

# Getting Started with an OpenStack-based Cloud Using SUSE® Cloud to Run SAP Applications

Written by B1 Systems GmbH

May 2014



Introduction to the Software Components .....	3
Introduction to SUSE Cloud.....	3
Installing SUSE Cloud.....	8
SUSE Linux Enterprise Server for SAP Applications .....	22
Installing the SAP Software .....	31
Outlook .....	39

## Introduction to the Software Components

This document provides a basic guide to the installation of a SUSE® Cloud environment and describes how you prepare a SUSE Linux Enterprise Server for SAP Applications image for use in the cloud. It outlines which aspects are to be taken into account and how an SAP instance is generated in the cloud.

The following products are included in this document:

- **SUSE Cloud:** A cloud based on OpenStack (the leading open source cloud project). SUSE Cloud allows you to create and manage virtual machines easily and without advanced system expertise.
- **SUSE Linux Enterprise Server for SAP Applications:** A specialized SUSE LINUX distribution with advanced support, system adjustment to SAP requirements and an integrated high-availability add-on and anti-virus product.
- **SAP NetWeaver:** A technology platform for particularly flexible and powerful business applications, such as ERP or business intelligence applications. SAP NetWeaver consists of a database layer, an application layer and a presentation layer.

## Introduction to SUSE Cloud

Powered by OpenStack, SUSE Cloud is an IaaS (infrastructure as a service) solution that enables the deployment and management of a private cloud within a company.

Unlike a native installation of OpenStack packages, which requires manual adjustments at various points in the individual components, SUSE Cloud offers the advantage of automated installation of the servers that form the cloud infrastructure. The individual servers can be rolled out centrally from here using either the CLI or a Web GUI. This allows a complete cloud to be deployed for productive use within a short space of time. Compared with the OpenStack project, the enterprise support provided makes SUSE Cloud the product of choice for productive use in the business environment.

SUSE Cloud enables multiple hypervisors to be run in the same cloud. Supported hypervisors include KVM, Xen, Microsoft Hyper-V and VMware ESXi through integration with VMware vCenter. The support for mixed hypervisor environments ensures increased flexibility in designing cloud deployments and helps optimize licensing costs.

This document also looks at the components that make up SUSE Cloud and describes how these are installed. It describes a minimal cloud only, which forms the basis for further scenarios. Depending on requirements, significantly larger scenarios can also be modeled.

### SUSE Cloud Components

SUSE Cloud is characterized in particular by its manageability, speedy and straightforward installation and easy configuration. Each of the SUSE Cloud components and their function are introduced briefly in the following paragraphs.

SUSE Cloud was developed based on OpenStack. OpenStack comprises a number of core projects. The incorporation of functionality that allows the infrastructure components of OpenStack to also be managed by means of configuration management means the cloud can be fully installed and configured from a server (known as the administration server). The SUSE Cloud packages are added to the administration server during the installation

process, allowing the control, compute and storage servers to then be installed and updated and the required services to be rolled out.

The administration server is used to configure all other components, and Crowbar is used to install additional servers. Initiated by Dell, the Crowbar open source project allows servers to be installed and configured centrally. SUSE and Dell are the driving forces behind this project. Crowbar uses what are known as barclamps (modules) to carry out server installation and configuration systematically.

The actual configuration settings as well as any configuration changes made are distributed further by means of Chef. Chef is configuration management software. The individual configurations are divided into barclamps, which means the administrator needs only to update the part to be changed rather than the entire configuration for all components. Chef is then used to roll out the new configuration. All machines in the infrastructure for SUSE Cloud obtain configuration settings in this way— from the central admin server.

When SUSE Cloud is installed, one or more control nodes and as many compute nodes as required are installed in addition to the admin server. The number of control and compute nodes depends largely on the applications that are to be run in the cloud. Only a minimal scenario from the required components is outlined in this document.

The control node is responsible for the network configuration of the virtual machines and provides additional server services required by the OpenStack components as well as a Web interface for managing the cloud. This Web interface is very different from that provided by Crowbar. The Web interface based on Crowbar is used to manage SUSE Cloud and the corresponding components (network and physical servers), whereas the Web interface provided on the control node is used “only” to manage the virtual infrastructure of the cloud. The database server that the OpenStack components use to save data also runs on the control node. A service called “AMQP” is also provided, which, as a centrally used messaging back end, provides the message queues required by OpenStack. The various OpenStack components use these message queues to communicate with one another. The user and identity management solution provided by OpenStack (Keystone) is also installed here. The dynamic network configuration required for the virtual machines is controlled by the Neutron service that runs here. Along with the publicly addressable OpenStack APIs, the central service for managing operating system images for virtual machines (Glance) rounds off the server services provided by the control node.

The role of the compute nodes is to host virtual machines. These machines can have different fundamental units (with respect to the vCPUs, main memory, network, etc.) and can, of course, be started using various operating system images as a basis. The service installed and configured there by OpenStack is the Nova service.

It is also possible to set up storage nodes, which can provide storage for the cloud should this be required. This is recommended for production environments. The block storage service (Cinder) is rolled out on the control node for the setup outlined in this document.

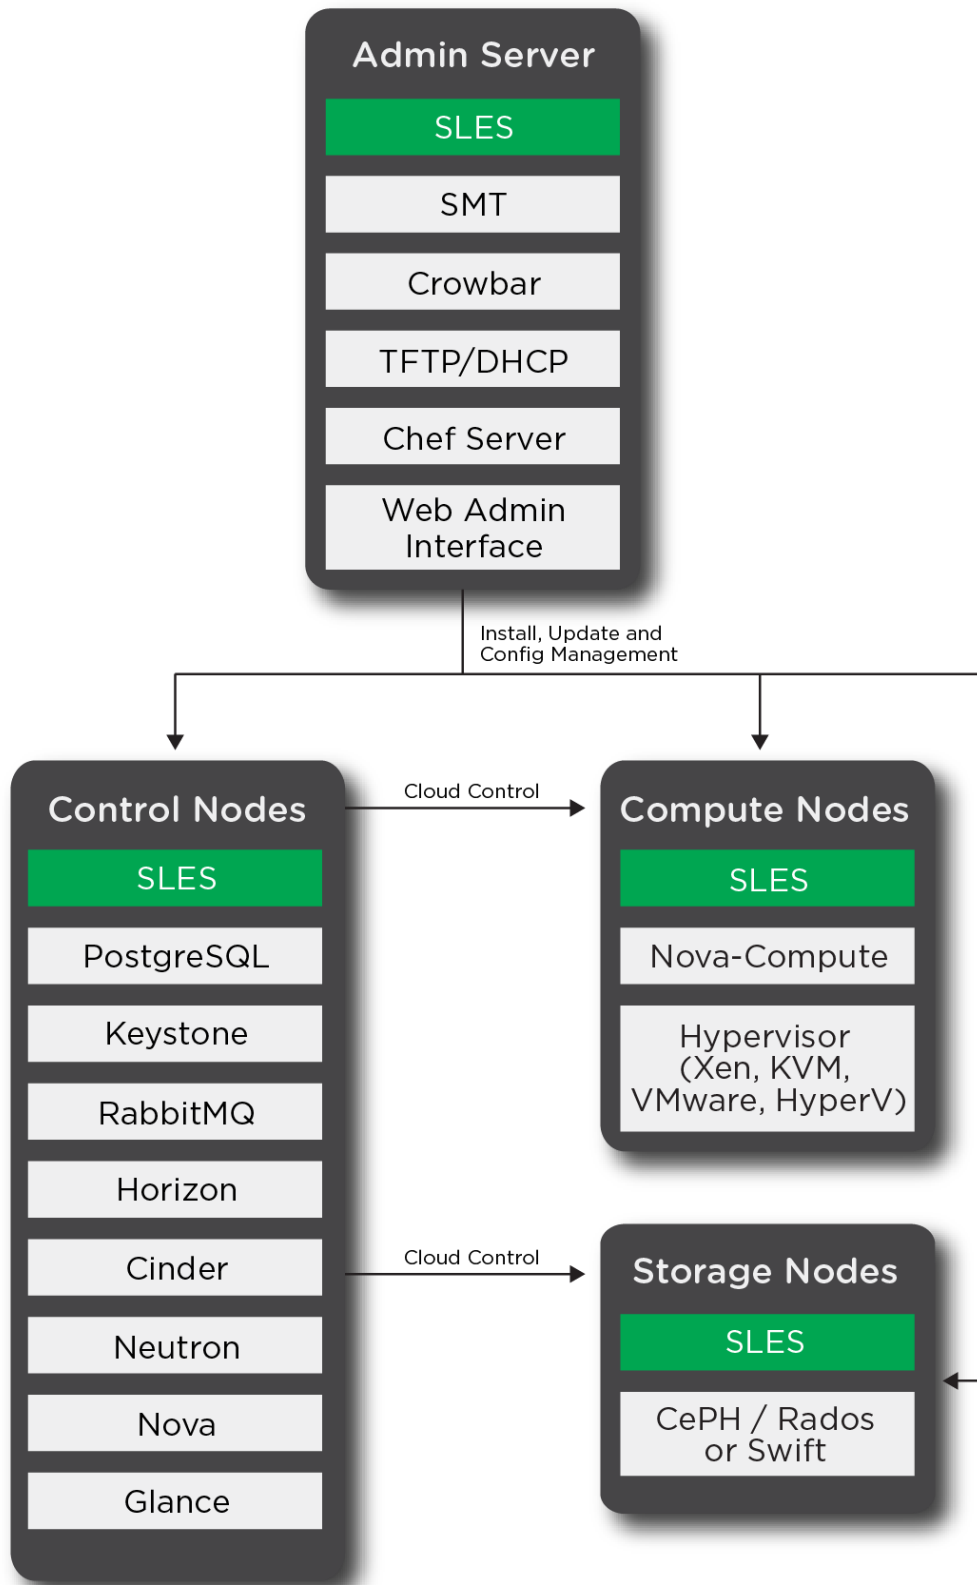


Figure 1: Overview of SUSE Cloud

## Admin Server

The admin server is responsible for installing and configuring all further nodes. Updates for both the operating system and SUSE Cloud are made available as is a PXE server. By using this, all other nodes can be booted and installed.

