

ProLiant for Private Cloud

Optimized for SUSE OpenStack Cloud



Table of contents

Introduction	2
Private cloud computing	2
OpenStack software	3
SUSE OpenStack Cloud	3
Solution architecture	3
Intended audience	3
Additional information	3
Considerations for using SUSE OpenStack Cloud and HP ProLiant hardware	4
Facility considerations.....	4
Physical networking	4
Computing platform.....	4
Software components	4
Hardware selection and server roles.....	6
Network topology and hardware selection.....	8
HP SUSE OpenStack Cloud pre-installation configuration	9
Pre-installation requirements.....	10
VLAN configuration.....	11
SUSE OpenStack Cloud installation.....	13
Deploying SUSE OpenStack Cloud.....	13
Conclusion.....	16
Appendix A: SUSE OpenStack Cloud rack layout	17
Appendix B: Bills of material	18
Resources	24

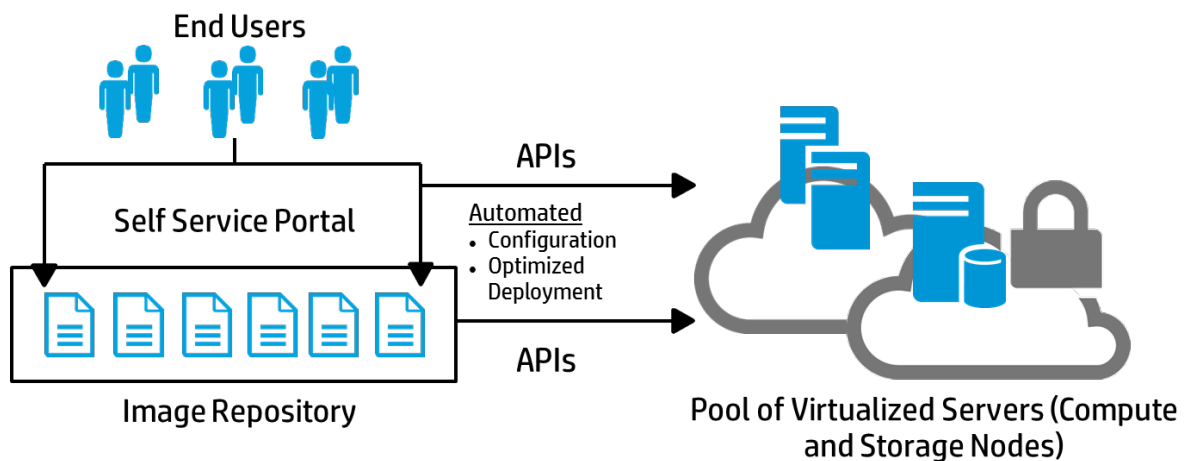
Executive summary

This document provides information about building a private cloud using specific HP ProLiant rack servers for both compute and storage nodes utilizing SUSE® OpenStack Cloud, with detailed hardware information and setup recommendations for implementing a functional OpenStack cloud. The document includes a solution architecture and discussion of design decisions that were made in arriving at this implementation.

Introduction

Cloud computing delivers access to a pool of processing, storage, networking and other virtualized resources over a network. Cloud deployment abstracts the complex, internal infrastructures from consumers. It allows individual virtual machines, or workloads, to be provisioned on demand from predefined templates with little concern about the underlying infrastructure or resources. Services, consisting of many workloads can be provisioned for use in minutes and then scaled appropriately to meet service demands for use by multiple tenants. Figure 1 illustrates all of these components and interactions to deliver an Infrastructure-as-a-Service (IaaS) private cloud deployment:

Figure 1. Infrastructure-as-a-Service (IaaS) private cloud deployment



For further reference, the [National Institute of Standards and Technology's \(NIST\) Definition of Cloud Computing](#) describes the important characteristics of cloud computing.

Private cloud computing

A private cloud provides Infrastructure as a Service (IaaS) by delivering network, compute, and storage resources. These resources can be integrated through private cloud software to provide the robust foundation for applications and services before adding deployed workloads. Networking is the unifying component tying the computing and storage resources together in this solution architecture. Given the complexity of such an infrastructure, the network must be fully configurable and adapt to change requests upon demand. High-speed, scalable devices like the HP FlexFabric series of network switches provide a perfect foundation. Compute resources are provided with physical machines or virtual machines running atop a hypervisor. Storage resources can be provided through either dedicated storage elements or distributed resilient file systems across commonly available computing systems, each with local storage. Modern, industry-leading x86-based servers like the HP ProLiant Gen 9 servers provide an ideal platform for private clouds because they balance performance and efficiency, and various models can fulfill both the compute and storage node roles. Finally, it is the private cloud software that provides the user interface for setup, configuration, maintenance, and long term operation of these three core components bonding them into a cohesive service offering.

OpenStack software

OpenStack software is an open source cloud operating system that controls these large pools of compute, storage, and networking resources—all managed through a dashboard that gives administrators control while empowering their users to provision resources through a Web interface. OpenStack software is designed for scalability, so you can easily add new compute and storage capacity to grow your cloud over time. Companies such as HP have built massive public clouds based on OpenStack technology.

SUSE OpenStack Cloud

SUSE OpenStack Cloud, based on SUSE Linux Enterprise Server, is a distribution of upstream OpenStack project deliverables and other value-add open-source components. SUSE OpenStack Cloud eases the management of your private cloud infrastructure by delivering enhanced networking flexibility and improved operational efficiency. It also provides “as-a-service” capabilities, so your application teams (this includes development, big data analytic, as well as other teams) can deliver business solutions rapidly. Integrated with SUSE High Availability Extensions and SUSE Enterprise Storage, SUSE OpenStack Cloud provides a scale-out, hardened, and enterprise-ready product. It’s designed to deliver the most innovative open source cloud computing technology in a resilient, maintainable platform. For more information about SUSE OpenStack Cloud, go to suse.com/cloud.

Solution architecture

This solution architecture provides guidance in the deployment of a scalable OpenStack-based private cloud using SUSE OpenStack Cloud software. The OpenStack control-plane services are provided in a highly available fashion to prevent outages of these core infrastructure components. Compute services are delivered via the KVM hypervisor technology with object and block storage services furnished through Ceph¹ technology. The entire solution relies on specific HP ProLiant DL Gen9 servers.

This document describes the rationale for choosing each of these components and the steps necessary to successfully install the complete SUSE OpenStack Cloud solution. When other technology choices are deployable as part of the overall framework, these will be noted but not described in detail.

Intended audience

To be successful with this guide, you should be:

- Familiar with OpenStack cloud technology.
- Familiar with the SUSE Linux Enterprise distribution of Linux, SQL databases, and virtualization.
- Comfortable administering and configuring multiple Linux machines for networking and perhaps other operating systems should you choose to include those hypervisor alternatives.
- Familiar with networking concepts such as DHCP, Linux bridges, VLANs, and IP tables.

You should also have access to a network hardware expert who can configure the switches and routers required in your OpenStack cloud.

Additional information

The primary source of documentation for installation of SUSE OpenStack Cloud is available in the “SUSE Cloud Deployment Guide” at suse.com/documentation/suse-cloud-5/.

For additional information on how to operate your cloud, you can read the “SUSE Cloud OpenStack Admin User Guide”, “SUSE Cloud OpenStack End User Guide”, or “SUSE Cloud Supplement to Admin User Guide and End User Guide” at suse.com/documentation/suse-cloud-5/. You can find more general information in the “OpenStack Operations Guide” available at docs.openstack.org/ops. The “Architecture” section contains useful information about how an OpenStack Cloud is constructed. However, SUSE OpenStack Cloud takes care of many of these details for you.

¹ Ceph is a massively scalable, open source, distributed storage system. It is comprised of an object store, block store, and a POSIX-compliant distributed file. Read more at <https://www.openstack.org/assets/presentation-media/cephopenstackthewinningpair.pdf>

Considerations for using SUSE OpenStack Cloud and HP ProLiant hardware

This section explores the necessary preparations for core components in a private cloud instance.

