SUSE Cloud

1.0

www.suse.com

December 03, 2012

Deployment Guide

Deployment Guide

List of Authors: Tanja Roth, Frank Sundermeyer

Contents

About This Guide v

1 Available Documentation v

2 Feedback vi

3 Documentation Conventions vii

4 About the Making of This Manual vii

1 The SUSE Cloud Architecture 1

1.1 The Administration Server 2

1.2 The Controller Node 3

1.3 The Compute Nodes 4

1.4 The Storage Nodes 4

2 Considerations and Requirements 5

2.1 Network 5

2.2 Product and Update Repositories 14

2.3 Storage 15

2.4 SSL Encryption 19

2.5 Hardware Requirements 20

2.6 Summary: Considerations and Requirements 22

3 Installing and Configuring the Administration Server 25

3.1 Operating System Installation 26

3.2 Post-Installation Configuration 32

4 Installing the OpenStack Nodes 41
4.1 Preparations 41
4.2 Node Installation 43
4.3 Post-Installation Configuration 44
4.4 Editing Allocated Nodes 46

5 Deploying the OpenStack Services 47
5.1 Barclamp 48
5.2 Deploying the Database 50
5.3 Deploying Keystone 50
5.4 Deploying Swift (optional) 51
5.5 Deploying Ceph (optional, unsupported) 52
5.6 Deploying Glance 54
5.7 Deploying Nova 55
5.8 Deploying the Nova Dashboard 56
5.9 How to Proceed 57

6 Troubleshooting and Support 59
6.1 FAQ 59
6.2 Support 62

A Log Files 63
A.1 On the Administration Server 63
A.2 On All Other Crowbar Nodes 64
A.3 On the Controller Node 64
A.4 On Compute Nodes 65
A.5 On Nodes with Ceph Barclamp 65

Terminology 67

About This Guide

SUSE® Cloud is an open source software solution that provides the fundamental capabilities to deploy and manage a cloud infrastructure based on SUSE Linux Enterprise. SUSE Cloud is powered by OpenStack, the leading community-driven, open source cloud infrastructure project. It seamlessly manages and provisions workloads across a heterogeneous cloud environment in a secure compliant, and fully-supported manner. The product tightly integrates with other SUSE technologies and with the SUSE maintenance and support infrastructure.

This guide provides cloud operators with the information needed to deploy and maintain SUSE Cloud administrative units, the Administration Server, and the Controller Node, as well as the Compute and Storage Nodes. The Administration Server provides all services needed to manage and deploy all other nodes in the cloud. The Controller Node hosts all OpenStack services needed to operate virtual machines deployed on the Compute Nodes in the SUSE Cloud. Each virtual machine (instance) started in the cloud will be hosted on one of the Compute Nodes. Object storage is managed by the Storage Nodes.

Many chapters in this manual contain links to additional documentation resources. These include additional documentation that is available on the system as well as documentation available on the Internet.

For an overview of the documentation available for your product and the latest documentation updates, refer to http://www.suse.com/documentation/suse_cloud10.

1 Available Documentation

The following manuals are available for this product:

Deployment Guide (page i)
: Gives an introduction to the SUSE® Cloud architecture and describes how to set up, deploy, and maintain the individual components.

User Guide for Administrators (↑*User Guide for Administrators*)
Guides you through management of projects and users, images, flavors, and quotas with SUSE Cloud Dashboard or the command line interface.

End User Guide (↑*End User Guide*)
Describes how to launch instances, manage volumes, and track usage.

HTML versions of the product manuals can be found in the installed system under `/usr/share/doc/manual`. Find the latest documentation updates at `http://www.suse.com/documentation` where you can download the manuals for your product in multiple formats.

2 Feedback

Several feedback channels are available:

Bugs and Enhancement Requests
For services and support options available for your product, refer to `http://www.suse.com/support/`.

To report bugs for a product component, log in to the Novell Customer Center from `http://www.suse.com/support/` and select *My Support* > *Service Request*.

User Comments
We want to hear your comments about and suggestions for this manual and the other documentation included with this product. Use the User Comments feature at the bottom of each page in the online documentation or go to `http://www.suse.com/documentation/feedback.html` and enter your comments there.

Mail
For feedback on the documentation of this product, you can also send a mail to `doc-team@suse.de`. Make sure to include the document title, the product version, and the publication date of the documentation. To report errors or suggest enhancements, provide a concise description of the problem and refer to the respective section number and page (or URL).

3 Documentation Conventions

The following typographical conventions are used in this manual:

- `/etc/passwd`: directory names and filenames
- *`placeholder`*: replace *`placeholder`* with the actual value
- `PATH`: the environment variable PATH
- `ls`, `--help`: commands, options, and parameters
- `user`: users or groups
- Alt, Alt + F1: a key to press or a key combination; keys are shown in uppercase as on a keyboard
- *File*, *File* > *Save As*: menu items, buttons
- This paragraph is only relevant for the architectures `amd64`, `em64t`, and `ipf`. The arrows mark the beginning and the end of the text block.

  This paragraph is only relevant for the architectures `System z` and `ipseries`. The arrows mark the beginning and the end of the text block.
- *Dancing Penguins* (Chapter *Penguins*, ↑Another Manual): This is a reference to a chapter in another manual.

4 About the Making of This Manual

This book is written in Novdoc, a subset of DocBook (see `http://www.docbook.org`). The XML source files were validated by `xmllint`, processed by `xsltproc`, and converted into XSL-FO using a customized version of Norman Walsh's stylesheets. The final PDF is formatted through XEP from RenderX.

1 The SUSE Cloud Architecture

SUSE Cloud is a cloud infrastructure solution that can easily be deployed and managed. It offers a cloud management solution that helps organizations to centralize virtual machine deployment. SUSE Cloud 1.0 provides the following features:

- Open source software that is based on the OpenStack Essex release.
- Centralized resource tracking providing insight into activities and capacity of the cloud infrastructure for optimized automated deployment of services.
- A self-service portal allowing end users to configure and deploy services as necessary, also offering the ability to track resource consumption (Nova Dashboard).
- An image repository allowing to publish standardized, pre-configured virtual machines (Glance).
- Automated installation processes via Crowbar utilizing predefined scripts for configuring and deploying Compute and Storage Nodes.
- Multi-tenant, role-based provisioning and access control enabling provisioning for multiple departments and users within your organization.
- APIs allowing to integrate third-party software, such as identity management and billing solutions.

SUSE Cloud is based on SUSE Linux Enterprise Server, OpenStack, Crowbar, and Chef. SUSE Linux Enterprise Server is used as the underlying operating system for all cloud infrastructure machines (also called nodes), whereas OpenStack, the cloud management layer, works as the “Cloud Operating System”. Crowbar and Chef are used to

automatically deploy and manage the OpenStack nodes from a central Administration Server.

Figure 1.1: *SUSE Cloud Infrastructure*

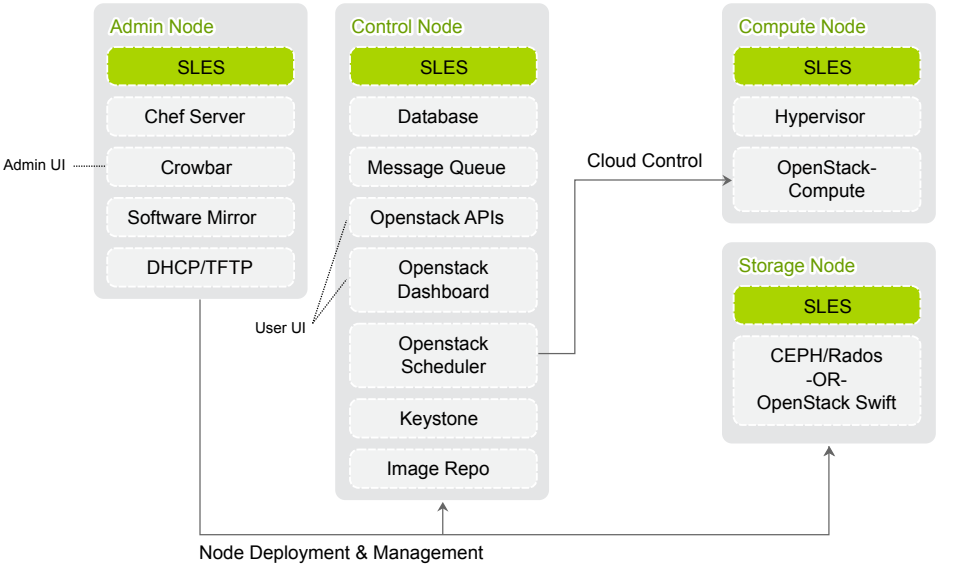

SUSE Cloud is deployed to four different types of machines:

- one Administration Server for node deployment and management
- one Controller Node hosting the cloud management services
- several Compute Nodes on which the instances are started
- several Storage Nodes for block and object storage

1.1 The Administration Server

The Administration Server provides all services needed to manage and deploy all other nodes in the cloud. Most of these services are provided by the Crowbar tool that automates in conjunction with Chef all the required installation and configuration tasks. Among the services provided by the server are DHCP, DNS, NTP, PXE, TFTP.

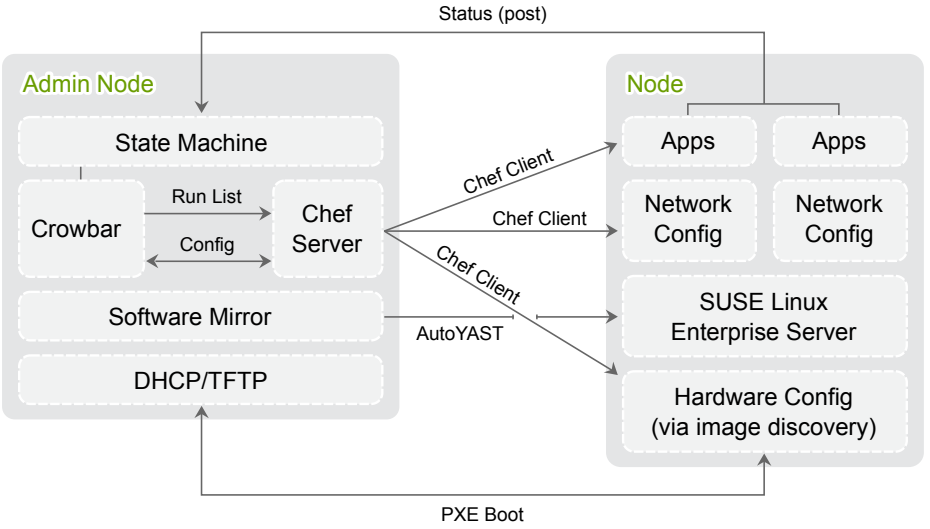

The Administration Server also hosts the software repositories for SUSE Linux Enterprise Server and SUSE Cloud, since they are needed for node deployment. Optionally (if no other sources for the software repositories are available) it can also host the Subscription Management Tool (SMT), providing up-to-date repositories with updates and patches for all nodes.

1.2 The Controller Node

The Controller Node hosts all OpenStack services needed to orchestrate virtual machines deployed on the Compute Nodes in the SUSE Cloud. OpenStack on SUSE Cloud uses a PostgreSQL database. It is managed and deployed through the Administration Server. The following OpenStack components and dependencies run on the Controller Node:

- PostgreSQL database
- Image (Glance) for managing virtual images
- Identity (Keystone), providing authentication and authorization for all OpenStack services

- Dashboard (Horizon), providing the Dashboard, which is a user Web interface for the OpenStack services
- Nova API and scheduler
- Message broker (RabbitMQ)

In addition to that, the management parts of the compute and storage services also run on the Controller Node.

1.3 The Compute Nodes

The Compute Nodes are the pool of machines on which the instances are running. These machines need to be equipped with a sufficient number of CPUs and enough RAM to start several instances. The Controller Node effectively distributes instances within the pool of Compute Nodes and provides the necessary network resources. The OpenStack service Compute (Nova) runs on the Compute Nodes and provides means for setting up, starting, and stopping virtual machines.

1.4 The Storage Nodes

The Storage Nodes are the pool of machines providing storage. SUSE Cloud offers two different types of storage: object and block storage. Object storage is provided by the OpenStack Swift component, while block storage is provided by Nova Volume.

Nova Volume can use different backends. By default it will use the LVM backend with iSCSI on the Controller Node. Alternatively, it can use Ceph's RADOS Block Device (RBD). As of SUSE Cloud 1.0, Ceph is not officially supported but rather included as a technical preview, so using Nova Volume without Ceph is recommended.

Deploying Swift is optional.

Considerations and Requirements

2

Before deploying SUSE Cloud, there are a few requirements to be met and considerations to be made. Make sure to thoroughly read this chapter—some decisions need to be made *before* deploying SUSE Cloud, since you cannot change them afterwards.

2.1 Network

SUSE Cloud requires a complex network setup consisting of several networks that are configured during installation. These networks are for exclusive cloud usage. In order to access them from an existing network, a router is needed.

The network configuration on the nodes in the SUSE Cloud network is entirely controlled by Crowbar. Any network configuration not done with Crowbar (e.g. with YaST) will automatically be overwritten. Once the cloud is deployed, network settings cannot be changed anymore!