### **Hardware requirements:**

- X86\_64
- At least 2 GB RAM (4 GB RAM recommended)
- Hard disk of at least 40 GB (60 GB preferable)

### **Software requirements:**

The following software components are required to install an administration server:

- SUSE Linux Enterprise Server (latest release)
- SMT Add-On (latest release)
- SUSE Cloud (latest release)

You can find the relevant download links on the following page: <https://download.suse.com/index.jsp>.

### **Services provided by the admin server:**

- Crowbar: an open source project developed by SUSE and Dell that allows server installation and configuration to be controlled centrally
- Chef server: the server responsible for configuration management that also updates the configuration files for the services rolled out by Crowbar
- SMT (Subscription Management Tool for SUSE Linux Enterprise): provides updates from the SCC (SUSE Customer Center) to keep its own software version and that of the nodes installed by the administration server up to date
- TFTP/DHCP server: installs the nodes and assigns them dynamic IP addresses, from which Crowbar later generates static addresses

## Control Node

A control node makes available all OpenStack services that are required to manage the virtual machines that are started on the compute nodes.

### **Hardware requirements:**

- X86\_64
- 2 GB RAM recommended, depending on the services required
- Hard disk: requirement varies depending on how many different images are to be made available and whether the images require additional memory as block storage. If block storage is made available using raw devices, at least one additional hard disk should be provided to make these available separately to the instances.

### **Services provided by the control node:**

- Database: PostgreSQL database; stores data for most OpenStack services
- Identity service (Keystone):
  - Provides authentication and authorization services for the OpenStack components
  - Manages the following objects:
    - User: users
    - Tenants: customers/projects
    - Role: assigns authorizations to users

- Message broker (RabbitMQ): the individual OpenStack services do not communicate directly with one another, but use the Advanced Message Queueing Protocol (AMQP). This functionality is made available by a standalone server service (RabbitMQ).
- Image service (Glance): registers and manages image files that can then be booted as an instance
- Block storage (Cinder): provides block storage to make additional storage space available to instances and save data persistently. In addition, block devices are made available as hard disks for the images, thereby improving the performance of virtual machines when compared with image files.
- Networking (Neutron): Network as a Service: Provides IP addresses to instances and if applicable, separates networks, thereby allowing different customers to be hosted. However, some configuration file names still contain references to “quantum,” the old name of the service.
- Compute (Nova):
  - Responsible for managing virtual machines and consists of a number of individual services:
    - Nova API: interface for commands
    - Nova conductor: controls database access
    - Nova scheduler: selects a suitable host to run an instance
- Dashboard (Horizon): Horizon provides a Web-based dashboard for OpenStack. This allows the OpenStack components to be administered and virtual machines to be managed.

## Compute Node

The virtual machines that contain the SAP application are subsequently deployed on the compute nodes.

The hardware must facilitate the overall requirements of the applications. You can find the various specifications for the SAP applications in the documentation provided or by using the online sizing tool from SAP.

### **Hardware requirements:**

- X86\_64
- CPU and RAM: As much as possible to run the virtual machines in parallel. Please use the SAP sizing tools for virtual machines or ask your hardware partner about suitable sizing for your SAP workload.

### **The following services are run on the compute nodes:**

- Nova compute:
  - Starts virtual machines
  - Capable of using various hypervisors to run virtual machines. These include KVM, Xen, Microsoft Hyper V, and VMware vSphere. Only the following virtualizations are supported for SAP installations: KVM, Xen, VMware ESX (see SAP Note 1122387).
- Neutron Linuxbridge Agent: facilitates network communication of the instances

## Storage Node

Storage nodes are one or more servers that provide either block storage or object storage. Block storage is storage that is offered to clients in the form of block devices. Clients can treat these as new (unpartitioned and unformatted) hard disks. Object storage, on the other hand, allows large files to be uploaded and downloaded using an HTTP(s) interface.

Block storage can be made available either using the OpenStack Cinder component or using Ceph. The provision of block storage using Ceph is offered only as a technical preview in the example used (SUSE Cloud 3), however.

Object storage is made available using the OpenStack Swift service, which offers automatic data replication. For larger setups, it is recommended to provision the Cinder block storage service on a separate node with sufficient memory. In the simple setup described in this document, the service is rolled out on the control node.

**Hardware requirements:**

- X86\_64
- The RAM and hard disk requirements depend on the type and amount of storage to be made available. Additional storage nodes are not required for the setup outlined in this document.

**Services provided on the storage node:**

- Ceph/RADOS: Ceph provides object storage that can be addressed as block storage using RADOS (technical preview, not yet supported).
- Swift: object storage for storing files, images, snapshots, and their automatic replication

## Installing SUSE Cloud

### Installing the Admin Server

To install the admin server, it must be booted by a SUSE Linux Enterprise Server image.

SMT installation is optional and can be skipped if an SMT or SUSE Manager for updating the nodes is available elsewhere. If this is not the case, an SMT has to be installed on the admin server.

If an SMT server is also installed, steps must be taken when partitioning during installation to ensure that the size of the /srv folder is at least 25 GB since this is where the SMT manages its updates.

The following add-on products should be added and installed during the installation phase:

- SMT (if required)
- SUSE Cloud

The following packages should be selected when the packages and package groups are selected:

- Base system
- Minimal system (appliances)
- SUSE Cloud admin server
- Subscription management tool (optional)
- Web and LAMP Server (only required if an SMT is installed)

Figure 2 Big Picture - SUSE Cloud for SAP on page 11 provides an overview of the planned setup.



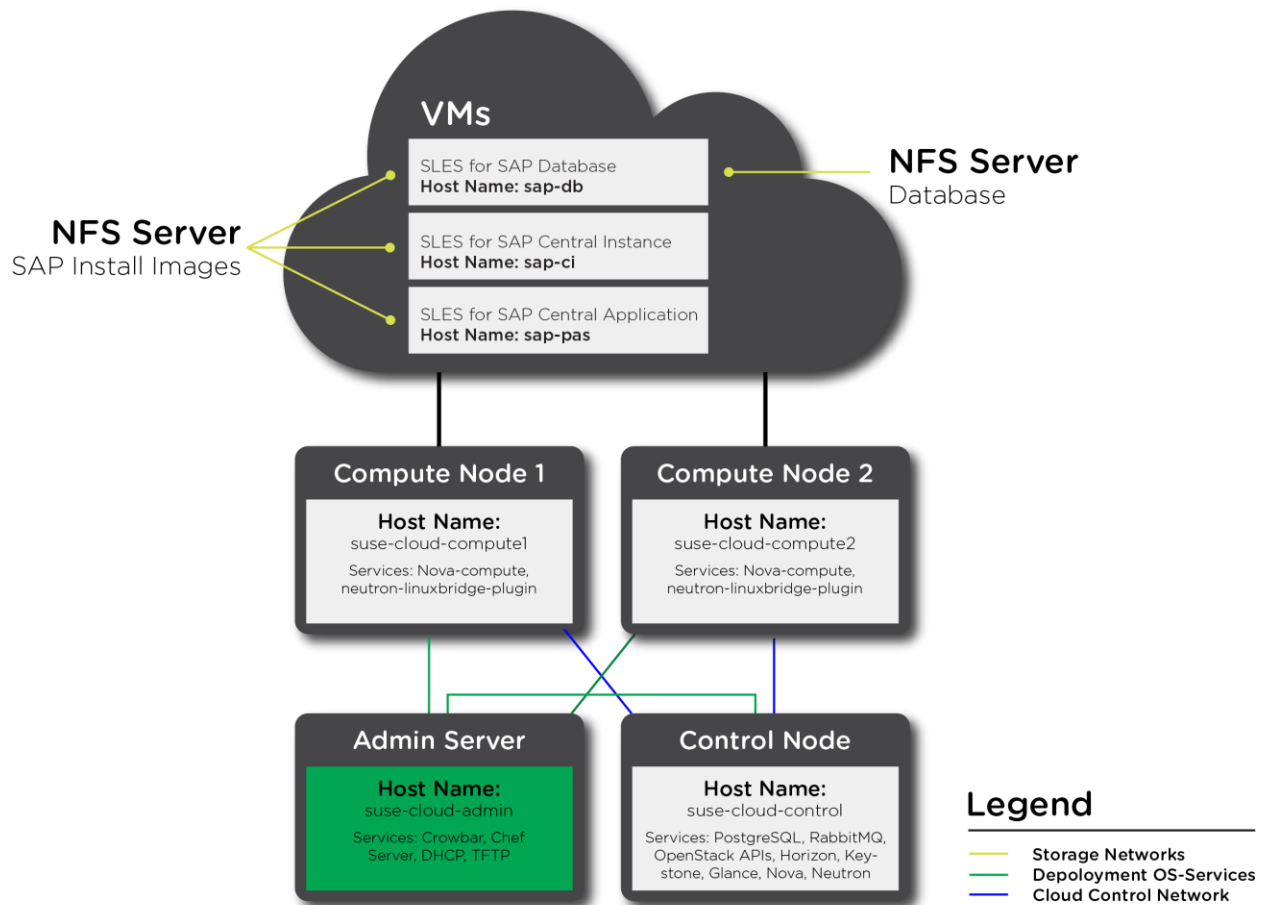


Figure 2: Big Picture - SUSE Cloud for SAP

Once installation of the minimal SUSE Cloud components has been completed, the network is the first thing that should be configured. SUSE Cloud requires a complex network setup comprising multiple networks that are configured during installation. These networks are for cloud use only. A router is required to access them from an existing network.

