REFERENCE ARCHITECTURE
Intel® Cloud Builders
Cloud Computing

SUSE® OpenStack Cloud Reference Architecture with Dell Hardware

Purpose

The intent of this document is to provide an overview of a complete, OpenStack-based, private cloud solution, including the core physical infrastructure of Dell PowerEdge servers and Dell Networking switches integrated with enterprise-class SUSE® OpenStack Cloud software. Deploying this fully supported private cloud solution helps you to increase business agility, maximize your existing data center investments and improve your resource utilization with scaling of both compute and storage resources.

Introduction

This white paper aims to help organizations create and deploy a private cloud instance within their own network space using computing and storage resources under their own control. It covers networking, computing hardware, software and key areas of integration.

This document also focuses on how to leverage the highly automated nature of this private cloud installation to quickly and confidently deploy the infrastructure, yielding a scalable, highly available solution for any type of workload whether for testing, development or production services.

The target audience is IT professionals responsible for setting up, configuring, administering and operating a private cloud infrastructure.

Table of Contents

Solution Focus

Through integrated, open source software, enterprises can establish an Infrastructure-as-a-Service (IaaS[1]) private cloud that delivers on-demand access to pools of compute, storage and networking resources for use within their organization. These resources are provisioned as necessary to meet the specific requirements of the applications and services deployed in the cloud. Compute resources are provided through the provisioning of physical servers or virtual machines running on top of a hypervisor. Providing storage resources can be accomplished with either dedicated storage elements or via distributed, resilient file systems across industry-standard hardware with local storage devices. Networking provides the unifying connectivity that ties the compute, storage and management elements of the private cloud together. Given the complexity of such an infrastructure, the network must be provided in a fully configurable fashion, adapting to change requests on demand. Finally, it is the private cloud software that provides the user interface for setup, configuration and maintenance as well as the long-term operation of the cloud by bonding these three core components into a cohesive service offering.

Background

Cloud computing delivers access to resources like hardware and software over a network, abstracting the complex, internal infrastructures from end users and, to an extent, from some development and information technology operational professionals (DevOps). 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 further reference, the National Institute of Standards and Technology's (NIST) Definition of Cloud Computing[2] describes the important characteristics of cloud computing.

While public cloud services have existed for some time, a private cloud enables companies to take advantage of the benefits of cloud computing with less risk of data exposure and more control of resources because it is operated behind the corporate firewall. However, compared to public clouds, private clouds impact capital cost saving profiles because they continue to require investment in locally maintained infrastructure. Yet, because private cloud resources can be shared via higher levels of automation, they offer potentially higher overall utilization with the total cost of ownership likely lower than traditional IT infrastructure deployment change requests on demand.

Approach

This white paper provides a complete private cloud solution including the core physical infrastructure and enterprise-ready versions of many open source offerings. With this framework, enterprises can take a secure approach to increase agility, enhance scalability and improve utilization of resources. Such a private cloud instance can even be integrated with or extended to public cloud resources with the Dell Cloud Manager[3] product, enabling the provisioning, automation and management of workloads across this hybrid cloud approach. Conceptually, the basic building blocks

[1] http://en.wikipedia.org/wiki/Infrastructure_as_a_service
[2] http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf
[3] http://software.dell.com/products/cloud-manager

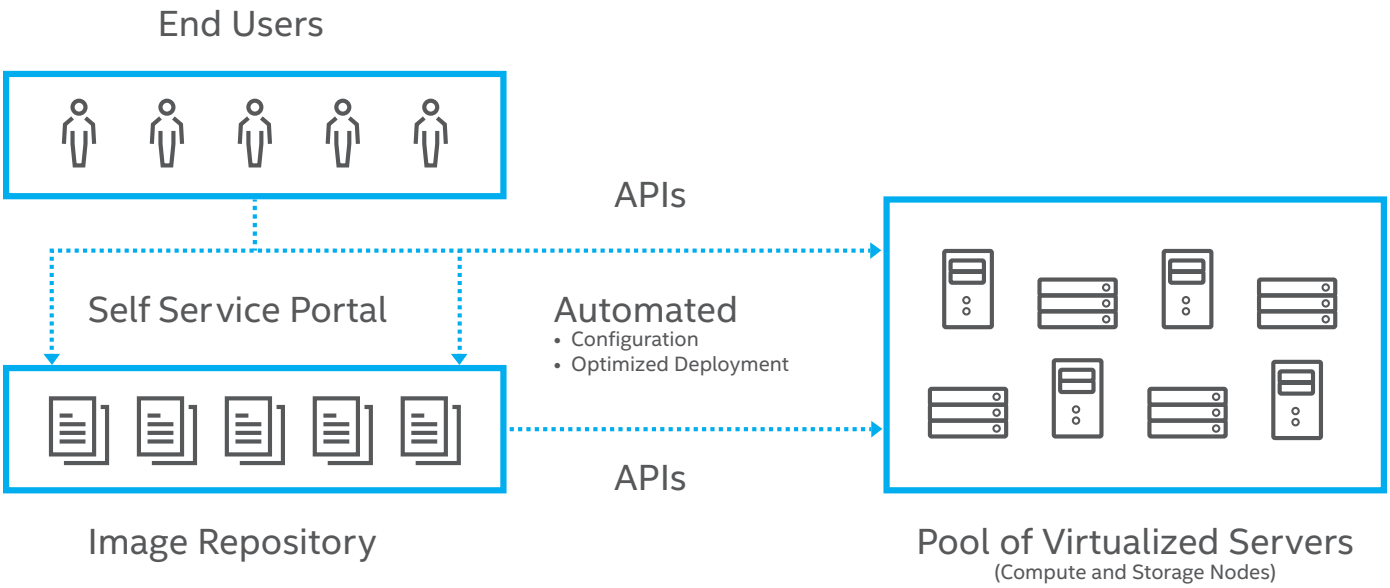

Figure 1: Infrastructure-as-a-Service (IaaS) private cloud deployment

for deploying a private cloud (i.e., networking infrastructure, computing hardware and software) are similar to any IT infrastructure.

With private clouds, delivering robust network resources for the core infrastructure and the needs of the workloads requires significant attention.

High-speed, scalable devices like the Dell Network Managed S-series of network switches and modern, industry-standard x86-based servers like the Dell PowerEdge R Servers provide an ideal platform for private clouds because they balance performance and efficiency. The emerging software standard for private cloud implementations is the open source Linux operating system. The current reference framework for the cloud itself, OpenStack[4], is also open source, as are many cloud sub-components. Examples include KVM or Xen virtualization technologies and Crowbar, an open source project founded by Dell, which facilitates system deployment and enables scaling with the private cloud instance. These open source components deliver an overall private cloud ecosystem that yields a completely functional, usable and stable installation today. In addition, the rapid innovation of open source development will continue to rapidly enhance current features and add capabilities. The figure below illustrates all of these components and interactions to deliver an Infrastructure-as-a-Service (IaaS) private cloud deployment: This white paper introduces the setup and configuration of networking services for a private cloud instance using ultra-low latency Dell Networking top-of-rack (ToR) switches as the networking infrastructure, Dell PowerEdge servers for both compute and storage resources and SUSE OpenStack Cloud[5] as the private cloud software.

For the remainder of this document, the user should refer to the documentation for the specific network and computing hardware and software products under discussion, using this white paper as supplemental information to augment, clarify and further refine the use of the other documentation for a successful deployment.

Site Preparation Needed for the Deployment

Facility Considerations

The heating, ventilation, air conditioning (HVAC) and power requirements for deployment can be estimated using the Dell Energy Smart Solution Advisor[6]. Using this tool, you can plan the needs for your solution, order the correct Power Distribution Unit (PDU) and have the proper HVAC ready for the installation.

This white paper does not specify any country-specific localization or PDU since power requirements vary depending on the exact configuration and the power configuration in the final installation location.

[4] www.openstack.org/

[5] www.suse.com/products/suse-cloud

[6] http://essa.us.dell.com/DellStarOnline/DCCP

Networking Considerations

Because networking is the technology component likely to take the most upfront planning, networking requirements for a private cloud instance have a fair amount of complexity, especially when integrated with an existing IT infrastructure. For the physical level, use a pair of Dell Networking S4048[7] top-of-rack (ToR) 10GbE with 40GbE uplinks connected together with Virtual Link Trunking (VLT). These Dell Networking switches offer an ultra-low-latency switch fabric providing non-blocking performance. To complete the fabric, a single Dell Networking S3048[8] switch is used to handle the 1GbE connections from the dedicated BMC/iDRAC ports of each resource node and is uplinked to both of the S4048 switches for redundancy. With each server having multiple NIC ports, forming a bonded link across the ports with at least one port on each switch in the stack provides both performance and redundancy across the multiple NIC and switch configuration. The appendices include the recommended configuration for each of these Dell Networking switches plus the types of cables to include in an order.

Computing Platform Considerations

One of the benefits of a private cloud implementation is that industry-standard servers can fulfill all the 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 popular Dell PowerEdge R430[9], R630[10], R730[11] and R730xd[12] series of servers with their powerful and balanced performance, advanced I/O capabilities and flexible, scalable networking options are ideally suited for all the various nodes in this private cloud installation reference configuration . The appendices contain the recommended configurations for each of these system platforms with a preferred model for each of the various private cloud node roles.