Figure 2.1: *SUSE Cloud Network: Overview*

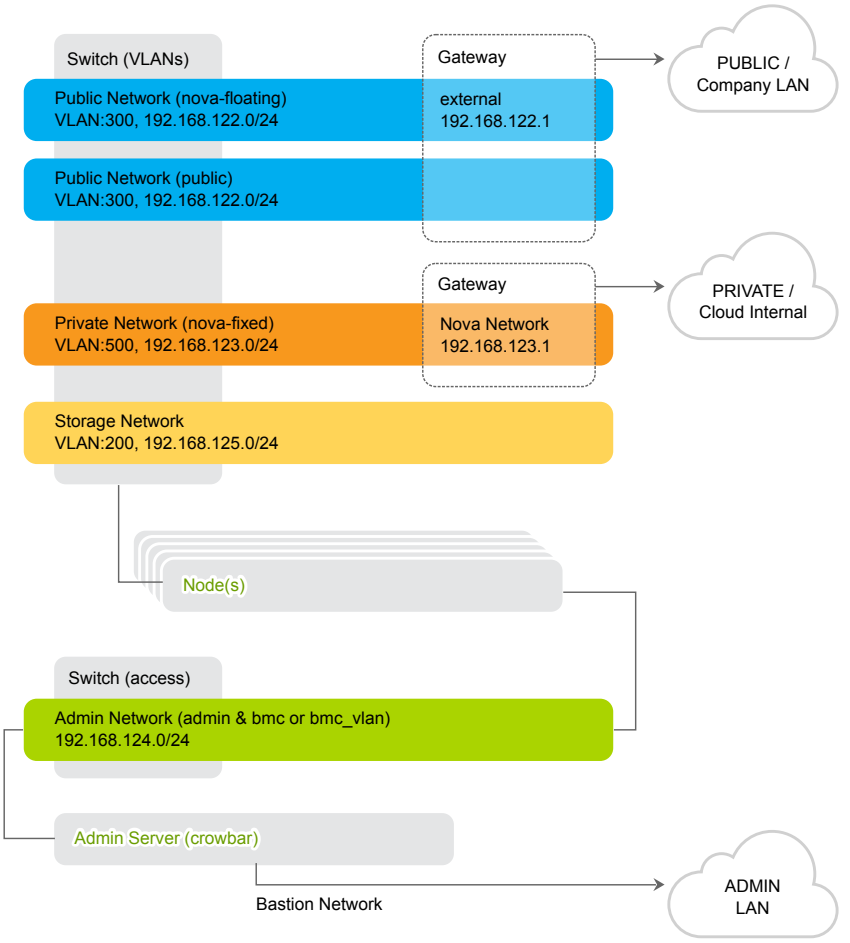

The following networks are defined when setting up SUSE Cloud. The IP addresses listed are the default addresses and can be changed using the YaST Crowbar module (see Section 3.1.9, “Crowbar Setup” (page 30)).

Admin Network (`192.168.124/24`)
: A private network to access the Administration Server and all nodes for administration purposes. The default setup lets you also access the BMC (Baseboard Management Controller) data via IPMI (Intelligent Platform Management Interface) from this network. If required, BMC access can be swapped to a separate network.

To access this network, you have the following options:

- do not allow access from the outside and keep the admin network completely separated
- allow access from a single network (e.g. your company's administration network) via the “bastion network” option configured on an additional network card with a fixed IP address
- allow access from one or more networks via a gateway

Storage Network (`192.168.125/24`)
: Private, SUSE Cloud internal virtual network. This network is used by Ceph, and Swift, only. It should not be accessed by users.

Private Network (nova-fixed, `192.168.123/24`)
: Private, SUSE Cloud internal virtual network. This network is used for inter-instance communication only. The gateway required is also automatically provided by SUSE Cloud.

Public Network (nova-floating, public, `192.168.122/24`)
: The only public network provided by SUSE Cloud. You can access the Nova Dashboard as well as instances (provided they have been equipped with a floating IP) on this network. This network can only be accessed via a gateway, which needs to be provided externally. All SUSE Cloud users and administrators need to be able to access the public network.

NOTE: No IPv6 support

As of SUSE Cloud 1.0, IPv6 is not supported. This applies to the cloud internal networks as well as to the instances.

The following diagram shows the SUSE Cloud network in more detail. It demonstrates how the OpenStack nodes and services use the different networks.

Figure 2.2: *SUSE Cloud Network: Details*

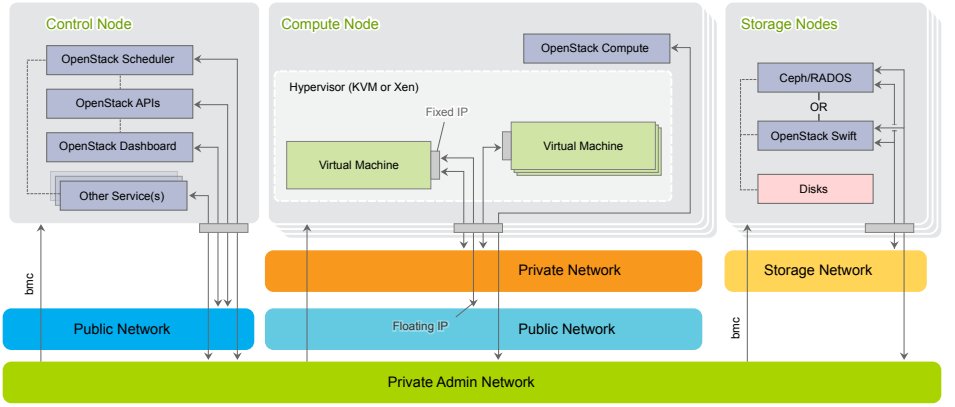

2.1.1 Network Address Allocation

The default networks set up in SUSE Cloud are class C networks with 256 IP addresses each. This limits the maximum number of instances that can be started simultaneously. Addresses within the networks are allocated as outlined in the following table. Use the YaST Crowbar module to customize (see Section 3.1.9, “Crowbar Setup” (page 30)). The `.255` address for each network is always reserved as the broadcast address. This assignment cannot be changed.

Table 2.1: *192.168.124.0/24 (Admin/BMC) Network Address Allocation*

Function	Address	Remark
router	`192.168.124.1`	Provided externally.
admin	`192.168.124.10` - `192.168.124.11`	Fixed addresses reserved for the Administration Server.
dhcp	`192.168.124.21` - `192.168.124.80`	Address range reserved for node allocation/installation. Determines the maximum number of parallel allocations/installations.

Function	Address	Remark
host	192.168.124.81 - 192.168.124.160	Fixed addresses for the OpenStack nodes. Determines the maximum number of OpenStack nodes that can be deployed.
bmc vlan host	192.168.124.161	Fixed address for the BMC VLAN.
bmc host	192.168.124.162 -192.168.124.240	Fixed addresses for the OpenStack nodes. Determines the maximum number of OpenStack nodes that can be deployed.
switch	192.168.124.241 -192.168.124.250	

Table 2.2: *192.168.125/24 (Storage) Network Address Allocation*

Function	Address	Remark
host	192.168.125.10 - 192.168.125.239	

Table 2.3: *192.168.123/24 (Private Network/nova-fixed) Network Address Allocation*

Function	Address	Remark
router	192.168.123.1 - 192.168.123.49	Each Compute Node also acts as a router for “it's” instances, getting an address from this range assigned. This effectively limits the maximum number of Compute Nodes that can be deployed with SUSE Cloud to 49.
dhcp	192.168.123.50 - 192.168.123.254	Address range for instances.

Table 2.4: *192.168.122/24 (Public Network nova-floating, public) Network Address Allocation*

Function	Address	Remark
public host	`192.168.122.2 - 192.168.122.49`	Public address range for external SUSE Cloud services such as the Dashboard or the API.

2.1.2 Network Modes

SUSE Cloud supports different network modes: single, dual, and teaming. As of SUSE Cloud 1.0 the networking mode is applied to all nodes as well as the Administration Server. That means that all machines need to meet the hardware requirements for the chosen mode. The network mode can be configured using the YaST Crowbar module (Section 3.1.9, “Crowbar Setup” (page 30)). The network mode cannot be changed once the cloud is deployed.

Other, more flexible network mode setups can be configured by manually editing the Crowbar network configuration files. See the documentation on the Crowbar wiki (`https://github.com/dellcloudedge/crowbar/wiki`) for more information. SUSE can assist you in creating a custom setup within the scope of a Level 3 support contract.

2.1.2.1 Single Network Mode

In single mode you just use one ethernet card for all the traffic:

Single Mode

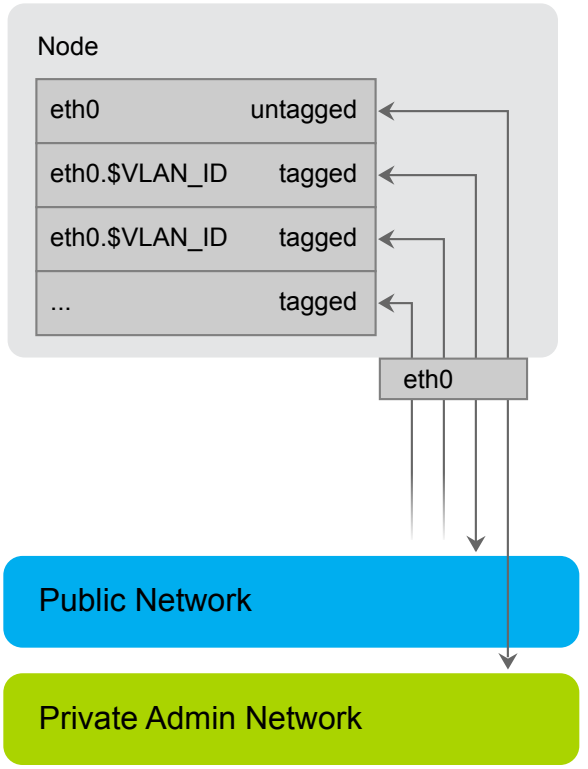

2.1.2.2 Dual Network Mode

Dual mode needs two ethernet cards and allows you to completely separate traffic to/from the Admin Network and to/from the public network:

Dual Mode

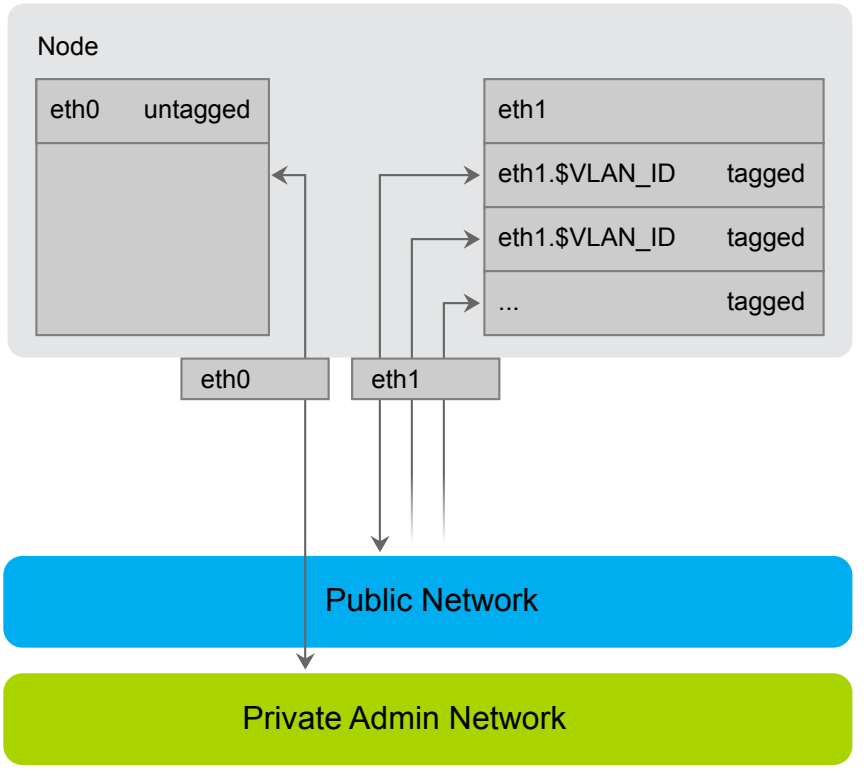

2.1.2.3 Teaming Network Mode

Teaming mode is almost identical to single mode, except for the fact that you combine several ethernet cards to a “bond” in order to increase the performance. Teaming mode needs two or more ethernet cards.