Facility considerations

While beyond the scope of this document, the heating, ventilation, air conditioning (HVAC) requirements of hosting such a private cloud solution should be carefully considered and planned. To aid in determining the power requirements for system deployment, use the HP Power Advisor as either an on-line version or with a downloadable application (both available in the [Resources](#) section of this document). Using this tool, you can plan the needs for your solution and order the correct Power Distribution Unit (PDU). An example PDU is cited in this document, however please review that configuration and adjust as needed for the local power conditions and connections in the final installation location.

Physical networking

Networking is the technology component likely to require the most advance planning. Requirements for a private cloud instance have a fair amount of network complexity, especially when integrated with an existing IT infrastructure. For the physical level, using pairs of HP FlexFabric Switch Series devices as top-of-rack (ToR) switches allows all servers, each having multiple 10Gb NIC ports, to form a link aggregation group (LAG) across the ports with at least one port on each switch in the stack. When used in conjunction with network interface bonding on the resource nodes, the LAG ideally offers switch redundancy within the rack and enables high availability throughout the infrastructure. However, for this basic setup, a single 10GB switch is used, having all resource nodes' 10GB NIC ports connected here, but could be expanded to the dual ToR configuration at a later date.

Computing platform

One of the benefits of a private cloud implementation is that HP industry-standard servers can fulfill all the hosting and computational needs. To reduce the time spent on hardware specification for an initial private cloud implementation, the hardware should be general purpose and allow for a wide range of configuration options. The industry leading HP ProLiant Gen9 DL360 (1U) and DL380 (2U) rack mount servers with their high performance, multitude of network options and low energy consumption characteristics are ideal virtualization and storage nodes for a private cloud deployment.

Software components

[SUSE® OpenStack Cloud](#) provides the enterprise-grade implementation of OpenStack components and includes other tools for a complete private cloud implementation. SUSE OpenStack Cloud forms the basis of the private cloud operating solution and overall framework to abstract the network, computing, and storage resources. By including an installation framework and some implemented choices for the software components, the setup and maintenance of a private cloud instance is less complex, repeatable, and scalable. SUSE OpenStack provides the following services:

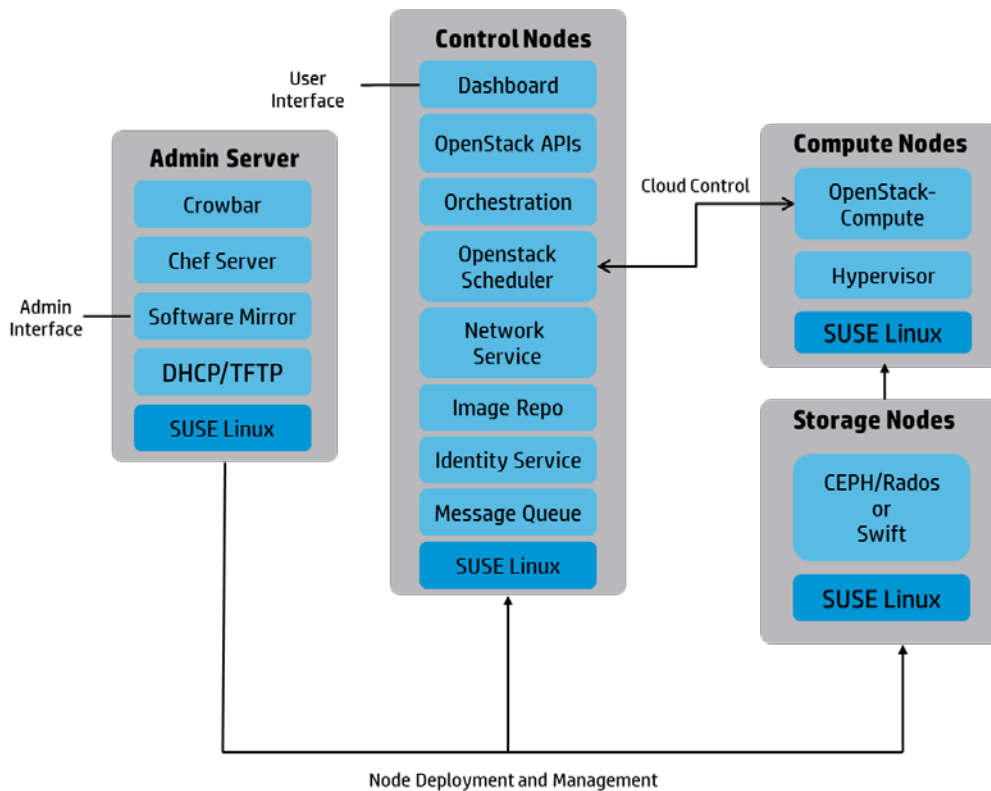
- **Installation Framework** - Provides the services and tools needed to quickly and easily deploy all the remaining nodes in your private cloud instance. All of these are implemented on an Administration Server and are setup, configured and accessed from a web interface that guides the installation process. A cloud administrator would utilize this interface to manage the resource nodes over their entire lifecycle in the private cloud instance. The services provided are
 - Dynamic Host Control Protocol (DHCP) to allow resource nodes to easily become part of the private cloud network
 - Trivial File Transfer Protocol (TFTP) to enable delivery of a simple boot kernel and initial ramdisk to new resource nodes
 - Pre-boot Execution Environment (PXE) to help provision the operating system on the various resource nodes
 - Domain Name Service(DNS) to provide mapping of hostname to IP for resource nodes within the private cloud
 - Provisioning tools, based upon Chef and Crowbar, for repeatable provisioning of the remaining software components to the resource nodes
 - Software repositories of all of the needed components to initially install the resource nodes, plus allow updates to be applied to the nodes over time
 - Pacemaker and associated tools, provide the process and ability to create clusters for the various OpenStack services, making the overall control plane highly available and resilient to outages
- **OpenStack Shared Services** – Provides those components utilized across the three pillars of compute, storage and networking, including:
 - Database is an underlying relational data store. SUSE OpenStack Cloud implements this with PostgreSQL²

² Open source database – more at www.postgresql.org/

- Message queue for peer-to-peer interprocess communication. SUSE OpenStack Cloud implements this with RabbitMQ³
- Ceph for distributed object and block store. SUSE OpenStack Cloud provides this functionality with SUSE Enterprise Storage components
- Identity (Keystone) is a central directory for authentication and authorization, mapping users to services and resources they can access
- Image (Glance) for discovery, registration and delivery of disk and server images
- Telemetry (Ceilometer) aggregates usage and performance data across services
- Orchestration (Heat) is a template-driven engine to describe and automate deployment of infrastructure
- **OpenStack Core Capabilities** – Provides the core services to specifically manage compute, network and storage resources, including:
 - Block Storage (Cinder) is pluggable management of persistent, block-level storage for running workload instances
 - Networking (Neutron) is pluggable, scalable, API-driven network and IP management
 - Compute (Nova) manages the lifecycle of workload instances
 - Dashboard (Horizon) is a web-based, self-service portal to access and manage the underlying private cloud resources
 - Database (Trove) provides an optional, scalable Database-as-a-Service functionality

All of these software components and services in the private cloud instance are associated with various resource nodes (See Figure 2), which are described in more detail later in this document.