NETWORK NAME	NETWORK DESCRIPTION/CONSIDERATIONS
Admin	**Private** network to access the Administration Server and all nodes for administration purposes. The default setup lets you also access and manage any available BMC (Baseboard Management Controller) data via IPMI (Intelligent Platform Management Interface) from this network. If required, BMC access can be utilized on a separate network. You have the following options for controlling access to this network: • Do not allow access from the outside and keep the admin network completely separated. • Allow access to the Administration Server from a single network (for example, your company's administration network) via the "bastion network" option configured on an additional network card with a fixed IP address. • Allow access from one or more networks via a gateway.
Storage VLAN:200	**Private** SUSE Cloud internal virtual network. This network is used by Ceph and OpenStack Object Store (Swift) only. It should not be accessed by users.
Public (nova-floating) VLAN:300	The only **public** network provided by SUSE OpenStack Cloud. You can access the OpenStack Dashboard (Horizon) and all workload instances (provided they have been equipped with a floating IP) via this network. You have the following options for controlling access to this network: • This network can be accessed only by a gateway, which needs to be provided externally. • All SUSE OpenStack Cloud users and administrators need to be able to access the public network. • The two functions can be split, if so desired, to address security concerns, so that the OpenStack Dashboard and APIs are assessable from on range and the floating IPs assigned to the workload are in another network range.
Software Defined Network (os_sdn) VLAN:400	**Private** SUSE Cloud internal virtual network. This network is used when OpenStack Networking (Neutron) is configured to use Open vSwitch with GRE tunnelling for the virtual networks. It chould not be accessed by users.
Private (nova-fixed) VLAN:500	**Private** SUSE OpenStack Cloud internal virtual network. This network is used for inter-instance communications and provides access to the outside world for the instances. The gateway required is also automatically provided by SUSE OpentStack Cloud.

[7] www.dell.com/us/business/p/open-networking-switches/pd
[8] www.dell.com/us/business/p/open-networking-switches/pd
[9] www.dell.com/us/business/p/poweredge-r430/pd
[10] www.dell.com/us/business/p/poweredge-r630/pd
[11] www.dell.com/us/business/p/poweredge-r730/pd
[12] www.dell.com/us/business/p/poweredge-r730xd/pd

Software Component Considerations

SUSE OpenStack Cloud provides the enterprise-grade implementation of OpenStack components and includes other tools for a complete private cloud implementation. Fundamental to a private cloud implementation, various networks are needed to manage, operate, access and provide isolation for all the components and services. SUSE OpenStack Cloud utilizes the following set of network namespaces:

SUSE OpenStack Cloud forms the basis of the IaaS private cloud operating solution and overall framework to abstract the network, computing and storage resources by including a number of OpenStack projects, components and services, as shown in the table below. Based upon the previously described network namespace model and the various OpenStack services, multiple physical hosts are required for the various server roles in a SUSE OpenStack Cloud implementation:

Administration Server

- Usually the first system installed, the SUSE OpenStack Cloud Administration Server is a dedicated system that provides all services needed to set up and deploy all other resource nodes in the private cloud installation. It runs on the SUSE Linux Enterprise Server operating system and provides the deployment framework with the Crowbar deployment engine along with other tools. Among the services provided to the cloud instance by this server are DHCP, DNS, NTP, PXE and TFTP. These components are set up, configured and accessed from a web interface that guides all the remaining installation and configuration tasks through the use of barclamps.
- The Administration Server also commonly hosts the necessary software repositories for SUSE Linux Enterprise Server and SUSE OpenStack Cloud, along with other software products, since they are needed to deploy the Control, Compute and Storage Nodes. The remaining Control, Compute and Storage Nodes are all provisioned by simply PXE network booting from the SUSE OpenStack Cloud Administration Server. The primary user of this node and its web interface is the cloud administrator.

Control Node(s)

- A SUSE OpenStack Cloud Control Node hosts the entire core OpenStack services needed to orchestrate virtual machines deployed on the Compute and utilizing volumes from the Storage Nodes in the private cloud installation. Like the SUSE

OPENSTACK PROJECTS: COMPONENTS & SERVICES

OPENSTACK SERVICE	PROJECT NAME	DESCRIPTION
Dashboard	Horizon	Provides a web-based, self-service portal to interact with underlying OpenStack services, such as launching an instance, assigning IP addresses and configuring access controls.
Compute	Nova	Manages the lifecycle of compute instances in an OpenStack environment. Responsibilities include spawning, scheduling and decommissioning of virtual machines on demand.
Networking	Neutron	Enables network-connectivity-as-a-service for other OpenStack services, such as OpenStack Compute. Provides an API for users to define networks and the attachments into them. Has a pluggable architecture that supports many popular networking vendors and technologies.
Object storage	Swift	Stores and retrieves arbitrary unstructured data objects via a RESTful, HTTP based API. It is highly fault tolerant with its data replication and scale-out architecture. Its implementation is not like a file server with mountable directories. In this case, it writes objects and files to multiple drives, ensuring the data is replicated across a server cluster.
Block storage	Cinder	Provides persistent block storage to running instances. Its pluggable driver architecture facilitates the creation and management of block storage devices.
Identity service	Keystone	Stores and retrieves virtual machine disk images. OpenStack Compute makes use of this during instance provisioning.
Image service	Glance	Stores and retrieves virtual machine disk images. OpenStack Compute makes use of this during instance provisioning.
Telemetry	Ceilometer	Monitors and meters the OpenStack cloud for billing, benchmarking, scalability and statistical purposes.
Orchestration	Heat	Orchestrates multiple composite cloud applications by using either the native HOT template format or the AWS Cloud Formation template format, through both an OpenStack-native REST API and a Cloud Formation-compatible Query API.
Application Programming Interfaces (API)	OpenStack API	Provides application programing interfaces for block storage, compute, identity, image services, networking and other OpenStack components.

OpenStack Cloud Administration Server, it runs on the SUSE Linux Enterprise Server operating system. It is managed and deployed through the SUSE OpenStack Cloud Administration Server web interface.

TIP: For high availability setups, creating multiple dedicated Control Nodes is required with at least one pair of nodes needed for a cluster, yet an odd number is preferred to establish and maintain a quorum. Once a cluster is established, via the Pacemaker barclamp, it can be assigned to many of the core OpenStack services. This document will describe the setup of a three-node cluster of Control Nodes, but other combinations are supported as noted in the deployment guide.

- A particularly noteworthy service, the OpenStack Dashboard (Horizon) usually runs on the Control Node and gives users and administrators of the private cloud the ability to interact with, deploy and manage their workloads and storage volumes.

Compute Node(s)

- 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 and services. Both KVM and Xen are supported hypervisors within SUSE OpenStack Cloud and are included with a Compute Node deployed with SUSE Linux Enterprise Server. More Compute Nodes can be added over time as needed to address increased workload hosting, and availability zones can be utilized to achieve logical groupings of system and then to allow selective spreading of workloads across these zone for availability concerns.

  TIP: In addition, one can completely provision and utilize a Microsoft Hyper-V Compute Node from the SUSE OpenStack Cloud Administration Server web interface by following the preparatory steps outlined in the SUSE OpenStack Cloud Deployment Guide and the respective appendix. Another option is to integrate a VMware vSphere node into the compute role through interaction and delegation of vCenter, as also described in the deployment guide and its respective appendix.

- Users do not typically interact directly with these nodes since the Control Node provides the self-service dashboard to deploy and manage workloads and services on these Compute Nodes.

Storage Node(s)

- A SUSE OpenStack Cloud Storage Node acts as a provider of persistent cloud-based storage elements and offers OpenStack Object Storage (Swift) as well as Ceph-based options. This storage can be presented in both object or block format and furnish volumes for a workload instance or even as the underlying format for some of the core OpenStack services. Multiple Storage Nodes should be deployed to provide data redundancy and resiliency due to component or system failures and as detailed in this document. Three Storage Nodes, based upon Ceph, are deployed, but this can be scaled by adding more nodes as needed to provide more storage.

  TIP: It is also possible to integrate one cloud instance, leveraging resources from another storage instance, for example, by using an existing external Ceph storage cloud as outlined in the deployment guide.

- As with the Compute Nodes, users typically interact with the Storage Nodes through the use of volumes via the self-service dashboard on the Control Node.

Figure 2 represents a logical diagram of the user's interaction with these various node types.

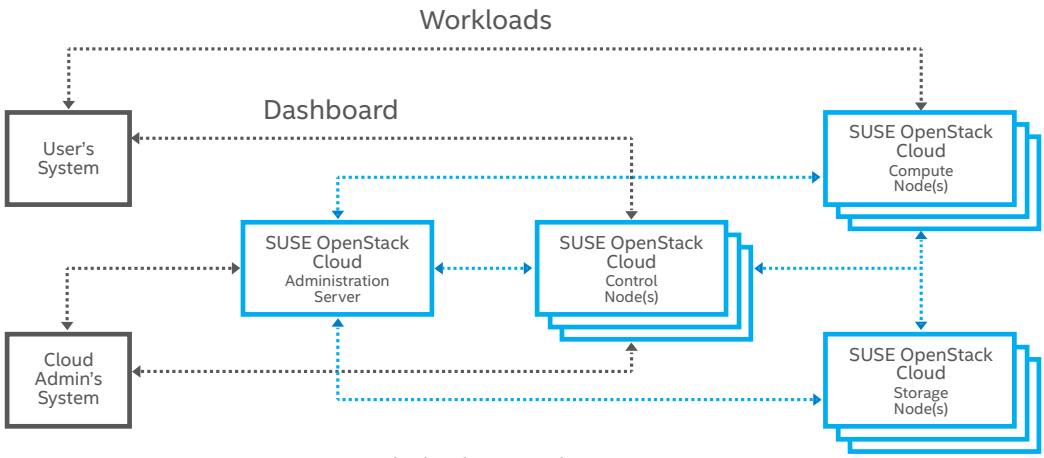

Figure 2: SUSE Openstack Cloud User Node Interaction Diagram

Target environments for private cloud are varied, ranging from small development and test instances to full production implementations that address both scale and availability. Some private cloud instances target only workload management while others are set up to provide storage clouds, yet mixed environments are quite common. Both Dell 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 configurations and sizing guidelines 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 all the features. In this scenario, workload lifecycle management is handled in a true cloud sense, in that already updated images are deployed

Small Starter Instance

This type of instance provides all of the features noted in this reference configuration but at the lowest levels of node count possible. As such, this is a starter configuration and 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 set to yield more scale across hypervisor and storage types and provides node count to address high availability concerns about key control plane services of the private cloud instance.

Instance Sizing Guidelines

As noted in the table, the remainder of this document focuses on a small starter instance that can be easily scaled with additional compute and storage resources, yet has a highly available control plane. All of the required software components are noted in the appendices.

