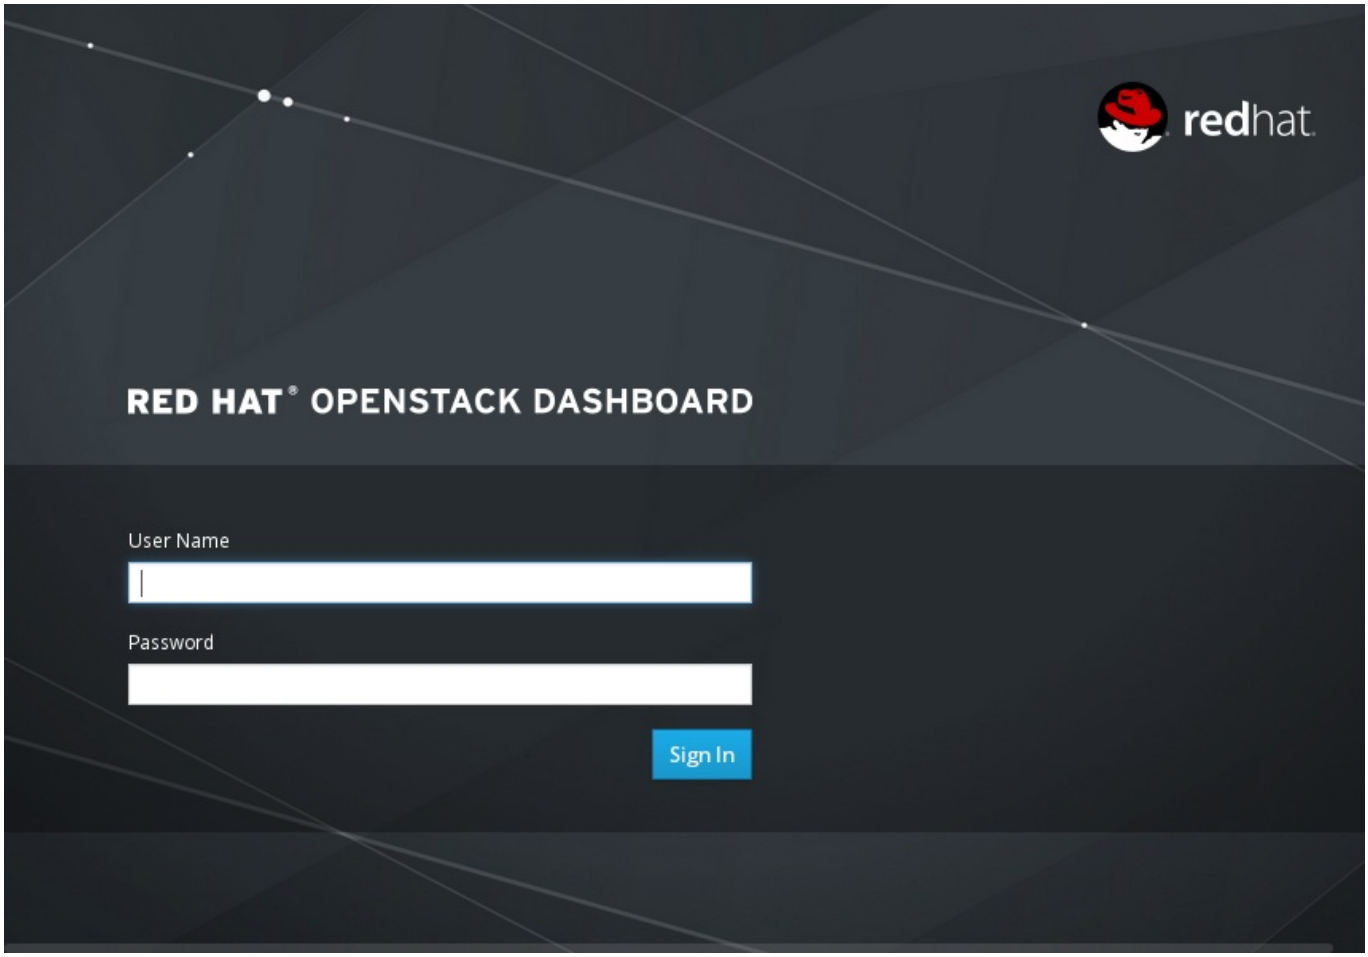

Figure 6.13. The Dashboard Login Screen

6.6.4. Next Steps

After you have provisioned your RHEL OpenStack Platform environment, all additional configuration of services in the environment must be performed individually on the nodes in the environment. See the following documentation in the RHEL OpenStack Platform documentation suite for information on how to get started working with your environment and performing additional configuration.

Important

If you change the values of settings that the installer configures when it provisions a host, the newly configured values are overwritten when a Puppet run is performed on the host. To prevent these changes from being overwritten, you must stop the **puppet** service on the host.

Administration User Guide

Provides an overview of how to create and manage resources in a RHEL OpenStack Platform environment using the Horizon dashboard or client commands for administrative users.

End User Guide

Provides an overview of how to create and manage resources in a RHEL OpenStack Platform environment using the Horizon dashboard or client commands for end users.

Cloud Administrator Guide

Provides an overview of the software that administrators can use to manage and troubleshoot a RHEL OpenStack Platform environment.