The network configurations on the nodes in SUSE Cloud are controlled completely by Crowbar. Any network configurations that do not take place using Crowbar (with YaST, for example) are automatically overwritten.

Network settings cannot be changed once the cloud has been provided.

For a detailed description of networks and options, see the [SUSE Cloud Deployment Guide](#).

Three different network types can be selected to meet the various requirements for applications in the cloud:

- Single = one network card for admin/public network
- Double = two network cards, one for public and one for the admin network
- Team = at least two network cards that are combined to create a bonding device, either for the purposes of load balancing or to safeguard against failure

(See [SUSE Cloud Deployment Guide](#) — under “Considerations and Requirements,” “Network.”)

The network type selected in this setup for SAP applications is “Team.” A Linux bonding device consisting of at least two network cards is created that then manages the network traffic within the cloud. The bonding mode selected determines how many network cards are required.

It is also important to ensure that the same number of network cards must be available in the bonding on all nodes that are installed using the administration server.

The bonding mode to be selected in this setup depends on various considerations, such as performance or safeguarding against failure. The mode is set for SUSE Cloud by specifying the relevant number (see Table 1: Bonding Modes) in YaST. If performance is considered particularly important, modes 0 or 5 are recommended whereas mode 1 is recommended to safeguard well against failure. The default value in SUSE Cloud is 5.

Number	Bonding Mode	Explanation
0	balance-rr	Round-robin rule: Transfer packages in sequential order from the first available slave until the last. This mode supports load balancing and fault tolerance.
1	active-backup	Only one slave is active in the bond. A second slave is only activated if the active slave is down. The MAC address of the bond is only visible externally via one port, which avoids confusion for the switch.
2	balance-xor	XOR rule: Transfer packets based on the result of an XOR operation of the MAC address of the sender with the MAC address of the receiver.
3	broadcast	Transfer everything using all slave interfaces. This mode supports fault tolerance.
4	802.3ad	This mode creates a bond device in accordance with IEEE standard 802.3ad – Link Aggregation Control Protocol (LACP). A bond with the same speed and duplex settings is created. It uses all slaves in the bond as per the 802.3ad specification. The switch must support this and needs to be configured accordingly.
5	balance-tlb	Adaptive transmit load balancing: Channel bonding that does not require any special switch support. The outgoing data traffic is distributed based on the current load on each slave. Incoming data traffic is received by the currently designated slave. If the slave fails to receive the traffic, another slave takes on the MAC address of the faulty slave.

Number	Bonding Mode	Explanation
6	balance-alb	Adaptive load balancing: Includes balance-tlb and additional load balancing. Special switch support is not required.

Table 1: Bonding Modes

Other hardware (switches) used in the network structure must, of course, also be taken into account.

The following default IP address ranges used by the admin network:

- gateway → 192.168.124.1
- admin → 192.168.124.10 – 192.168.124.11
- dhcp → 192.168.124.21 – 192.168.124.80
- host → 192.168.124.81 -192.168.124.160

These network configurations are rolled out to the other nodes at a later stage.

When the control and compute nodes are being installed, dynamic IP addresses are initially assigned by the DHCP server to install the nodes. Fixed IPs from the “host” range are subsequently assigned to the nodes.

An IP address from the “admin” range must be assigned to the admin server as a fixed address. In our setup, this IP address is 192.168.124.10/24.

If you want to change the IP addresses for the admin network and VMs, you can use the following YaST Crowbar module:

```
# yast → Miscellaneous → Crowbar.
```

**Note: These IP addresses cannot be changed once SUSE Cloud has been installed!**

Since the IP addresses are used only within the cloud, an external network interface must exist to facilitate external access to the admin server. This “bastion network” still needs to be defined. An additional network card should be used for this purpose. If three network cards are used, for example, it is recommended that you define the last of these as a bastion network card to prevent issues when rolling out the network configuration.

It is not currently possible to set up a bastion network using YaST, so this must be done manually.

To do this, edit the /etc/crowbar/network.json file, and insert the following statements below the “admin” : { ... }” area (see also SUSE Cloud Deployment Guide — 3.1.9.3.1. Setting Up a Bastion Network):

```
"bastion" : {
"add_bridge" : false, "vlan" : 50,
"router" : "172.19.134.1",
"ranges" : {
"admin" : {
"start" : "172.19.134.45",
"end" : "172.19.134.45"
}
}
}
```

```
},  
"broadcast" : "172.19.255.255",  
"netmask" : "255.255.0.0",  
"conduit" : "bastion1", "subnet" : "172.19.0.0",  
"use_vlan" : false
```

You then also need to specify on which network card the bastion network is set up and on which network cards the bonding device is to be created.

You also specify this in the `/etc/crowbar/network.json` file, this time in the following area: "pattern":

```
"team/././.*",  
"conduit_list" : { "intf0" : {  
  "if_list" : [  
    "1g1", "1g2"  
  ]  
},  
"bastion1" : {  
  "if_list" : [  
    "1g3"  
  ]  
},  
},
```

Since the name assigned using `"conduit" : "bastion1"` was selected for the bastion network, this must also be reflected in the network card assignment.

In this example, the third network card (eth2) with the IP address 172.19.134.45 is assigned to the bastion network.

Network cards eth0 and eth1, which are configured as device "bond0," are assigned to the admin network. This assignment looks as follows:

1g1 = eth0

1g2 = eth1

1g3 = eth2

"1g" here stands for "1 Gigabit network cards." The number at the end indicates the sequence. To validate the syntax, run the following command again:

```
# yast → Miscellaneous → Crowbar
```

Your bastion network should then appear on the "Network" tab.

## SMT Setup

If your setup already includes an SMT server, you can skip the following section and proceed as described in “Providing Repositories when an SMT Server Already Exists” on page 17.

The easiest way of configuring the SMT server on the admin server is using YaST:

Invoke:

```
# yast2 smt-server
```

Then enter your mirroring credentials and specify a password for database access. You can find the mirroring credentials in the SUSE Customer Center (SCC).

The following channels from the SUSE Customer Center, for example, are required for the SUSE Cloud installation:

- SLES11-SP3-Pool
- SLES11-SP3-Updates
- SUSE-Cloud-3.0-Pool
- SUSE-Cloud-3.0-Updates

These reflect the latest operating system and SUSE Cloud versions available when this document was created. You should always use the latest version available.

To set up and synchronize the relevant software channels, use the following commands (specifying the latest version number):

```
# for REPO in SLES11-SP3-{Pool,Updates} SUSE-Cloud-3.0-{Pool,Updates}; do \ smt-repos $REPO sle-11-x86_64 -e \
done
# smt-mirror -L /var/log/smt/smt-mirror.log
```

The last command synchronizes the pools from the SCC, which means this may take some time to run.

### ***Providing Update Repositories for Deployment***

The repositories must then be linked from SMT to /srv/tftpboot so that they are subsequently available to the nodes to be updated (as before, please use the latest version numbers):

```
# for REPO in SLES11-SP3-{Pool,Updates} SUSE-Cloud-3.0-{Pool,Updates}; do \
ln -s /srv/www/htdocs/repo/\$RCE/$REPO/sle-11-x86_64 /srv/tftpboot/repos/$REPO \ done
```

To update the actual administration server (and ensure these updates are also available to the servers installed by the administration server), repositories are set up. Any repositories that reference installation sources that have been removed must first be deactivated, however. These repositories would otherwise be relayed but would not be accessible from the control and compute nodes (as before, ensure you use the latest version numbers):

```
# zypper modifyrepo --disable --no-refresh --remote
```

```
# for REPO in SLES11-SP{3-Pool,3-Updates} SUSE-Cloud-3.0-{Pool,Updates}; do \ zypper addrepo -f
/srv/tftpboot/repos/$REPO $REPO \
done
# zypper addrepo /srv/tftpboot/repos/Cloud "SUSE-Cloud-3.0"
# zypper addrepo /srv/tftpboot/suse-11.3/install/ "SLES 11 SP3"
```

### ***Providing Repositories when an SMT Server Already Exists***

If an SMT server is already available in your network, you can skip the previous section and simply need to execute the following commands:

```
# for REPO in SLES11-SP3-{Pool,Updates} SUSE-Cloud-3.0-{Pool,Updates}; do ATTR_REPO=$(echo $REPO
| sed "s:\.:\\\.:g")
/opt/dell/bin/json-edit \
-a "attributes.provisioner.suse.autoyast.repos.$ATTR_REPO.url" \
-v "http://smt-server.com/repo/\$RCE/$REPO/sle-11-x86_64/" /etc/crowbar/provisioner.json
done
```

Replace the name “smt-server.com” with the IP address or the FQDN of your SMT server and specify current version numbers.

### ***Product Repositories***

To enable SUSE Cloud to install other nodes using PXE, you still need to add the installation media for SUSE Linux Enterprise Server and SUSE Cloud on the admin server.

If the images are stored on a DVD, insert this DVD into the DVD drive. If the images are stored on an NFS server, mount the shared directory.

### ***Example of an SUSE Linux Enterprise Server 11 SP3 image:***

```
# mkdir -p /srv/tftpboot/suse-11.3/install/
```

DVD:

```
# mount /dev/dvd /media
```

Or Image on an NFS Server:

```
# mount ip-of-nfs-server:/shared directory /media
# mount -o loop /media/SLES-11-SP3_y86_64.img /media
```

### Copying the Medium:

```
# rsync -avP /media/ /srv/tftpboot/suse-11.3/install/  
# umount /media
```

### SUSE Cloud image:

```
# mkdir -p /srv/tftpboot/repos/Cloud/
```

### DVD:

```
# mount /dev/dvd /media
```

### Or Image on an NFS Server:

```
# mount ip-of-nfs-server:/shared directory /media  
# mount -o loop /media/SUSE-Cloud-2.0-DVD1.img /media
```

### Copying the Medium:

```
# rsync -avP /media/ /srv/tftpboot/repos/Cloud/  
# umount /media
```

## Installing SUSE Cloud Software

Before you proceed with the actual SUSE Cloud installation, you need to run some final commands:

- # hostname -f Check whether the FQDN of your administration server is displayed. If this is not the case, enter it in the /etc/HOSTNAME and /etc/hosts files. This is important because a DNS server is also installed during the SUSE Cloud installation, and this needs to know the correct naming schema.
- # zypper patch Install all updates to ensure the administration server is up-to-date.