INSTANCE SIZING GUIDLINES

	DEVELOPMENT	STARTER[a]	SCALABLE
Administration Server	1	1	1
Control Node(s)	1	3	2–9[c]
Compute Node(s)	2	3	8+
Storage Node(s)	2[b]	3[b]	4+

Notes:

[a] The focus of this reference implementation

[b] For Swift, minimum of two nodes, for Ceph, a minimum of three or more nodes

[c] At a minimum, a two-node cluster is required, but multiple two- and three-node clusters are recommended for scalability and performance of core services.

Solution

Based upon the previous section's considerations, an overall deployment diagram can now be formulated as shown in the accompanying figure.

For each of the technology layers, networking, computing platform and software, the following sections details the specific implementation details to complete the private cloud instance

Network

Configuration

From a physical level, it is recommended to be attentive and consistent with all network cabling, methodically wiring all the nodes in a very similar fashion to ease configuration and later troubleshooting. For the Dell PowerEdge servers and Dell Networking switches referenced in this document, the following drawing may be a useful guide: Further detailed network switch configuration settings may be found in the Appendix.

Notes:

- The physical configuration, as ordered from the factory, included a built-in LAN on Motherboard (LOM) with a 10 GbE em1, a second 10GbE em2 and a pair of 1GbE NIC ports (em3, em4, respectively), plus a pair of 10GbE ports on an add-on card.
- To take advantage of this performance-oriented configuration, em1 was the default PXE-interface, so was a natural choice for the Admin network, with the remaining 10GbE NIC ports becoming part of a LACP-bonded set where all the cloud networking was deployed. For consistency, default settings were used for MTUs on all NICs with higher values allowed on bounded switch ports.

TIP: In this reference architecture, the 1GbE NICs are basically unused, except on the Administration Server, but scenarios to include them are easy to rationalize, especially if the add-on NIC ports were not present. In that case the two 1GbE ports could

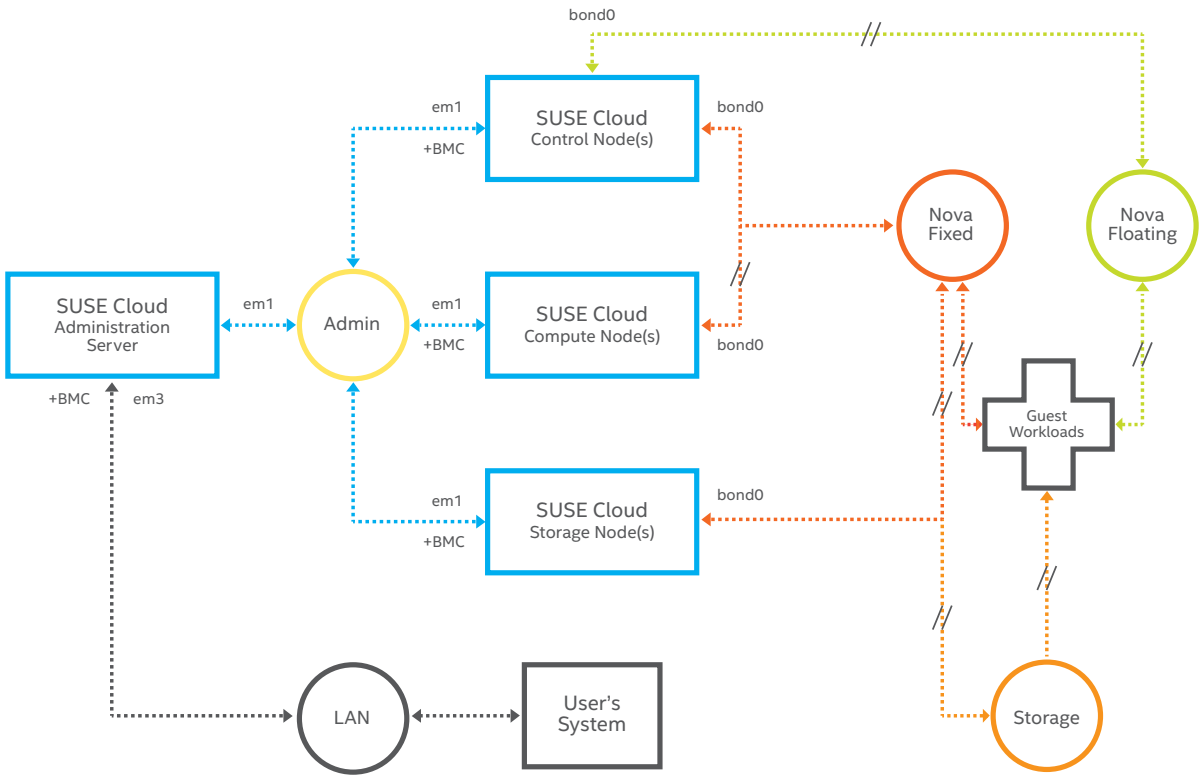

Notes / Legend:

- // boded (across 3 10Gbe NICs) for HA implementation
- Admin network must be untagged VLAN to allow PXE booting (includes BMC/iDRAC of all resource nodes)

be used singly or bonded together for the Admin Network, and the two 10GbE ports could be bonded together for the remaining cloud networks.

As previously described, multiple networks are required for the proper operation of a SUSE OpenStack Cloud instance. Further information can be found in the SUSE OpenStack Cloud Deployment Guide[13] along with other important considerations and limitations. From a software configuration standpoint, matching the physical cable and switching infrastructure is accomplished by considering the respective networking modes supported by the various Compute Node and hypervisor combinations and must be correctly established before proceeding with the private cloud instance setup because later changes cannot be accommodated without a re-installation.

Notes:

- Unlike the external and internal VLAN segments, the administrative VLAN does not use 802.1q VLAN tagging. Using flat or LAG-bonded interfaces and software controlled network configurations for the various VLAN segments is recommended.
- With this default network configuration cited in the deployment guide, up to 49 Compute Nodes may be accommodated and up to 61 public IP addresses, with about 200 workload instances. Adjust as needed to fit within the local network configuration and for the desired scale of the deployment.
- There is no IPv6 support for SUSE OpenStack Cloud implementations at this time.
- In addition, a network gateway providing access to the public network is required and in this configuration was implemented on the switches.

In order to facilitate setting up and administering the various Dell PowerEdge servers, the embedded BMC/iDRAC network interface for the system hosting the SUSE OpenStack Cloud Administration Server should be connected to a subnet typically outside the private cloud network infrastructure. This allows easy administration access while all other node being used in the private cloud instance should have their BMC/iDRAC network interfaces connected to and managed on the internal private cloud admin network.

Note: The SUSE OpenStack Cloud Administration Server will configure and manage the BMC/iDRAC for the Control, Compute and Storage Nodes during set up, and the operating system install will use all the network interfaces presented to it.

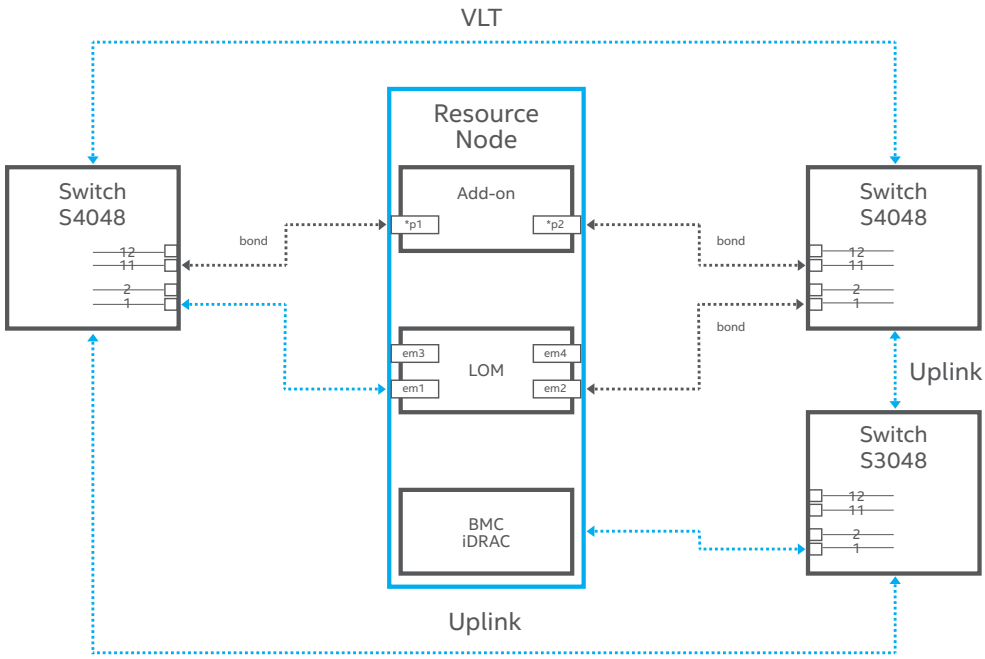

[13] www.suse.com/documentation/suse-cloud-5/

Computing Hardware

Configuration

Depending upon how the physical Dell PowerEdge servers were ordered or their current state, it is wise to start from a known, consistent configuration for both firmware and storage settings. If the servers were previously used, returning to factory firmware defaults is recommended. A Dell Service engineer may also be able to quickly configure settings in the environment to ensure a compatible, consistent state of settings. At a minimum, certain settings, beyond the defaults, are critical for their proper operation within the cloud installation, as shown below and in the respective sections for each node's role.

Software

Configuration

As preparation, download the SUSE Linux Enterprise Server and SUSE OpenStack Cloud products from the SUSE Downloads11 site with a free 60-day evaluation subscription or directly order them from Dell with your hardware order .The remainder of the process is described in a suggested deployment order.