Configuration Reference Guide

Provides a reference for administrative users for looking up configuration options. This guide contains lists of the configuration options available with RHEL OpenStack Platform and uses auto-generation to generate options and the descriptions from the code for each project.

Chapter 7. Monitoring Hosts

The Red Hat Enterprise Linux OpenStack Platform installer user interface provides a number of functions for monitoring the hosts it manages. These functions are accessed via the **Monitor** menu in the main title bar.

7.1. Dashboard

The **Dashboard** page is the main screen displayed when you log in to the user interface, and presents an overview of the current status of all hosts that the installer manages. You can access the dashboard at any time from anywhere in the user interface by clicking **Monitor → Dashboard**.

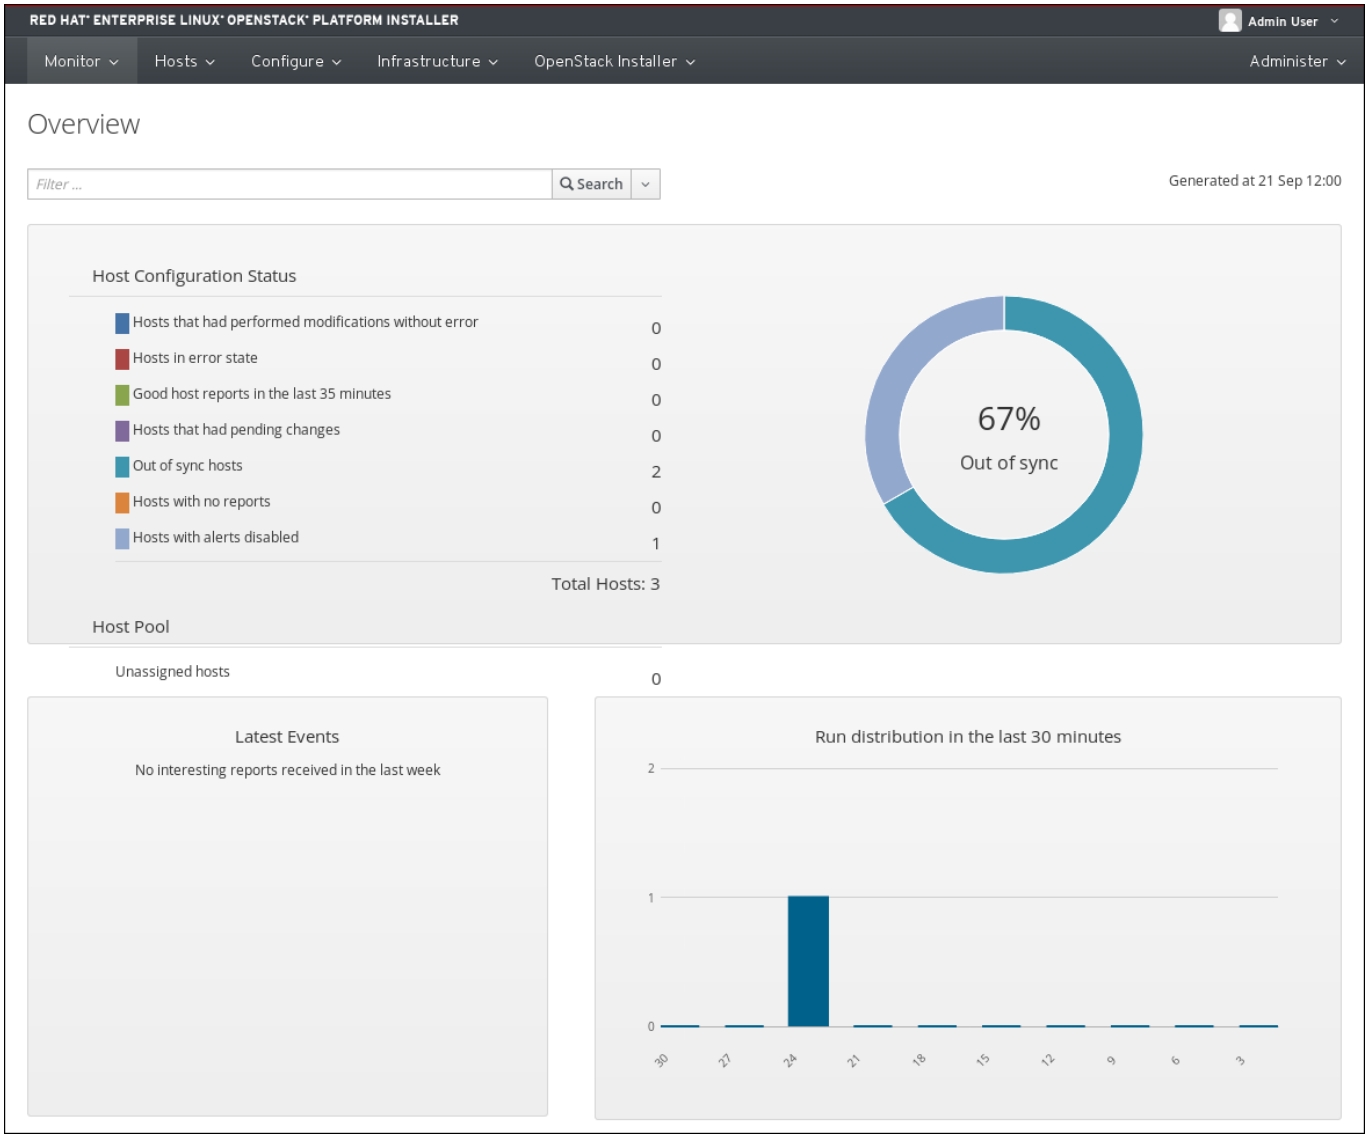

Figure 7.1. The Dashboard

There are three key elements to the dashboard:

Host Configuration Status

This section contains a table and a pie graph that outline the total number of hosts that the installer manages and the current state of each host. Hosts are categorized according to status, and you can click the title of a status type to display a list of all hosts currently reporting that status.