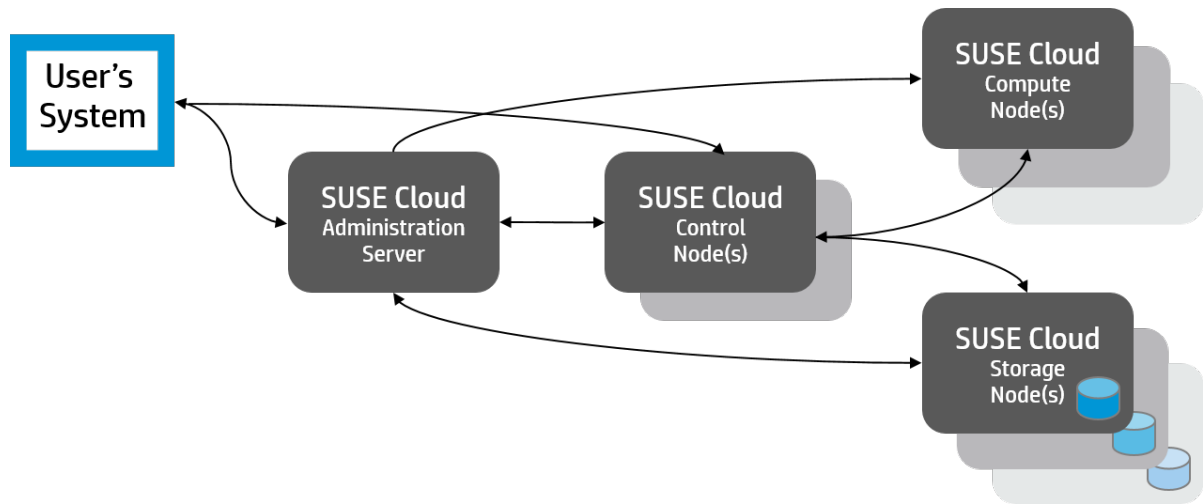
Figure 2. SUSE OpenStack Cloud service assignment to resource nodes



³ Open source message broker software – more at www.rabbitmq.com/

Figure 3 represents a logical diagram of the interaction of these various software components.

Figure 3. SUSE Open Cloud software components



Target environments for private clouds are varied, ranging from small developments and test instances to full production implementations that address both scale and high availability. Some private cloud instances target only workload management while others are set up to provide storage clouds, yet mixed environments are quite common. HP and SUSE provide professional services and technical pre-sales teams that can size a private cloud instance according to a customer's needs. The following sizing considerations can be used as a starting point:

- **Minimal Development Instance** - This minimal instance provides only the necessary components to cover private cloud usage modes in a functional sense so that developers of next-generation services have access to the basic features. In this scenario, workload life cycle management is handled in a true cloud sense, meaning that updated images are deployed.
- **Small Starter Instance** - This small instance provides all of the features noted in this reference configuration, but at the lowest levels of node count possible. As such, this starter configuration can be scaled over time to meet increasing usage and availability concerns.
- **Scalable and Highly Available Instance** - This instance expands upon the small starter node and is set to yield more scale across hypervisor and storage types. Utilizing these considerations, an instance of this scope also provides additional node counts for each role to address high availability concerns about key control plane services of the private cloud instance.

This solution architecture documents the core components of the “Scalable and Highly Available” instance which can grow as needed to address the required scale and performance.

Hardware selection and server roles

The HP ProLiant DL360 and DL380 Gen9 Servers are excellent platforms to implement a SUSE OpenStack Cloud configuration. The DL360 Gen9 Server is compact in size, taking just 1U in rack space while still allowing up to 12 TB of internal SAS storage to be configured. The DL380 Gen9 Server provides greater storage capacity, allowing up to 90 TB with LFF drives, or up to 31.2 TB with SFF drives. Both server models are highly configurable and can provide optimized performance for various OpenStack roles. You can find more information on these servers at hp.com/go/proliant.

Administration Server

The administration server is the first system installed and is used for provisioning the remaining resource nodes. It also provides the administration web interface, allowing you to assign service roles to each node used in the cloud. The Administration Server runs SUSE Linux Enterprise Server plus the SUSE OpenStack Cloud install framework. This is a single dedicated system, but does not have a high workload and only requires about 2 TB of storage for its OS and repositories used during the deployment of the other resource nodes. A DL360 Gen9 Server with one 6-core Intel® Xeon® CPU, 32 GB memory, and 4 LFF drive bays provides a good match of performance at reasonable cost for this role. RAID 10 is used for protection and performance.

Control nodes

A control node runs many of the core services for the OpenStack cloud environment, including Nova, Keystone, Glance, Cinder, Heat, Neutron, and Horizon. Like the SUSE OpenStack Cloud Administration Server, the control node runs on the SUSE Linux Enterprise Server operating system. Management and deployment of any remaining software is performed through the SUSE OpenStack Cloud Administration Server web interface. High availability setups require the use of SUSE High Availability Extensions, allowing the creation of multiple dedicated control nodes in a clustered configuration. For this solution architecture, a three-node cluster of control nodes are deployed serving all core OpenStack services.

Once the cluster is established through the Pacemaker⁴ barclamp⁵, it can be assigned as a unit to any of the core OpenStack services. These control nodes must provide high performance because all key OpenStack services run on these nodes, and they service the entire production cloud. Each of the control nodes is a DL360 Gen9 server with 2 12-core 2.6 GHz Intel Xeon CPUs, 64 GB memory, and 4 LFF drive bays with fast 600 GB 15k drives setup as RAID 10 for protection and performance.

Storage nodes

A SUSE OpenStack Cloud storage node acts as a controller for persistent cloud-based storage elements. All of these nodes run SUSE Linux Enterprise Server and are managed with the necessary services from the SUSE OpenStack Cloud Administration Server. The storage volumes for the data stores are managed from the OpenStack Dashboard. A SUSE OpenStack Cloud implementation can contain multiple dedicated storage nodes and types of storage, but for this solution architecture the focus is on Ceph-based options through the included SUSE Enterprise Storage product.

A Ceph-based cluster is implemented on at least three servers for availability, where the number of data drives should come in multiples of four, or instead, on a four-node cluster. Using HP ProLiant hardware, Ceph enables provisioning of persistent block storage at the virtual machine level for quick retrieval and fast processing along with the construction of a resilient image and object storage option. Each of these nodes runs the full complement of Ceph roles and by default Ceph makes three copies of the data spread over the available nodes. This storage model allows for redundancy so that the storage can survive failure of one of these nodes. Additional Ceph nodes can be added, one or more at a time. The Ceph replica count is configurable and a value of four copies of the data is suggested for maximum resiliency to failures if you have that multiple of drive volumes or nodes.

Capacity is the primary objective for storage nodes, but drives don't need to be particularly fast for most object storage use cases. To achieve the highest capacity, use LFF drives that currently offer the largest capacity per drive of 6 TB. For this scalable solution architecture, three DL380 Gen9 Servers, offering up to 15 LFF drive bays are recommended. Two of the drive bays have smaller but faster 15k 600 GB drives installed. These two drives are configured as a RAID 1 volume for the operating system. The remaining drive bays are installed with 7.2K rpm 6TB drives. These data drives are each configured as a RAID 0 JBOD.

Compute nodes

A SUSE OpenStack Cloud compute node is a physical, dedicated server running a hypervisor and serves as the host for the eventual deployment of workloads. While multiple hypervisor technology alternatives are available, this solution architecture focuses on the use of KVM. Compute nodes also provide ephemeral storage for the VMs. The number of VMs that can be run on a compute node is directly proportional to the number of CPU cores, memory, and drive space available. Use the maximum number of CPU cores for density, which is 18 cores per processor for the currently available DL360 Gen9 Servers. Configured with 256 GB memory, this provides an average of 7 GB memory per physical core, matching all but the largest typical flavors of cloud workloads. In this solution architecture, four KVM Compute nodes are implemented to provide compute host targets for VM instances. Users can easily scale up by adding 50 to 100 additional compute nodes.

Drive speed is very important for compute nodes, but you also want the best capacity you can get with fast drives. To achieve this, a 4 SFF drive configuration is used. Each compute node has 4x1.2 TB 10k SFF drives, configured as a RAID 10 volume for protection and performance, which provides 2.4 TB of usable storage.

The ratio of compute node memory to disk storage capacity is approximately 10 to 1, which matches very closely with the default VM flavors' memory to disk ratios of 10 to 1. To best utilize the compute nodes' available resources and capabilities while providing for the maximum number of VMs, you should use SUSE OpenStack Cloud-provided default VM flavors. If creating new VM flavors, utilize ratios similar to this same Compute node memory/configured disk ratio.