To install SUSE Cloud, run the following command:

**Note: Changes cannot be made once the installation has been completed!**

```
# screen install-suse-cloud
```

This command starts the SUSE Cloud installation.

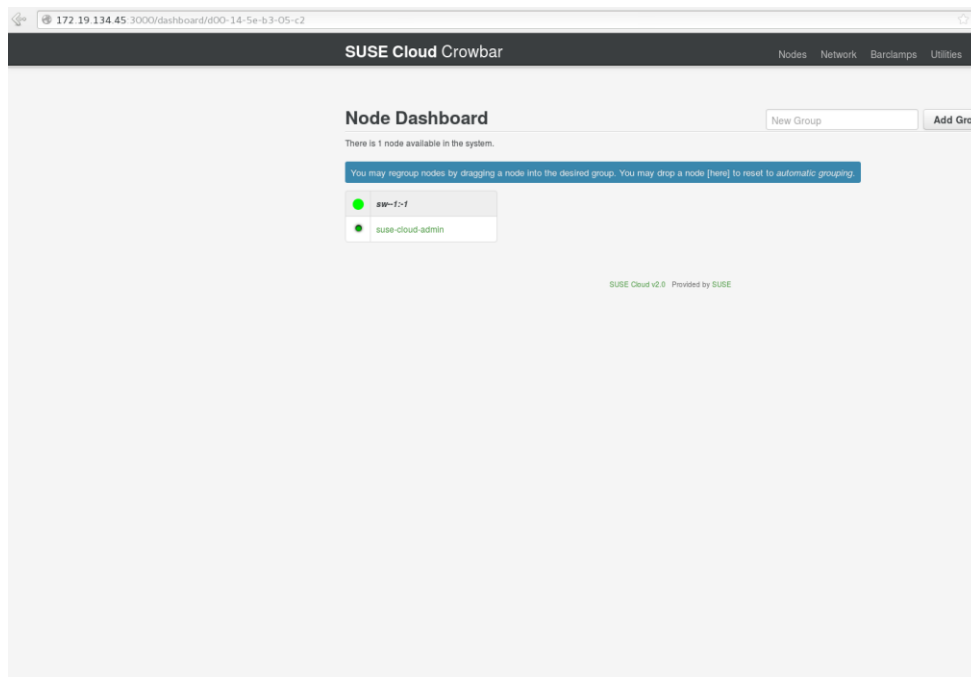
The screen command can be used to run the installation command at this point. The reasons for this are that it allows you to monitor the installation progress directly and that the connection may drop when setting up the network during installation. When you start the installation with the screen command, this will still continue to run.

The installation progress is indicated in the log file /var/log/crowbar/install.log.

Once installation has completed successfully, a message appears indicating that the Crowbar Web interface on port 3000 is to be opened. Both the username and password correspond with those set up automatically during the installation: `http://192.168.124.10:3000`

**User name: crowbar**

**Password: crowbar**



*Figure 3: Crowbar Web Interface*

### **Installing Additional Nodes**

To install the control node, you can simply have it booted by the network.

For information about how to configure the network boot settings for your servers, please see the documentation



provided by your hardware manufacturer. As soon as the image provided by the admin server has been booted, the node appears on the Web interface of the admin server. The node is assigned a color reflecting the stage of the installation (see Figure 4).

To complete the installation, click Nodes → Bulk Edit.

Then assign an alias name and a public name and choose “Allocate” (see Figure 6).

Figure 4: Complete List of All Nodes with Grouping Option

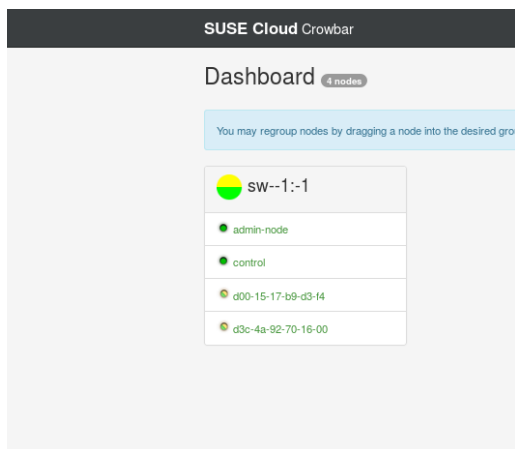
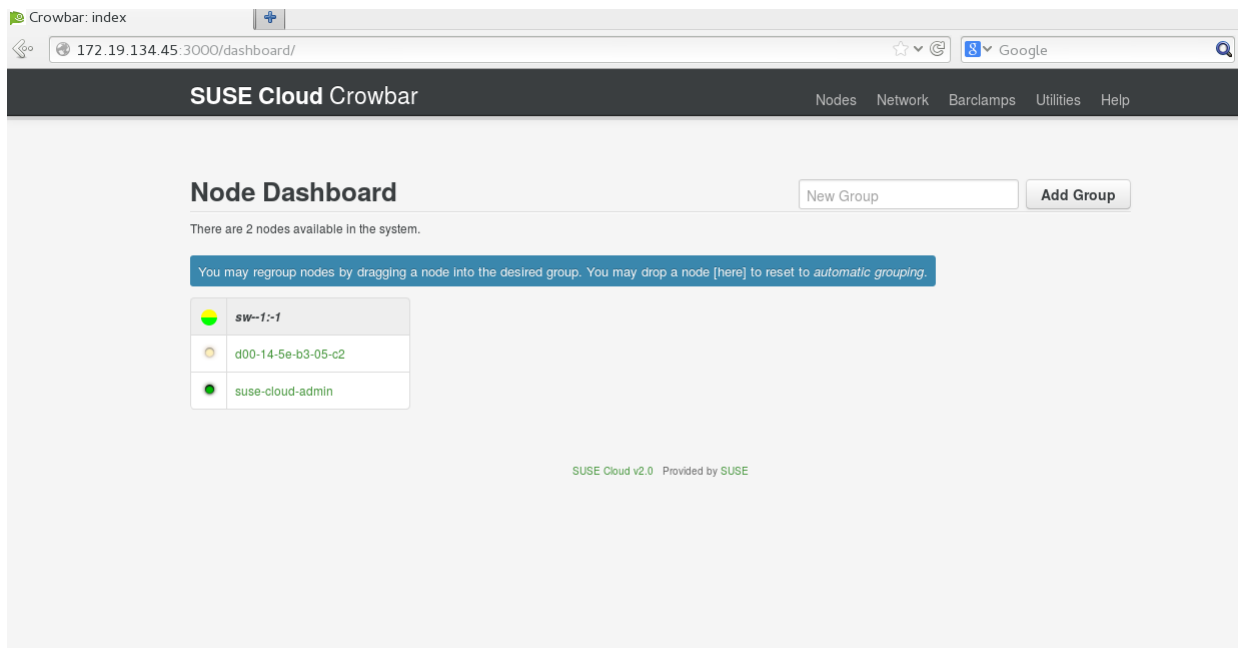


Figure 5: Overview of all nodes; nodes prefixed with “d” are currently being installed and have not yet been assigned new names

SUSE Cloud Crowbar

Nodes Network Barclamps Utilities Help

## Edit node

d00-15-17-b9-d3-f4.site Cancel Save

**Target Platform**  
SLES 11 SP3

**Alias**  
compute1

**Public Name**  
compute1

The public name is the hostname that users will use to access services on this node. Re-apply proposals on the node to have this setting taken into account immediately, or wait for an automatic update on the node. Any name specified here should already exist in the upstream DNS zones.

**Description**

**Group**  
Using default: sw-1-1

**Override group?**  
sw-1-1

**Intended Role**  
Undecided

Intended Role is the way you intend to use the node in your cloud infrastructure. The value is used to propose the initial node deployment for barclamps.

**Availability Zone**

Availability zones allow to arrange sets of either OpenStack Compute or OpenStack Block Storage hosts into logical groups. If empty, the default availability zone will be used.

Cancel Save

Figure 6: Editing Node Details (Name, Zones, etc.)

To ensure all nodes also receive the relevant updates, these should be distributed centrally using Crowbar.

Crowbar uses what are known as barclamps for this purpose. These are modules that specify what is to be executed or installed on which node. The Updater is then configured.

Barclamps → Crowbar → Updater

SUSE Cloud Crowbar

Nodes Network Barclamps Utilities Help

## • Edit Proposal

Updater Save Apply Delete Cancel

**Attributes** Raw

**Use zypper**  
patch

**Enable GPG checks**  
true

**Automatically agree with licenses**  
false

**Include patches that need reboots (Kernel)**  
false

**Reboot nodes if needed**  
false

Figure 7: Updater

Here, you specify all settings for updating packages on the systems involved. Configure the behavior of the system as required for kernel patches and restarts that may be required due to software updates.

You specify the assignment, that is, for which machines the updater is to be configured, at the end of the page in the “Deployment” section. Drag a node that has not been configured (“Available Nodes”) to the column on the right (“Updater”).

All barclamps are assigned in the same way.