Latest Events

This section displays a list of any recent events that have been recorded in the installer over the most recent one-week period.

Run distribution in the last 30 minutes

This section contains a time series that displays the number of Puppet runs executed over the most recent report interval. By default, this interval is 30 minutes. The time series is broken into sub-intervals that represent 10% of the overall report interval length, each of which display the number of runs executed during that sub-interval.

7.2. Reports

The **Reports** page displays a list of reports of Puppet runs executed on the hosts that the installer manages. You can access this page at any time from anywhere in the user interface by clicking **Monitor → Reports**.

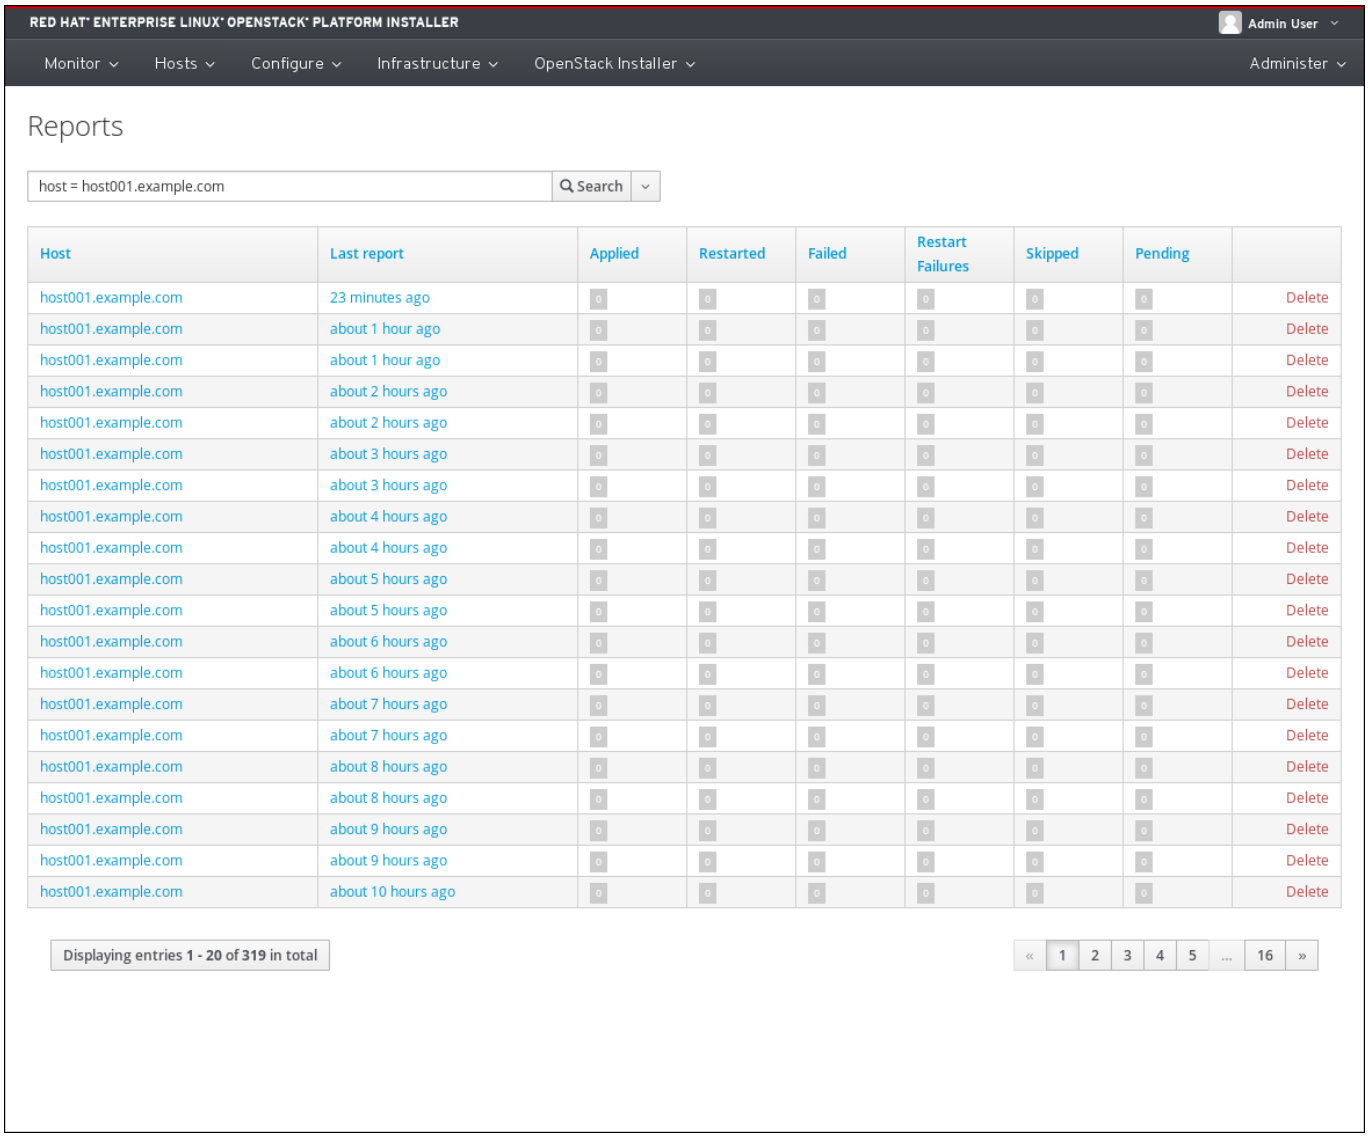

Figure 7.2. The Reports page

The reports table displays a list of the number of actions performed on individual hosts during a report interval. The time at which a report was collected is displayed in the **Last report** column, and actions are divided into a number of categories as outlined below:

- **Applied**: The number of changes that were applied to resources on the host, such as files, directories, or user accounts.
- **Restarted**: The number of resources, such as services, that were restarted.
- **Failed**: The number of actions that were executed but did not complete successfully.
- **Restart Failures**: The number of resources, such as services, that were scheduled to restart but did not restart successfully.
- **Skipped**: The number of actions that were scheduled, but were skipped due to a problem scheduling the Puppet run.
- **Pending**: The number of actions that were scheduled to be applied, but had not yet been performed at the time the report was collected.

7.2.1. Changing the Report Interval

Change the interval by which events are reported in the user interface.

Procedure 7.1. Changing the Report Interval

1. From the main title bar, click **Administer** → **Settings**.
2. Click the **Puppet** tab.
3. Click the edit icon next to the value of the **puppet_interval** setting.
4. Enter a new report interval in minutes.
5. Click **Save**.

You have changed the interval by which events are reported in the user interface, and information on hosts that the installer manages is collected in accordance with the newly specified interval.

7.2.2. Deleting Reports in a Batch

You can delete all reports older than a specific number of days if those reports are no longer needed.

Procedure 7.2. Deleting Reports

1. Log in to the machine on which the user interface is installed as the root user.
2. Delete all reports more than *[days]* number of days old:

```
# foreman-rake reports:expire days=[days]
```

All reports older than the specified number of days are deleted.

7.2.3. Deleting Reports Individually

You can manually delete individual reports that are no longer needed.

Procedure 7.3. Deleting a Report

1. From the title bar in the main screen of the user interface, click **Monitor** → **Reports**.
2. Click **Delete** for the report to delete.
3. Click **OK** when prompted.

7.3. Facts

The **Facts** page displays a list of details about each host that the installer manages, such as the MAC address, current IP address and the version of the operating system installed. These details are current to the time of the last report created for a given host. You can access the **Facts** page at any time from anywhere in the user interface by clicking **Monitor** → **Facts**.