For reference, this white paper relies on the following core design decisions and settings, summarized in the following table:

Software Deployment Decisions

Administration Server Setup and Configuration

Using the designated physical server to act as the Administration Server, you have a couple of deployment methods at your disposal. You can utilize either:

- A ready-made SUSE OpenStack Cloud Admin Appliance that has all the necessary components and soft-

SETTING		NODE ROLE			
NAME	DEFAULT VALUE	ADMINISTRATION SERVER	CONTROL NODE	COMPUTE NODE	STORAGE NODE
iDRAC					
iDRAC.IPMILan.Enable	Enabled				
iDRAC.IPMILan.PrivLimit	4				
iDRAC.IPv4.Enable	Enabled				
iDRAC.Users.2.Enable	Enabled				
iDRAC.Users.2.IpmiLanPrivilege	4				
iDRAC.Users.2.Privilege	0x1ff				
iDRAC.WebServer.Enable	Enabled				
BIOS					
BootMode	UEFI	BIOS	BIOS	BIOS	BIOS
BootSeqRetry	Enabled				
DcuIpPrefetcher	Enabled				
DcuStreamerPrefetcher	Enabled				
DynamicCoreAllocation	Disabled				
IntegratedRaid	Enabled				
InternalSdCard	Off				
IoatEngine	Enabled				
LogicalProc	Enabled				
MemOpMode	OptimizerMode				
MemTest	Disabled				
NodeInterleave	Disabled				
OsWatchdogTimer	Disabled				
ProcAdjCacheLine	Enabled				
ProcCores	All				
ProcExecuteDisable	Enabled				
ProcHwPrefetcher	Enabled				
ProcPwrPerf	MaxPerf				
ProcTurboMode	Enabled				
ProcVirtualization	Disabled	Enabled		Enabled	
QpiSpeed	MaxDataRate				
RtidSetting	Disabled				
SriovGlobalEnable	Enabled		Disabled		Disabled

Note: The plus (+) denotes those services with roles deployed on the HA (high availability) cluster of Control Nodes
www.suse.com/documentation/suse-cloud3/book_cloud_ deploy/data/book_cloud_deploy.html

For storage settings, use the following recommended configurations for each type of node:

SETTING	NODE ROLE			
VOLUME PURPOSE	ADMINISTRATION SERVER	CONTROL NODE	COMPUTE NODE	STORAGE NODE
Operating System	RAID 10 (6 * 500GB)	RAID 10 (6 * 500GB)	RAID 10 (6 * 500GB)	RAID 1 (2 * 300GB)
Data—Journal	N/A	N/A		RAID 0 (3 @ 200GB SSD)
Data—Storage	N/A	N/A		RAID 0 (13 @ 4TB)

ware pre-packaged in an easy-to-use preload format

- The documented process to build an Administration Server

While the recommended approach is to utilize this physical server solely for the Administration Server function, another viable alternative is to create this system as a virtual machine host or solution admin host. In this way, the SUSE OpenStack Cloud Administration Server can be a virtual machine guest, and other guests may also be created for other functions of your private cloud such as Calamari (Ceph master cluster monitor), a vCenter Compute Node proxy or other miscellaneous guests. If you deploy a solution admin host with virtual machines, be vigilant to not oversubscribe the memory available. For the Administration Server, there is actually a low footprint of system requirements, and as noted in the deployment guide the runtime dependencies on this node for the private cloud instance are simply the few services in use and the provisioning aspect of resource nodes and role assignment.

Suse Openstack Cloud Administration Server (setup via appliance)

The appliance image incorporates the process from SUSE OpenStack Cloud 5 Deployment Guide and is pre-populated with all of the necessary SUSE Linux Enterprise media and software repositories in order to rapidly deploy SUSE OpenStack Cloud, The SUSE OpenStack Cloud Administration Appliance should not be used if the target physical node has:

- An installation boot hard disk drive larger than 2TB
- UEFI Secure Boot is enabled/ required

The SUSE OpenStack Cloud Administration Appliance is available via SUSE Studio™[14] Download the Preload ISO or USB Image and boot the system from this appliance image, either directly or through virtual mounts of the media via the BMC/iDRAC.

- Follow the install instruction[15] prompts to localize your installation.
- Jump to the next section, at the "Crowbar Setup" step.

SUSE Openstack Cloud Administration Server (setup via deployment guide process)

Using either direct console access with an external media drive or the virtual media option of the BMC/iDRAC and the ISO image file, perform an x86_64 architecture installation of:

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

Notes:

- Either a default install or, minimally, the patterns cited in the SUSE OpenStack Cloud deployment guide are required.
- If the available capacity of the primary LUN is greater than 2TB, ensure that a multi-partition setup with a smaller "/boot" partition is used that is completely contained in the first 2TB to be compatible with the recommended legacy BIOS setting. The remaining free space then can be allocated for the "swap" and "/" (root) partitions plus any other disk volume partitions that are desired.

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 the IPv6 are disabled, since these operating modes are not

[14] https://susestudio.com/a/Mrr6vv/suse-openstack-cloud-5-admin

[15] https://github.com/cseader/suse-cloud-appliances/blob/suse-open- stack-cloud-5/docs/SUSE-Cloud-AA-Guide.pdf

SOFTWARE DEPLOYMENT DECISIONS

ATTRIBUTE	DESCRIPTION
OpenStack Release	SUSE OpenStack Cloud 5 based on OpenStack Juno
Host Operating System(s)	SUSE Linux Enterprise Server 11 SP3 (Administration Server, Control Nodes) SUSE Linux Enterprise Server 12 (Compute and Storage Nodes)
Hypervisor	KVM
Network Mode	Dual (single 10GbE NIC for Admin Network, multiple 10GbE NICs bonded for remaining)
Control Node (HA Cluster)	Pacemaker (cluster across 3 Control Nodes)
Database (+)	SUSE OpenStack Cloud utilizes PostgreSQL
Message Queue (+)	SUSE OpenStack Cloud utilizes RabbitMQ
OpenStack Identity Service (+)	Keystone
Ceph	Mon/OSD across 3 Storage Nodes with SSD journals
OpenStack Object Storage	Swift, not used
OpenStack Image Service (+)	Glance (configured with Rados backend)
OpenStack Block Storage (+)	Cinder (with Ceph/RADOS plug-in)
OpenStack Networking (+)	Neutron (Open vSwitch, gre)
OpenStack Compute (+)	Nova
OpenStack Dashboard (+)	Horizon
OpenStack Orchestration (+)	Heat
OpenStack Telemetry (+)	Ceilometer
OpenStack Database	Trove, not used

VOLUME PURPOSE	ADIMINSTRATION SERVER	CONTROL NODE	COMPUTE NODE	STORAGE NODE
Operating System	RAID 10 (6* 500GB)	RAID 10 (6* 500GB)	RAID 10 (6* 500GB)	RAID 10 (2* 500GB)
Data- Journal	N/A	N/A	N/A	RAID 0 (3* @ 200GB SSD)
Data-Storage	N/A	N/A	N/A	RAID 0 (13 @ 4 TB

currently supported, and the OpenStack networking model encompasses good network separation and isolation practices)

- SMT configuration (optional) At this point, carefully review the Crowbar Setup section, selecting the appropriate network mode for this high availability, performant configuration which matches the physical system and network cabling and switching infrastructure:
- The “dual” network mode selection permits the use of two distinct network interfaces with the first 10GbE NICE used for the Admin network and the other, via a bonded-mode of the remaining three 10GbE NICs, used for all remaining networks and their corresponding VLAN designations.

In addition:

- Manually ensure that the “conduit map” section pertaining to the “team” configuration section of /etc/crowbar/network.json file correctly references “10g2”, “10g3”, 10g4” across the respective “intf1” interface listings. See the Network Conduits section of the deployment guide for more details.
- Further, to match the switch settings, ensure that the “teaming mode” is set to “4” to match the LACP settings of the respective switch ports.
- Perform the setup step for all the necessary software repositories, as noted in the SUSE OpenStack Cloud Deployment Guide.

  Note: 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.

At this point, carefully review the network configuration to ensure your setup complies with the check points noted in the SUSE OpenStack Cloud Deployment Guide. 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 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 all nodes in the private cloud instance. As such, ensure a valid name server is configured on this specific system; to forward host resolution requests to that are beyond the private cloud environment. A similar configuration is also required for NTP as well.
- For the remaining nodes, if any of the primary LUN exceeds 2TB, you should configure a multi-partition setup with a smaller /boot partition. This can be accomplished for all subsequent node deployments by performing the following steps on the SUSE OpenStack Cloud Administration Server:
- Edit /opt/dell/chef/cookbooks/provisioner/templates/ default/autoyast.xml.erb (other settings 13 can also be adjusted to local preferences)
- In the @raid_type == “single” section after the CT_DISK line, add the following:

```
<partitions config:type="list">
  <partition>
    <filesystem
    config:type="symbol">ext3</
    filesystem>
    <mount>/boot</mount>
    <size>1gb</size>
  </partition>
  <partition>
    <mount>swap</mount>
    <size>2gb</size>
  </partition>
  <partition>
    <filesystem
    config:type="symbol">xfs</
    filesystem>
    <mount>/</mount>
    <size>auto</size>
  </partition>
</partitions>
```

- **knife cookbook upload -o /opt/dell/chef/cookbooks/ provisioner**
- **chef-client**

  The remaining required nodes of the SUSE OpenStack Cloud infrastructure can be set up, configured, and managed from the Administration Server web interface. The only prerequisite is to be able to PXE boot each of the systems against the Administration Server via the admin VLAN subnet. This includes the Control, Compute and Storage Nodes.

Other physical servers which become the resource nodes can be set up (assuming the default network setup, adjust as necessary), by executing the following commands:

Notes:

- Ensure that all the remaining target systems have the BMC/iDRAC configured and available on the admin subnet as this will get managed and integrated into the private cloud instance for basic power on/off and reboot control along with having a link from the Admin web interface for direct access.
- Ensure that all the remaining systems are able to PXE boot from the Administration Server via the admin subnet. The em1 10GbE interfaces must be specifically enabled to PXE boot from their respective device settings menu.

Each resource node will become visible to the Administration Server after PXE booting, at which point 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 Control, Compute, or Storage Node. Follow the SUSE OpenStack Cloud Deployment Guide, paying careful attention to the respective operating system and order listed.