## Configuring the Control Node

### *Rolling out the OpenStack services:*

The following services are rolled out with the default configuration on the control node:

- Database
- Keystone (set the administration password to “admin”)
- RabbitMQ
- Nova
- Horizon
- Cinder

The screenshot shows the SUSE Cloud Crowbar interface. At the top, there is a navigation bar with 'Nodes', 'Network', 'Barclamps', 'Utilities', and 'Help'. Below this is a yellow notification bar that says 'Successfully created proposal [clear]'. The main content area is titled 'Edit Proposal' and includes a 'Cancel' button. The 'Database' service is selected, with 'Delete', 'Apply', and 'Save' buttons. Under 'Attributes', there is a 'PostgreSQL Options' section with a 'Global Connection Limit (max\_connections)' field set to '1000'. The 'Node Deployment' section has a 'Raw' link and a instruction: 'Drag nodes for deployment from Available Nodes into the selected Role'. On the left, under 'Available Nodes', there is a search box and a list of nodes: 'suse-cloud-admin', 'suse-cloud-compute-01', and 'suse-cloud-control-01'. On the right, under 'database-server', there is a search box and a list of nodes: 'suse-cloud-control-01'. At the bottom right of the deployment section, there are 'Delete', 'Apply', and 'Save' buttons.

Figure 8: Deploying the OpenStack Services

We recommend that you configure the settings as follows for the “Neutron” barclamp to allow you to use the more powerful network configuration within the cloud:

Networking Plugin → linuxbridge

Using Linux bridges, the virtual machines and the real network cards are combined in a bridge and can therefore communicate. Any projects that differ from each other are separated using vLANs.

When a virtual machine is started, it is assigned an IP address from a configured network in exactly the same way as any other virtual machine. This is provided for individual virtual machines via DHCP. In this way, virtual machines of a project can communicate with each other provided they are located within the same configured network (Neutron network).

Networks configured in Neutron can also be configured as external networks. Addresses can then be mapped from these network areas to specific virtual machines. The virtual machine does not know its external address; accessibility is achieved using the Neutron L3 Agent.

These assigned IP addresses can be released at any time (during virtual machine runtime), and other virtual machines can be made available. IP addresses that are used in this way are known as “floating IPs.”

For the Nova barclamp, the nova-multi-controller is rolled out on the “suse-cloud-control-01” node and the nova-multi-compute-kvm on the “suse-cloud-compute-01” node. These names were configured when the node details were defined (see Figure 6).

The screenshot displays the 'Node Deployment' interface. On the left, under 'Available Nodes', there is a search bar and a list of three nodes: 'suse-cloud-admin', 'suse-cloud-compute-01', and 'suse-cloud-control-01'. On the right, several roles are listed with corresponding node assignments in input boxes: 'nova-multi-controller' is assigned to 'suse-cloud-control-01'; 'nova-multi-compute-hyperv' is empty; 'nova-multi-compute-kvm' is assigned to 'suse-cloud-compute-01'; 'nova-multi-compute-qemu' is empty; 'nova-multi-compute-xen' is empty; and 'nova-multi-compute-esxi' is empty. A 'Raw' link is visible in the top right corner.

Figure 9: Deploying the OpenStack Nova Service

Figure 10: Deploying the OpenStack Cinder Service

The following values should be configured for the Cinder barclamp to ensure that high performance and persistent storage are available to your virtual disks:

“Type of Volume” → Raw devices

This is the recommended value for performance-relevant applications. This setting is used to present the hard disks still available as raw devices to guests.

“Disk selection method” → All Available

This setting ensures that all free hard disks that are available are made available as a backing device for the VMs.

It should be noted, however, that the first hard disk (/dev/sda) is not generally free since your operating is installed on this.

“Name of Volume” →

Here, you can define a freely selectable name, to which prefixes will be added at a subsequent stage to ensure the volumes for each of the VMs are named consistently.

## Installing the Compute Nodes

Compute nodes are the servers on which the virtual machines are later started.

They are installed in the same way as the control node, but only Nova is rolled out on the compute nodes by the OpenStack services. It must be ensured here that the compute nodes are listed under “nova-multi-compute-\*”. The strong recommendation here is to use XEN or KVM. At this point, however, the SAP documentation must be consulted to determine which hypervisors are certified. When this document was created, these were XEN, KVM, and VMWare. In a cloud infrastructure, it is also possible to use different virtualization solutions on different nodes.

## SUSE Linux Enterprise Server for SAP Applications

SUSE Linux Enterprise Server for SAP Applications is an optimized operating system platform that is tailored specifically to the needs of SAP users.

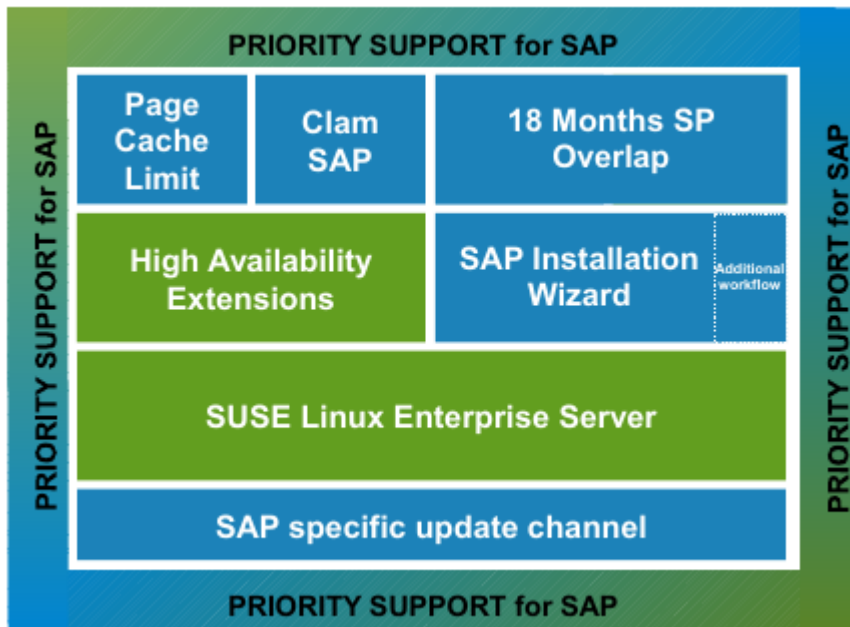


Figure 11: SLES Components for SAP Applications

Software:

- SUSE Linux Enterprise Server
- SUSE Linux Enterprise High Availability Extension
- An installation wizard for installing SAP applications and operating systems in one step in an automated workflow
- PageCacheLimit: optimizes the paging behavior of Linux for large databases and applications that are stored in the memory
- ClamSAP as a more cost-effective alternative for cross-platform virus protection using ClamAV and the SAP NetWeaver virus scanning API

Service:

- Support for the latest service pack is extended to 18 months (from 6 months).
- SUSE Linux Enterprise Priority Support: offers seamless 24x7 support from SAP and SUSE using the SAP ticket system.

In the upcoming sections of this document, we use SUSE Linux Enterprise Server for SAP Applications as the operating system for SAP applications.

### Preparing an Installation Server

An installation server can be used to make available the packages to be installed. This exports CDs/DVDs read using NFS for clients that can use this data for the installation.

A storage node can be used as the source for the installation server.

Install and configure the installation server with the following command:

```
# yast2 instserver
```

Then choose “Configure NFS repository”. Next, select a directory to which the data read is to be saved and exported using NFS (referred to as /install in the upcoming text). In the next dialog box, confirm the default settings for wild cards and options. Then add a new repository by choosing “Add” and rename it; in the example, this is SLES11SP3. Finally, make a DVD or an ISO image with the required files available to the wizard. The wizard will read and copy the data.

### Providing an NFS Server

Additional software packages can be made available easily using NFS. To do this, install an NFS server on a node (again, a storage node can be used for this purpose). Then generate a shared directory and copy the files to be made available to this directory:

```
# yast2 nfs_server
```

Choose “Add directory” to add a directory to the list of directories to be exported and specify the sharing options (host wild card, e.g. \* for all hosts) and export options.

When configuration has been completed in the wizard, copy the data to be shared to the directory that was just created. Files such as the SAP installation files can be stored here.

### Installing the Base Image

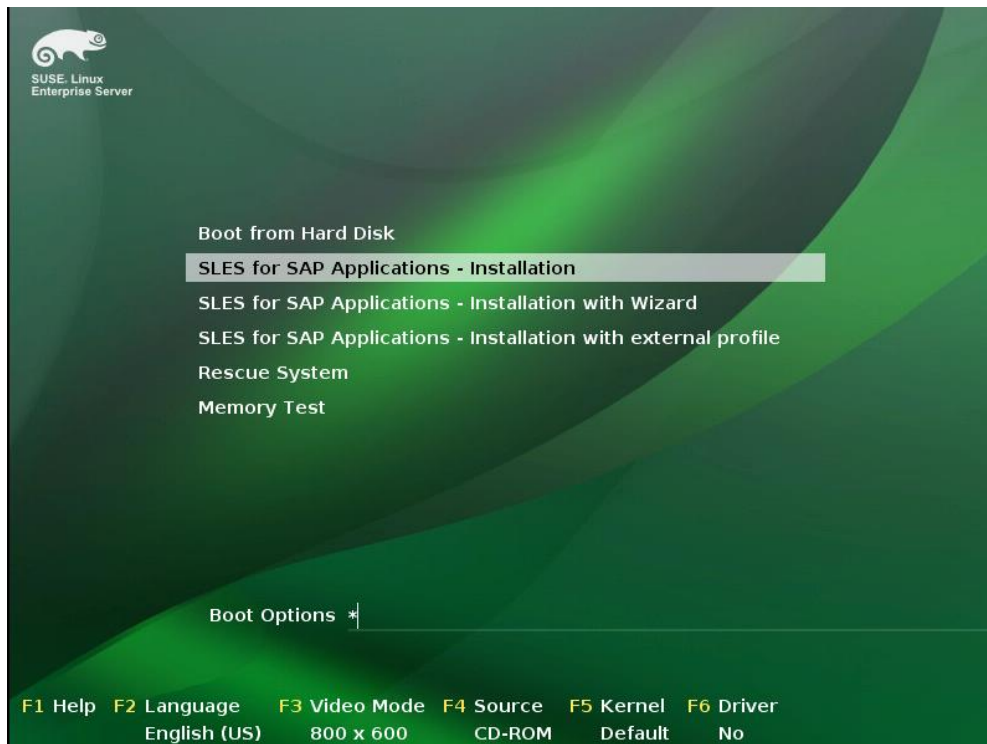


Figure 12: Options for Installing SLES for SAP Applications

When you install SUSE Linux Enterprise Server for SAP Applications, there are three options that you can use (see Figure 12):

- SLES for SAP Applications — Installation:  
The “basic” installation variant offers the same installation process as that for SUSE Linux Enterprise Server. This allows you to carry out the SAP installation manually or at a later stage using the installation wizard. Any packages required for SAP need to be selected manually.
- SLES for SAP Applications — Installation with wizard:  
With this installation variant, the operating system and SAP applications are installed by means of a wizard. The operating system and all applicable packages and optimizations as well as the SAP application are installed automatically.
- SLES for SAP Applications — Installation with external profile:  
With this type of installation, a distinct AutoYaST profile can be specified.