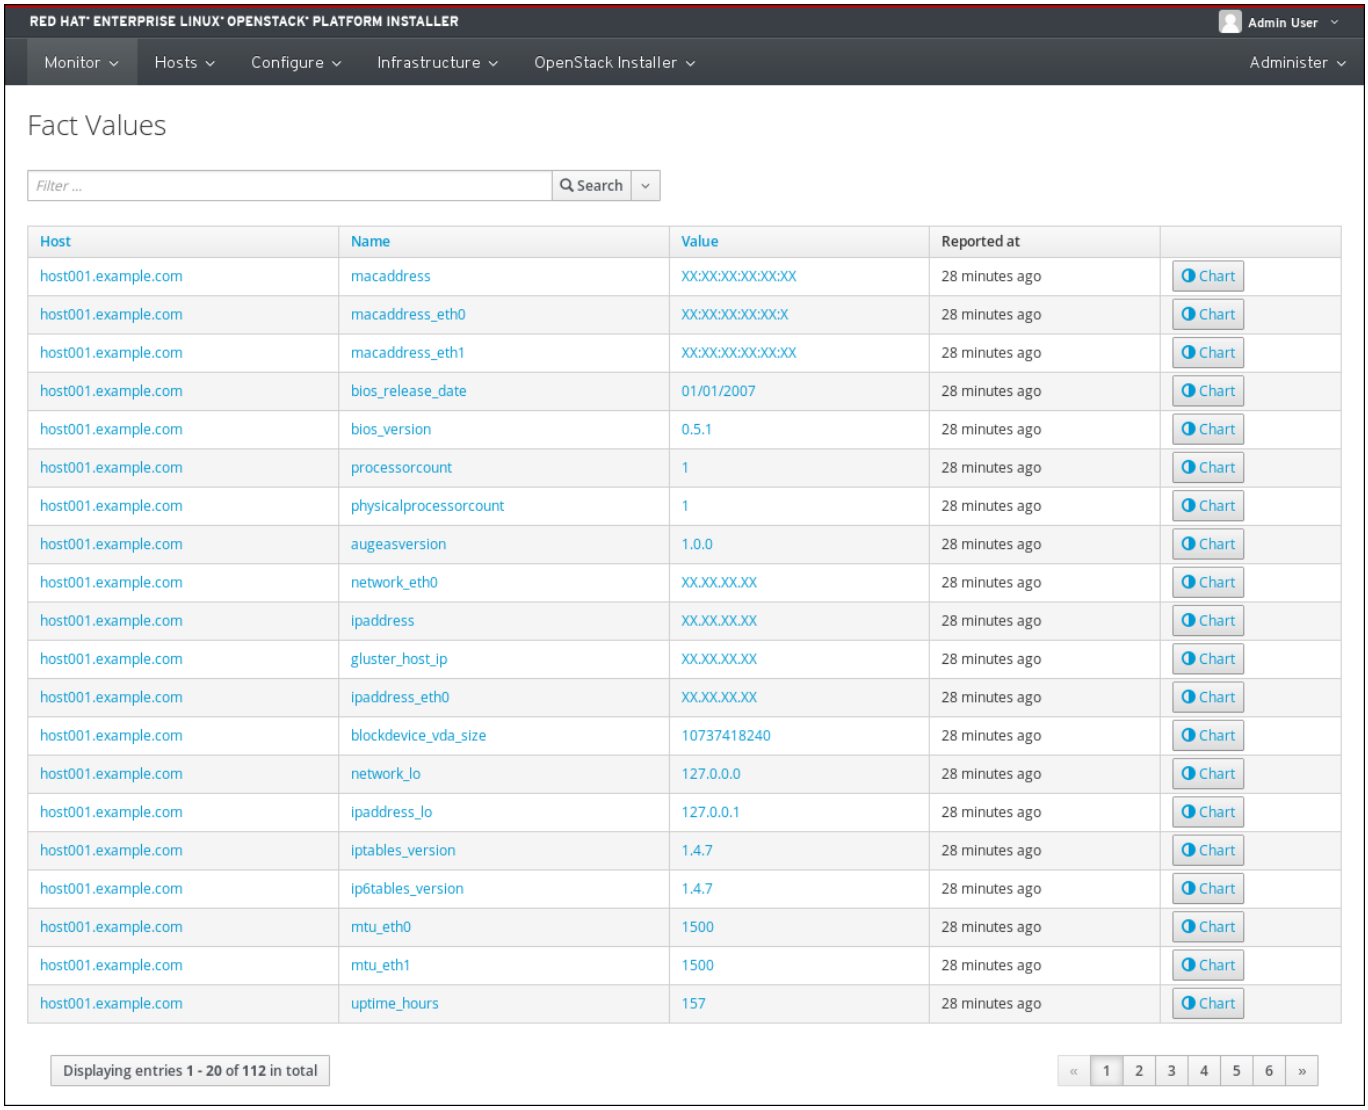

Figure 7.3. The Facts page

The **Facts** page includes the following columns:

- **Host**: The host that the fact describes.
- **Name**: The name of the fact.

- **Value**: The value of the fact.
- **Reported at**: The time at which the fact was collected.
- **Chart**: A button that displays a pie chart representing the relative frequency of the value of the given fact across all hosts that the installer manages.

7.4. Tasks