VM flavors and configuration

To further illustrate the scaling capability of this cloud instance, in terms of the end goal of running virtual machine workloads, consider that each compute node is configured with 256 GB of physical memory, 2.4 TB of formatted disk storage, and a total of 36 processor cores with two threads/core yielding 72 total vCPUs. These compute node resources are directly available for allocation by specifying a specific VM flavor when creating it. With the default VM flavors, shown in

⁴ Pacemaker is a scalable high availability cluster resource manager. You can read more about Pacemaker at this [Clusterlabs](#) link.

⁵ OpenStack services are automatically installed on the nodes by using so-called barclamps—a set of recipes, templates, and installation instructions.

the Horizon **Admin->Flavors** tab, all have an approximate 10 to 1 memory to disk allocation ratio as seen in the following table:

Table 1. VM flavor definitions

Flavor name	vCPUs	RAM	Root disk
m1.tiny	1	512 MB	1 GB
m1.small	1	2048 MB	20 GB
m1.medium	2	4096 MB	40 GB
m1.large	4	8192 MB	80 GB
m1.xlarge	8	16385 MB	160 GB

Any arbitrary combination of the above VM sizes may be assigned to compute nodes (subject to a user's quota), but each VM created will consume vCPU, memory, and formatted disk storage on its assigned compute node. You can view the currently available hypervisor resources, which can be assigned to VMs in the Horizon **Admin->Hypervisors** tab. When the available RAM or available root disk space is insufficient to instantiate a VM of a given flavor, the VM creation will fail. On the other hand, vCPUs can often be oversubscribed. For example, 26 of the m1.large flavor VMs will fit on a single compute node. The resource constraint which limits the instance count to 26 is the lack of available root disk space, not physical memory. Note that these 26 VMs will consume 104 vCPUs, so available hardware CPU threads will be oversubscribed by a factor 1.44.

Table 2 lists hardware used in this solution configuration.

Table 2. Solution configuration hardware

Server role and quantity	Server model	Processor	Memory	Storage	Network
Administration Server (1)	DL360 Gen9 4 LFF configuration	1P 6-core, 2.4 GHz Intel Xeon	32 GB	2x2 TB, 7.2k SAS LFF	10 Gb dual port FlexibleLOM
Control Node (3)	DL360 Gen9 4 LFF configuration	2P 6-core, 2.4 GHz Intel Xeon	64 GB	4x600 GB 15k SAS LFF	10 Gb dual port FlexibleLOM
Storage Node (3)	DL380 Gen9 15 LFF configuration	2P 12-core, 2.6 GHz Intel Xeon	64 GB	2x600 GB 15k SAS LFF + up to 13x6 TB 7.2k SAS LFF	10 Gb dual port FlexibleLOM
Compute Node (4)	DL360 Gen9 8 SFF configuration	2P 18-core, 2.3 GHz Intel Xeon	256 GB	4x1.2 TB 10k SAS SFF	10 Gb dual port FlexibleLOM

Network topology and hardware selection

To accommodate high availability, the network fabric must have at least two physical links which can be trunked together for all of the networking and VLAN functionality. For this solution architecture, two of the 10 Gb NIC interfaces are used in this fashion. On HP ProLiant servers, the iLO-designated network connects to the HP iLO management port on each server. This network is used by the Administration Server to control the server state, performing activities such as remotely powering servers on and off during bare-metal deployment.

In order to facilitate setup up and administration of the affected servers, the embedded iLO network interface of the SUSE OpenStack Cloud Administration Server should be connected to a traditional management subnet which typically resides outside the private cloud network infrastructure to allow easy administration access. For simplicity, the remaining resource nodes should have their iLO connected to the Admin network.

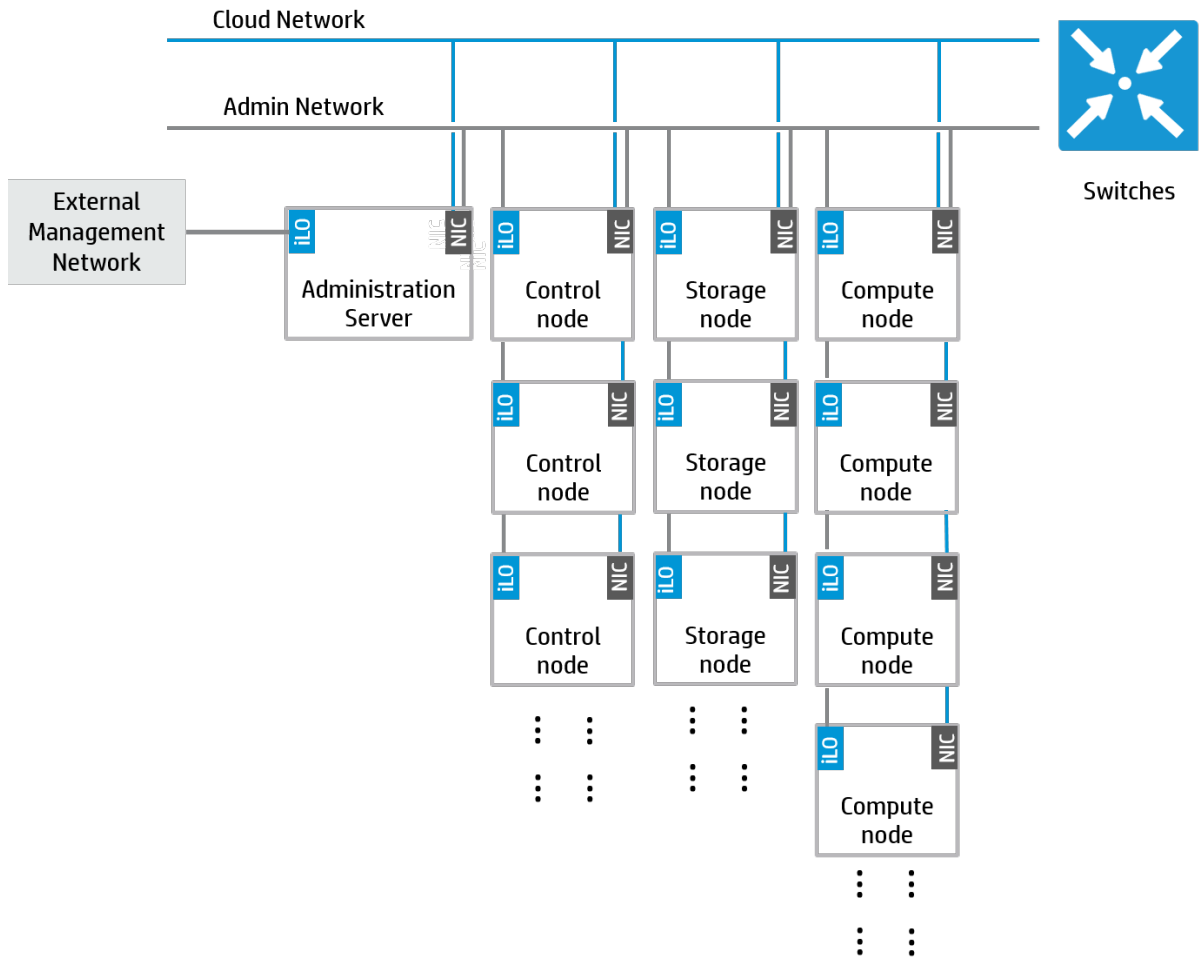
This subnet typically resides outside the private cloud network infrastructure to allow easy administration access. All other systems should have their iLO network interfaces connected to the internal private cloud admin iLO network.

Note

The SUSE OpenStack Cloud Administration Server will configure the iLO for the Control, Compute, and Storage nodes during set up. The operating system install process will bond the network interfaces presented into a teamed configuration.

In this solution configuration, use one HP FF 5700-48G-4XG-2QSFP+ Switch (JG894A) to implement a 1 GB Admin (IPMI) network on all cloud resource node servers with the iLO ports connected to this switch. For the 10 GB NIC on each node, connect both ports to the HP FlexFabric 5700-40XG-2QSFP+ Switch (JG896A) with each switch port configured as “untagged” for the Admin Network VLAN ID.