### ***Preparing SUSE Linux Enterprise Server for SAP Applications Image for Cloud Use***

For subsequent use, you need a base image, which is later copied twice and used as a basis for all three virtual machines that are to make SAP services available.

You install the operating system on a virtual machine on a separate host (or on a compute node) for this purpose. You use a disk image as the basis for the hard disk of the virtual machine.

Start the installation process for a virtual machine on the compute node:

```
# vm-install -c 2 -m 2048 -n BaseImage -s nfs://<IP-installation-  
server>:/install/SLES11SP3/CD1 \  
-d /disk0.qcow2,0,disk,w,20000
```

Adjust the parameters passed to your specific needs. -c stands for the number of CPUs and -m for the minimum main memory of the virtual machine. -n indicates the name and -s indicates the source for the files required for the installation. The data you must specify is that indicated in the “Preparing an Installation Server” section. You use -d to specify the details for the hard disk or its image. 20 GB are specified for the hard disk in this example. The wizard prompts for further data, including the network card required. A window with the console for the virtual machine then opens. Perform an installation in this window in the usual way. Note the data provided in the upcoming sections and use the SAP documentation as a reference.

When the installation has completed, you copy the resulting image file, which contains the fully installed operating system, and upload the image to your SUSE Cloud.

As an alternative to the manual installation, Studio Online (<http://susestudio.com>) and the SUSE Studio™ Onsite version offer an easy method of creating images. The SUSE Image Builder KIWI, which is also the basis for SUSE Studio, can also be used. KIWI is a command line tool provided free of charge by SUSE.

Since SAP applications have high demands with respect to main memory and hard disk space, the server on which the image is prepared must have sufficient hard disk space. This also applies later to the compute nodes. Sufficient main memory is also required here to be able run the SAP applications.



A further recommendation is to use two network cards in the image to subsequently obtain better performance and allow network traffic to be better controlled.

Exactly how much main memory is required depends on the SAP application, its purpose and the number of instances being used. For information about the resource requirements (such as the hard disk space and main memory required), see the Installation Guide for the SAP application. Alternatively, you can also use the online sizing tool from SAP to calculate the resources required.

Experience has generally shown that over-provisioning virtual CPUs can have a negative impact on performance for virtual machines with SAP applications. It is therefore recommended that you refrain from over-provisioning.

To ensure the best possible performance, an installation on raw devices is recommended. Once installation to an image file has been completed successfully, the corresponding image is converted into a volume using the dashboard.

### ***Creating a Generic SAP Image***

To ensure maximum flexibility when it comes to the type of SAP application to be installed, choose the installation option “SLES for SAP Applications – Installation.” This installs a SUSE Linux Enterprise Server for SAP Applications and along with the installation wizard, provides an easy means of performing a subsequent NetWeaver installation.

### ***Partitioning***

- The size of the swap space depends on the available main memory. The following values have been taken from SAP Note 1597355 (<http://service.sap.com/sap/support/notes/1597355>). Compare these values with those for your installation to verify that they are up to date.

<b>Physical Memory (RAM)</b>	<b>Recommended Swap-Space</b>
< 32 GByte	2 x RAM
32 – 63 GByte	64 GByte
64 – 127 GByte	96 GByte
128 – 255 GByte	128 GByte
256 – 511 GByte	160 GByte
512 – 1023 GByte	192 GByte
1024 – 2047 GByte	224 GByte
2048 – 4095 GByte	256 GByte
4096 – 8191 GByte	288 GByte
> 8192 GByte	320 GByte

*Table 2: Size of the Swap Space Based on the Main Memory*

- /boot 512 MB
- The rest of the hard disk is assigned to the LVM:
  - The VG (volume group) should have the name “system.”
  - Create LV (logical volume): lv\_root no smaller than 5 GB

### **Software**

Packages to be installed in addition to the basic system:

- sap\_installation\_wizard (optional)
- sap-locale
- sapconf
- sap\_nw pattern (optional if SAP NetWeaver is to be installed)
- cloud-init (extended configurations of the virtual machine)

### **Network Configuration:**

- Hostname change using DHCP
- Deactivate firewall (with the standard setting, your virtual machines are visible only within your virtual network)
- Set network cards to DHCP

Once SUSE Linux Enterprise Server for SAP Applications has been installed on a virtual machine, another installation source must be added for the SAP software to be installed.

In this case, we recommend installing from an NFS server. To do this, the SAP software must first be copied to the shared directory of the NFS server.

Of course, the NFS server must later also be accessible from the IP range of SUSE Cloud assigned to the instances.

### **Setting up an NFS mount for SAP installation products**

Enter the NFS server to be mounted in the /etc/fstab file below the last line. Example:

```
# vi /etc/fstab
192.168.124.10:/SAP-install /mnt nfs defaults,wsiz=32768,rsiz=32768,soft,bg,intr 0 0
```

Replace the IP address “192.168.124.10” with the IP address or name of your NFS server and replace the folder shared using NFS (in this example, “/SAP-Install”) with the folder that contains your SAP installation media.

It is important that all of the systems involved have the same time. You can achieve this by using NTP (Network Time Protocol). Install and configure automatic time synchronization using

```
# yast2 timezone
```

In addition to NTP, DNS is another important aspect to be considered. It should be possible to resolve all systems using DNS, including the names of virtual machines for the SAP system.

If a DNS is not available in the landscape, the local file /etc/hosts can also be used.

Note: The DNS server of SUSE Cloud on the admin server manages only the nodes within the SUSE Cloud infrastructure.

You obtain the best possible block IO throughput when you explicitly deactivate guest IO scheduling ("elevator=noop"), thereby transferring the guest block IO directly to the hypervisor and allowing it to perform scheduling. You can define this kernel parameter using YaST in the boot manager settings, for example, or you can edit the GRUB configuration file directly and add the "elevator=noop" entry to the "kernel" line.

### **Completing Image Creation**

To ensure that the image to be started in the cloud also obtains the correct network card name each time, the `/etc/udev/rules.d/70-persistent-net.rules` file must be deleted.

Also, remove the key files of the ssh daemon under `/etc/ssh/` (\*key, \*key.pub).

Ensure that the hostname of the virtual machine can be accepted and defined using DHCP. To do this, select the "Change hostname via DHCP" option on the "Hostname" tab of the network configuration (accessible using "yast2 lan").

Finally, delete the user history using "history -c". Then shut down the VM.

Now copy the image twice — a total of three images is required:

```
# cp /disk0.qcow2 /disk_pas.qcow2
# cp /disk0.qcow2 /disk_ci.qcow2
# mv /disk0.qcow2 /disk_db.qcow2
```

The first copy is intended for subsequent installation of the SAP primary application server (SAP-PAS); the second is for the SAP central instance (SAP-CI); and the third is for the SAP database instance (SAP-DB).

### **Creating the SAP Database Image**

For the image for the SAP database (SAP-DB), we recommend that you set up another NFS export for the database on a high-performance NFS server. This is later installed not in the image but on the NFS export outside the cloud. This increases database performance and makes it easier to implement subsequent changes. The recommended size of the NFS export is at least 40 GB (depending on the database).

Boot the third image and create another entry in `/etc/fstab`:

```
192.168.124.11:/DB-Export /DB-Export nfs defaults,wsize=32768,rspace=32768,soft,bg,intr 0 0
```

Replace the IP address "192.168.124.11" with the IP address or name of your NFS server and replace the folder shared using NFS (in this example: "/DB-Export") with the folder in which the database is to be installed.

Create a directory "/DB-Export" and delete the `/etc/udev/rules.d/70-persistent-net.rules` file before you shut down the virtual machine.

Now copy the three images to the control node. In a later step, the image will be converted to ensure you obtain the best possible performance.

## Loading SUSE Linux Enterprise Server for SAP Images to SUSE Cloud

Once the images have been created, they need to be uploaded to the cloud. To do this, you require information about the current image format. You determine this using:

```
# qemu-img info /path/to/image/imagefile
```

The output (qcow2 or raw) is needed in the following steps.

You can load the images either by using the dashboard (Horizon) or by copying them to the control node by means of a CLI. The following section outlines how you do this using the dashboard.

Access the dashboard using the browser (HTTP), e.g., using <http://192.168.124.81>. The credentials used here correspond with those set up automatically during the installation:

**Login: admin/admin**




Figure 13: SUSE Cloud Dashboard

Choose the “Admin” project, then the “Images” item on the right side, and finally “Create Image” (see Figure 14).



Figure 14: Uploading Images to SUSE Cloud

Then enter the URL from which the image can be uploaded, or choose the “Browse” button to upload it from your local computer. Finally, you must define the “Image Name” and “Format” fields (e.g., container format: bare; disk format: in this case, qcow2) and select the “Public” checkbox.