The **Tasks** page displays a list of details about current tasks being run by the installer. You can access the **Tasks** page at any time from anywhere in the user interface by clicking **Monitor** → **Tasks**.

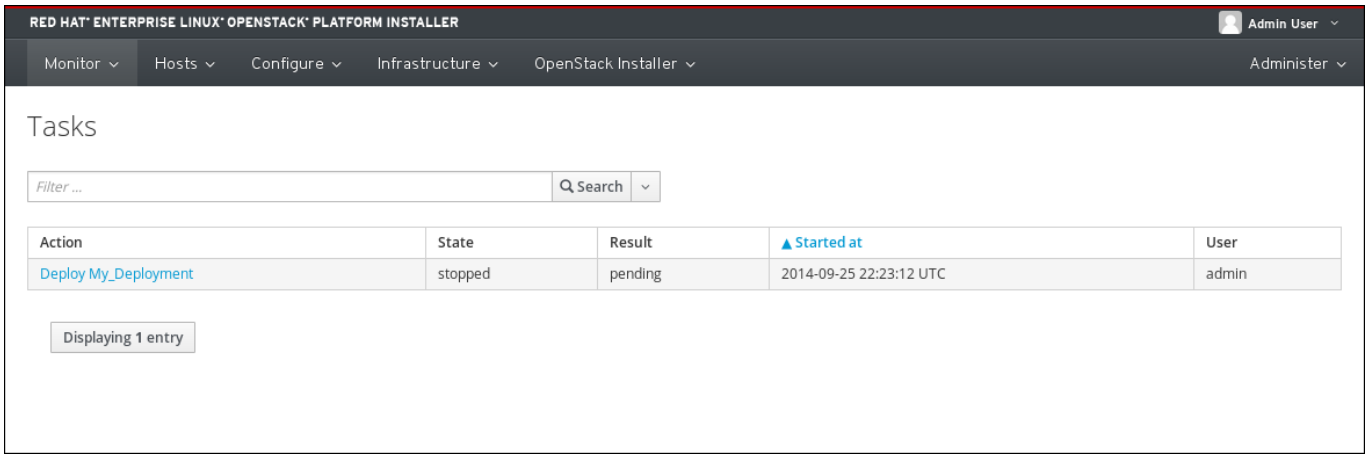

Figure 7.4. The Tasks page

The **Tasks** page includes the following columns:

- **Action**: The action being performed.
- **State**: The current state of the task.
- **Result**: The outcome of the task, if any.
- **Started at**: The time at which the task was initiated.
- **User**: The user in the user interface that initiated the task.