Notes:

- Validate that the respective system’s network, hardware and storage devices match the requirements for the intended role.
- As each node is “Discovered,” it is advisable to group the nodes by functional type, enter a meaningful description.
- Create a public name to provide a handle to arrange and manage all your nodes. Alias names can also be used to make identifying the node easier, especially on the Crowbar web interface.

- Set the intended role for each node to aid later association of roles to a given node so that the subsequent barclamps auto-populate nodes appropriately.
- Where appropriate, availability zones, which signify differing power, network or other attributes, may also be entered for the Compute Nodes.

Control Node(S) Setup and Configuration

The Control Node hosts all OpenStack services needed to orchestrate virtual machines deployed on the Compute Nodes in SUSE OpenStack Cloud. It is managed and deployed through the SUSE OpenStack Cloud Administration Server web interface, as follows:

- PXE boot the designated systems to act as the Control Nodes, ensuring that they meet the necessary hardware requirements.
- Allocate the Control Node, selecting SUSE Linux Enterprise Server 11 SP3 as the target platform.
- In this high availability setup, creating multiple Control Nodes is required. All three Control Nodes will be used in the creation of a single cluster.

To utilize the high availability feature set, first apply the Pacemaker barclamp to the respective cluster nodes:

- Create a recognizable proposal name (e.g., ControlHA)
- For STONITH, select "Configured with one resource per node" with the respective parameters:
  - » Fencing agent: external/ipmi
  - » Parameters for each agent (respectively):
  - »ipaddr=<BMCIPAddress> userid=<BMCLogin> passwd=<BMCPassword> interface=<AdminNetworkNIC>
- Setup non-web GUI (hb gui) as "true"
- Include all three nodes into "pacemaker-cluster-member" and "hawk-server."

Then for the remaining OpenStack core services, assign this cluster to the role instead of an individual Control Node.

Compute Node(S) Setup and Configuration

- The Compute Node is a physical server running a hypervisor, serving as a host for workload guest virtual machines. Generally speaking, Compute Nodes also house the root and ephemeral disk images of the running workload guests, but this space is reclaimed when a guest is terminated.
- PXE boot the designated systems to act as Compute Nodes, ensuring that each meets the necessary hardware requirements.

  TIP: If desired, the backing store of the Compute Nodes can reside on shared storage, like an NFS mount or Ceph, to facilitate live migration of workload instances. For the former, this can be set up via barclamps from the web interface while for Ceph this requires changes to the OpenStack Compute services configuration

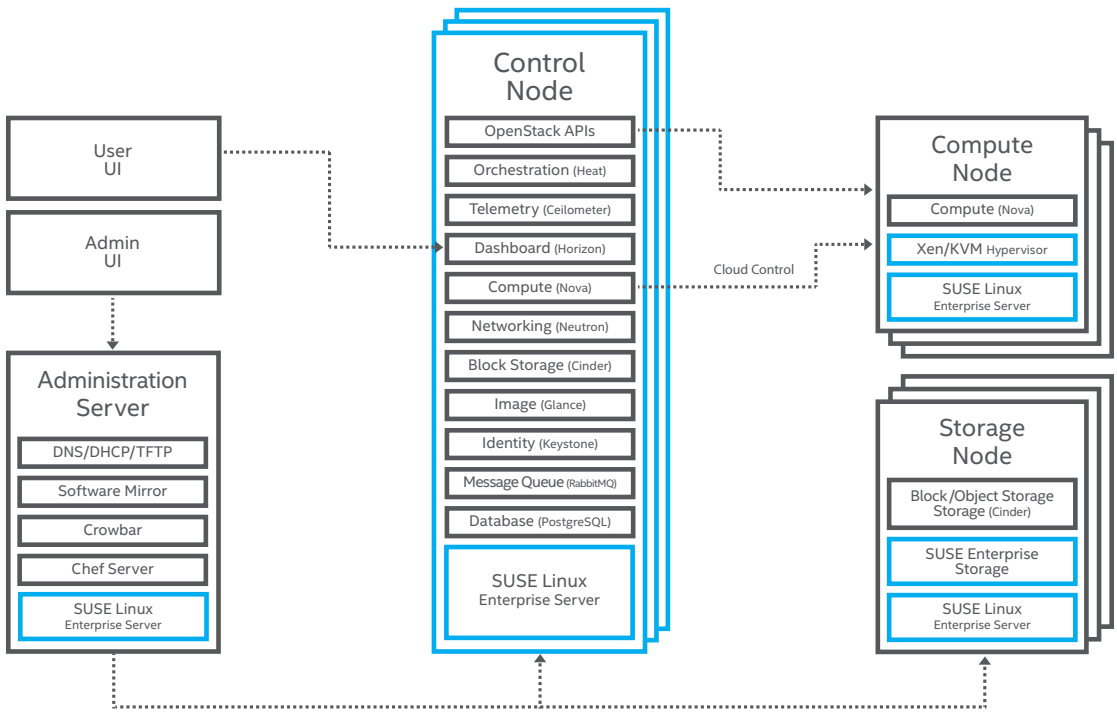

parameters ("images_rbd_pool" and "images_rbd_ceph_conf") accessible within the raw mode of the Nova barclamp. An easy approach is to leverage the "volumes" pool created during the Ceph setup and the stock Ceph configuration file location.

For KVM-based compute nodes, select SUSE Linux Enterprise Server 12 as the target platform.

Storage Node(S) Setup and Configuration

The Storage Node acts as a controller for persistent cloud-based storage elements. A SUSE OpenStack Cloud infrastructure should contain multiple Storage Nodes and provides object (via Swift and/or Ceph) and block storage elements (via Cinder plug-ins).

PXE boot the designated systems to act as the Storage Nodes, ensuring that each meets the necessary hardware requirements.

Notes:

- By default, SUSE OpenStack Cloud is configured to always use all unused disks on a node for Storage Nodes; it is advisable to configure the storage as noted in the storage setup tables.
- For this implementation, Ceph will be deployed on all Storage Nodes with a minimum of three nodes.

Openstack Services Setup and Configuration

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 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.

Notes:

- 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**: assign the ControlHA cluster to the designated controller role(s); then provide either an SBD or shared volume to house the data for this active/passive role.
- **Keystone, Neutron, 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.

  TIP: as your Ceph cluster is scaled further, it is only necessary to have the first three nodes (for small-large collections of storage nodes) or five nodes (only for very large installations) assigned to the "ceph-mon" role

  - To utilize any SSD devices as journals for the OSD drives for increased performance, modify the Ceph barclamp, in Raw mode, and insert the respective SSD device designation (e.g., /dev/sdb) into the "journal devices" section. If you have multiple SSD devices, these will automatically be split across the OSDs.
  - If a spare host (or virtual machine from the solution admin host) is available, assign that to the "ceph-calamari" role to view and monitor the Ceph functions
- **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 the ControlHA cluster to the designated controller role(s), and assign all of the compute nodes to the "nova-multi-compute-kvm" role.

  TIP: If desired, the backing store of the Compute Nodes can reside on shared storage, such as an NFS mount or Ceph, to facilitate live migration of workload instances. For the former, this can be set up via barclamps from the web interface. For Ceph this requires changes to the OpenStack Compute services configuration parameters accessible within the Raw mode of the Nova barclamp ("images_rbd_pool" and "images_rbd_ceph_conf"). An easy approach is to leverage the "volumes" pool created during the Ceph setup.

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 or via the "Utilities -> Exported" action on the Crowbar web interface before contacting your support organization

At this point, you should be able to utilize your private cloud instance, according to the SUSE OpenStack Cloud Admin, User and Supplement Guides. 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 as follows:

- Install the Tempest barclamp on your Administration Server, via **zypper install crowbar-barclamp-tempest**.

- From the SUSE OpenStack Cloud Crowbar web interface, assign the "tempest" role, typically to a control node or another dedicated resource node.
- Logged into that node, execute the following command: **/var/lib/openstack-tempest-test/run_tempest.sh –N –t –s –L /etc/tempest/logging.conf**

Summary

After completing the steps described in this white paper along with the steps in the respective network, hardware and software products documentation, you should have a fully functional private cloud installation.

Support

Both Dell and SUSE can provide guidance, training, support and services for more sophisticated deployments than the one described in this white paper; however, they are beyond the scope of this document.

Appendices

Rack and Power

Rack and power distribution units are not listed, as they are usually site-specific. The physical dimensions and power requirements need to be reviewed. The Dell PowerEdge R430, R630, R730 and R730xd all require rear cable management and power distribution.

Networking Switches

Bill of Materials—Dell Networking S4048-ON Switch

Quantities shown are for a single switch, but two devices are required as documented in this reference architecture.