Team Mode

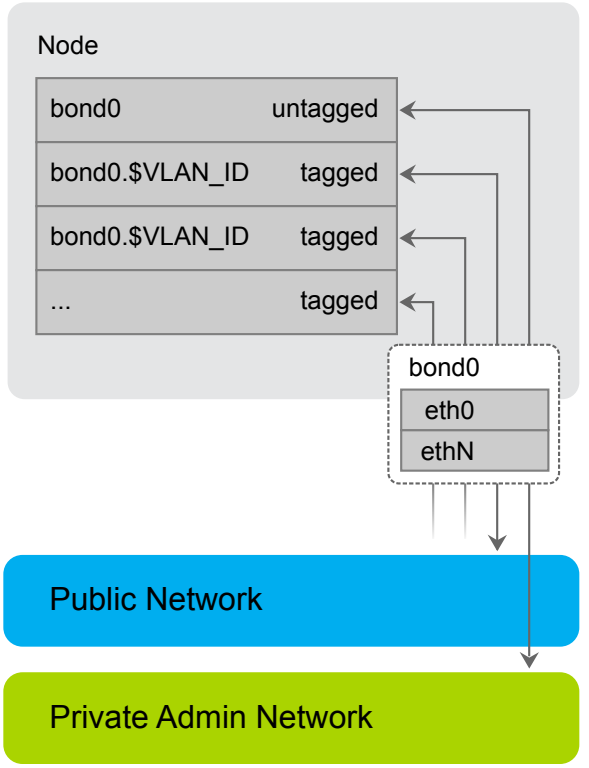

2.1.3 Accessing the Admin Network via a Bastion Network

If you want to allow to access the cloud's admin network from another network, you can do so by providing an external gateway. This option offers maximum flexibility, but requires additional machines and may be less secure than you require. Therefore SUSE Cloud offers a second option for accessing a single external network (e.g. a dedicated server administration network): the bastion network. You just need a dedicated ethernet card and a static IP address from the external network to set it up.

2.1.4 DNS and Hostnames

The Administration Server acts as a name server for all nodes in the cloud. If you allow access to the admin network from outside, you may want to add additional name servers to your network setup prior to deploying SUSE Cloud. If additional name servers are found on cloud deployment, the name server on the Administration Server will automatically be configured to forward requests for non-local records to those servers.

The Administration Server needs to be configured to have a full qualified hostname. This hostname must not be changed after SUSE Cloud has been deployed. The OpenStack nodes will be named after their MAC address by default, but you can provide aliases which are easier to remember when allocating the nodes. The aliases for the OpenStack nodes can be changed any time. It is useful to have a list of MAC addresses and the intended use of the corresponding host at hand when deploying the OpenStack nodes.

2.2 Product and Update Repositories

The Administration Server as well as the OpenStack nodes need to get security updates and patches for the operating system (SUSE Linux Enterprise Server) as well as for SUSE Cloud itself. Furthermore product repositories for SUSE Linux Enterprise Server and SUSE Cloud are needed as an installation source for the OpenStack nodes. In SUSE Cloud the Administration Server is designed to work as the single source for all repositories.

While the product repositories do not change, the update repositories need to be regularly updated. Depending on your network setup there are several possibilities to provide up-to-date repositories on the Administration Server:

"Sneakernet"
: If you choose to completely seal off your admin network from all other networks, you need to manually update the repositories from removable media.

Installing a Subscription Management Tool (SMT) Server on the Administration Server
: The SMT server, a free add-on product for SUSE Linux Enterprise Server, regularly synchronizes repository data from Novell Customer Center with your local host. Installing the SMT server on the Administration Server is recommended if you do

not have access to update repositories from elsewhere within your organization. This option requires the Administration Server to be able to access the Internet. Subscription Management Tool 11 SP2 is available from http://www.novell.com/linux/smt/.

Utilizing Existing Repositories
: If you can access existing repositories from within your company network from the Administration Server, you can either mount or sync these repositories to the required locations on the Administration Server.

As of SUSE Cloud 1.0, the following update repositories need to be mirrored:

- SLES11-SP2-Core
- SLES11-SP2-Updates
- SLES11-SP1-Pool
- SLES11-SP1-Updates
- SUSE-Cloud-1.0-Pool
- SUSE-Cloud-1.0-Updates
- SLES11-SMT-SP2-Pool (only needed when the SMT is installed)
- SLES11-SMT-SP2-Updates (only needed when the SMT is installed)

In addition to the update repositories you also need to mirror the contents of the product media (SUSE Linux Enterprise Server 11 SP2 and SUSE Cloud 1.0) to your local disk.

2.3 Storage

When talking about “storage” on SUSE Cloud, there are two completely different aspects to discuss: the block and object storage services SUSE Cloud offers on the one hand and the hardware related storage aspects on the different node types.

2.3.1 Cloud Storage Services

As mentioned above, SUSE Cloud offers two different types of storage services: object and block storage. Object storage lets you upload and download files (similar to an FTP server), whereas a block storage provides mountable devices (similar to a hard-disk partition). Furthermore SUSE Cloud provides a repository to store the virtual disk images used to start instances.

Object Storage with Swift
: The object OpenStack storage service is called Swift. Swift needs to be deployed on dedicated nodes where no other cloud services run. In order to be able to store the objects redundantly, it is required to deploy at least two Swift nodes. SUSE Cloud is configured to always use all unused disks on a node for storage. Offering object storage with Swift is optional.

Block Storage
: Block storage on SUSE Cloud is provided by Nova Volume. By default Nova Volume uses an LVM backend with iSCSI. This default setup utilizes a single device on the Controller Node—using a RAID for this purpose is strongly recommended.

  Alternatively, Nova Volume can use Ceph RBD as a backend. Ceph offers data security and speed by storing the devices redundantly on different servers. Ceph needs to be deployed on dedicated nodes where no other cloud services run. In order to be able to store the objects redundantly, it is required to deploy at least two Ceph nodes. You can configure which devices Ceph uses for storage.

  IMPORTANT: Ceph not Supported

  As of SUSE Cloud 1.0, Ceph is not officially supported but rather included as a technical preview, so using Nova Volume instead is recommended.

The Glance Image Repository
: Glance provides a catalog and repository for virtual disk images used to start the instances. Glance is usually installed on the Controller Node. The image repository resides in a directory on the file system by default—it is recommended to mount a partition or volume to that directory.

2.3.2 Storage Hardware Requirements

Apart from sufficient disk space to install the SUSE Linux Enterprise Server operating system, each node in SUSE Cloud has to store additional data.Requirements and recommendations for the various node types are listed below.

IMPORTANT: Choose a Hard Disk for the Operating System Installation

The operating system will always be installed on the *first* hard disk, the one that is recognized as `/dev/sda`. This is the disk that is listed *first* in the BIOS, the one from which the machine will boot. If you have nodes with a certain hard disk you want the operating system to be installed on, make sure it will be recognized as the first disk.

2.3.2.1 Administration Server

If you store the update repositories directly on the Administration Server (see Section 2.2, "Product and Update Repositories" (page 14) for details), it is recommended to mount `/srv` to a separate partition or volume with a minimum of 30 GB space.

2.3.2.2 Controller Node

The virtual disk image repository resides under `/var/lib/glance/images` by default. It is recommended to mount a separate partition or volume into this directory that provides enough space to host all virtual disk images needed.

Unless deploying the Ceph RDB service as a backend for Nova Volume (which is currently not supported, but included as a technical preview), it uses LVM with iSCSI on the Controller Node. This setup allows to use only one device, therefore it is highly recommended to provide a RAID with sufficient disk space.

2.3.2.3 Compute Nodes

Unless an instance is started via "Boot from Volume", it is started with at least one disk—a copy of the image from which it has been started. Depending on the flavor you start, the instance may also have a second, so-called "ephemeral" disk. The size of the root disk depends on the image itself, while ephemeral disks are always created as

sparse image files that grow (up to a defined size) when being “filled”. By default ephemeral disks have a size of 10 GB.

Both disks, root images and ephemeral disk, are directly bound to the instance and are deleted when the instance is terminated. Therefore these disks are bound to the Compute Node on which the instance has been started. The disks are created under `/var/lib/nova` on the Compute Node. Your Compute Nodes should be equipped with enough disk space to store the root images and ephemeral disks.

NOTE: Ephemeral Disks vs. Block Storage

Do not confuse ephemeral disks with persistent block storage. In addition to an ephemeral disk, which is automatically provided with most instance flavors, you can optionally add a persistent storage device provided by Nova Volume. Ephemeral disks are deleted when the instance terminates, while persistent storage devices can be reused in another instance.

The maximum disk space required on a compute node depends on the available flavors. A flavor specifies the number of CPUs, as well as RAM and disk size of an instance. Several flavors ranging from *tiny* (1 CPU, 2512 MB RAM, no ephemeral disk) to *xlarge* (8 CPUs, 8 GB RAM, 10 GB ephemeral disk) are available by default. Adding custom flavors as well as editing and deleting existing flavors is also supported.

To calculate the minimum disk space needed on a compute node, you need to determine the highest "RAM to disk space" ratio from your flavors. Example:

Flavor small: 2 GB RAM, 100 GB ephemeral disk => 50 GB disk /1 GB RAM
Flavor large: 8 GB RAM, 200 GB ephemeral disk => 25 GB disk /1 GB RAM

So, 50 GB disk /1 GB RAM is the ratio that matters. If you multiply that value by the amount of RAM in GB available on your compute node, you have the minimum disk space required by ephemeral disks. Pad that value with sufficient space for the root disks plus a buffer that enables you to create flavors with a higher RAM to disk ratio in the future.

WARNING: Overcommitting Disk Space

The scheduler that decides in which node an instance is started does not check for available disk space. If there is no disk space left on a compute node, this will not only cause data loss on the instances, but the compute node itself will

also stop operating. Therefore you must make sure all compute nodes are equipped with enough hard disk space!

2.3.2.4 Storage Nodes

The block-storage service Ceph and the object storage service Swift need to be deployed onto dedicated nodes—it is not possible to mix these services. Each storage service requires at least two machines (more are recommended) to be able to store data redundantly.

Each Ceph/Swift Storage Node needs at least two hard disks. The first one (`/dev/sda`) will be used for the operating system installation, while the others can be used for storage purposes. While you can configure which devices Ceph uses for storage, Swift always uses all devices.

IMPORTANT: Ceph not Supported

As of SUSE Cloud 1.0, Ceph is not officially supported but rather included as a technical preview, so using Nova Volume instead is recommended.

2.4 SSL Encryption

Whenever non-public data travels over a network it needs to be encrypted. Encryption protects the integrity and confidentiality of data. Therefore you should enable SSL support when deploying SUSE Cloud to production (it is not enabled by default). The following services (and their APIs if available) can make use of SSL:

- Keystone
- Glance
- Nova
- VNC
- Nova Dashboard

Each service requires valid certificates signed by a trusted third party. You may either use the same certificates for all services or use dedicated certificates for each service. See http://www.suse.com/documentation/sles11/book_sle_admin/data/sec_apache2_ssl.html for instructions on how to create certificates and get them signed by a trusted organization.

2.5 Hardware Requirements

Precise hardware requirements can only be listed for the Administration Server and the OpenStack Controller Node. The requirements of the OpenStack Compute and Storage Nodes depends on the number of concurrent instances and their virtual hardware equipment.

The minimum number of machines required for a SUSE Cloud setup featuring all services is seven: one Administration Server, one Controller Node, one Compute Node, and four Storage Nodes. In addition to that, a gateway providing access to the public network is required.

IMPORTANT: Physical Machines and Architecture

All SUSE Cloud nodes need to be physical machines. Although the Administration Server and the Controller Node can be virtualized in test environments, this is not supported for production systems.

SUSE Cloud currently only runs on `x86_64` hardware.

2.5.1 Administration Server

- Architecture: x86_64
- RAM: at least 2 GB, 4 GB recommended
- Hard disk: at least 40 GB. It is recommended to put `/srv` on a separate partition with at least 30 GB space, unless you mount the update repositories from another server (see Section 2.2, “Product and Update Repositories” (page 14) for details).

- Number of network cards: 1 for single mode, 2 for dual mode, 2 or more for team mode. Additional networks such as the bastion network and/or a separate BMC network each need an additional network card. See Section 2.1, “Network” (page 5) for details.

2.5.2 Controller Node

- Architecture: x86_64
- RAM: at least 1 GB, 2 GB recommended
- Number of network cards: 1 for single mode, 2 for dual mode, 2 or more for team mode. See Section 2.1, “Network” (page 5) for details.
- Hard disk: See Section 2.3.2.2, “Controller Node” (page 17).

2.5.3 Compute Node

The Compute Nodes need to be equipped with a sufficient amount of RAM and CPUs, matching the numbers required by the maximum number of instances running concurrently. An instance started in SUSE Cloud cannot share resources from several physical nodes, but rather uses the resources of the node on which it was started. So if you offer a flavor (see Flavor (page 68) for a definition) with 8 CPUs and 12 GB RAM, at least one of your nodes should be able to provide these resources.

See Section 2.3.2.3, “Compute Nodes” (page 17) for storage requirements.

2.5.4 Storage Node

The Storage Nodes are sufficiently equipped with a single CPU and 1 or 2 GB of RAM. See Section 2.3.2.4, “Storage Nodes” (page 19) for storage requirements.

2.5.5 Software Requirements

The following software requirements need to be met in order to install SUSE Cloud:

- SUSE Linux Enterprise Server 11 SP2 installation media (ISO image, included in the SUSE Cloud Administration Server subscription)
- Access to the SUSE Linux Enterprise Server 11 SP2 Update repositories (either by registering SUSE Linux Enterprise Server 11 SP2 or via an existing SMT server).
- SUSE Cloud installation media (ISO image).
- A SUSE/Novell account (for product registration and SMT setup). If you do not already have one, go to http://www.suse.com/login to create it.
- Optional: Subscription Management Tool 11 SP2 installation media. A free download is available on http://www.novell.com/linux/smt/. See Section 2.2, “Product and Update Repositories” (page 14).