Appendix A. Firewall Rules

The following table outlines the firewall rules that the **`rhel-osp-installer`** command configures when you install the user interface. The installer uses these ports to communicate with and control other machines in the environment. This table is provided for your information; no further configuration is necessary beyond that provided by **`rhel-osp-installer`**.

Table A.1. Red Hat Enterprise Linux OpenStack Platform Installer Firewall Rules

Ports	Protocols	Service	Purpose
22	TCP	SSH	Connecting to other machines on the private network that the installer defines.
53	TCP, UDP	DNS	Resolving the host names and addresses of machines on the private network that the installer defines.
67	TCP	DHCP	Assigning IP addresses to machines on the private network that the installer defines.
69	TCP	TFTP	Enabling the PXE booting of machines on the private network that the installer defines.
80, 443	TCP	HTTP, HTTPS	The Apache web server for hosting the user interface for the installer.
8140	TCP	Puppet	Communication between Puppet clients and the Puppet master.

Appendix B. Configuring a Gateway

Configure the machine on which the Red Hat Enterprise Linux OpenStack Platform installer user interface is installed to act as a gateway so that traffic from the private provisioning network that the installer defines can be forwarded to a network interface with external connectivity.

Procedure B.1. Configuring a Gateway

1. Log in to the machine on which you will install the user interface as the root user.
2. Edit **/etc/sysctl.conf** and change the value of **net.ipv4.ip_forward** to **1**:

```
net.ipv4.ip_forward = 1
```

3. Load the new value:

```
# sysctl -p
```

4. Enable IP masquerading:

```
# iptables -t nat -A POSTROUTING -o [if_name] -j MASQUERADE
# iptables -A FORWARD -s [XX.XX.XX.XX/XX] -j ACCEPT
# iptables -A FORWARD -d [XX.XX.XX.XX/XX] -j ACCEPT
# iptables -A FORWARD ! -s [XX.XX.XX.XX/XX] -j DROP
```

   - *[if_name]*: The name of the network interface to which to forward network traffic. You must specify the name of the network interface that will not be used for the private provisioning network.
   - *[XX.XX.XX.XX/XX]*: The network address of the private provisioning network that the installer defines. You must specify this address using CIDR notation. For example, **XX.0.0.0/8**, **XX.XX.0.0/16**, or **XX.XX.XX.00/24**.

5. Save the changes to the firewall:

```
# service iptables save
```

6. Restart networking:

```
# service network restart
```

Traffic from the private provisioning network can now be forwarded to a network interface with external network connectivity, allowing machines on that network to access resources on other networks.

Appendix C. Configuring the NTP Server

The nodes in your Red Hat Enterprise Linux OpenStack Platform environment might require NTP configuration to ensure they all have correctly synchronized time. Perform the following operation after deploying a Red Hat Enterprise Linux OpenStack Platform environment.

Procedure C.1. Configuring the NTP Server

1. Log into the user interface.
2. Navigate to **Configure** → **Puppet Classes**.
3. In **Filter**, type **ntp** and click **Search**.
4. Click on **ntp** and navigate to **Smart Class Parameter** → **Servers**
5. Check **Override** and enter the NTP server's hostname or IP address in the **Default Value** field.
6. Click **Submit**.
7. Navigate to **Configure** → **Host Groups**.
8. For each Host Group:
    a. Click on the Host Group's name.
    b. Navigate to **Puppet Classes**.
    c. Navigate to the **ntp** Puppet class group and expand it.
    d. Select the **ntp** class. This class moves to the **Included Classes** list.
    e. Click **Submit**.

Each node configures its NTP service the next time it connects to the Installer.

Appendix D. Manual Procedures

The following procedures outline methods for manually configuring the user interface that have been replaced by new user interface elements or are not required to provision Red Hat Enterprise Linux OpenStack Platform environments using the installer, and are provided here for reference only.

D.1. Configuring Fencing on High-Availability Nodes

If you provision Red Hat Enterprise Linux OpenStack Platform 5.0 on Red Hat Enterprise Linux 7.2 with high availability, you must configure fencing on each of the controller nodes in the environment. In principle, you must configure fencing in the user interface before you provision your environment. However, if necessary, you can manually configure fencing on hosts after they have been provisioned via the following procedure.

Note

For more information about configuring high availability in Red Hat Enterprise Linux 7.2, see High Availability Add-On Reference.

Important

This procedure assumes you have a fencing device in your Red Hat Enterprise Linux OpenStack Platform environment that you can use to fence the nodes in that environment. All steps in the procedure must be performed on each controller node.

Procedure D.1. Configuring Fencing on High-Availability Nodes

1. Stop and disable the **puppet** service:

```
# systemctl stop puppet.service
# systemctl disable puppet.service
```

Note

Disabling the **puppet** service prevents future changes from being applied to the controller nodes. To apply any changes, you must enable the **puppet** service and apply the changes. You must then disable the **puppet** service again and repeat the final step in this procedure to reactivate fencing.