Figure 4. SUSE solution architecture server and network configuration



HP SUSE OpenStack Cloud pre-installation configuration

To ensure a successful HP SUSE Cloud installation, your HP Network and ProLiant server platforms should meet these pre-installation and VLAN requirements.

Pre-installation requirements

To ensure a successful installation of SUSE OpenStack Cloud, reboot each system and press **F9** to run **System Utilities** to make sure each system has the following configuration options set up correctly in BIOS:

1. All systems must be set to the correct date and time.
 - A. In **System Configuration**, select **BIOS/Platform Configuration (RBSU)→Date and Time**. Set the correct Date, Time, and Time Zone.
2. All systems servers should have Boot Option set to “UEFI Mode” in BIOS.
 - A. In **System Configuration**, select **BIOS/Platform Configuration→Boot Option→Boot Mode**, and make sure Boot Mode is set to **UEFI Mode**.
 UEFI mode is the default boot option for Gen9 servers. If you need to change from another boot option to UEFI, the server needs to reboot once for the new UEFI boot mode to take effect.
3. Make sure that each system is configured in the BIOS to stay powered off in the event of being shutdown rather than automatically restarting.
 - A. In **System Configuration**, select **BIOS/Platform Configuration→Server Availability→Automatic Power-On**, and set it to Restore Last Power State.
4. Disable network cards not being used.
 The DL360 and DL380 Servers used in this solution configuration are equipped with two network cards. In this solution configuration, SUSE OpenStack Cloud does not use the embedded HP Ethernet 1Gb 4-port 331i Adapter. By disabling this unused network device, the first port of the 10 Gb dual port FlexibleLOM that connects to the hypervisor or OS network will show up as “eth0” on each cloud node after SUSE OpenStack Cloud deployment. This consistent Ethernet interface name across all cloud nodes greatly simplifies overall configuration. The second 10 Gb port is used in the bonding pair automatically setup on each node.
 To disable the unused network card:
 - A. In **System Configuration**, select **BIOS/Platform Configuration (RBSU)→PCI Device Enable/Disable**, and set “Embedded LOM 1: HP Ethernet 1Gb 4-port 331i Adapter—NIC” to **Disabled**.
5. All cloud servers should be configured in BIOS such that only the desired network interfaces are enabled for PXE and Network Boot. In this solution configuration, the ports of the 10 Gb dual port FlexibleLOM is the NIC interface connected to the “Cloud network” as shown in Figure 5.
 - A. In **System Configuration**, select **BIOS/Platform Configuration→Network Options→Network Boot Options**.
 - B. Set “Embedded FlexibleLOM 1 Port 1” to **Network Boot** and repeat for Port2
 - C. For all other NIC interfaces, set their network boot option to **Disabled**.
6. The boot order on each cloud node should be configured with Network/PXE Boot as the first option.
 - A. In **System Configuration**, select **BIOS/Platform Configuration→Boot Options→UEFI Boot Order**. Move “Embedded FlexibleLOM port 1” up to be the first one in the boot order list.
7. For cloud servers, the local storage controller and disk drives should be configured to match their respective role, utilizing these steps:
 - A. Reboot each server, and press **F9** to get into **System Utilities**
 - B. elect **System Configuration→Embedded RAID**: <embedded Smart Array controller model>
 - C. Select **Exit** and Launch **HP Smart Storage Administrator (HPSSA)**
 - D. Select **Smart Storage Administrator**
 - E. Follow the GUI to select proper storage controller and hard drives, then create arrays and logical volumes according to the following table. For detailed instruction on how to use HP Smart Storage Administrator, refer to: www8.hp.com/us/en/products/server-software/product-detail.html?oid=5409020.

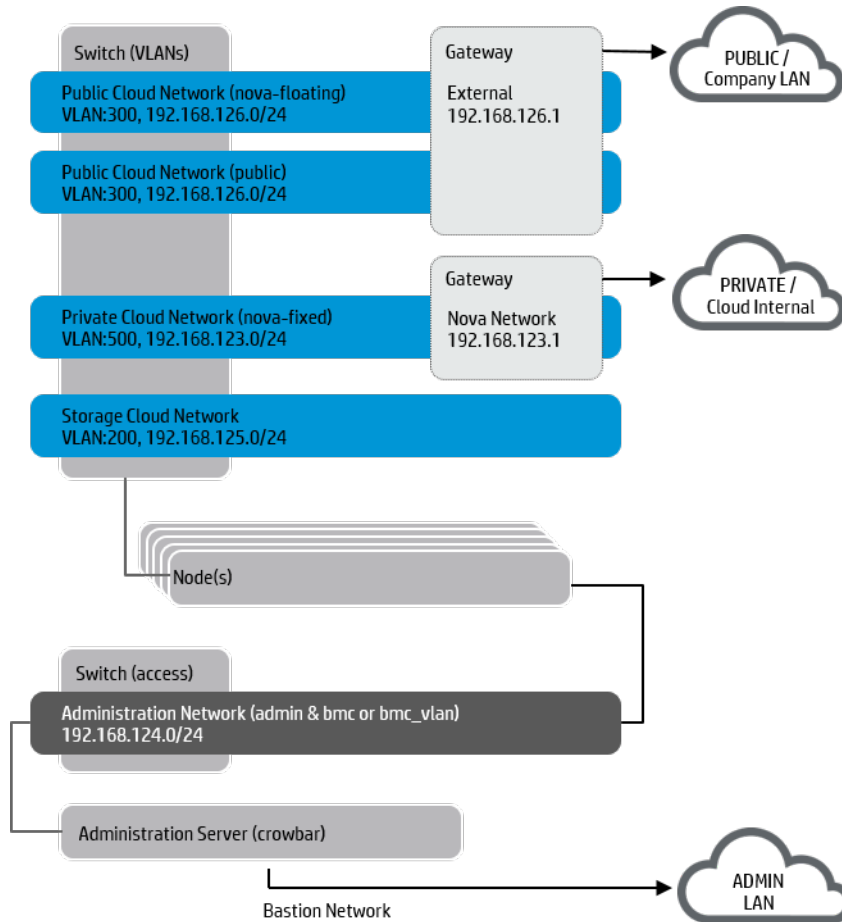
Table 3. Storage configuration

Role	OS Boot Drive Logical Volume	Data Drive Logical Volume
Administration Server	RAID10 – 2x2TB LFF drives (1863GB)	None
Control node	RAID10 – 4x600GB LFF drives (1117GB)	None
Storage node	RAID10 -- 2x600GB LFF drives (558 GB)	RAID0 – 5x6TB LFF drives (29.29TB)
Compute node	RAID10 – 4x1.2TB SFF drives (2235 GB)	None

VLAN configuration

This solution for configuring VLANs for the SUSE OpenStack Cloud networks uses an HP FlexFabric 5700 Switch as an aggregation router to provide routing among the networks displayed in Figure 5.

Figure 5. Network VLAN configuration and topology for the solution configuration



As shown in Figure 5, use one 5700 Switch (HP FlexFabric 5700-48G-4XG-2QSFP+ Switch) to implement a 1 Gb network for the Admin network. The first 10 ports of this switch are wired to each bare metal server's iLO port (except for the SUSE OpenStack Cloud Administration Server).