2.6 Summary: Considerations and Requirements

As outlined above, there are some important considerations to be made before deploying SUSE Cloud. The following briefly summarizes what was discussed in detail in this chapter. Keep in mind that as of SUSE Cloud 1.0 it is not possible to change some aspects such as the network setup once SUSE Cloud is deployed!

Network

- If you do not want to stick with the default networks and addresses, define custom networks and addresses. You need four different networks, at least three of them VLANs. If you need to separate the admin and the BMC network, a fifth network is required. Class C networks are sufficient. See Section 2.1, “Network” (page 5) for details.
- Determine how to allocate addresses from your network. Make sure not to allocate IP addresses twice. See Section 2.1.1, “Network Address Allocation” (page 8) for the default allocation scheme.
- Define which network mode to use. Keep in mind that all machines within the cloud (including the Administration Server) will be set up with the chosen mode and

therefore need to meet the hardware requirements. See Section 2.1.2, “Network Modes” (page 10) for details.

- Define how to access the admin and BMC network(s): no access from the outside (no action is required), via an external gateway (gateway needs to be provided), or via bastion network. See Section 2.1.3, “Accessing the Admin Network via a Bastion Network” (page 13) for details.

- Provide a gateway to access the public network (public, nova-floating).

- Make sure the admin server's hostname is correctly configured (`hostname -f` needs to return a full qualified hostname).

- Prepare a list of MAC addresses and the intended use of the corresponding host for all OpenStack nodes.

Update Repositories

- Depending on your network setup you have different options on how to provide up-to-date update repositories for SUSE Linux Enterprise Server and SUSE Cloud on the Administration Server: Sneakernet, installing Subscription Management Tool, syncing data with an existing repository, or mounting remote repositories. Choose the option that best matches your needs.

Storage

- Decide whether you want to deploy the object storage service Swift. If so, you need to deploy at least two nodes with sufficient disk space exclusively dedicated to Swift.

- Decide whether to use Nova Volume with Ceph as backend for block storage (not supported). If deploying Ceph, you need to deploy at least two nodes with sufficient disk space exclusively dedicated to it. If you choose not to deploy Ceph and use the default setup for Nova Volume (recommended), your Controller Node needs to be equipped with additional disk space (a RAID is strongly recommended).

IMPORTANT: Ceph not Supported

As of SUSE Cloud 1.0, Ceph is not officially supported but rather included as a technical preview, so using Nova Volume instead is recommended.

- Optionally, provide a volume for storing the Glance image repository. Doing so is recommended.
- Make sure all nodes are equipped with sufficient hard disk space.

SSL Encryption

- Decide whether to use different SSL certificates for the services and the API or whether to use a single certificate.
- Get one or more SSL certificates certified by a trusted third party source.

Hardware and Software Requirements

- Make sure the hardware requirements for the different node types are met.
- Make sure to have all required software at hand.

Installing and Configuring the Administration Server

3

Deploying and installing SUSE Cloud is a multi-step process, starting by deploying a basic SUSE Linux Enterprise Server installation and the SUSE Cloud add-on product to the Administration Server. Now the product and update repositories need to be set up and the SUSE Cloud network needs to be configured. Next the Administration Server setup will be finished. Once the Administration Server is ready, you can start deploying and configuring the OpenStack nodes. The complete node deployment is done automatically via Crowbar and Chef from the Administration Server. All you need to do is to PXE boot the nodes and to deploy the OpenStack services to them.

Procedure 3.1: *High Level Overview of the SUSE Cloud Installation*

1. Install SUSE Linux Enterprise Server 11 SP2 on the Administration Server with the Add-On products Subscription Management Tool (optional) and SUSE Cloud. See below.

2. Once the Administration Server is set up, PXE boot all nodes onto which the OpenStack components should be deployed and allocate them in the Crowbar Web interface to start the automatic SUSE Linux Enterprise Server installation. See Chapter 4, *Installing the OpenStack Nodes* (page 41).

3. Configure and deploy the OpenStack services via the Crowbar Web interface or command line tools. See Chapter 5, *Deploying the OpenStack Services* (page 47).

4. When all OpenStack services are up and running, SUSE Cloud is ready. The cloud admin can now upload images to enable users to start deploying instances. See *User Guide for Administrators* (↑*User Guide for Administrators*).

In this chapter you will learn how to install and set up the Administration Server from bare metal. As a result, the Administration Server will be ready to deploy OpenStack nodes and services. It will run on SUSE Linux Enterprise Server 11 SP2 and will include the add-on products SUSE Cloud and SMT (optional). Installing the Administration Server involves the following basic steps:

3.1 Operating System Installation

Start the installation by booting from the SUSE Linux Enterprise Server 11 SP2 installation medium.

NOTE: Differences from the Default Installation Process

For an overview of a default SUSE Linux Enterprise Server installation, refer to the SUSE Linux Enterprise Server *Installation Quick Start*. Detailed installation instructions are available in the SUSE Linux Enterprise Server *Deployment Guide*. Both documents are available at http://www.suse.com/documentation/sles11/.

The following sections will only cover the differences from the default installation process.

3.1.1 Add-On Product Selection

Installing the Add-On products SUSE Cloud and SMT (optional) during the SUSE Linux Enterprise Server installation is recommended. Make sure to be able to access the installation media (DVD or ISO image). Alternatively, install the add-on products after the SUSE Linux Enterprise Server installation.

If you have access to remote update repositories for SUSE Linux Enterprise Server and SUSE Cloud from the cloud's admin network, you may want to skip the SMT Add-On product installation. Please refer to Section 2.2, “Product and Update Repositories” (page 14) for details.

On the *Installation Mode* screen, click *Include Add-On products from Separate Media*. Proceed with *Next* to the Add-On product installation dialog. If you have direct access to the installation media (for example, via DVD or USB stick), skip the network instal-

lation dialog. Otherwise configure the network as described in Section 3.1.7, “Basic Network Configuration” (page 29). Add SUSE Cloud and SMT (optional) as add-on products and proceed with the installation. Consult the SUSE Linux Enterprise Server *Deployment Guide* at http://www.suse.com/documentation/sles11/book_sle_deployment/data/sec_i_yast2_inst_mode.html for detailed instructions.

3.1.2 Partitioning

Currently, Crowbar requires `/opt` to be writable. Apart from that, SUSE Cloud has no special requirements in regards of partitioning. However, it is recommended to create a separate partition or volume for `/srv`. /srv will host all update and product repositories for SUSE Linux Enterprise Server and SUSE Cloud. A size of at least 25 GB is required. Help on using the partitioning tool is available at http://www.suse.com/documentation/sles11/book_sle_deployment/data/sec_yast2_i_y2_part_expert.html.

3.1.3 Software Selection

Installing a minimal base system is sufficient to set up the Administration Server. The following patterns are the minimum requirement:

- *Base System*
- *Minimal System (Appliances)*
- *Subscription Management Tool* (optional)
- *SUSE Cloud Admin Node*
- *Web and LAMP Server* (only needed when installing SMT)

3.1.4 Product Registration

Although you can also register your products at any time after the installation, it is recommended to register SUSE Linux Enterprise Server and SUSE Cloud now, because it will give you immediate access to the update channels. If you have installed the SMT

Add-On product, you *must* register your SUSE Linux Enterprise Server version at the Novell Customer Center *now*, otherwise you will not be able to configure the SMT server. You have received registration keys with the SUSE Cloud Administration Server subscription. See http://www.suse.com/documentation/sles11/book_sle_deployment/data/sec_i_yast2_conf.html for details on the Novell Customer Center registration.

NOTE: SUSE Login Required

In order to register a product, you need to have a SUSE/Novell login. If you do not have such a login, create it at http://www.suse.com/login.

3.1.5 Online Update

A successful product registration will add update repositories for SUSE Linux Enterprise Server and all add-on products. After having successfully registered you will be asked to perform an online update, which will update the system and the add-on products. It is strongly recommended to perform the update at this point in time. If you choose to skip the update now, you must perform it later, before running the Cloud installation script.

3.1.6 CA Setup

In case you have installed SMT you need to provide a certification authority (CA). If you already have a CA certificate in your organization, import it. Otherwise generate all certificates in the Administration Server itself by accepting the YaST proposal. See http://www.suse.com/documentation/sles11/book_security/data/cha_security_yast_ca.html for more information.

If SMT is not installed, click on the *CA Management* link and choose to not set up a CA.

3.1.7 Basic Network Configuration

Only the first interface (`eth0`) on the Administration Server needs to be configured during the installation. Other interfaces will be automatically configured by the cloud installation script.

`eth0` needs to be given a fixed IP address from the admin network—when sticking with the default network addresses this would be `192.168.124.10`. The address you need to enter for the *Default Gateway* depends on whether you have provided an external gateway for the admin network (use the address of that gateway) or not (use *`xxx.xxx.xxx`*.1, e.g. `192.168.124.1`). Using a custom IP address or more than one network interfaces requires to adjust the Crowbar configuration in a later step as described in Section 3.1.9, “Crowbar Setup” (page 30).

If you allow to access the admin network from another network (via gateway or bastion network), you can also add one or more name servers. The Administration Server's name server will automatically be configured by the cloud installation script to forward requests for non-local records to those server(s).

You also need to assign a hostname and a full qualified domain name (FQDN) such as *`admin.cloud.example.com`* to `eth0`.

Last, the firewall need to be disabled for all interfaces.

IMPORTANT: Administration Server Domain Name and Hostname

Setting up the SUSE Cloud will also install a DNS server for all nodes in the cloud. The domainname you specify for the Administration Server will be used for the DNS zone. It is recommended to use a sub-domain such as *`cloud.example.com`*.

The hostname and the FQDN need to be resolvable with `hostname -f`. Double-check whether `/etc/hosts` contains an appropriate entry for the Administration Server. It should look like the following:

```
191.168.124.10 admin.cloud.example.com admin
```

It is *not* possible to change the Administration Server hostname or the FQDN once the cloud installation script has been run.

3.1.8 SMT Configuration (optional)

Skip this step if you have not installed the SMT add-on product. In case you have installed it, you will be asked to configure it. Configuring the SMT server requires you to have your mirroring credentials and your registration e-mail address at hand. To access them, log in to the Novell Customer Center at http://www.novell.com/center/. Get the mirror credentials by selecting *My Products* > *Mirror Credentials* in the left navigation. Obtain your registration e-mail address from *My Profile* > *Login Profile*.

Enter this data at the *SMT Configuration Wizard Step 1/2* into the fields *User*, *Password*, and *NCC E-mail Used for Registration*. Accept the pre-filled defaults for the other input fields. Make sure to *Test* the credentials.

In step two of the SMT configuration you need to enter a database password and specify an e-mail address for getting reports. Refer to http://www.suse.com/documentation/smt11/ for the complete *SMT for SUSE Linux Enterprise 11 Guide*.

3.1.9 Crowbar Setup

This YaST module allows you to basically configure all networks within the cloud and set the network mode for all networks. Furthermore you can also change the username and password for the Crowbar Web interface with which you can manage the OpenStack nodes.

Start YaST and choose *Miscellaneous* > *Crowbar* to start the YaST Crowbar module. The *Administration Settings* tab lets you change the username and password for the Crowbar Web interface.

On the *Network Mode* tab you can choose between *single*, *dual*, and *team* mode. When choosing *team*, you also need to set the *Bonding Policy*. See Section 2.1.2, “Network Modes” (page 10) for details on SUSE Cloud and network modes. In-depth information about the *Bonding Policy* (also known as bonding modes) is available at https://www.kernel.org/doc/Documentation/networking/bonding.txt in section 2, *Bonding Driver Options*, under *mode*.

If you do not want to use the default IP addresses and the default address allocation, change these settings on the *Networks* tab. See Section 2.1, “Network” (page 5) for details on the cloud network. You can also change the Bridge and VLAN allocation on the *Networks* tab. Only change them if you really know what you require, sticking with the defaults is recommended.

If you want to separate the admin and the BMC network, you must change the settings for the networks *bmc* and *bmc_vlan*. The *bmc_vlan* is used to used to generate a VLAN tagged interface on the Administration Server that can access the *bmc* network. The *bmc_vlan* needs to be in the same ranges as *bmc*, and *bmc* has to have *VLAN* enabled.

Table 3.1: *Separate BMC Network Example Configuration*

	bmc	**bmc_vlan**
Subnet	`192.168.126.0`	
Netmask	`255.255.255.0`	
Router	`192.168.126.1`	
Broadcast	`192.168.126.255`	
Host Range	`192.168.126.10 - 192.168.126.100`	`192.168.126.101 - 192.168.126.101`
VLAN	yes	
VLAN ID	100	
Bridge	no	

IMPORTANT: No Network Changes after Having Run the Cloud Installation Script

As of SUSE Cloud 1.0 it is not possible to change the network setup after having run the cloud installation script. Allowing such changes is planned for future releases of SUSE Cloud.