## Converting the Images

You now use the dashboard to convert the available images into volumes. By using central storage, you improve the availability. If you want to improve performance instead, however, use the images previously available, and during the subsequent installation steps ensure you store the user data in the network, such as on an NFS server.

In the “Images & Snapshots” section of the dashboard, you will now see the images that are currently installed. By choosing the “More” button next to each image, you can convert the image to a volume by choosing “Create Volume.” This creates a volume on the configured block storage of Cinder. In the wizard that then appears, you can influence the size of the volume and provide a description. You then initiate the volume creation process by choosing “Create Volume” (see Figure 15).

The screenshot shows a 'Create Volume' wizard with the following fields and values:

- Volume Name \***: SAP-Application-Server
- Description**: Volume for SAP Application Server
- Type**: (empty dropdown)
- Size (GB) \***: 10
- Use image as a source**: SAP-Application-Server (10.0 GB)

**Description:** Volumes are block devices that can be attached to instances.

**Volume Limits**

- Total Gigabytes (0 GB)**: 1,000 GB Available
- Number of Volumes (0)**: 10 Available

Buttons: Cancel, Create Volume

Figure 15: Converting an Image to a Volume

Perform this operation for all three installed images. As a result, any virtual machines that start up access the corresponding volume instead of an image file.

## Creating Flavors

Flavors are predefined templates for virtual machines that specify which hardware resources (that is, CPUs, main memory, and hard disk memory) a system obtains. They can be freely defined and adapted to the environment.

Appropriate flavors for an SAP installation are set up based on SAP specifications.

The values indicated below are minimal examples. To determine your actual requirements, consult the SAP documentation and the online sizing tool from SAP. You will find the relevant information and links in your SAP documentation.

The minimum requirements for an SAP system are as follows:

- At least one dual-core or two single-core CPUs
- At least 7.5 GB RAM for DB, CI and AS

- At least 72 GB hard disk memory for DB, CI and AS
- SWAP space — at least twice the size of the main memory. When divided among the individual applications, the following values result:

**SAP Central Instance:**

- 2 vCPUs
- 2 GB RAM
- 2 GB HDD

**SAP Database:**

- 2 vCPUs
- 2 GB RAM
- 40 GB HDD

**SAP Primary Application Server:**

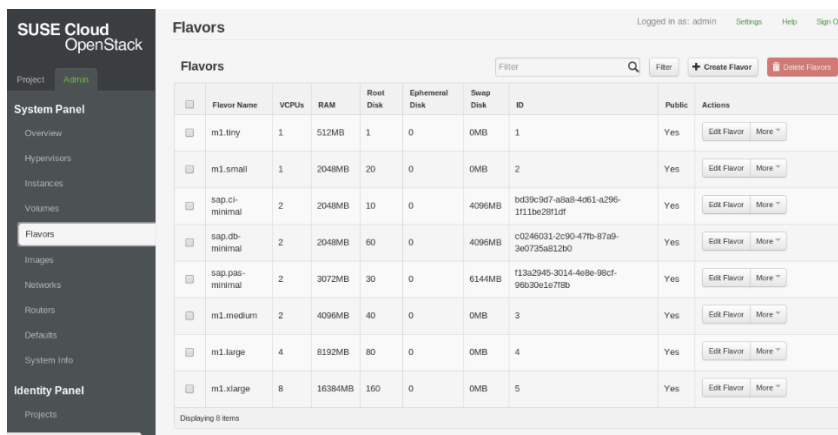
- 2 vCPUs
- 3 GB RAM
- 20 GB HDD

These values are used as the basis for a number of minimum flavors (see Table 3 “Minimum Flavors Created” below):

Name	Number of vCPUs	RAM	Root Disk	SWAP
sap.db-minimal	4	6 GB	100 GB	12 GB
sap.ci-minimal	2	4 GB	40 GB	8 GB
sap.pas-minimal	2	4 GB	50 GB	8 GB

*Table 3: Minimum Flavors Created*

The flavors are created using the SUSE Cloud OpenStack dashboard. “System Panel” contains the “Flavors” item, where the existing flavors can be edited or new flavors created (see Figure 16 on page 30).



*Figure 16: Managing Flavors in the Dashboard*

### ***Final Steps before Booting an Instance***

Within a project, the instances can communicate with one another with no restrictions. Network packages must not be able to obtain internal access from outside, however. The package filter is used for this purpose. You can access it by choosing “Access & Security” in the “Security Groups” area. This allows you to permit port 22 on the TCP protocol from outside, for example. For more detailed information, see the “SUSE Cloud User Guide” (“Configuring Access to the Instances” section).

SSH access to the instances is required for further access. You permit this by first opening the “Access & Security” area in the dashboard. Then select “Edit Rules” on the “Security Groups” tab. In the window that then opens, click “Add Rule” and select “ssh” in the “Rule” field. To permit SSH access from all addresses to the instances in this security group, leave the “Remote” and “CIDR” fields unchanged.

### **Starting Instances Using the Dashboard**

- Log on to the dashboard
- Select the project
- Instances → Launch Instance
- Select the image → “Boot from Volume”, select the volume
- Select a name for the instance
- Select the flavor
- Select the network
- Start the instance using “Launch”

## **Installing the SAP Software**

You can install all of the SAP components on one single VM if required. This is known as a “two-tier installation.” It is sufficient for small to medium-sized SAP installations, test environments, training, sandboxes and development purposes.

The three components can also be distributed across various VMs to obtain better utilization of the compute nodes in SUSE Cloud. This is known as a “three-tier installation” or a “distributed installation.” This is used to implement large SAP NetWeaver installations with a large number of users. Both variants are conceivable in SUSE Cloud. Distributed installation is described here.

The SAP host names sap-ci, sap-db and sap-pas used here are only examples. Names can be defined as required.

### ***Installing ASCS***

To install the ASCS (ABAP SAP Central Services) components, boot the SAP central instance volume with the sap.ci-minimal flavor using the dashboard and assign the virtual machine the name SAP ASCS Server (see Figure 17).

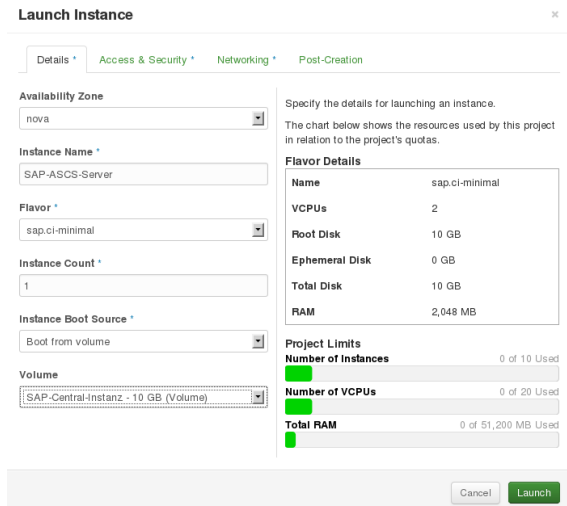


Figure 17: Deploying the SAP ASCS Server Virtual Machine

Depending on the database used (Oracle, MAXDB, DB2), you will find details about installing the ABAP stack in the corresponding SAP installation document (see <http://service.sap.com/instguidesNW73>).

Log on to the SAP ASCS server virtual machine using `ssh -X`.

You can determine the relevant IP address either in the dashboard or using the command:

```
# nova list
```

Go to the `/SAP-install` directory. Here you should now see the exported SAP installation archives. Create the following users and groups, depending on the SID used:

User	Primary Group	Secondary Groups	Description
<sapsid>adm	sapsys	oper, sapinst	SAP system administrator
<daside>adm	sapsys	sapinst	Diagnostics agent administrator
sapadm	sapsys	sapinst	Host agent administrator

Table 4: Users for the ASCS Installation

Create and change to the `/sapinst/ascs` directory:

```
# mkdir /sapinst/ascs
# cd /sapinst/ascs
```

Start the SAP installer:

```
# /SAP-install/51042312_MASTER/DATA_UNITS/NW730_IM_LINUX_X86_64_ORA/sapinst \
SAPINST_USE_HOSTNAME=sap-ascs
```

The SAP installer is now started (see Figure 18).



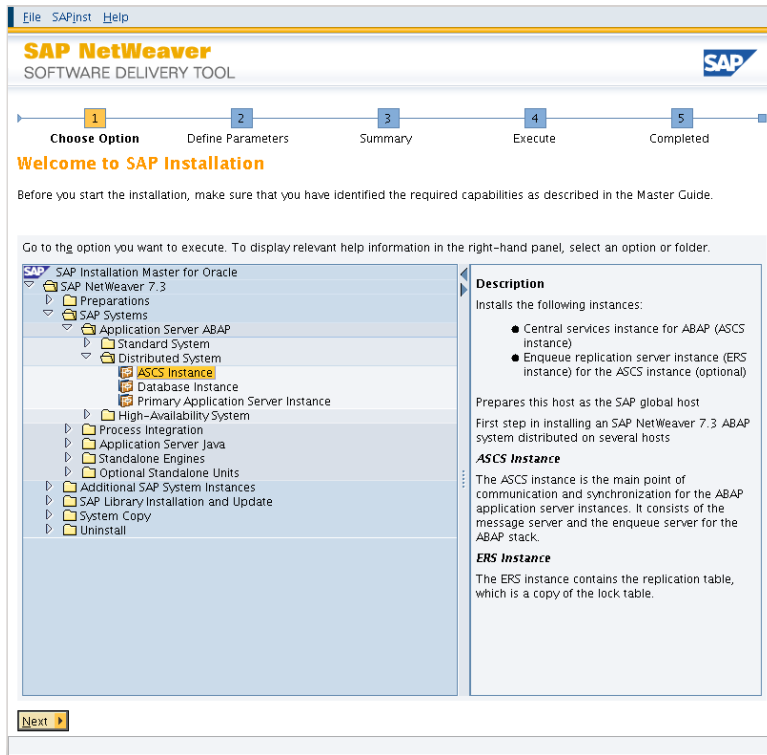


Figure 18: SAP Installer – ASCS Instance

You can find the subsequent installation steps in your corresponding SAP installation document.