The first SFP+ port of this switch is connected to the second 5700 Switch to connect the ILO network to the other network. Configuration of this 5700 Switch and ILO network is done with following steps:

```
# Create ILO VLAN
<switch5700> system-view
[switch5700] vlan 100
[switch5700-vlan100] name iLO_vlan
[switch5700-vlan100] quit
[switch5700] quit
<switch5700>
```

Port 1-10 will be connected to each resource node's iLO. These ports are configured as untagged for ILO
 # VLAN. The following example shows how port #1 is configured. The rest of the ports can be configured the

```

# same way.
<switch5700> system-view
[switch5700] interface GigabitEthernet 1/0/1
[switch5700-GigabitEthernet1/0/1] port link-type access
[switch5700-GigabitEthernet1/0/1] port access vlan 100
[switch5700-GigabitEthernet1/0/1] quit
[switch5700] quit
<switch5700>

# Port 49 is connected to the second 5700 Switch to connect ILO network to
other networks.
<switch5700> system-view
[switch5700] interface Ten-GigabitEthernet 1/0/49
[switch5700-Ten-GigabitEthernet1/0/49] port link-type access
[switch5700-Ten-GigabitEthernet1/0/49] port access vlan 100
[switch5700-Ten-GigabitEthernet1/0/49] quit
[switch5700] quit
<switch5700>

```

The second 5700 Switch (HP FlexFabric 5700-40XG-2QSFP) in Figure 5 is used to create a 10 Gb Cloud network. The first 11 ports of this switch are connected to each baremetal server's 10 Gb dual port FlexibleLOM port 1, while the next 11 ports connect to port 2, providing a bonded teamed interface. The last SFP+ port of this switch is connected to the first 5700 Switch for IPMI network access. The configuration of this second 5700 Switch is done with following steps:

```

# Create mgmt., external, service, IPMI, and internet vlan
<switch5700> system-view
[switch5700] vlan 500
 [switch5700] name vlan_fixed
[switch5700] quit
[switch5700] vlan 300
[switch5700] name vlan_public
[switch5700] quit
[switch5700] vlan 400
 [switch5700] name vlan_sdn
[switch5700] quit
[switch5700] vlan 100
[switch5700] name iLO_vlan
[switch5700] quit
[switch5700] vlan 200
[switch5700] name vlan_storage
[switch5700] quit
[switch5700] quit

# Configure routing interface on each vlan
<switch5700> system-view
[switch5700] interface vlan 500
[switch5700] ip address 192.168.123.1 255.255.255.0
[switch5700] quit
[switch5700] interface vlan 300
[switch5700] ip address 192.168.126.1 255.255.255.0
[switch5700] quit
[switch5700] interface vlan 400
[switch5700] ip address 192.168.130.1 255.255.255.0
[switch5700] quit
[switch5700] interface vlan 100

```

```
[switch5700] ip address 192.168.124.1 255.255.255.0
[switch5700] quit
[switch5700] interface vlan 200
[switch5700] ip address 192.168.125.1 255.255.255.0
[switch5700] quit
[switch5700] quit
```

Switch port 1-11 will be connected to each resource node's port 1 of dual ports 10 Gb FlexibleLOM (and ports 12-22 connect to port 2).
 # These switch ports need to be configured as untagged for management VLAN 500, then tagged for external

network 300 and service network 400. Following is an example of how to configure port 1. Other ports can be configured with similar steps.

```
<switch5700> system-view
[switch5700] interface Ten-GigabitEthernet 1/0/1
[switch5700-Ten-GigabitEthernet 1/0/1] jumboframe enable
[switch5700-Ten-GigabitEthernet 1/0/1] port link-type trunk
[switch5700-Ten-GigabitEthernet 1/0/1] port trunk permit vlan all
[switch5700-Ten-GigabitEthernet 1/0/1] port trunk pvid vlan 500
[switch5700-Ten-GigabitEthernet 1/0/1] quit
[switch5700] quit
```

Port 40 is connected to IPMI network switch for ILO access.

```
<switch5700> system-view
[switch5700] interface Ten-GigabitEthernet 1/0/40
[switch5700-Ten-GigabitEthernet1/0/40] port link-type access
[switch5700-Ten-GigabitEthernet1/0/40] port access vlan 100
[switch5700-Ten-GigabitEthernet1/0/40] quit
[switch5700] quit
<switch5700>
```

SUSE OpenStack Cloud installation

For the remainder of this software installation section, consult the respective product documentation and use the following information to augment and highlight important configuration guidance.

Deploying SUSE OpenStack Cloud

With the HP Network and ProLiant server platforms prepared, the SUSE OpenStack Cloud software can be deployed to create the private cloud instance.

Administration Server setup and configuration

With the designated HP ProLiant DL360 server, use direct console access and an external media drive or the virtual media option of the iLO and the ISO image file, perform an x86_64 architecture installation of the following components:

- SUSE Linux Enterprise Server
- SUSE OpenStack Cloud
- SUSE Subscription Management Tool or SMT (needed only if you don't already have access to such repositories)

Additional guidance-

- Currently, SUSE OpenStack Cloud 5 requires SLES 11 SP3. Going forward, you can find the latest updates at: downloads.suse.com. Utilize the respective kISO or Bootable Driver Kit (drivers.suse.com/hp/HP-ProLiant-Gen9/) as the initial boot media before proceeding on to the main operating system media.
- Only the minimal operating system patterns cited in the SUSE OpenStack Cloud deployment guide are required.

Complete the next set of steps as described in the “SUSE OpenStack Cloud Deployment Guide”, including:

- Product Registration
- Online Update
- CA Setup
- Basic Network Configuration (ensuring that the firewall for all network interfaces and IPv6 are disabled)
- SMT Configuration (optional)

Carefully review the **Crowbar Setup** section, selecting the appropriate network mode for this high availability configuration. The “team” network mode selection with a bonding policy of “4” allows multiple 10G network interfaces to bond together and validate that the remaining resource nodes are all similar with regard to network connectivity, adapter configuration, and overall configuration, as per the previous pre-installation requirements. In addition, manually ensure that the “conduit map” section pertaining to the “team” configuration section of /etc/crowbar/network.json file correctly references “10g1”, “10g2” across all the respective interface listings. See the Network Conduits section of the deployment guide for more details.

Perform the setup step for all the necessary software repositories, as noted in the “SUSE OpenStack Cloud Deployment Guide”.

Additional guidance-

- Some time is required to set up all of this properly and to mirror all content. Before proceeding, it is imperative that all the software repositories are mirrored and available on the SUSE OpenStack Cloud Administration Server.
- As the compute and storage nodes will be deployed with SUSE Linux Enterprise 12, ensure that these repositories are also populated and available before proceeding.

Then, using a terminal multiplexer like the **screen** utility, complete the following process:

- Execute the **screen install-suse-cloud** setup script.
- When it completes, review the noted log file for any issues.

Note

The network setup is fixed at this point and cannot be modified. To make changes, you must restart the installation of the Administration Server and associated setup process.

Some additional integration may be undertaken at this point to ease later setups and provide a more convenient working environment for the administrator and users. Remember that the SUSE OpenStack Cloud Administration Server is your default DNS server for the private cloud instance.

Many of the remaining required nodes in the SUSE OpenStack Cloud infrastructure can be set up, configured, and managed from the administration server web interface, which is found on port 3000 of the IP address of your administration server (by default located at IP address http://192.168.124.10:3000). PXE boot each of the systems (including the Compute and Storage Nodes) against the administration server via the Admin Network VLAN subnet. Details for the provisioning of the Control Nodes is found in that respective section later in this document. As each resource node becomes visible to the administration server after PXE booting, basic system information is returned to the administration web interface in a “Discovered” mode. At this point, the new node can be “Allocated” for use as a Compute or Storage Node.

Additional guidance-

As each node is “Discovered,” it is advisable to:

- Enter a meaningful description
- Create a public name (optional) to provide a handle to manage each of your nodes.
- Use alias names to make identifying the node easier (instead of the default MAC address designation)
- Create functional groups to visually organize the nodes on the interface
- Set the intended role for each node to aid later association of roles to a given node, and so that the subsequent actions auto-populate nodes appropriately

It is recommended to allocate all the nodes (Control, Compute and Storage) with their designated operating system (as described in the following sections) before starting to deploy the OpenStack services and roles.

Control Node(s) setup and configuration

Access the designated resource nodes and prepare to install the operating system. Since the control nodes are currently required to run SUSE Linux Enterprise 11 SP3 (you can find the latest updates at: downloads.suse.com), install each using the kISO/Bootable Driver Kit media and then the operating system media, via direct console access or using iLO Virtual Media, with the same minimal patterns as noted for the Administration Server. Then follow the Deployment Guide information, utilizing “crowbar_register”, to convert existing SLES 11 SP3 systems into SUSE OpenStack Cloud Nodes.