NOTE: Setting up a Bastion Network

As of SUSE Cloud 1.0 it is not possible to set up a bastion network with YaST Crowbar. It needs to be configured manually—see Section 3.2.3, “Setting Up a Bastion Network” (page 37).

Other, more flexible network mode setups can be configured by manually editing the Crowbar network configuration files. See the documentation on the Crowbar wiki (https://github.com/dellcloudedge/crowbar/wiki) for more information. SUSE can assist you in creating a custom setup within the scope of a Level 3 support contract.

3.2 Post-Installation Configuration

After the installation has finished, you need to set up product and update repositories and, optionally, configure the bastion network. Once the Administration Server host is fully configured, start the cloud installation script.

3.2.1 Setting up the SMT Repositories (optional)

Skip this step if you have not installed the SMT add-on product. In case you have installed it, the SMT server was set up to be able to communicate with the Novell Customer Center during the installation. In this step we will add and mirror the update repositories for SUSE Linux Enterprise Server and for SUSE Cloud. They will serve as the update source for the OpenStack nodes. Run the following commands as user `root`:

```
for REPO in SLES11-SP{2-Core,2-Updates,1-Updates,1-Pool} SUSE-Cloud-1.0-{Pool,Updates}; do
  smt repos $REPO sle-11-x86_64 -e
done
smt mirror -L /var/log/smt/smt-mirror.log
```

The `smt mirror` command will download approximately 14 GB of patches. This process may last up to several hours. A log file is written to `/var/log/smt/smt-mirror.log`.

3.2.2 Setting Up Local Repositories

In order to deploy the OpenStack nodes and to provide update repositories for them, product and update repositories for SUSE Linux Enterprise Server and SUSE Cloud must be locally available at `/srv/tftpboot`. The source of the repositories can either be an SMT server installed on the Administration Server or your company's network. Please refer to Section 2.2, “Product and Update Repositories” (page 14) for details. The following table lists all repositories, their file system location on the SMT server, and the location at which they need to be made available on the Administration Server:

Repository Names and Locations

SLES11-SP1-Pool
SMT dir: `/srv/www/htdocs/repo/$RCE/SLES11-SP1-Pool/slee-11-x86_64`

Local dir: `/srv/tftpboot/repos/SLES11-SP1-Pool`

SLES11-SP1-Updates
SMT dir: `/srv/www/htdocs/repo/$RCE/SLES11-SP1-Updates/slee-11-x86_64`

Local dir: `/srv/tftpboot/repos/SLES11-SP1-Updates`

SLES11-SP2-Core
SMT dir: `/srv/www/htdocs/repo/$RCE/SLES11-SP2-Core/slee-11-x86_64`

Local dir: `/srv/tftpboot/repos/SLES11-SP2-Core`

SLES11-SP2-Updates
SMT dir: `/srv/www/htdocs/repo/$RCE/SLES11-SP2-Updates/slee-11-x86_64`

Local dir: `/srv/tftpboot/repos/SLES11-SP2-Updates`

SUSE-Cloud-1.0-Pool
SMT dir: `/srv/www/htdocs/repo/$RCE/SUSE-Cloud-1.0-Pool/slee-11-x86_64`

Local dir: `/srv/tftpboot/repos/SUSE-Cloud-1.0-Pool`

SUSE-Cloud-1.0-Updates

SMT dir: `/srv/www/htdocs/repo/$RCE/SUSE-Cloud-1.0-Updates/slee-11-x86_64`

Local dir: `/srv/tftpboot/repos/SUSE-Cloud-1.0-Updates`

SUSE Linux Enterprise Server 11 SP2 Product

SMT dir: n/a

Local dir: `/srv/tftpboot/suse-11.2`

3.2.2.1 Update Repositories

The update repositories for SUSE Linux Enterprise Server and SUSE Cloud not only need to be available locally on the Administration Server, they also need to be kept in sync with the official update repositories provided by Novell Customer Center. It is highly recommended to install an SMT server either on the Administration Server or within your company network. An SMT server automatically synchronizes the repositories with the Novell Customer Center. There are several possibilities to make the repositories locally available on the Administration Server.

SMT Server installed on the Administration Server

Link the repositories mirrored by SMT to `/srv/tftpboot`:

```
for REPO in SLES11-SP{2-Core,2-Updates,1-Updates,1-Pool} SUSE-Cloud-1.0-{Pool,Updates};
 do
  ln -s /srv/www/htdocs/repo/\$RCE/$REPO/sle-11-x86_64 /srv/tftpboot/repos/$REPO
done
```

SMT Server installed on a Remote Host

If the SMT server is installed on a remote host that can be accessed from the Administration Server you can either mount the update repositories, for example via `NFS`, or regularly `rsync` them.

To `NFS`-mount the repositories from a remote host, either use the YaST *NFS Client* module or edit `/etc/fstab`. The local mount point should be `/srv/tftpboot/repos/<REPOSITORY_NAME>`.

To `rsync` the repositories from a remote host, create a daily cron job running the following command on the Administration Server. This command will *pull* the files from a host named SMT.example.com:

```
for REPO in SLES11-SP{2-Core,2-Updates,1-Updates,1-Pool} SUSE-Cloud-1.0-{Pool,Updates};
 do
  rsync -avPz SMT.example.com:/srv/www/htdocs/repo/\\\$RCE/$REPO/sle-11-x86_64/ \
  /srv/tftpboot/repos/$REPO/
done
```

Alternatively you may set up the cron job on the remote host and *push* the file to the Administration Server (which has the IP address `192.168.124.10` in the following example):

```
for REPO in SLES11-SP{2-Core,2-Updates,1-Updates,1-Pool} SUSE-Cloud-1.0-{Pool,Updates};
 do
  rsync -avPz /srv/www/htdocs/repo/\\\$RCE/$REPO/sle-11-x86_64/ \
  192.168.124.10:/srv/tftpboot/repos/$REPO/ \
done
```

NOTE: Mind the Trailing Slash

The `rsync` command must be used with trailing slashes in the directory names as shown above. Otherwise rsync would copy the repositories into the wrong directory.

Sneakernet

If your admin network is isolated from other networks, you need to manually sync the update repositories from removable media. To do so you can either use `rsync` (see above for an example) or `cp -axu`.

3.2.2.2 Product Repositories

The files in the product repositories for SUSE Linux Enterprise Server and SUSE Cloud do not change, therefore they do not need to be synced with a remote source. It is sufficient to copy the data once, either from a remote host or directly from the installation media. Alternatively you may mount the product repository from a remote server via `NFS`. Please note that the data *must* be directly available from the local directories listed in Repository Names and Locations (page 33). It is not possible to use links.

If copying, it is recommended to use `rsync`. If the installation data is located on a removable device, make sure to mount it first (for example, after inserting the DVD in the Administration Server and waiting for the device to become ready):

```
mkdir -p /srv/tftpboot/suse-11.2/install/
mount /dev/dvd /mnt
rsync -avP /mnt/ /srv/tftpboot/suse-11.2/install/
umount /mnt
```

If the installation data is provided by a remote machine, log in to that machine and push the data to the Administration Server (which has the IP address `192.168.124.10` in the following example):

```
rsync -avPz /data/sles11sp2/ 192.168.124.10:/srv/tftpboot/suse-11.2/install/
```

Also make the contents of the SUSE Cloud product repository available at `/srv/tftpboot/repos/Cloud/` using one of the techniques described in the previous step.

3.2.2.3 Software Repository Sources

Now that the product and update repositories for SUSE Linux Enterprise Server and SUSE Cloud are available locally, the OpenStack nodes can be installed and updated from these sources. However, it also makes sense to use these repositories as resources to install and update packages on the Administration Server as well. Therefore you need to replace the existing remote repositories that have been added automatically during the product registration by the local ones. You can either use zypper or the YaST module *Software Repositories* to do so.

One way to do so, would be to disable all existing remote services and repositories and to add the local repositories afterwards:

```
zypper ms -dR --remote
zypper mr -dR --remote
for REPO in SLES11-SP{2-Core,2-Updates,1-Updates,1-Pool} SUSE-Cloud-1.0-{Pool,Updates}; do
  zypper ar -f /srv/tftpboot/repos/$REPO $REPO
done
zypper ar /srv/tftpboot/repos/Cloud "SUSE-Cloud-1.0"
zypper ar /srv/tftpboot/suse-11.2/install/ "SLES 11 SP2"
```

IMPORTANT: Remote Repositories and Services Need to be Disabled

The cloud installation script will refresh all active repositories and services. In case repositories cannot be refreshed, the script will fail. Even if you have a

permanent Internet connection on the Administration Server, it may temporarily not be available during the run of the cloud installation script, since this script also reconfigures the network.

Therefore all remote repositories and services need to be disabled prior to running the cloud installation script. This is archived by running the `zypper ms` and `zypper mr` commands listed above. In case you need remote repositories (such as the SMT update repositories), re-enable them after the cloud installation script has run.

3.2.3 Setting Up a Bastion Network

As outlined in Section 2.1, "Network" (page 5), one way to access the admin network from a defined external network is via a Bastion network and a second network card (as opposed to providing an external gateway).

To set up the Bastion network, you need to have a static IP address for the Administration Server from the external network. You need to adjust the network template file `/opt/dell/chef/data_bags/crowbar/bc-template-network.json`. The example configuration used below assumes that the external network from which to access the admin network has the following addresses. You need to adjust them according to your needs.

Subnet: `10.10.1.0`
Netmask: `255.255.0.0`
Broadcast: `10.10.1.255`
Gateway: `10.10.1.1`
Static Administration Server address: `10.10.1.125`

Adjust `/opt/dell/chef/data_bags/crowbar/bc-template-network.json` according to the following patch (it only directly matches if you have not changed the default network configuration). Once the bastion network configuration has been added to `bc-template-network.json`, it can be adjusted using the YaST Crowbar module.

```
--- /opt/dell/chef/data_bags/crowbar/bc-template-network.json
+++ /opt/dell/chef/data_bags/crowbar/bc-template-network.json
@@ -86,6 +86,11 @@
                           "1g1"
```

```
                    ]
                },
+               "bastion1" : {
+                  "if_list" : [
+                     "1g2"
+                  ]
+               },
                "intf1" : {
                   "if_list" : [
                      "1g1"
@@ -209,6 +214,23 @@
                "subnet" : "192.168.122.128",
                "use_vlan" : true
            },
+           "bastion" : {
+              "add_bridge" : false,
+              "vlan" : 50,
+              "router" : "10.10.1.1",
+              "ranges" : {
+                 "admin" : {
+                    "start" : "10.10.1.125",
+                    "end" : "10.10.1.125"
+                 }
+              },
+              "broadcast" : "10.10.1.255",
+              "netmask" : "255.255.255.0",
+              "use_vlan" : false,
+              "conduit" : "bastion1",
+              "subnet" : "10.10.1.0",
+              "router_pref" : 5
+           },
            "public" : {
               "add_bridge" : false,
               "vlan" : 300,
```

3.2.4 Running the Cloud Installation Script

Before running the cloud installation script to finish the configuration of the Administration Server make sure to double-check the following items.

Final Check Points

- Make sure the network configuration is correct. Run *YaST > Crowbar* to review/change the config. See Section 3.1.9, “Crowbar Setup” (page 30) for further instructions.

- Make sure `hostname -f` returns a full qualified hostname. See Section 3.1.7, “Basic Network Configuration” (page 29) for further instructions.
- Make sure all update and product repositories are available locally. See Section 3.2.2, “Setting Up Local Repositories” (page 33) for further instructions.
- Make sure the operating system and SUSE Cloud are up-to-date and have the latest patches installed. Run `zypper patch` to install them.

Now everything is in place to finally configure the Administration Server. This is done by running the script `/opt/dell/bin/install-chef-suse.sh`. This command will install and configure Chef, and use it to complete the installation of Crowbar and all required Barclamps. It will take several minutes to complete.

```
screen /opt/dell/bin/install-chef-suse.sh
```

IMPORTANT: Use a Terminal Multiplexer to run the Cloud Installation Script

Run the installation script `install-chef-suse.sh` inside of a terminal multiplexer like GNU Screen (provided by the `screen` package).

During the run of this script the network will be reconfigured. This may result in interrupting the script when being run from a network connection (like SSH). Using `screen` will continue running the script in a session to which you can reconnect via `screen -r` if you lose the connection.

`install-chef-suse.sh` will produce a lot of output that gets written to a log file located at `/var/log/chef/install.log`. Check this log file in case something goes wrong. You can run `install-chef-suse.sh` multiple times as long as you have not started to deploy the OpenStack services.

If the script has successfully finished, you will see a message telling you how to log in to the Crowbar Web interface.

WARNING: No Network Changes After Having Run the Cloud Installation Script

Once you have run the cloud installation script, you cannot change the network setup anymore. If doing so, you would have to completely set up the Administration Server again.

3.2.4.1 Activating the Bastion Network

In case you have configured a Bastion Network, you need to activate its network interface by running the following commands:

```
/opt/dell/bin/crowbar network -U crowbar -P crowbar allocate_ip \
  default $(hostname -f) bastion admin
chef-client
```

This command needs to be executed directly on the admin server, so you either need direct access to the machine or a serial console.