2. Configure IPMI:

   a. Ensure the **ipmi** service is running:

```
# systemctl start ipmi.service
```

   b. Configure and activate an IPMI user with administrative privileges:

```
# ipmitool user set name 2 user_name
# ipmitool user set password 2 password
# ipmitool user priv 1 4
# ipmitool user enable 2
```

3. Create a fencing agent:

```
# pcs stonith create fence_agent_name fence_ipmilan
login=user_name passwd=password ipaddr=ip_address
pcmk_host_list="host_name"
```

   - **fence_agent_name**: A human-readable name for the fencing agent.
   - **user_name**: The name of the IPMI user on the fencing device.
   - **password**: The password for the IPMI user on the fencing device.
   - **ip_address**: The IP address of the fencing device.
   - **host_names**: The host name of the node.

4. Confirm that the fencing agent has been correctly created:

```
# pcs status
```

5. Enable fencing:

```
# pcs property set stonith-enabled=true
```

You have created a fencing agent that uses a fencing device to fence the nodes in your Red Hat Enterprise Linux OpenStack Platform environment.

D.2. Adding a Host Manually

Add a host to the user interface by manually entering the details of that host. This procedure assumes the host is connected to the private network on which the installer provides the DHCP service.

Note

The following procedure only outlines the configuration of the basic options required to add a host manually to the user interface. All options are then automatically configured when you assign the host to a deployment role.

Procedure D.2. Adding a Host Manually

1. Click **Hosts** → **New host**.
2. Configure host details:
   a. Enter a name to represent the host in the user interface in the **Name** field.
   b. Leave the **Host Group** list blank.

c. Select **Bare Metal** from the **Deploy on** list.

d. Select **Discovery** from the **Environment** list.

e. Select the server to act as the Puppet CA for the host from the **Puppet CA** list.

f. Select a server to act as the initial puppet server for the host from the **Puppet Master** list.

3. Configure networking options:

a. Click the **Network** tab.

b. Select the domain to which the host is to belong from the **Domain** list.

c. Leave the **Realm** list blank.

d. Enter the MAC address of the host in the **MAC address** field.

e. Select the subnet to which the host belongs from the **Subnet** list. The **IP address** field is automatically populated based on the subnet you select.

4. Configure the Operating System

a. Click the **Operating System** tab.

b. Select the architecture of the host from the **Architecture** list.

c. Select the operating system to install on the host from the **Operating system** list.

d. Select the **Build mode** check box.

e. Select the installation media entry by which to install the selected operating system from the **Media** list.

f. Select the partitioning scheme to be applied to the host from the **Partition table** list.

g. Enter a password for the **root** user on the host in the **Root password** list.

5. Click **Submit**.

Appendix E. Revision History

Revision 5.0.0-32 **Thu Mar 17 2016** **Dan Macpherson**
Adding NTP configuration instructions

Revision 5.0.0-31 **Tue Nov 17 2015** **Dan Macpherson**
Updated 7.1 to 7.2

Revision 5.0.0-30 **Thu Apr 16 2015** **Dan Macpherson**
Updated 7.0 to 7.1

Revision 5.0.0-29 **Wed Feb 11 2015** **Andrew Dahms**
Corrected a description of the product version.

Revision 5.0.0-28 **Wed Feb 11 2015** **Andrew Dahms**
Updated the link to the article outlining supported guest operating systems.

Revision 5.0.0-27 **Mon Dec 15 2014** **Andrew Dahms**
BZ#1154528 - Added a sample deployment configuration.

Revision 5.0.0-26 **Wed Nov 12 2014** **Andrew Dahms**
BZ#1145375 - Added information about configuring a Red Hat Enterprise Linux OpenStack Platform after provisioning.

Revision 5.0.0-25 **Tue Nov 11 2014** **Andrew Dahms**
BZ#1161978 - Updated the procedure for downloading a Red Hat Enterprise Linux 7.0 ISO file.
BZ#1149070 - Updated the reference to the supported virtual machine operating systems.

Revision 5.0.0-24 **Thu Nov 6 2014** **Andrew Dahms**
BZ#1160918 - Added information on how to manually add a BMC network interface controller.

Revision 5.0.0-23 **Thu Nov 6 2014** **Andrew Dahms**
BZ#1157916 - Added further information on the options available when installing the user interface for the installer.

Revision 5.0.0-22 **Fri Oct 31 2014** **Andrew Dahms**
BZ#1157915 - Added a procedure outlining how to configure the default installation media entry.
BZ#1155999 - Added information on how to configure bonding on host network interfaces.
BZ#1155950 - Updated the procedure for creating deployments.
BZ#1154463 - Added instructions on creating subnets when creating a deployment.
BZ#1150362 - Updated the minimum operating systems requirements for the installer.
BZ#1149070 - Updated the description of supported virtual machine operating systems.

Revision 5.0.0-21 **Mon Oct 27 2014** **Andrew Dahms**
BZ#1156285 - Added a note regarding the welcome page for the installer user interface.
BZ#1154527 - Added a section on the host overview page.
BZ#1145381 - Added a procedure outlining how to edit host group parameters.
BZ#1145380 - Added information on hardware requirements for nodes.
BZ#1145377 - Added further information about the options available when creating a deployment.
BZ#1145376 - Added further information on adding hosts via discovery.