Additional guidance-

Both the kISO/Bootable Driver Kit and SUSE Linux Enterprise Server media can be hosted on the Administration Server’s SMT web server for ready access by the Control Node’s iLO virtual media interface.

With all three Control Nodes installed and integrated into the private cloud instance, apply the Pacemaker role from the SUSE OpenStack Cloud Crowbar web interface to the respective cluster nodes with the following changes to the supplied defaults:

- Create a recognizable proposal name (e.g. ControlHA)
- For STONITH, select “Configured with IPMI data from IPMI barclamp”
- Setup non-web GUI (hb_gui) as “true”
- Include all three nodes into “pacemaker-cluster-member” and “hawk-server”

Compute node(s) setup and configuration

For each of the four designated resource nodes, integrate each into the private cloud instance by:

- PXE booting the respective systems to act as Compute Nodes
- After successful discovery, verify all of the Compute Node’s resource are accurately reflected in the inventory before allocation of the SUSE Linux Enterprise Server 12 target platform

Storage node(s) setup and configuration

Access the three respective resource nodes to be used as storage by:

- PXE booting the designated systems.
- For storage nodes, select SUSE Linux Enterprise Server 12 as the target platform during the allocate phase.

Additional guidance-

As a reminder, SUSE OpenStack Cloud is configured to always use all unused disks (beyond the first disk which is assigned to the operating system)

Deploying the OpenStack Services

With all the resource nodes discovered, allocated and now ready, follow the respective section of the SUSE OpenStack Cloud Deployment Guide as a reference to apply each of the OpenStack barclamps. Step through each service (see the [Figure 2](#) earlier in this document showing the services in the private cloud instance) by creating a proposal, adjusting configuration parameters to your local situation and assigning nodes (or the control node cluster) to each role, finally applying the barclamp.

Additional guidance-

Beyond any localized login, password, or credential configuration settings, the changes from default settings for each of the OpenStack barclamps are noted below:

- Database, RabbitMQ, Keystone, Neutron, Nova, Horizon, Heat, Ceilometer – assign the ControlHA cluster to the designated controller role(s)
- Ceph – assign all of the Storage Nodes to the “ceph-mon” and “ceph-osd” roles
- Glance – select “Rados” as the “Default Storage Store” to use the Storage Nodes from the Ceph setup
- Cinder – select “rbd” as the “Type of Volume” for the “Add new Cinder Backend” and delete the “Backend: default” Cinder Backend. Also assign the ControlHA cluster to the “cinder-controller” role.
- Nova – assign all of the compute nodes to the “nova-multi-compute-kvm” role
- If any errors arise, address these first before proceeding onto the next service. If you cannot resolve the issue quickly, check if the troubleshooting section of the deployment guide might help or generate a report on the suspect nodes using the supportconfig tool, before contacting your support organization.

At this point, you should be able to utilize your private cloud instance, according to the Admin, User and Supplement guides found at suse.com/documentation/suse-cloud-5/. You can also optionally perform a basic functionality and API test of your overall installation, utilizing the OpenStack Integration Test Suite (Tempest) to exercise the “smoke” tests that are part of the OpenStack gate process by:

- Installing the Tempest barclamp on your Administration Server, via “zypper install crowbar-barclamp-tempest”
- From the SUSE OpenStack Cloud Crowbar web interface, assign one of your Control Nodes to the “tempest” role
- Then login to the Control Node and run (these should all pass)
- `/var/lib/openstack-tempest-test/run_tempest.sh -N -t -s -L /etc/tempest/logging.conf`
- Refer to docs.openstack.org/developer/tempest/overview.html for more information on Tempest tests


Conclusion

After understanding and working through the steps described in this document, you should have a working cloud that is scalable through the addition of even more compute and storage nodes, as needed. SUSE OpenStack Cloud is a complete suite of software and may be configured in many different ways. This solution architecture provides a baseline for a private cloud implementation with highly available control plane services, using HP ProLiant rack servers and can serve as a functional environment for many workloads.

HP ProLiant DL Gen9 Server Series hardware is an excellent platform for implementation of OpenStack software. It provides powerful, dense compute, and storage capabilities via the servers we’ve selected for this solution architecture. And the HP iLO management capability is indispensable in managing a cloud cluster of this kind.

Appendix A: SUSE OpenStack Cloud rack layout

HP ProLiant for Private Cloud with SUSE OpenStack Cloud Rack Layout



Appendix B: Bills of material

Important

Some components in the following bills of material are for use in North America and Japan only. For deployments outside of these regions, you should consider localized versions of these components.

Table 4. Administration Server - DL360p Gen9 4LFF— 1 of Intel Xeon E5-2620v3 -- 6 Core 2.4Ghz - Quantity 1

Quantity	Product number	Name	Description
1	755259-B21	ProLiant DL360 Gen9	HP ProLiant DL360 Gen9 4LFF Configure-to-order Server HP Dynamic Smart Array B140i 4-Bay LFF Drive Cage
1	755382-L21	Processor	Intel Xeon E5-2620v3 (2.4GHz/6-core/15MB/85W)
4	726718-B21	Memory	8GB (1x8GB) Single Rank x4 DDR4-2133 CAS-15-15-15 Registered Memory Kit
1	665243-B21	Flexible LOM	HP Ethernet 10Gb 2-port 560FLR-SFP+ FIO Adapter
2	720478-B21	Power Supplies	HP 500W Flex Slot Platinum Hot Plug Power supply Kit
1	749974-B21	Storage Controller	HP Smart Array P440ar/2GB FBWC 12Gb 2-ports Int FIO SAS Controller
1	766211-B21	Storage cable	HP DL360 Gen9 LFF P440ar/H240ar SAS Cbl
2	652757-B21	Hard Drives	HP 2TB 6G SAS 7.2K rpm LFF (3.5-inch) SC Midline 1yr Warranty Hard Drive
1	663202-B21	Rail Kit	HP 1U Large Form Factor Ball Bearing Rail Kit
1	663203-B21	Cable Management Arm	HP 1U Cable Management Arm for Ball Bearing Rail Kit

Table 5. Control Nodes - DL360 Gen9 4LFF 2 of Intel Xeon E5-2690 -
2 of 12 core at 2.6Ghz -Quantity 3

Quantity	Product number	Name	Description
1	755259-B21	ProLiant DL360 Gen9	HP ProLiant DL360 Gen9 4LFF Configure-to-order Server HP Dynamic Smart Array B140i 4-Bay LFF Drive Cage
1	755382-L21	Processor 1	Intel® Xeon® E5-2690v3 (2.6GHz/12-core/30MB/135W) FIO Processor Kit
1	755382-B21	Processor 2	Intel® Xeon® E5-2690v3 (2.6GHz/12-core/30MB/135W) FIO Processor Kit
4	726718-B21	Memory for 1st processor	8GB (1x8GB) Single Rank x4 DDR4-2133 CAS-15-15-15 Registered Memory Kit
4	726718-B21	Memory for 2nd processor	8GB (1x8GB) Single Rank x4 DDR4-2133 CAS-15-15-15 Registered Memory Kit
1	665243-B21	FlexibleLOM	HP Ethernet 10Gb 2-port 560FLR-SFP+ FIO Adapter
2	720478-B21	Power Supplies	HP 800W Flex Slot Platinum Hot Plug Power Supply Kit
4	765424-B21	Hard Drives	HP 600GB 12G SAS 15K rpm LFF (3.5-inch) SC Enterprise 3yr Warranty Hard Drive
1	749974-B21	Storage Controller	HP Smart Array P440ar/2GB FBWC 12Gb 2-ports Int FIO SAS Controller
1	766211-B21	Storage Cable	HP DL360 Gen9 LFF P440ar/H240ar SAS Cbl
1	663202-B21	Rail Kit	HP 1U Large Form Factor Ball Bearing Rail Kit
1	663203-B21	Cable Management Arm	HP 1U Cable Management Arm for Ball Bearing Rail Kit

Table 6. Storage Nodes - DL380 Gen9 12LFF, 2 of Intel Xeon E5-2690
 - 2 of 12 core at 2.6Ghz-Quantity 3