3.2.4.2 Re-enabling Remote Repositories

Prior to running the cloud installation script, all required update repositories have been made available locally and all remote repositories and services have been disabled (see Section 3.2.2.3, “Software Repository Sources” (page 36)). In case you still need remote repositories that have been disabled (e.g. SLE11-SMT-SP2-Pool and SLE11-SMT-SP2-Updates if you have installed SMT), you may re-enable them now using YaST or zypper.

Installing the OpenStack Nodes

4

The OpenStack nodes represent the actual cloud infrastructure. Node installation and service deployment is done automatically from the Administration Server. Before deploying the OpenStack services, you need to install SUSE Linux Enterprise Server on every node. In order to do so, each node needs to be PXE booted using the `tftp` server from the Administration Server. Afterwards you can allocate the nodes and trigger the operating system installation. There are three different types of nodes:

Controller Node: The central management node interacting with all other nodes.
Compute Nodes: The nodes on which the instances are started.
Storage Nodes: Nodes providing object or block storage.

4.1 Preparations

Meaningful Node names
: Make a note of the MAC address and the purpose of each node (for example, controller, storage Ceph, storage Swift, compute). This will make deploying the OpenStack services a lot easier and less error-prone, since it allows you to assign meaningful names (aliases) to the nodes, which are otherwise listed with the MAC address by default.

BIOS Boot Settings
: Make sure PXE-booting (booting from the network) is enabled and configured as the *primary* boot-option for each node. The nodes will boot twice from the network during the allocation and installation phase.

Custom Node Configuration

All nodes are installed using AutoYaST with the same configuration located at `/opt/dell/chef/cookbooks/provisioner/templates/default/autoyast.xml.erb`. If this configuration does not match your needs (for example if you need special third party drivers) you need to make adjustments to this file. An AutoYaST manual can be found at http://www.suse.com/documentation/sles11/book_autoyast/data/book_autoyast.html. Having change the AutoYaST config file, you need to re-upload it to Chef, using the following command:

```
knife cookbook upload -o /opt/dell/chef/cookbooks/ provisioner
```

Direct `root` Login

By default, the `root` account on the nodes has no password assigned, so a direct `root` login is not possible. Logging in on the nodes as `root` is only possible via SSH public keys (for example, from the Administration Server).

If you want to allow direct `root` login, you can set a password that will be used for the `root` account on all OpenStack nodes before deploying the nodes. This must be done before the nodes are deployed; setting a `root` password at a later stage is not possible.

Setting a `root` Password for the OpenStack Nodes

1. Create an md5-hashed `root`-password, for example by using `mkpasswd --method=md5` (`mkpasswd` is provided by the package `whois`, which is not installed by default).

2. Open a browser and point it to the Crowbar Web interface available at port `3000` of the Administration Server, for example http://192.168.124.10:3000/. Log in as user `crowbar`. The password defaults to `crowbar`, if you have not changed it during the installation.

3. Open the Barclamp menu by clicking *Barclamps* > *All Barclamps*. Click the *Provisioner* Barclamp entry and *Edit* the *Default* proposal.

4. Click *Raw* to edit the configuration file.

5. Add the following line within the *Provisioner* section of the file:

   ```
   "root_password_hash": "HASHED_PASSWORD"
   ```

replacing "*HASHED_PASSWORD*" with the password you generated in the first step.

4.2 Node Installation

To install a node, you need to PXE boot it first. It will be booted with an image that allows the Administration Server to discover the node and make it available for installation. Once you have allocated the node, it will PXE boot again and the automatic installation will start.

1. PXE-boot all nodes you want to deploy. Although it is possible to allocate nodes one-by-one, doing this in bulk-mode is recommended, because it is much faster. The nodes will boot into the "SLEShammer" image, which performs initial hardware discovery.

2. Open a browser and point it to the Crowbar Web interface available at port `3000` of the Administration Server, for example `http://192.168.124.10:3000/`. Log in as user `crowbar`. The password defaults to `crowbar`, if you have not changed it.

   Click *Nodes* > *Dashboard* to open the *Node Dashboard*.

3. Each node that has successfully booted will be listed as being in state `Discovered`, indicated by a yellow bullet. The nodes will be listed with their MAC address as a name. Wait until all nodes are listed as being `Discovered` before proceeding.

4. Although this step is optional, it is recommended to properly group your nodes at this stage, since it allows you to clearly arrange all nodes. Grouping the nodes by role would be one option, for example control, compute, object storage (Swift), and block storage (Ceph) .

   4a Enter the name of a new group into the *New Group* input field and click *Add Group*.

   4b Drag and drop a node onto the title of the newly created group. Repeat this step for each node you would like to put into the group.

5 To allocate the nodes click on *Nodes* > *Bulk Edit*. If you prefer to allocate the nodes one-by-one, click a node's name followed by a click on *Edit* instead.

6 Provide a meaningful *Alias* and a *Description* for each node and check the *Allocate* box. The entries for *BIOS* and *RAID* are currently not used.

TIP: Alias Names

Providing an alias name will change the default node names (MAC address) to the name you provided, making it easier to identify the node. Furthermore, this alias will also be used as a DNS `CNAME` for the node in the admin network. As a result, you will be able to access the node via this alias when, for example, logging in via SSH.

---

7 Once you have filled in the data for all nodes, click *Save*. The nodes will reboot and commence the AutoYaST-based SUSE Linux Enterprise Server installation via a second PXE boot. Click *Nodes* > *Dashboard* to return to the *Node Dashboard.*

8 Nodes that are being installed are listed with the status `Installing` (yellow/green bullet). Once the installation of a node has finished, it is listed as being `Ready`, indicated by a green bullet. Wait until all nodes are listed as being `Ready` before proceeding.

4.3 Post-Installation Configuration

The following lists some *optional* configuration steps like configuring node access and SSL-enablement. You may entirely skip the following steps or perform the steps necessary for accessing the nodes or the SSL enablement at any later stage.

4.3.1 Providing a Volume or Separate Partition for the Glance Image Repository

If you plan to host the Glance Image Repository on a separate volume (recommended) or partition, you need to prepare the Controller Node before deploying the Glance service.

Log in to the Controller Node as `root` via SSH from the Administration Server (see Section 6.1.2, “OpenStack Node Deployment” (page 60) for detailed instructions). Set up the volume or format the partition and mount it to `/var/lib/glance/images` (if you do not use YaST for this tasks, you need to create the directory prior to mounting).

4.3.2 Accessing the Nodes

By default, the `root` account on the nodes has no password assigned, so `root` login is only possible via SSH. The default setup allows to execute the `ssh` command as user `root` from the Administration Server (see “How can I log in to a node as `root`?” (page 60)). In order to be able to execute the `ssh` command as a different user, you need to add this user's public SSH keys to `root`'s `authorized_keys` file on all nodes. Proceed as follows:

Procedure 4.1: *Copying SSH Keys to all Nodes*

1. Log in to the Crowbar Web interface available at port `3000` of the Administration Server, for example `http://192.168.124.10:3000/` (username and default password: `crowbar`).

2. Open the Barclamp menu by clicking *Barclamps* > *All Barclamps*. Click the *Provisioner* Barclamp entry and *Edit* the *Default* proposal.

3. Copy and paste the SSH keys into the *Additional SSH Keys* input field. Each key needs to be placed on a new line.

4. Click *Apply* to deploy the keys and save your changes to the proposal.

4.3.3 Enabling SSL

In order to enable SSL to encrypt communication within the cloud (see Section 2.4, “SSL Encryption” (page 19) for details), the respective certificates need to be available on the nodes.

The certificate file and the key file need to be copied to the Controller Node, into the following locations:

SSL Certificate File

`/etc/apache2/ssl.crt/`

SSL Key File

`/etc/apache2/ssl.key/`

4.4 Editing Allocated Nodes

All nodes that have been allocated can be decommissioned or re-installed. Click a node's name in the *Node Dashboard* and then click *Edit*. The following options are available:

Forget

Deletes a node from the pool. If you want to re-use this node again, it needs to be reallocated and re-installed from scratch.

Deallocate

Temporarily removes the node from the pool of nodes. Once you reallocate the node it will take its former role. Useful for adding additional machines in times of high load or fir decommissioning machines in times of low load.

Reinstall

Triggers a reinstallation. The machine stays allocated.

WARNING: Editing Nodes in a Production System

When deallocating nodes that provide essential services, the complete cloud will become unusable. While it is uncritical to disable single storage nodes (provided you have not disabled redundancy) or single compute nodes, disabling the Controller Node will “kill” the complete cloud. You should also not disable nodes providing Ceph monitoring services or the nodes providing swift ring and proxy services.

Deploying the OpenStack Services

5

Once the nodes are installed and configured you can start deploying the OpenStack services in order to finalize the installation. The services need to be deployed in a given order, because they depend on one another. Deployment is done from the Crowbar Web interface through recipes, so-called "Barclamps".

The services controlling the cloud (including storage management and control services) need to be installed on the Controller Node. However, you may *not* use your Controller Node as a compute or storage host. Here is a list with services that may *not* be installed on the Controller Node: *Swift-storage*, *Ceph-store*, *Nova-multi-compute*. These services need to be installed on dedicated nodes.

The OpenStack services need to be deployed in the following order. For general instructions on how to edit and deploy Barclamp, refer to Section 5.1, "Barclamp" (page 48). Deploying Swift and Ceph is optional; all other services must be deployed.

1. Deploying the Database

2. Deploying Keystone

3. Deploying Swift (optional)

4. Deploying Ceph (optional, unsupported)

IMPORTANT: Ceph not Supported

As of SUSE Cloud 1.0, Ceph is not officially supported but rather included as a technical preview, so using Nova Volume instead is recommended.

5.1 Barclamp

The OpenStack services are automatically installed on the nodes by using so-called Barclamps—a set of recipes, templates, and installation instructions. All existing Barclamps can be accessed from the Crowbar Web interface by clicking on *Barclamps*. To edit a Barclamp, proceed as follows:

1. Open a browser and point it to the Crowbar Web interface available at port `3000` of the Administration Server, for example `http://192.168.124.10:3000/`. Log in as user `crowbar`. The password defaults to `crowbar`, if you have not changed it.

   Click *Barclamps* to open the *All Barclamps* menu. Alternatively you may filter the list to *Crowbar* or *OpenStack* Barclamps by choosing the respective option from *Barclamps*. The *Crowbar* Barclamps contain general recipes for setting up and configuring all nodes, while the *OpenStack* are dedicated to OpenStack service deployment and configuration.

2. Click a Barclamp's name. You can either *Create* a proposal or *Edit* an existing one.

   When creating a new proposal, give it a meaningful name and description. You can create several proposals, for example, for testing purposes, but only one at a time can be deployed.

   Most OpenStack Barclamps consist of two sections: the *Attributes* section lets you change the configuration, and the *Node Deployment* section lets you choose onto which nodes to deploy the Barclamp.

3. To edit the *Attributes* section, change the values via the Web form. Alternatively you can directly edit the configuration file by clicking *Raw*.

WARNING: Raw Mode

Only use the *Raw* mode in case an option cannot be changed via Web form. Raw mode does not perform any syntax checks.

If you switch between *Raw* mode and Web form (*Custom* mode), make sure to *Save* your changes before switching, otherwise they will be lost.

In the *Node Deployment* section of the OpenStack Barclamp you can drag and drop nodes from the *Available Nodes* column to the desired role. You need to drop the node onto the role name. Do *not* drop a node onto the input field—this is rather used to filter the list of *Available Nodes*!

One or more nodes are usually automatically pre-selected for available roles. If this pre-selection does not meet your requirements, remove it *before* dragging new nodes to the role. To remove a node from a role, click the respective *Remove* icon.

4 To save and deploy your edits, click *Apply*. To just save your changes without deploying them, click *Save*. To remove the complete proposal, click *Delete*. A proposal that already has been deployed can only be deleted manually, see Section 5.1.1, "Deactivate a Proposal that Already has been Deployed" (page 50) for details.

If you deploy a proposal onto a node where a previous one is still active, the new proposal will overwrite the old one.

NOTE: Wait Until a Proposal has been Deployed

Deploying a proposal might take some time (up to several minutes). It is strongly recommended to always wait until you see the note "Successfully applied the proposal" before proceeding on to the next proposal.

WARNING: Barclamp Deployment Failure

In case the deployment of a Barclamp fails, make sure to fix the reason that has caused the failure and deploy the Barclamp again. Refer to the respective troubleshooting section at Section 6.1.2, "OpenStack Node Deployment" (page 60) for help. A deployment failure may leave your node in an inconsistent state.

5.1.1 Deactivate a Proposal that Already has been Deployed

To finally deactivate a proposal that already has been deployed, you first need to *Deactivate* it in the Crowbar Web interface. Run the following commands as `root` on the Administration Server afterwards:

```
P_NAME="proposal_name_to_delete"
SERVICE="service_name"
crowbar $SERVICE proposal delete $P_NAME
crowbar $SERVICE delete $P_NAME
```

`proposal_name_to_delete` needs to be replaced by the name of the proposal you want to delete. *`service_name`* needs to be replaced by one of the following strings representing the OpenStack services: `ceph`, `database`, `glance`, `keystone`, `nova_dashboard`, `nova`, `swift`.