DELL NETWORKING S4048-ON

QUANTITY	PART NUMBER	PRODUCT DESCRIPTION
1	210-ADUZ	Dell Networking S4048-ON, 48x 10GbE SFP+ and 6x 40GbE QSFP
1	634-BCWZ	Dell Networking OS9, S4048-ON
1	450-AASX	Dell Networking, Jumper Cord, 250V, 12A, 2 Meters, C13/C14
1	634-BCWX	Dell Networking S4048-ON User Guide
1	997-6304	Dell H Limited Warranty Extended Years
1	997-6179	ProSupport: Next Business Day Onsite Service
1	997-6186	ProSupport: 7x24 HW/SW Tech Support
1	997-6174	ProSupport: Next Business Day Onsite Service
1	997-6305	Dell Hardware Limited Warranty Intial Year
1	997-6306	Info 3rd Party Software Warranty
1	989-3439	Dell ProSupport
1	332-1286	Us Order On-Site Installation Declined
1	900-9997	Onsite Installation Declined

DELL S4048 NETWORK SWITCH

SETTINGS NAME(S)	SWITCH #1 PORT RANGE	SWITCH #2 PORT RANGE
no ip address portmode hybrid switchport spanning-tree rstp edge-port no shutdown interface Vlan 100 ip address <Unused Admin Network IP> untagged TenGigabitEthernet tagged Port-channel 3 ip helper-address <Admin Server IP> no shutdown	1-10	
no ip address portmode hybrid switchport spanning-tree rstp edge-port no shutdown		1
no ip address portmode hybrid switchport spanning-tree rstp edge-port no shutdown	11, 21-22, 31	11, 21-22, 31

DELL S4048 NETWORK SWITCH

SETTINGS NAME(S)	SWITCH #1 PORT RANGE	SWITCH #2 PORT RANGE
1no ip address		
port-channel-protocol LACP		
port-channel <PortNumber> mode active		
no shutdown		
Interface Port-channel <Respective 12-20>		
no ip address		
Mtu 12000		
portmode hybrid		
switchport		
spanning-tree rstp edge-port		
vlt-peer-lag port-channel <Respective 12-20>		
no shutdown		
Interface Vlan 200		
ip address <Unused Storage Net IP>		
tagged Port-channel 12-20		
no shutdown		
interface Vlan 300		
ip address <Router for Public/Floating Net IP>		
tagged Port-channel <Respective 12-20>		
no shutdown		
interface Vlan 400		
no ip address		
tagged Port-channel 12-20		
no shutdown		
interface Vlan 500		
no ip address		
tagged Port-channel 12-20		
no shutdown		
no ip address		
Mtu 12000		
portmode hybrid		
switchport		
spanning-tree rstp edge-port		
vlt-peer-lag port-channel <Corresponding 12-20>		
no shutdown	12-20	2-10, 12-20

DELL S4048 NETWORK SWITCH

SETTINGS NAME(S)	SWITCH #1 PORT RANGE	SWITCH #2 PORT RANGE
no ip address shutdown	23-30, 32-46, 49-54	23-30, 32-46, 49-54
no ip address port-channel-protocol LACP port-channel 3 mode active no shutdown interface Port-channel 3 description Uplink_R5-S3048 no ip address mtu 12000 portmode hybrid switchport vlt-peer-lag port-channel 3 no shutdown	47	47
no ip address port-channel-protocol LA port-channel 5 mode active no shutdown interface Port-channel 5 description Uplink_Core no ip address mtu 12000 portmode hybrid switchport vlt-peer-lag port-channel 5 no shutdown	48	48
no ip address portmode hybrid switchport no shutdown	50	50
no ip address port-channel-protocol LACP port-channel 3 mode activemode no shutdown Interface Port-channel 3 description Uplink_R5-S4048 no ip address mtu 120 portmode hybrid switchport no shutdown	51-52	51-52

DELL S4048 NETWORK SWITCH		
SETTINGS NAME(S)	**SWITCH #1 PORT RANGE**	**SWITCH #2 PORT RANGE**
interface port-channel 1		
description VLTi		
no ip address		
channel-member fortyGigE		
switchport		
vlt-peer-lag port-channel 3		
no shutdown		
interface port-channel 3		
description uplink S3048		
no ip address		
mtu 12000		
portmode hybrid	53-54	53-54
no shutdown		

Bill of Materials—Dell Networking S3048-ON Switch

For reference, the per-port configuration of this switch follows:

DELL NETWORKING S3048-ON		
QUANTITY	**PART NUMBER**	**PRODUCT DESCRIPTION**
1	210-AEDP	Dell Networking S3048-ON, 48x 1GbE, 4x SFP+ 10GbE ports
1	634-BCXS	Dell Networking OS9, S3048-ON
1	450-AASX	Dell Networking, Jumper Cord, 250V, 12A, 2 Meters, C13/C14
1	634-BCXR	Dell Networking S3048-ON User Guide
1	997-6306	Info 3rd party SW Warranty Provided by Vendor
1	802-7475	Basic HW Svcs, Business Hours (5x10) NBD on Site HW Warranty Repair
1	802-7389	Dell HW limited Warranty
1	996-8029	Declined recommended ProSupport Service
1	802-7391	SW Warranty, NW SW, 90 Day and 60 days media
1	900-9997	Onsite Installation Declined
1	332-1286	US Order

DELL S4048 NETWORK SWITCH

SETTINGS NAME(S)	PORT RANGE
no ip address switchport spanning-tree rstp edge-port no shutdown interface Vlan 100 ip address <Unused Admin Network IP> untagged GigabitEthernet tagged Port-channel 3 untagged TenGigabitEthernet !untagged Port-channel 3	1-10
shutdown interface Vlan 1 !untagged GigabitEthernet !untagged TenGigabitEthernet !untagged Port-channel 3	11-49
no ip address portmode hybrid switchport no shutdown	50
no ip address port-channel-protocol LACP Port-channel 3 mode activemode no shutdown Interface Port-channel 3 description Uplink_R5-S4048 no ip address mtu 120 portmode hybrid switchport no shutdown	51-52

Bill of Materials—Dell Networking Cables

Quantities shown are per system; order a set of these for each system.

DELL NETWORKING CABLES (PER DELL POWEREDGE SYSTEM)

QUANTITY	PART NUMBER	PRODUCT DESCRIPTION
4	470-AAGN	Dell Networking S3048-ON, 48x 1GbE, 4x SFP+ 10GbE ports
2	A7459287	Dell Networking OS9, S3048-ON
1	A0375239	Dell Networking, Jumper Cord, 250V, 12A, 2 Meters, C13/C14

Computing Platform

The computing hardware listed should be used as initial guidance only as it represents what was tested. Additional configurations are possible and will likely be required as each customer's environment and use case are unique. Common parameters that could differ include:

Processors—Higher frequencies and core counts may improve performance while lower voltage/TDP processors can improve power efficiency.

Local Storage—Disk capacity, drive technology and spindle speed can be matched to budget and performance requirements as necessary.

Memory—Depending on the usage of various services, more or less memory may be necessary on various nodes according to their role.

Bill of Materials—Solution Admin Hosts (Administration Server)

Quantities shown are for the single Administration Server required. Select one of the following models, with Dell PowerEdge R630 as the preferred choice or one of the alternatives.

DELL POWEREDGE R630 SERVER (PREFERRED)

QUANTITY	PART NUMBER	PRODUCT DESCRIPTION
1	210-ACXS	PowerEdge R630 Server
1	329-BCIY	PowerEdge R630 Motherboard
1	321-BBKM	Chassis with up to 10, 2.5" Hard Drives, 3 PCIe Slots
1	340-AKPR	PowerEdge R630 Shipping—10/24 Drive Chassis
1	338-BFFU	Intel Xeon E5-2630 v3 2.4GHz,20M Cache,8.00GT/s_x000D_QPI,Turbo,HT,8C/16T (85W) Max Mem 1866MHz
	374-BBHD	Upgrade to Two Intel Xeon E5-2630 v3 2.4GHz,20M_x000D_Cache,8.00GT/s QPI,Turbo,HT,8C/ 16T (85W)
1	412-AAEE	120W Heatsink for PowerEdge R630
2	370-ABWE	DIMM Blanks for System with 2 Processors
1	370-ABUF	2133MT/s RDIMMs
1	370-AAIP	Performance Optimized
1	370-ABUI	4GB RDIMM, 2133MT/s, Single Rank, x8 Data Width
8	780-BBJN	RAID 10 for H330/H730/H730P (4-24 HDDs or SSDs in pairs)
1	405-AAEG	PERC H730 Integrated RAID Controller, 1GB Cache
1	400-AEEN	500GB 7.2K RPM NLSAS 6Gbps 2.5in Hot-plug Hard Drive,13G
6	540-BBBB	Intel X520 DP 10Gb DA/SFP+, + I350 DP 1Gb Ethernet, Network_x000D_Daughter Card
1	540-BBHY	Intel X520 DP 10Gb DA/SFP+ Server Adapter, Low Profile
1	634-BBWU	OpenManage Essentials, Server Configuration Management
1	385-BBHO	iDRAC8 Enterprise, integrated Dell Remote Access Controller,_x000D_Enterprise
1	429-AAIQ	No Internal Optical Drive
1	325-BBIL	Quck Sync Bezel 10/24 Drive Chassis
1	770-BBBC	ReadyRails Slideing Rails without Cable Management Arm
1	384-BBBL	Performanc BIOS Settings
1	450-ADWQ	Dual, Hot-plug, Redundant Power Supply (1+1), 495W
1	492-BBDH	C13 to C14, PDU Style, 12 AMP, 2 Ft Power Cord, North America
2	343-BBDK	Electronic System Documenation and OpenManage DVD Kit, Power Edge R630
1	619-ABVR	No Operating System
1	421-5736	No Media Required
1	989-3439	Dell ProSupport, For tech support
1	976-7648	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, 3 year
1	976-7728	Dell HW Limited Warranty Plus On Site Service
1	978-7657	ProSupport: 7x24 HW/SW Tech Support and Assistance, 3 Year
1	900-9997	On-Ste Installation Declined
1	332-1286	US Order

DELL POWEREDGE R430 SERVER (ALTERNATE CHOICE)

QUANTITY	PART NUMBER	PRODUCT DESCRIPTION
1	210-ADLO	PowerEdge R430 Server
1	329-BCBR	PowerEdge R430/R530 Motherboard
1	340-AMJF	PowerEdge R430 Shipping
1	338-BFFU	E5-2630 v3 2.4GHz,20M C,85W
1	374-BBHD	E5-2630 v3 2.4GHz,20M C,85W
1	370-ABXP	DIMM Blanks for System w/ 2 Processors
1	370-ABXV	Cooling Fan
1	374-BBIJ	135W Heatsink
1	374-BBIJ	135W Heatsink
1	330-BBEF	Riser, 2LP, R430
1	370-ABUF	2133MT/s RDIMMs
1	370-AAIP	Performance Optimized
1	405-AAEG	PH730 Intg RD CTL,1GB Cache
1	384-BBBL	Performance BIOS Settings
1	450-AEGZ	Dual Hot Plug Pwr Sply 550W
1	343-BBDT	EDOCS for R430
1	619-ABVR	No Operating System
1	421-5736	No Media Required
1	632-BBDC	SanDisk DAS Cache, 90 Day Trial License
1	951-2015	INFO, PSP TECH SPT CONTACT, ENTERPRISE
1	997-2924	HW WRTY + SVC,PE R430,UNY
1	997-2983	PSP NBD OS,PE R430,UNY,3YR
1	997-2992	PSP TECH SPT,PE R430,3YR
1	909-0259	Dell Proactive Systems Mgmt,Declined,Ent
1	900-9997	ONSITE INSTL DECLINED
1	973-2426	INFO Declined Remote Consulting Service
1	332-1286	US Order
1	321-BBNK	2.5” Chas up to 8HDs,HP
1	429-AAQM	DVD ROM, SATA, Internal
8	370-ABUG	16GB RDIMM,2133MT/s,DR,x4
1	780-BBPK	No RAID,H330/H730/H730P
4	400-AEFO	HDD,1.2TB 10K SAS,6G,2.5,HP,13G
1	330-BBDX	iDRAC Port Card
1	385-BBHO	iDRAC8, Enterprise
2	540-BBHY	X520 DP 10Gb DA/SFP+,SA LP
1	350-BBBW	No Bezel
1	770-BBBC	Slide RdyRL,No CMA
2	492-BBDH	C13-C14,PDU,12A,2 ft,0.6m,NA