Quantity	Product number	Name	Description
1	719061-B21	ProLiant DL380 Gen9	HP ProLiant DL380 Gen9 12LFF Configure-to-order Server HP Dynamic Smart Array B140i 12-Bay LFF Drive Cage
1	719044-L21	Processor 1	Intel® Xeon® E5-2690v3 (2.6GHz/12-core/30MB/135W) FIO Processor Kit
1	719044-B21	Processor 2	Intel® Xeon® E5-2690v3 (2.6GHz/12-core/30MB/135W) FIO Processor Kit
4	726718-B21	Memory for 1st processor	8GB (1x8GB) Single Rank x4 DDR4-2133 CAS-15-15-15 Registered Memory Kit
4	726718-B21	Memory for 2nd processor	8GB (1x8GB) Single Rank x4 DDR4-2133 CAS-15-15-15 Registered Memory Kit
1	665243-B21	Flexible LOM	HP Ethernet 10Gb 2-port 560FLR-SFP+ FIO Adapter
2	720479-B21	Power Supply	HP 800W Flex Slot Platinum Hot Plug Power Supply Kit
1	724864-B21	Additional Disk Bay	HP DL380 Gen9 2SFF Bay Kit
2	765424-B21	Hard Drives	HP 600GB 12G SAS 15K rpm LFF (3.5-inch) SC Enterprise 3yr Warranty Hard Drive
5	761477-B21	Hard Drives	HP 6TB 6G SAS 7.2k rpm LFF (3.5-inch) SC Midline 1yr Warranty Hard Drive
1	761874-B21	Storage Controller	HP Smart Array P840/4G FIO Controller
1	785991-B21	Storage Cable	HP DL380 Gen9 12LFF SAS Cable Kit
1	783007-B21	Storage Cable	HP DL380 Gen9 P840/440 SAS Cable Kit
1	665249-B21	Network	HP Ethernet 10Gb 2-port 560SFP+ Adapter
1	720864-B21	Rail Kit	HP 2U Large Form Factor Ball Bearing Rail Kit
1	720865-B21	Cable Management Arm	HP 2U Cable Management Arm for Ball Bearing Rail Kit

Table 7. Compute Nodes - DL360p Gen9 8SFF 2 of Intel Xeon E5-2699
 -- 2 of 18 core at 2.3Ghz-Quantity 4

Quantity	Product number	Name	Description
1	755258-B21	ProLiant DL360 Gen9	HP ProLiant DL360 Gen9 8SFF Configure-to-order Server HP Dynamic Smart Array B140i 8-Bay SFF Drive Cage
1	780003-L21	Processor 1	Intel® Xeon® E5-2699v3 (2.6GHz/18-core/45MB/145W) FIO Processor Kit
1	780003-B21	Processor 2	Intel® Xeon® E5-2699v3 (2.6GHz/18-core/45MB/145W) FIO Processor Kit
8	726719-B21	Memory for 1st processor	16GB (1x16GB) Dual Rank x4 DDR4-2133 CAS-15-15-15 Registered Memory Kit
8	726719-B21	Memory for 2nd processor	16GB (1x16GB) Dual Rank x4 DDR4-2133 CAS-15-15-15 Registered Memory Kit
1	665243-B21	FlexibleLOM	HP Ethernet 10Gb 2-port 560FLR-SFP+ FIO Adapter
2	720479-B21	Power Supply	HP 800W Flex Slot Platinum Hot Plug Power Supply Kit
6	759212-B21	Hard Drives	HP 600GB 12G SAS 15K rpm SFF (2.5-inch) SC Enterprise 3yr Warranty Hard Drive
1	749974-B21	Storage Controller	HP Smart Array P440ar/2GB FBWC 12Gb 2-ports Int FIO SAS Controller
1	663201-B21	Rail Kit	HP 1U SFF Ball Bearing Rail Kit
1	663203-B21	Cable Management Arm	HP 1U Cable Management Arm for Ball Bearing Rail Kit

Table 8. Rack Infrastructure - Quantity 1

Quantity	Product number	Name	Description
1	H6J68A	Rack	HP 11642 1200mm Shock Universal Rack
1	AF630A	KVM console and monitor	HP LCD8500 1U US Rackmount Console Kit
1	AF617A	KVM switch	HP 0x2x16 KVM Svr Cnsl G2 SW
4	H5M58A	PDU	NA/JPN - HP 4.9kVA 208V 20out NA/JP bPDU
1	JG896A	Production Switch	HP FF 5700-40XG-2QSFP+ Switch
1	JG894A	Management Switch	HP FF 5700-48G-4XG-2QSFP+ Switch
2	ZU717A	Factory Express	Network Custom Cable Kit (10/kit, add more with scaling)
2	H1K92A3 TF4	Proactive Care 24x7	HP FF 5700 Switch Support
4	JG900A	Switch power supplies	HP 58x0AF 300W AC Power Supply
4	JC682A	Switch fan trays	HP 58x0AF Bck(pwr)-Frt(ports) Fan Tray
1	JD097C	SFP+ cables	HP X240 10G SFP+ SFP+ 3m DAC Cable (add more with scaling of compute or storage)
22	JG081C	SFP+ cables	HP X240 10G SFP+ SFP+ 5m DAC Cable
11	AF628A	Monitor/KVM adapters	HP KVM USB Adapter (add more with scaling of compute or storage)
1	BW932A	Rack Stabilizer	HP Rack Stabilizer Kit 600mm
1	BW930A	Airflow Optimization	HP Airflow Optimization Kit
1	BW909A	Side Panel Kit	HP 42U 1200mm Side Panel Kit

Table 9. Software Subscription, 1-year

Quantity	SKU	Description
1	874-007031	SUSE OpenStack Cloud Control Node plus Admin Server, x 86-64, 1 Instance, L3-Priority Subscription, 1 Year
2	874-007033	SUSE OpenStack Cloud Control Node, x86-64, 1 Instance, L3-Priority Subscription, 1 Year
4	874-007035	SUSE OpenStack Cloud Compute Node, x86-64, 1-2 Sockets, L3-Priority Subscription, 1 Year
4	874-006878	SUSE Linux Enterprise Server, x86 & x86-64, 1-2 Sockets with Unlimited Virtual Machines, L3-Priority Subscription, 1 Year
1	874-007043	SUSE Enterprise Storage Base Configuration, x86-64, 4 OSD Nodes with 1-2 Sockets, L3-Priority Subscription, 1 Year

Table 10. Software Subscription, 3-Year

Quantity	SKU	Description
1	874-007032	SUSE OpenStack Cloud Control Node plus Admin Server, x 86-64, 1 Instance, L3-Priority Subscription, 3 Year
2	874-007034	SUSE OpenStack Cloud Control Node, x86-64, 1 Instance, L3-Priority Subscription, 3 Year
4	874-007036	SUSE OpenStack Cloud Compute Node, x86-64, 1-2 Sockets, L3-Priority Subscription, 3 Year
4	874-006886	SUSE Linux Enterprise Server, x86 & x86-64, 1-2 Sockets with Unlimited Virtual Machines, L3-Priority Subscription, 3 Year
1	874-007044	SUSE Enterprise Storage Base Configuration, x86-64, 4 OSD Nodes with 1-2 Sockets, L3-Priority Subscription, 3 Year

Resources

HP ProLiant product information

hp.com/go/proliant

HP ProLiant Linux product information

hp.com/go/proliantlinux

SUSE Cloud OpenStack User Guides

suse.com/documentation/suse-cloud-5/

OpenStack Operations Guide

docs.openstack.org/ops

HP white paper library

hp.com/servers/technology

Sign up for updates

hp.com/go/getupdated



Share with colleagues



Rate this document

© Copyright 2015 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. Intel Xeon is a trademark of Intel Corporation in the U.S. and other countries. Windows is a trademark of the Microsoft group of companies. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. The OpenStack word mark is a trademark or registered trademark of OpenStack Foundation in the United States and other countries, and is used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community. Cloud Foundry is a trademark and/or registered trademark of Pivotal Software, Inc. in the United States and/or other countries.