5.2 Deploying the Database

The very first service that needs to be deployed is the *Database*. The database service is used by all other services. It must be installed on the Controller Node.

SUSE Cloud only supports *PostgreSQL* as an *SQL Engine*, so this value must not be changed.

5.3 Deploying Keystone

Keystone is another core component that is used by all other OpenStack services. It provides authentication and authorization services. *Keystone* needs to be installed on the Controller Node. You can configure the following parameters of this Barclamp:

SQL Engine
: Must be set to *PostgreSQL/MySQL*. This is the default.

SQL Instance
: Name of the proposal you deployed in the previous step (see Section 5.2, “Deploying the Database” (page 50)). The default value should be correct.

Default Tenant
: Tenant for the users. Do not change the default value of `openstack`.

Regular User/Administrator Username/Password
: Username and password for the regular user and the administrator. Both accounts can be used to log in to the SUSE Cloud Dashboard to manage Keystone users and access.

Administrator Token (long-lived)
: The permanent administrator token (random string).

Administrator Token Expiration
: Expiration Date for the administrator token. Must be entered in the form *YYYY-MM-DD*T*HH:MM*, e.g. 2015-03-31T12:00.

Security Attributes
: When sticking with the default value `HTTP`, public communication will not be encrypted. Choose `HTTPS` to use SSL for encryption and specify the path to the certificate files. Note that you need to create and copy the certificate prior to deploying Keystone, see Section 4.3.3, “Enabling SSL” (page 45) for instructions.

5.4 Deploying Swift (optional)

Swift adds an object storage service to SUSE Cloud that lets you store single files such as images or snapshots. It offers high data security by storing the data redundantly on a pool of Storage Nodes—therefore Swift needs to be installed on at least two dedicated nodes.

It is recommended not to change the defaults in the Barclamp proposal, unless you know exactly what you requir. However you should change the *Cluster Admin Password*. If you plan to change the *Zone* value, it is important to know that you need at least as many Storage Nodes as *Zones*.

The Swift service consists of three different roles:

Swift-ring-compute
: The ring maintains the information about the location of objects, replicas, and devices. It can be compared to an index, that is used by various OpenStack services

to look up the physical location of objects. *Swift-ring-compute* must only be installed on a single node; it is recommended to use the Controller Node.

Swift-proxy-acct
: The Swift proxy server takes care of routing requests to Swift. Installing a single instance of *Swift-proxy-acct* on the Controller Node is recommended.

Swift-storage
: The virtual object storage service. Install this role on all dedicated Swift Storage Nodes (at least two), but not on any other node.

WARNING: Swift-storage Needs Dedicated Machines

Never install the Swift-storage service on a node that runs other OpenStack services.

5.5 Deploying Ceph (optional, unsupported)

Ceph adds a redundant block storage service to SUSE Cloud. It lets you store persistent devices that can be mounted from instances. It offers high data security by storing the data redundantly on a pool of Storage Nodes—therefore Ceph needs to be installed on at least two dedicated nodes.

For more information on the Ceph project, vist `http://ceph.com/`.

IMPORTANT: Ceph not Supported

As of SUSE Cloud 1.0, Ceph is not officially supported but rather included as a technical preview, so using Nova Volume instead is recommended.

The Ceph Barclamp only has one configuration option: telling Ceph which devices to use on the nodes. Edit the Barclamp in *Raw* and search for the following lines

```
"devices": [

],
```

Add a comma-separated list of devices that should be used by Ceph. For example:

```
"devices": [
  "/dev/sdb", "/dev/sdc", "/dev/sdd"
],
```

IMPORTANT: Devices

Not all of the devices used for Ceph need to exist on all nodes. All devices from a node matching the list will be used. They must *not* be mounted prior to deploying Ceph. Any data stored on these devices will be lost.

The Ceph service consists of three different roles:

Ceph-mon-master
: Master cluster monitor daemon for the Ceph distributed file system. *Ceph-mon-master* must only be installed on a single node; it is recommended to use the Controller Node.

Ceph-mon
: Cluster monitor daemon for the Ceph distributed file system. *Ceph-mon* needs to be installed on two or four Storage Nodes.

IMPORTANT: Number of Ceph Monitor Nodes

In addition to the node running the Ceph-mon-master service an additional two or four nodes also need to run the *Ceph-mon* service. The sum of the *Ceph-mon-master* and the *Ceph-mon* nodes must always be an odd number (either three or five).

Nodes running *Ceph-mon* cannot be deleted or temporarily be disabled.

Ceph-store
: The virtual block storage service. Install this role on all dedicated Ceph Storage Nodes (at least two), but not on any other node.

WARNING: *Ceph*-store Needs Dedicated Machines

Never deploy *Ceph-store* on a node that runs other non-Ceph OpenStack services. The only service that may be deployed together with it is *Ceph-mon*.

Deploying Ceph requires to perform the steps in a given order:

1. Edit the Barclamp proposal to specify the devices to be used by Ceph as described above.

2. Drag and drop a node (for example, the Controller Node) to the *Ceph-mon-master* role.

3. Drag and drop two or four nodes to the *Ceph-mon* role. Note that the maximum number of *Ceph-mon* nodes cannot exceed four and that the sum of *Ceph-mon-master* and *Ceph-mon* nodes must be odd.

4. Drag and drop all dedicated Ceph Storage Nodes to the *Ceph-store* (at least two). You may also use the nodes with the *Ceph-mon* roles, but not the *Ceph-mon-master* node (you can add that one later).

5. Click *Apply* to deploy your proposal. This can take some time.

6. If you also want to use the *Ceph-mon-master* as a Storage Node, drag and drop it to the *Ceph-store* role and click *Apply* again. Note that it is not recommended to use the Controller Node for non-management purposes such as storage or compute.

5.6 Deploying Glance

Glance provides discovery, registration, and delivery services for virtual disk images. An image is needed to start an instance—it is its pre-installed root-partition. All images you want to use in your cloud to boot instances from, are provided by Glance.

Glance should be deployed onto the Controller Node. There are a lot of options to configure Glance. The most important ones are explained below—for a complete reference refer to `http://github.com/dellcloudedge/crowbar/wiki/Glance--barclamp`.

Image Store Directory
: Directory in which all images uploaded to Glance are stored. If you want to put the images onto a separate partition or volume, you need to mount this partition or volume on the Controller Node prior to deploying the Glance proposal (see Section 4.3.1, “ Providing a Volume or Separate Partition for the Glance Image Repository ” (page 44)). Specify the mount point of the partition or volume here.

Security Attributes

When sticking with the default value `HTTP`, public communication will not be encrypted. Choose `HTTPS` to use SSL for encryption and specify the path to the certificate files. Note that you need to create and copy the certificate prior to deploying Glance, see Section 4.3.3, “Enabling SSL” (page 45) for instructions.

API/Registry *Bind to All Addresses*

Set these two options to `true` to enable users to upload images to Glance. If unset, only the operator will be able to upload images.

Caching

Enable and configure image caching in this section. By default, image caching is disabled. Learn more about Glance's caching feature at `http://docs.openstack.org/developer/glance/cache.html`.

5.7 Deploying Nova

Nova provides key services for managing the SUSE Cloud, sets up the Compute Nodes, and provides a block storage service that either makes use of Ceph (if deployed) or uses local disks. There are a lot of options to configure Nova. The most important ones are explained below—for a complete reference refer to `https://github.com/dellcloudedge/crowbar/wiki/Nova--barclamp`. You will also find details about Nova's network modes on that page.

Hypervisor

Choose between the *KVM* and *Xen* hypervisors (other available options are currently not supported by SUSE). As of SUSE Cloud 1.0 choosing more than one hypervisor is not supported.

Choose device for nova-volume storage volume group

In case you have not deployed Ceph, Nova Volume will use a single device on the contrnode; to provide block storage. Specify which device to use with this option. If Ceph has been deployed, it is used automatically and you will not be able to choose any disks here.

NOTE: Nova Volume

Nova Volume only runs on the host onto which *Nova-multi-controller* is deployed. Make sure the chosen device is equipped with enough disk space (a RAID is strongly recommended). See Section 2.3.1, “Cloud Storage Services” (page 16) for more information.

Security Attributes

When sticking with the default value `HTTP`, public communication will not be encrypted. Choose `HTTPS` to use SSL for encryption and specify the path to the certificate files. Note that you need to create and copy the certificate prior to deploying Nova, see Section 4.3.3, “Enabling SSL” (page 45) for instructions.

noVNC Security Attributes

After having started an instance you can display its VNC console in the Nova Dashboard via browser using the `noVNC` implementation. By default this connection is not encrypted and can potentially be eavesdropped. To encrypt it, you can make use of SSL by setting *noVNC via SSL* to `true`. If you do not specify any additional certificates, the same ones as for Nova will be used.

The Nova service consists of two different roles:

Nova-multi-controller

Distributing and scheduling the instances is managed by the *Nova-multi-controller*. It also provides networking and messaging services. *Nova-multi-controller* needs to be installed on the Controller Node.

Nova-multi-compute

Provides the hypervisor (KVM or Xen) and tools needed to manage the instances. *Nova-multi-compute* needs to be installed on every Compute Node. The Compute Nodes are the “workhorses” of the cloud—each instance is started on a Compute Node.

5.8 Deploying the Nova Dashboard

The last service that needs to be deployed is the Nova Dashboard. It provides a Web interface for users to start and stop instances and for administrators to manage users,

groups, roles, etc. Nova Dashboard should be installed on the Controller Node. The following attributes can be configured:

SQL Engine
: Must be set to *PostgreSQL/MySQL*. This is the default.

SQL/Keystone Instance
: Name of the proposal for *Database* and *Keystone* you deployed in the previous steps. The default value should be correct.

Disable SSL Certification Verification
: Usually SSL certificates are checked for whether they have been signed by a trusted organization. Use this option to turn off checks. Useful in testing environments when using self-signed certificates. In production environments you should always use signed certificates and set this option to *false* (default).

Apache Attributes
: When sticking with the default value *HTTP* equals `true`, public communication will not be encrypted. Set `HTTPS` to `true` to use SSL for encryption and specify the path to the certificate files. Note that you need to create and copy the certificate prior to deploying Nova Dashboard, see Section 4.3.3, “Enabling SSL” (page 45) for instructions.

5.9 How to Proceed

With a successful deployment of the Nova Dashboard, the SUSE Cloud installation is finished. In order to be able to test your setup by starting an instance one last step remains to be done—uploading an image to the Glance service. Refer to Section “Managing Images” (Chapter 2, *Using OpenStack Command Line Interfaces*, ↑*User Guide for Administrators*) for instructions. Images for SUSE Cloud can be built in SUSE Studio—see this blog post for details: `http://blog.susestudio.com/2012/10/kvm-build-format-suse-cloud-support.html`.

Now you can hand over to the cloud administrator to set up users, roles, flavors, etc.—refer to the *User Guide for Administrators* (↑*User Guide for Administrators*) for details.

6 Troubleshooting and Support

6.1 FAQ

Find solutions for the most common pitfalls here. If your problem is not mentioned here, checking the log files on either the Administration Server or the OpenStack nodes may help. A list of log files is available at Appendix A, *Log Files* (page 63).

6.1.1 Admin Node Deployment

`/opt/dell/bin/install-chef-suse.sh` fails

Please check the script's log file at `/var/log/chef/install.log` for error messages.

I have configured a bastion network, but `ifconfig` does not show the second NIC after having run `install-chef-suse.sh`.

In order to activate the bastion network, you need to run two additional commands after having run `install-chef-suse.sh`:

```
crowbar network -U crowbar -P crowbar allocate_ip default $(hostname -f)
 bastion admin
chef-client
```

I have successfully set up a bastion network but cannot reach the Administration Server from outside the admin network. `route -n` shows no gateway for the bastion network.

Make sure the value for the bastion network's `"router_pref":` entry in `/opt/dell/chef/data_bags/crowbar/bc-template-network.json` is set to a *lower* value than the `"router_pref":` entry for the admin network.

Can I change the hostname of the Administration Server?

No, once you have run `install-chef-suse.sh` you cannot change the hostname anymore. Services like Crowbar, Chef, and the RabbitMQ will fail when having changed the hostname.

Browsing to the Chef Web UI gives a `Tampered with cookie` error:

You probably have an old cookie in your browser from a previous Chef installation on the same IP. Remove the cookie named `_chef_server_session_id` and try again.

6.1.2 OpenStack Node Deployment

How can I log in to a node as `root`?

By default you cannot directly log in to a node as `root`, because the nodes were set up without a `root` password. You can only log in via SSH from the Administration Server. You should be able to log in to a node with `ssh root@`*`NAME`* where *`NAME`* is the name (alias) of the node.

If name resolution does not work, go to the Crowbar Web interface and open the *Node Dashboard*. Click on the name of the node and look for its *admin (eth0) IP Address*. Log in to that IP address via SSH as user `root`.

A node refuses to boot or boots into a previous installation.

Make sure to change the boot order in the BIOS of the node, so that the first boot option is to boot from the network/PXE boot.

A node hangs during hardware discovery after the very first PXE boot into the “SLEShammer” image.