Revision 5.0.0-20 **Thu Oct 16 2014** **Andrew Dahms**

BZ#1151302 - Added a section outlining browser requirements for accessing the user interface.
BZ#1145387 - Added a chapter on how to use the Monitor menu.

Revision 5.0.0-19 Thu Oct 9 2014 Andrew Dahms

Added a note outlining the current status of support for upgrading the installer.

Revision 5.0.0-18 Tue Oct 7 2014 Andrew Dahms

BZ#1147729 - Added a description of network traffic types.

Revision 5.0.0-17 Tue Sep 30 2014 Andrew Dahms

BZ#1147729 - Updated the procedure for creating deployments and configuring host networking.
BZ#1147730 - Updated the procedure for configuring fencing.
BZ#1147731 - Updated the procedure for assigning hosts to deployment roles.
BZ#1134413 - Added a procedure outlining how to configure a gateway.

Revision 5.0.0-16 Fri Sep 26 2014 Andrew Dahms

BZ#1145396 - Added a description of common operating system entry parameters.
BZ#1145384 - Updated the installation media example and references to user interface elements.
BZ#1145383 - Added a procedure outlining how to use the New Host button.
BZ#1145381 - Added a procedure outlining how to edit service parameters.
BZ#1134413 - Added a procedure outlining how to configure a machine to act as a gateway.

Revision 5.0.0-15 Mon Sep 22 2014 Andrew Dahms

BZ#1135644 - Updated the procedure for preparing an installation medium.

Revision 5.0.0-14 Tue Sep 16 2014 Andrew Dahms

BZ#1141984 - Outlined the new randomly generated password for the default administrative user account.
BZ#1139308 - Added a note that specifying a Subscription Manager entitlement pool is recommended.
BZ#1135644 - Added further details on preparing an installation medium.
BZ#1133722 - Updated the Subscription Manager parameters that can be edited when installing the user interface.

Revision 5.0.0-13 Tue Sep 16 2014 Andrew Dahms

BZ#1139450 - Updated the content on logging in to the user interface.
BZ#1139123 - Updated the topic outlining how to add hosts via discovery.
BZ#1136156 - Added further detail to the procedure for installing the user interface.
BZ#1136134 - Updated the topic on provisioning Red Hat Enterprise Linux OpenStack Platform.

Revision 5.0.0-12 Mon Aug 25 2014 Andrew Dahms

BZ#1138097 - Updated the list of Subscription Manager parameters that users can edit.

Revision 5.0.0-11 Mon Aug 25 2014 Andrew Dahms

BZ#1127044 - Updated the subtitle and abstract.
BZ#1125040 - Collapsed the list of commands for enabling channels into a single screen section.
BZ#1122178 - Updated the note regarding domain name requirements.
BZ#1113270 - Updated the note regarding operating system requirements.

Revision 5.0.0-10 Thu Aug 21 2014 Andrew Dahms

Final revision for publication.

Revision 5.0.0-9 Thu Aug 21 2014 Andrew Dahms

BZ#1132713 - Updated notes regarding features and the procedure for logging in to the web user interface.
BZ#1130139 - Added a topic outlining how to configure fencing on high availability nodes.
BZ#1127997 - Updated the content on general service details.
BZ#1127209 - Added a section on how to create and remove partition tables.

Revision 5.0.0-8 Thu Aug 14 2014 Andrew Dahms

BZ#1130340 - Added notes to the installation procedure to clarify the steps required to enable provisioning.
BZ#1122178 - Added a note regarding fully qualified domain name requirements.

Revision 5.0.0-7 Fri Aug 8 2014 Andrew Dahms

BZ#1120304 - Added a note outlining domain membership requirements.
BZ#1091512 - Updated the name of the channel that provides Foreman packages.

Revision 5.0.0-6 Thu Aug 7 2014 Andrew Dahms

BZ#1127484 - Added a section outlining how to add hosts using discovery.
BZ#1125040 - Updated the list of channels required to install Red Hat Enterprise Linux OpenStack Platform Installer.

Revision 5.0.0-5 Wed Aug 6 2014 Andrew Dahms

BZ#1127044 - Updated the titles of all cited OpenStack material to reflect the latest titles.
BZ#1126247 - Added a section outlining how to add and work with operating system entries.
BZ#1126238 - Added a section outlining firewall requirements.
BZ#1125041 - Updated the steps involved in running the rhel-osp-installer command.

Revision 5.0.0-4 Tue Jul 22 2014 Andrew Dahms

BZ#1125126 - Changed references to Foreman to Red Hat Enterprise Linux OpenStack Platform Installer.
BZ#1125106 - Added sections outlining how to set up users, user groups, and roles.
BZ#1125028 - Removed references to the method of installing the provisioning portal using a script
BZ#1061978 - Added a chapter outlining how to provision Red Hat Enterprise Linux OpenStack Platform.

Revision 5.0.0-3 Fri Jul 11 2014 Scott Radvan

Prepare guide for Technical Review.

Revision 5.0.0-2 Wed Jul 9 2014 Andrew Dahms

BZ#1113270 - Added an introduction to the Foreman installer and a note on support for Red Hat Enterprise Linux 7.0.
BZ#1081201 - Updated the procedure on configuring installation media.

Revision 5.0.0-1 Thu Jun 26 2014 Don Domingo

Red Hat Enterprise Linux OpenStack 5.0 test build.