DELL POWEREDGE R730 SERVER (ALTERNATE CHOICE)

QUANTITY	PART NUMBER	PRODUCT DESCRIPTION
1	210-ACXU	PowerEdge R730 Server
1	591-BBCH	PowerEdge R730/xd Motherboard
1	340-AKKB	PowerEdge R730 Shipping
1	330-BBCO	R730/xd PCIe Riser 2, Center
1	330-BBCQ	R730 PCIe Riser 3, Left
1	330-BBCR	R730/xd PCIe Riser 1, Right
1	370-ABUF	2133MT/s RDIMMs
1	370-AAIP	Performance Optimized
1	384-BBBL	Performance BIOS Settings
1	631-AAJG	Edocs and OpenManage DVD, R730/xd
1	619-ABVR	No Operating System
1	421-5736	No Media Required
1	370-ABWE	DIMM Blanks for System with 2 Processors
1	374-BBHM	Standard Heatsink for PE R730/R730xd
1	374-BBHM	Standard Heatsink for PE R730/R730xd
1	332-1286	US Order
1	951-2015	INFO,PSP TECH SPT CONTACT,ENTERPRISE
1	976-8706	HW WRTY + SVC,PE R730,UNY
1	976-8728	PSP NBD OS,PE R730,UNY,3YR
1	976-8729	PSP TECH SPT,PE R730,3YR
1	900-9997	ONSITE INSTL DECLINED
1	338-BFFF	E5-2650 v3 2.3GHz,25M C,105W
1	374-BBGM	E5-2650 v3 2.3GHz,25M C,105W
8	370-ABUG	16GB RDIMM,2133MT/s,DR,x4
1	540-BBBB	X520 DP,10G,DA + I350 DP,1G,DC
1	540-BBCT	X520 DP 10Gb DA/SFP+ Svr Adpt
1	429-AAPU	DVD ROM, SATA
1	350-BBBW	No Bezel
1	770-BBBQ	Slide RdyRL,No CMA
1	450-ADWM	Dual,Redundant,Hot-plug PS,1100W
2	492-BBDH	C13-C14,PDU,12A,2 ft,0.6m,NA
1	385-BBHO	iDRAC8, Enterprise
1	350-BBEO	Chassis with up to 8, 3.5 HDs
1	780-BBJX	R10,H330/H730/H730P
1	405-AAEG	PH730 Intg RD CTL,1GB Cache
4	400-AJOV	HDD,1.2TB 10K SAS,12G,2.5,HYB
1	800-BBDM	UEFI BIOS with GPT Partition
1	387-BBIB	Energy Star, PowerEdge R730

Bill of Materials—Control Node

Quantities shown are for a single Control Node, but three are required as documented in this reference architecture. Select one of the following models, with Dell PowerEdge R630 as the preferred choice or one of the alternatives.

DELL POWEREDGE R630 SERVER (PREFERRED)

QUANTITY	PART NUMBER	PRODUCT DESCRIPTION
1	210-ACXS	PowerEdge R630 Server
1	329-BCIY	PowerEdge R630 Motherboard
1	321-BBKM	Chassis with up to 10, 2.5" Hard Drives, 3 PCIe Slots
1	340-AKPR	PowerEdge R630 Shipping—10/24 Drive Chassis
1	338-BFFF	Intel Xeon E5-2650 v3 2.3GHz,25M Cache,9.6GT/s_x000D_QPI,Turbo,HT,8C/16T (85W) Max Mem 1866MHz
1	374-BBGM	Upgrade to Two Intel Xeon E5-2650 v3 2.3GHz,25M_x000D_Cache,8.00GT/s QPI,Turbo,HT,8C/16T (85W)
2	412-AAEE	120W Heatsink for PowerEdge R630
1	370-ABWE	DIMM Blanks for System with 2 Processors
1	370-ABUF	2133MT/s RDIMMs
1	370-AAIP	Performance Optimized
8	370-ABUG	16GB RDIMM, 2133MT/s, Single Rank, x4 Data Width
1	780-BBJN	RAID 10 for H330/H730/H730P (4-24 HDDs or SSDs in pairs)
1	405-AAEG	PERC H730 Integrated RAID Controller, 1GB Cache
6	400-AEER	600GB 10K RPM SAS 6Gbps 2.5in Hot-plug Hard Drive,13G
1	540-BBBB	Intel X520 DP 10Gb DA/SFP+, + I350 DP 1Gb Ethernet, Network_x000D_Daughter Card
1	540-BBHY	Intel X520 DP 10Gb DA/SFP+ Server Adapter, Low Profile
1	634-BBWU	OpenManage Essentials, Server Configuration Management
1	385-BBHO	iDRAC8 Enterprise, integrated Dell Remote Access Controller,_x000D_Enterprise
1	429-AAIQ	No Internal Optical Drive
1	325-BBIL	Quck Sync Bezel 10/24 Drive Chassis
1	770-BBBC	ReadyRails Slideing Rails without Cable Management Arm
1	384-BBBL	Performanc BIOS Settings
1	450-ADWQ	Dual, Hot-plug, Redundant Power Supply (1+1), 495W
2	492-BBDH	C13 to C14, PDU Style, 12 AMP, 2 Ft Power Cord, North America
1	343-BBDK	Electronic System Documenation and OpenManage DVD Kit, Power Edge R630
1	619-ABVR	No Operating System
1	421-5736	No Media Required
1	989-3439	Dell ProSupport, For tech support
1	976-7648	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, 3 year
1	976-7728	Dell HW Limited Warranty Plus On Site Service
1	978-7657	ProSupport: 7x24 HW/SW Tech Support and Assistance, 3 Year
1	900-9997	On-Ste Installation Declined
1	332-1286	US Order

Bill of Materials—Compute Node

Quantities shown are for a single Compute Node, but three are required as documented in this reference architecture. Select one of the following models, with Dell PowerEdge R630 as the preferred choice or one of the alternatives.

DELL POWEREDGE R630 SERVER (PREFERRED)

QUANTITY	PART NUMBER	PRODUCT DESCRIPTION
1	210-ACXS	PowerEdge R630 Server
1	329-BCIY	PowerEdge R630 Motherboard
1	321-BBKM	Chassis with up to 10, 2.5" Hard Drives, 3 PCIe Slots
1	340-AKPR	PowerEdge R630 Shipping—10/24 Drive Chassis
1	338-BFFF	Intel Xeon E5-2650 v3 2.3GHz,25M Cache,9.6GT/s_x000D_QPI,Turbo,HT,8C/16T (85W) Max Mem 1866MHz
1	374-BBGM	Upgrade to Two Intel Xeon E5-2650 v3 2.3GHz,25M_x000D_Cache,8.00GT/s QPI,Turbo,HT,8C/ 16T (85W)
2	412-AAEE	120W Heatsink for PowerEdge R630
1	370-ABWE	DIMM Blanks for System with 2 Processors
1	370-ABUF	2133MT/s RDIMMs
1	370-AAIP	Performance Optimized
8	370-ABUG	16GB RDIMM, 2133MT/s, Single Rank, x4 Data Width
1	780-BBJN	RAID 10 for H330/H730/H730P (4-24 HDDs or SSDs in pairs)
1	405-AAEG	PERC H730 Integrated RAID Controller, 1GB Cache
6	400-AEER	600GB 10K RPM SAS 6Gbps 2.5in Hot-plug Hard Drive,13G
1	540-BBBB	Intel X520 DP 10Gb DA/SFP+, + I350 DP 1Gb Ethernet, Network_x000D_Daughter Card
1	540-BBHY	Intel X520 DP 10Gb DA/SFP+ Server Adapter, Low Profile
1	634-BBWU	OpenManage Essentials, Server Configuration Management
1	385-BBHO	iDRAC8 Enterprise, integrated Dell Remote Access Controller,_x000D_Enterprise
1	429-AAIQ	No Internal Optical Drive
1	325-BBIL	Quck Sync Bezel 10/24 Drive Chassis
1	770-BBBC	ReadyRails Slideing Rails without Cable Management Arm
1	384-BBBL	Performanc BIOS Settings
1	450-ADWQ	Dual, Hot-plug, Redundant Power Supply (1+1), 495W
2	492-BBDH	C13 to C14, PDU Style, 12 AMP, 2 Ft Power Cord, North America
1	343-BBDK	Electronic System Documenation and OpenManage DVD Kit, Power Edge R630
1	619-ABVR	No Operating System
1	421-5736	No Media Required
1	989-3439	Dell ProSupport, For tech support
1	976-7648	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, 3 year
1	976-7728	Dell HW Limited Warranty Plus On Site Service
1	978-7657	ProSupport: 7x24 HW/SW Tech Support and Assistance, 3 Year
1	900-9997	On-Ste Installation Declined
1	332-1286	US Order

DELL POWEREDGE R430 SERVER (ALTERNATE CHOICE)