The SLEShammer image has no `root` password set, so log in in for debugging purposes is not possible. To set a `root` password for the SLEShammer for the node that hangs, proceed as follows:

1. Make sure you know the MAC address of the hanging node

2. Log in to the Administration Server as `root`

3. Create a directory named

   ```
   /updates/d<HOSTNAME>
   ```

   where *<HOSTNAME>* is the MAC address of the hanging node with lowercase letters (e.g. `de-9a-88-bd-ff-c1`). Note that the MAC address is always prefixed by the lowercase letter `d`.

   Create a hook script `/updates/<HOSTNAME>/discovery-pre.hook` with the following content:

   ```
   #!/bin/bash
   echo "linux" | passwd --stdin root
   ```

   In this example, the `root` password is set to *`linux`*. You may want to change it to a more secure password.

4. Log in to the Crowbar Web interface and delete the hanging node from the pool as described in Section 4.4, “Editing Allocated Nodes” (page 46).

5. Reboot the hanging node to restart the hardware detection. Now you will be able to log in to the SLEShammer image as `root` using the password supplied by the hook script. In order to find out why the node hangs, you may want to look at `/var/log/chef/client.log` first.

When deploying a node after having allocated it, it fails to PXE boot with the following error message: `Could not find kernel image: ../suse-11.2/install/boot/x86_64/loader/linux`

The installation repository at `/srv/tftpboot/suse-11.2/install` on the Administration Server has not been set up correctly to contain the SUSE Linux Enterprise Server 11 SP2 installation media. Please review the instructions at Section 3.2.2, “Setting Up Local Repositories” (page 33).

When deploying a node after having allocated it, it hangs at `Unpacking initramfs` during PXE boot.

The node probably does not have enough RAM. You need at least 2 GB RAM.

After the installation of a node has finished, it hangs at `Executing AutoYast script: /var/adm/autoinstall/init.d/crowbar_join:`

Log in to the node as `root` and check the log files at `/var/log/crowbar-join*` for errors.

Applying a Barclamp proposal fails.

Check the Chef client logs located on the node(s) affected by the proposal (`/var/log/chef/client.log`), and also the logs of the service that failed to be deployed. Additional information may be gained from the Crowbar Web UI logs on the Administration Server. For a list of log file locations refer to Appendix A, *Log Files* (page 63).

6.2 Support

Whenever you contact support to help you with a problem on SUSE Cloud, it is strongly recommended that you gather as much information about your system and the problem as possible. For this purpose SUSE Cloud ships with a tool called `supportconfig`. It gathers system information such as the current kernel version being used, the hardware, RPM database, partitions, and other items. `supportconfig` also collects the most important log files, making it easier for the supporters to identify and solve your problem.

It is recommended to always run `supportconfig` on the Administration Server as well as on the Controller Node. If a Compute Node or a Storage Node is part of the problem, run `supportconfig` on the affected node as well. For details on how to run `supportconfig`, please refer to http://www.suse.com/documentation/sles11/book_sle_admin/data/cha_adm_support.html.

A Log Files

Find a list of log files below, sorted according to the nodes where they can be found.

A.1 On the Administration Server

- Crowbar Web Interface: `/opt/dell/crowbar_framework/log/production.log`
- Chef Web Interface: `/var/log/chef/webui.log`
- Chef server: `/var/log/chef/server.log`
- Chef expander: `/var/log/chef/expander.log`
- Chef client (for the Administration Server only): `/var/log/chef/client.log`
- Apache SOLR (Chef's search server): `/var/log/chef/solr.log`
- HTTP (AutoYaST) installation server for provisioner Barclamp: `/var/log/apache2/provisioner-{access,error}_log`
- Default SUSE log files: `/var/log/messages`, `/var/log/zypper.log` etc.
- Syslogs for all nodes: `/var/log/nodes/*.log` (these are collected via remote syslogging)

- Log file from mirroring SMT repositories (optional): `/var/log/smt/smt-mirror.log`
- Other client node log files saved on the Administration Server:
  - `/install-logs/h*.log`: Initial Chef client run on PXE-booted nodes prior to discovery by Crowbar.
  - `/opt/dell/crowbar_framework/log/d*.log`: Output from Chef client when proposals are applied to nodes. This is the first place to look if a Barclamp proposal fails to apply.

A.2 On All Other Crowbar Nodes

Logs for when the node registers with the Administration Server:

- `/var/log/crowbar-join.errlog`
- `/var/log/crowbar-join-$TOPIC.{log,err}`: STDOUT/STDERR from running commands associated with $TOPIC when the node joins the Crowbar cluster. $TOPIC can be:
  - `zypper`: package management activity
  - `ifup`: network configuration activity
  - `Chef`: Chef client activity
  - `time`: starting of ntp client
- Chef client log: `/var/log/chef/client.log`
- Default SUSE log files: `/var/log/messages`, `/var/log/zypper.log` etc.

A.3 On the Controller Node

- `/var/log/keystone/keystone.log`: OpenStack authentication, etc.

- `/var/log/rabbitmq/*`: logs for RabbitMQ, used by OpenStack for handling message queues
- `/var/log/nova/`: various logs relating to Nova services:
    - `api.log`
    - `consoleauth.log`
    - `network.log`
    - `nova-manage.log`
    - `scheduler.log`
    - `volume.log`
- `/var/log/apache2/openstack-dashboard-*`: Logs for the SUSE Cloud Dashboard

A.4 On Compute Nodes

`/var/log/nova/`: various logs relating to Nova services:

- `compute.log`
- `nova-manage.log`

A.5 On Nodes with Ceph Barclamp

`/var/log/ceph/*.log`

Terminology

Administration Server
: Also called Crowbar Administration Node. Manages all other nodes. It assigns IP addresses to them, PXE boots them, configures them, and provides them the necessary software for their roles. To provide these services, the Administration Server runs Crowbar, Chef, DHCP, TFTP, NTP, and other services.

Amazon Elastic Block Store (EBS)
: Block-level storage volumes for use with Amazon EC2 instances. One or more storage volumes can be attached to running instances and persist after the instance has been shut down. It is possible to boot from an EBS volume.

Amazon Elastic Compute Cloud (EC2)
: Amazon's virtual computing environment.

Amazon Machine Image (AMI)
: A virtual machine that can be created and customized by a user. AMIs can be identified by an ID prefixed with `ami-`.

Amazon Simple Storage Service (S3)
: A storage for the Internet that can be used to store and retrieve data on the Web.

Amazon Web Services (AWS)
: A collection of remote computing services (including Amazon EC2, Amazon S3, and others) that together make up Amazon's cloud computing platform.

Barclamp
: A set of Chef cookbooks, templates, and other logic. Used to apply a particular role to individual nodes or a set of nodes.

Chef
: An automated configuration management platform for deployment of your entire cloud infrastructure. The Chef server manages many of the software packages and allows the easy changing of nodes.

Ceph
: A massively scalable, open source, distributed storage system. It consists of an object store, a block store, and a POSIX-compliant distributed file system.

Compute Node
: Node within a SUSE Cloud. A physical server running a Hypervisor. A Compute Node is a host for guest virtual machines (VMs) that are deployed in the cloud. It starts virtual machines on demand using `nova-compute`. To split VM load across more than one server, a cloud should contain multiple Compute Nodes.

Controller Node
: Node within a SUSE Cloud. The Controller Node is configured through the Administration Server and registers with the Administration Server for all required software. Hosts the OpenStack API endpoints and the OpenStack scheduler and runs the `nova` services—except for `nova-compute`, which is run on the Compute Nodes. The Controller Node coordinates everything about cloud VMs: like a central communication center it receives all requests (for example, if a user wants to start or stop a VM) and communicates with the Compute Nodes to coordinate fulfillment of the request. A cloud can contain multiple Controller Nodes.

Cookbook
: A collection of Chef recipes which deploy a software stack or functionality. The unit of distribution for Chef.

Crowbar
: Bare-metal installer and an extension of Chef server. The primary function of Crowbar is to get new hardware into a state where it can be managed by Chef. That means: Setting up BIOS and RAID, network, installing a basic operating system, and setting up services like DNS, NTP, and DHCP. The Crowbar server manages all nodes, supplying configuration of hardware and software.

Ephemeral Disk
: Ephemeral disks offer machine local disk storage linked to the lifecycle of a virtual machine instance. When a virtual machine is terminated, all data on the ephemeral disk is lost. Ephemeral disks are not included in any snapshots.

Flavor
: The compute, memory, and storage capacity of `nova` computing instances (in terms of virtual CPUs, RAM, etc.).

Hybrid Cloud
: One of several deployment models for a cloud infrastructure. A composition of both public and private clouds that remain unique entities, but are bound together by standardized technology for enabling data and application portability.

Infrastructure-as-a-Service (IaaS)
: A service model of cloud computing where processing, storage, networks, and other fundamental computing resources are rented over the Internet. It allows the customer to deploy and run arbitrary software, including operating systems and applications. The customer has control over operating systems, storage, and deployed applications but does not control the underlying cloud infrastructure. Housing and maintaining it is in the responsibility of the service provider.

Instance
: A virtual machine that runs inside the cloud.

Node
: A (physical) server that is managed by Crowbar.

OpenStack
: A collection of open source software to build and manage public and private clouds. Its components are designed to work together to provide Infrastructure as a Service and massively scalable cloud computing software.

  At the same time, OpenStack is also a community and a project.

OpenStack Compute (Nova)
: One of the core OpenStack components and services. It is a cloud computing fabric controller and as such, the main part of an IaaS system. It provides virtual machines on demand.

OpenStack Dashboard (Horizon)
: One of the core OpenStack components or services. It provides a modular Web interface for OpenStack services and allows end users and administrators to interact with each OpenStack service through the service's API.

OpenStack Identity (Keystone)
: One of the core OpenStack components or services. It provides authentication and authorization for all OpenStack services.

OpenStack Image (Glance)
: One of the core OpenStack components or services. It provides discovery, registration, and delivery services for virtual disk images.

OpenStack Network (Quantum)
: One of the core OpenStack components or services. It provides "network connectivity as a service" between interface devices (for example, vNICs) managed by other OpenStack services (for example, `nova`). Allows users to create their own networks and attach interfaces to them.

OpenStack Object Store (Swift)
: One of the core OpenStack components or services. Allows to store and retrieve files while providing built-in redundancy and fail-over. Can be used for backing up and archiving data, streaming data to a user's Web browser, or developing new applications with data storage integration.

OpenStack Service
: A collection of Linux services (or daemons) that work together to provide core functionality within the OpenStack project, like storing objects, providing virtual servers, or authentication and authorization. All services have code names (noted in brackets), which are also used in configuration files and command line programs that belong to the service.

Platform-as-a-Service (PaaS)
: A service model of cloud computing where a computing platform and cloud-based application development tools are rented over the Internet. The customer controls software deployment and configuration settings, but not the underlying cloud infrastructure including network, servers, operating systems, or storage.

Project
: A concept in OpenStack Identity. Used to identify a group, an organization, or a project. Also called `tenant`. The term `tenant` is primarily used in the OpenStack command line tools, but occasionally also appears in the SUSE Cloud Dashboard.

Proposal
: Special configuration for a Barclamp. It includes Barclamp-specific settings, and a list of nodes to which the proposal should be applied.

Private Cloud
: One of several deployment models for a cloud infrastructure. The infrastructure is operated exclusively for a single organization and may exist on or off premises. The cloud is owned and managed by the organization itself, by a third party or a combination of both.

Public Cloud
: One of several deployment models for a cloud infrastructure. The cloud infrastructure is designed for use by the general public and exists on the premises of the cloud provider. Services like applications, storage, and other resources are made available to the general public for free or are offered on a pay-per-use model. The infrastructure is owned and managed by a business, academic or government organization, or some combination of these.

qcow (QEMU Copy on Write)
: A disk image format supported by the QEMU virtual machine manager. A `qcow2` image helps to optimize disk space as it consumes disk space only when contents are written on it and grows as data is added.

  `qcow2` is a more recent version of the `qcow` format where a read-only base image is used, and all writes are stored to the `qcow2` image.

Quota
: Restriction of resources to prevent overconsumption within a cloud. In OpenStack, quotas are defined per project and contain multiple parameters, such as amount of RAM, number of instances, or number of floating IP addresses.

Recipe
: A group of Chef scripts and templates. Recipes are used by Chef to deploy a unit of functionality.

Role
: In the Crowbar/Chef context: an instance of a Proposal (page 70) that is active on a node.

  In the OpenStack Identity (Keystone) (page 69) context: concept of controlling the actions that a user is allowed to perform.

Software-as-a-Service (SaaS)
: A service model of cloud computing where applications are hosted by a service provider and made available to customers remotely as a Web-based service.

Storage Node
: Node within a SUSE Cloud. Acts as the controller for cloud-based storage. A cloud can contain multiple Storage Nodes.

Tenant
: See Project (page 70).

Security Group
: Concept in OpenStack Compute. A collection of network access rules, like firewall policies. The access rules specify which incoming network traffic should be delivered to all virtual machines in the group. All other incoming traffic is discarded.

Volume
: Detachable block storage device. Unlike a SAN, it can only be attached to one instance at a time.