## Installing the DB

Boot the SAP database volume at minimum with the sap.db-minimal flavor and assign the virtual machine the name SAP-Database-Server (see Figure 19 on page 39).

Log on to the “SAP Database Server” virtual machine using ssh -X.

You can determine the relevant IP address either in the dashboard or using the command:

```
# nova list
```

Run the mount command. Two NFS mounts should be shown: /SAP-install, which includes the installation media for database installation, and /DB-Export, where the database is installed.

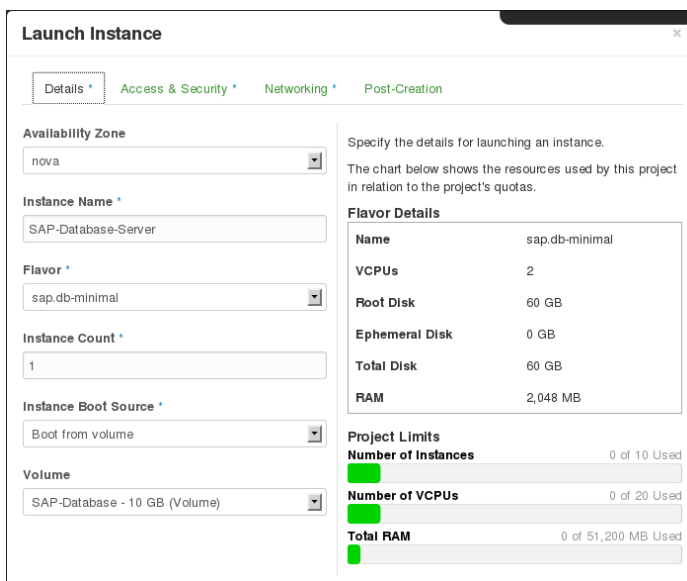


Figure 19: Deploying the SAP Database Server Virtual Machine

Create the following users, depending on the SID used:

User	Primary Group	Secondary Groups	Description
<sapsid>adm	Sapsys	oper, dba, sapinst	SAP system administrator
<dasid>adm	sapsys	sapinst	Diagnostics agent administrator
sapadm	sapsys	sapinst	Host agent administrator
ora<dbsid>	dba	oper, sapinst	

Table 5: Users for the Database Installation

Create and change to the /sapinst/db directory:

```
# mkdir /sapinst/db
# cd /sapinst/db
```

Now start the SAP installer for the database instance:

```
# /sap-medien/51042312_MASTER/DATA_UNITS/NW730_IM_LINUX_X86_64_ORA/sapinst \
SAPINST_USE_HOSTNAME=sap-db
```

The SAP installer is now started (see Figure 20).



Figure 20: SAP Installer – Database Instance

You can find the subsequent installation steps in your corresponding SAP installation document.

## Installing PAS

Boot the SAP Application Server image with the sap.pas-minimal flavor and assign the virtual machine the name SAP-Application-Server (see Figure 21 on page 42).

Log on to the SAP Application Server virtual machine using ssh -X.

You can determine the relevant IP address either in the dashboard or using the command:

```
# nova list
```

In the /SAP-install directory, you should then find the corresponding SAP archives for installing your primary application server.

**Launch Instance**

Details \* Access & Security \* Networking \* Post-Creation

Availability Zone: nova

Instance Name: SAP-Application-Server

Flavor: sap.pas-minimal

Instance Count: 1

Instance Boot Source: Boot from volume

Volume: SAP-Application-Server - 10 GB (Volume)

Specify the details for launching an instance. The chart below shows the resources used by this project in relation to the project's quotas.

**Flavor Details**

Name	sap.pas-min...
VCPUs	2
Root Disk	30 GB
Ephemeral Disk	0 GB
Total Disk	30 GB
RAM	3,072 MB

**Project Limits**

Number of Instances	1 of 10 Used
Number of VCPUs	2 of 20 Used
Total RAM	2,048 of 51,200 MB Used

Cancel Launch

Figure 21: Deploying the SAP Application Server Virtual Machine

Create the following users and groups with the relevant SID:

User	Primary Group	Secondary Groups	Description
<sapsid>adm	sapsys	oper, sapinst	SAP system administrator
<dasid>adm	sapsys	sapinst	Diagnostics agent administrator
sapadm	sapsys	sapinst	Host agent administrator

Table 6: Users for the Application Server Installation

Create and change to the /sapinst/pas directory:

```
# mkdir /sapinst/pas
```

```
# cd /sapinst/pas
```

Start the SAP installer for the PAS component:

```
# /sap-medien/51042312_MASTER/DATA_UNITS/NW730_IM_LINUX_X86_64_ORA/sapinst \
SAPINST_USE_HOSTNAME=sap-pas
```

The SAP installer is now started (see Figure 22).

You can find the subsequent installation steps in your corresponding SAP installation document.



Figure 22: SAP Installer – Application Server

Now that all relevant SAP components have been installed, the SAP client can be used to access the applications in question using the SAPNet set up.

### Automating the Installation Process

To further automate the process of installing the SAP components, a separate image can be provided for each component.

The users, groups and local directories must already have been created in this image, and the necessary NFS mounts must also already exist after the boot process.

The process of starting the SAP installer can be initiated automatically using cloud-init after the image has been booted. The information for the SAP installer is then transferred directly to the SAP installer using an inifile.xml file without any user interaction. For further information about this procedure, see SAP Note 950619:

<http://service.sap.com/sap/support/notes/950619>

The SUSE SAP installation wizard for SUSE Linux Enterprise Server for SAP Applications offers another option. The SAP software is also installed automatically in this case, but a distribution SAP installation is not possible.

The benefit of this approach lies in the interactive, wizard-based collection of the data and SAP media that are needed for subsequent automatic installation.

To use the SUSE SAP installation wizard for subsequent automatic installation, you must change the "RUN\_SAP\_INSTALL" variable in the /etc/sysconfig/sap-installation-wizard file from "yes" to "no." This prevents direct automatic installation and deletion of the collected data and SAP media.

The installation wizard writes all data and media to the /data/SAP\_CDs directory by default. This path can also be changed as required in /etc/sysconfig/sap-installation-wizard.

Once the "RUN\_SAP\_INSTALL" variables have been changed, the wizard can be run as described in the product documentation (see [www.suse.com/documentation/sles\\_for\\_sap](http://www.suse.com/documentation/sles_for_sap)) However, an SAP installation is not initiated, but the collected data is stored in the /data/SAP\_CDs/x/ directory or the directory you specified.

0 is the first installation, 1 is the second, 2 is the third, and so on. Example:

```
/data/SAP_CDs/1
├─ ay_q_db_logmode
├─ ay_q_masterpass
├─ ay_q_sapcrypto
├─ ay_q_sapinstnr
├─ ay_q_sid
├─ ay_q_tshirt
├─ ay_q_virt_hostname
├─ ay_q_virt_ip_addr
├─ ay_q_virt_ip_netmask
├─ ay_q_virt_network
├─ EXPORT1
├─ EXPORT2
├─ EXPORT3
├─ Instmaster
├─ product.data
├─ RDBMS-ADA
├─ SAPLUP
```

```
└─ scripts.data
└─ UKERNEL
```

To apply adjustments, such as a different SAP SID or IP address, the corresponding “ay\_q\_XXXXX” files must be changed.

The necessary SAP media are stored in the corresponding subdirectories, such as “UKERNEL.” To be able to create an installation request, further information is required from the /data/SAP\_Cds/1/product.data file.

Example of a product.data file:

```
"DATABASE"      : "ADA",
"INSTALL"       : "PD",
"PARTITIONING"  : "/usr/share/YaST2/include/sap-installation-wizard/base_partitioning",
"PRODUCT_ID"   : "NW_ABAP_OneHost:NW740.ADA.PI",
"SAPLUP"        : true,
"SCRIPT_NAME"   : "/usr/share/YaST2/include/sap-installation-wizard/sap_inst.sh",
"SEARCH"        : "NW_ABAP_OneHost",
"STACK"         : "AS-ABAP",
"TYPE"          : "CENTRAL",
"chosenProduct" : "NW740",
"instMaster"    : "/data/SAP_Cds/1/Instmaster"
```

The PRODUCT\_ID, the internal SAP ID for the SAP product to be installed, the SCRIPT\_NAME (called the "instMaster" data path) and, if partitioning is required, the path of the partitioning template are needed.

To be able to use the data and media obtained for an installation, the following must be entered in the command line:

If the SAP data partition (/sapdata) is to be partitioned:

```
# yast sap_create_storage <Partitioning Template>
```

To install the SAP product:

```
# <script-name> <Parameter: -i <Product_id> -m <Path to Media> -t <database type> -y <product type>
```

Example:

```
# /usr/share/YaST2/include/sap-installation-wizard/sap_inst.sh \
-m /data/SAP_Cds/1/Instmaster -i NW_ABAP_OneHost:NW740.ADA.PI \
-t ADA -y CENTRAL
```

It is recommended that you also redirect STDOUT and STDERR to log files.

The SAP installation uses the standard /usr/sap and /sapdata SAP directories.

## Outlook

You have now fully completed the installation of your SUSE Cloud. This document serves as a guide for a basic SUSE Cloud installation for providing SAP applications.

Please note: There are a number of settings that you need to configure based on your requirements — always in conjunction with the prerequisites specified by SAP. The ability of SUSE Cloud to deploy your Cloud infrastructure allows you to enlarge the structure represented easily and at any time.