QUANTITY	PART NUMBER	PRODUCT DESCRIPTION
1	210-ADLO	PowerEdge R430 Server
1	329-BCBR	PowerEdge R430/R530 Motherboard
1	340-AMJF	PowerEdge R430 Shipping
1	338-BFFU	E5-2630 v3 2.4GHz,20M C,85W
1	374-BBHD	E5-2630 v3 2.4GHz,20M C,85W
1	370-ABXP	DIMM Blanks for System w/ 2 Processors
1	370-ABXV	Cooling Fan
1	374-BBIJ	135W Heatsink
1	374-BBIJ	135W Heatsink
1	330-BBEF	Riser, 2LP, R430
1	370-ABUF	2133MT/s RDIMMs
1	370-AAIP	Performance Optimized
1	405-AAEG	PH730 Intg RD CTL,1GB Cache
1	384-BBBL	Performance BIOS Settings
1	450-AEGZ	Dual Hot Plug Pwr Sply 550W
1	343-BBDT	EDOCS for R430
1	619-ABVR	No Operating System
1	421-5736	No Media Required
1	632-BBDC	SanDisk DAS Cache, 90 Day Trial License
1	951-2015	INFO,PSP TECH SPT CONTACT,ENTERPRISE
1	997-2924	HW WRTY + SVC,PE R430,UNY
1	997-2983	PSP NBD OS,PE R430,UNY,3YR
1	997-2992	PSP TECH SPT,PE R430,3YR
1	909-0259	Dell Proactive Systems Mgmt,Declined,Ent
1	900-9997	ONSITE INSTL DECLINED
1	973-2426	INFO Declined Remote Consulting Service
1	332-1286	US Order
1	321-BBNK	2.5” Chas up to 8HDs,HP
8	370-ABUG	16GB RDIMM,2133MT/s,DR,x4
1	780-BBPK	No RAID,H330/H730/H730P
4	400-AEFO	HDD,1.2TB 10K SAS,6G,2.5,HP,13G
1	330-BBDX	iDRAC Port Card
1	385-BBHO	iDRAC8, Enterprise
2	540-BBHY	X520 DP 10Gb DA/SFP+,SA LP
1	350-BBBW	No Bezel
1	770-BBBC	Slide RdyRL,No CMA
2	492-BBDH	C13-C14,PDU,12A,2 ft,0.6m,NA

DELL POWEREDGE R730 SERVER (ALTERNATE CHOICE)

QUANTITY	PART NUMBER	PRODUCT DESCRIPTION
1	210-ACXU	PowerEdge R730 Server
1	591-BBCH	PowerEdge R730/xd Motherboard
1	340-AKKB	PowerEdge R730 Shipping
1	330-BBCO	R730/xd PCIe Riser 2, Center
1	330-BBCQ	R730 PCIe Riser 3, Left
1	330-BBCR	R730/xd PCIe Riser 1, Right
1	370-ABUF	2133MT/s RDIMMs
1	370-AAIP	Performance Optimized
1	384-BBBL	Performance BIOS Settings
1	631-AAJG	Edocs and OpenManage DVD, R730/xd
1	619-ABVR	No Operating System
1	421-5736	No Media Required
1	370-ABWE	DIMM Blanks for System with 2 Processors
1	374-BBHM	Standard Heatsink for PE R730/R730xd
1	374-BBHM	Standard Heatsink for PE R730/R730xd
1	332-1286	US Order
1	951-2015	INFO,PSP TECH SPT CONTACT,ENTERPRISE
1	976-8706	HW WRTY + SVC,PE R730,UNY
1	976-8728	PSP NBD OS,PE R730,UNY,3YR
1	976-8729	PSP TECH SPT,PE R730,3YR
1	900-9997	ONSITE INSTL DECLINED
1	338-BFFF	E5-2650 v3 2.3GHz,25M C,105W
1	374-BBGM	E5-2650 v3 2.3GHz,25M C,105W
8	370-ABUG	16GB RDIMM,2133MT/s,DR,x4
1	540-BBBB	X520 DP,10G,DA + I350 DP,1G,DC
1	540-BBCT	X520 DP 10Gb DA/SFP+ Svr Adpt
1	429-AAPU	DVD ROM, SATA
1	350-BBBW	No Bezel
1	770-BBBQ	Slide RdyRL,No CMA
1	450-ADWM	Dual,Redundant,Hot-plug PS,1100W
2	492-BBDH	C13-C14,PDU,12A,2 ft,0.6m,NA
1	385-BBHO	iDRAC8, Enterprise
1	350-BBEO	Chassis with up to 8, 3.5 HDs
1	780-BBJX	R10,H330/H730/H730P
1	405-AAEG	PH730 Intg RD CTL,1GB Cache
4	400-AJOV	HDD,1.2TB 10K SAS,12G,2.5,HYB
1	800-BBDM	UEFI BIOS with GPT Partition
1	387-BBIB	Energy Star, PowerEdge R730

Bill of Materials—Storage Node

Quantities shown are for a single Storage Node, but three are required as documented in this reference architecture.

DELL POWEREDGE R730XD SERVER

QUANTITY	PART NUMBER	PRODUCT DESCRIPTION
1	210-ADBC	PowerEdge R730xd Server
1	591-BBCH	PowerEdge R730/R730xd Motherboard
1	338-BFFF	Intel Xeon E5-2650 v3 2.3GHz,25M Cache,9.60GT/s_x000D_QPI,Turbo,HT,10C/20T (105W) Max Mem 2133MHz
1	350-BBEX	Chassis with up to 12 + 4 Internal, 3.5" Hard Drives and 2, 2.5"_x000D_Flex Bay Hard Drives
1	340-AKPM	PowerEdge R730xd Shipping
1	374-BBGM	Upgrade to Two Intel Xeon E5-2650 v3 2.3GHz,25M_x000D_Cache,9.60GT/s QPI,Turbo,HT,10C/20T (105W)
1	370-ABUF	2133MT/s RDIMMs
1	370-AAIP	Performance Optimized
16	370-ABUI	4GB RDIMM, 2133MT/s, Single Rank, x8 Data Width
1	619-ABVR	No Operating System
1	421-5736	No Media Required
1	780-BBLR	RAID 1+Unconfigured RAID forH330/H730/H730P (2 + 1-20 HDDs
1	405-AAEG	"PERC H730 Integrated RAID Controller, 1GB Cache"
3	400-AEIL	200GB Solid State Drive SAS Write Intensive 12Gbps 2.5in Hot-plug
9	400-AEGH	4TB 7.2K RPM NLSAS 6Gbps 3.5in Hot-plug Hard Drive,13G
2	400-AEOC	300GB 10K RPM SAS 6Gbps 2.5in Flex Bay Hard Drive,13G
4	400-AENM	4TB 7.2K RPM NLSAS 6Gbps 3.5in Internal Bay Hard Drive,13G
1	540-BBBB	Intel X520 DP 10Gb DA/SFP+, + I350 DP 1Gb Ethernet, Network
1	330-BBCO	R730/xd PCIe Riser 2, Center
1	374-BBHT	R730xd PCIe Riser 1 Filler Blank, Right
1	540-BBCT	Intel X520 DP 10Gb DA/SFP+ Server Adapter
1	450-ADWS	Dual, Hot-plug, Redundant Power Supply (1+1), 750W
2	492-BBDH	C13 to C14, PDU Style, 12 AMP, 2 Feet (.6m) Power Cord
1	384-BBBL	Performance BIOS Settings
1	770-BBBQ	ReadyRails Sliding Rails Without Cable Management Arm
1	350-BBER	Quick Sync Bezel
1	631-AACK	No Systems Documentation, NoOpenManage DVD Kit"
1	800-BBDM	UEFI BIOS
2	374-BBHR	Heatsink for 12 + 4 Chassis PowerEdge R730xd
1	370-ABWE	DIMM Blanks for System with 2 Processors
1	634-BBWU	OpenManage Essentials, Server Configuration Management
1	385-BBHO	iDRAC8 Enterprise, integrated Dell Remote Access Controller,_x000D_Enterprise
1	332-1286	US Order
1	976-9009	ProSupport: Next Business Day Onsite Service After_x000D_ProblemDiagnosis, 3 Year
1	976-9007	Dell Hardware Limited Warranty Plus On Site Service
1	976-9008	ProSupport: 7x24 HW / SW Tech Support and Assistance, 3 Year
1	989-3439	Dell ProSupport. For tech support,
1	900-9997	On-Site Installation Declines

Software

Bill of Materials—Software

Quantities shown cover the documented configuration (one Administration Server, three Control Nodes, three Compute Nodes, three Storage Nodes). See Notes below for guidelines on what adjustments are needed to scale beyond this initial configuration.

QUANTITY	SKU	DESCRIPTION
1	A7650089	SUSE OpenStack Cloud Control Node plus Admin Server, x 86-64, 1 Instance, Priority Subscription, 3 Year
2(a)	A7650090	SUSE OpenStack Cloud Control Node, x86-64, 1 Instance, Priority Subscription, 3 Year
3(b)	A7648810	SUSE OpenStack Cloud Compute Node, x86-64, 1-2 Sockets, Priority Subscription, 3 Year
3(b)	A8344781	SUSE Linux Enterprise Server, x86 & x86-64, 1-2 Sockets with Unlimited Virtual Machines, Priority Subscription, 3 Year
1	A8703778	SUSE Enterprise Storage Base Configuration, x86-64, 4 OSD Nodes with 1-2 Sockets, Priority Subscription, 3 Year
(c)	A8703779	SUSE Enterprise Storage Expansion Node, x86-64, 1 OSD Node with 1-2 Sockets, L3-Priority Subscription, 3 Year

Note: Possible quantity adjustments:

(a) *Increase quantity, if needed, for each additional Control Node, beyond the total of three cited.*

(b) *Increase quantity, if needed, for each additional Compute Node, beyond the total of three cited.*

(c) *Incrementally increase quantity, for each Storage Node beyond the first four allowed in SUSE Enterprise Storage Base.*

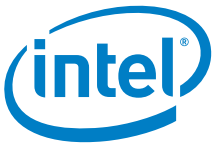

Disclaimers

Printed in USA 0316/HDW/MMcQ/PDF ♻ Please Recycle