SUSE® OpenStack Cloud Reference Implementation with Dell Hardware

Table of Contents

Introduction

This white paper is written to help an organization create and deploy an on-premise private cloud instance contained within its own network space, using computing and storage resources under its direct control. The document discusses key areas of integrating Dell networking and computing hardware with SUSE® OpenStack Cloud software while including specific information to leverage the highly automated nature of the software. Following these recommendations and best practices will allow an organization to quickly and confidently deploy the required infrastructure that yields a secure, scalable, performant, highly available private cloud instance for workloads, whether targeting testing, development or production usage of the services.

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

Solution Overview

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 local corporate firewall. Private cloud resources can be used across many groups, delivering access to IT resources and abstracting the complex, internal infrastructures from end users. This abstraction also applies to development and information technology operational professionals (DevOps). It allows individual virtual machines or workloads to be provisioned on demand, subject to allowed quotas, 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 describes the important characteristics of cloud computing.

Through integrated, open source software frameworks like OpenStack, enterprises can establish an Infrastructure-as-a-Service (IaaS) private cloud that delivers on-demand access to pools of compute, storage and networking resources for use within their organization, as shown in the following illustration:

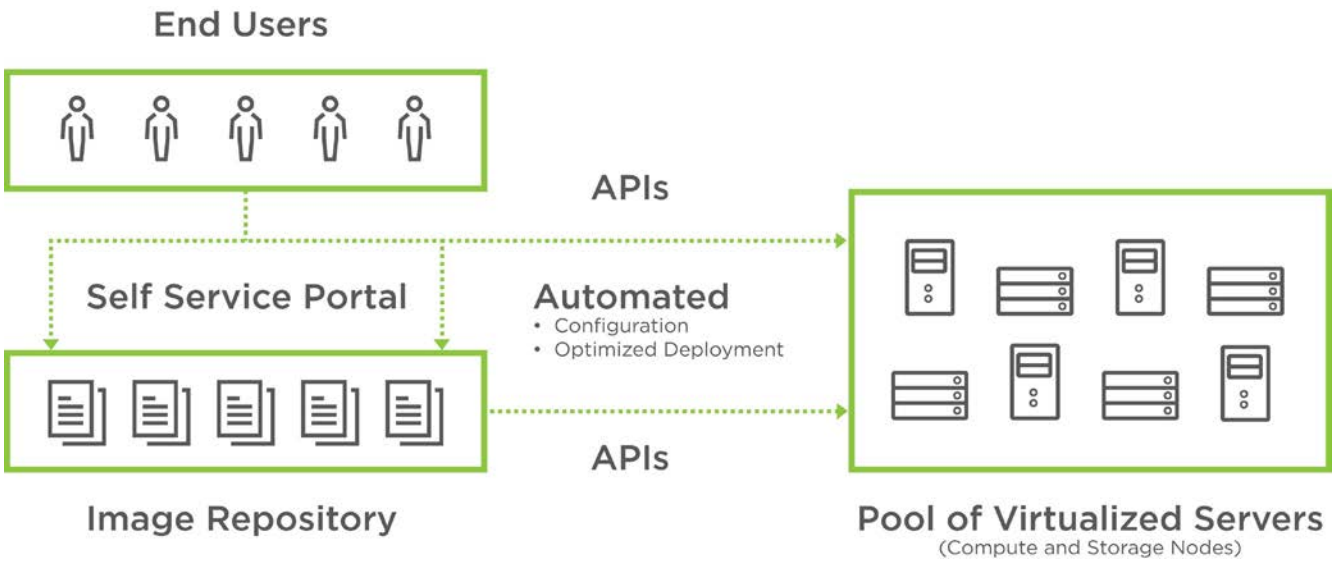

The underlying network, compute and storage resources can essentially be pre-provisioned as necessary to meet the overall requirements of the applications and services to be deployed later in the cloud instance. Compute resources are provided via physical servers by utilizing a hypervisor to run virtual machine workloads. Providing storage resources is accomplished 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. Finally, the private cloud software provides the user interface for set up, configuration and maintenance as well as the long-term operation of the cloud by bonding these three core components into a cohesive service offering.

The remainder of this white paper will present a complete, ready-to-use private cloud solution. First, it will discuss preparations and considerations across the various infrastructure elements common to the overall solution. Then it will detail the steps required for private cloud integration, starting with the setup of the control plane and followed by the storage and compute resources.

Preparation

Facility Considerations

The heating, ventilation, air conditioning (HVAC) and power requirements for a private cloud deployment need to be carefully assessed. As a reminder, ensure that calculated power consumption and cooling take into account the requirements of a fully utilized infrastructure since a private cloud instance is a shared resource subject to the needs and use of potentially many users.

You can estimate the heating, ventilation, air conditioning (HVAC) and power requirements for deploying the various network and hardware components by using the Dell Energy Smart Solution Advisor. With this tool, you can plan for the needs of your solution, to help order the correct Power Distribution Unit (PDU) and have the required HVAC for the installation.

> *Note: In the appendices, there are specific component listings for USA country-specific localizations, so adjust those items depending on your exact locale and the power configuration in the final installation location.*

Network Considerations

Networking is the technology component likely to take the most upfront planning because networking requirements for a private cloud instance have a fair amount of complexity, especially when it is being integrated into an existing IT infrastructure and network topology. This white paper provides specific recommendations for the networking topology and equipment that meet both scale and performance requirements, with later sections describing the specific network subnets and virtual LAN designations needed to satisfy the design decisions.

For the physical switching layer, use a pair of Dell Networking S4048-ON top-of-rack (ToR) 10GbE switches with 40GbE up-links 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-ON switch is used to handle the 1GbE connections from the dedicated BMC/iDRAC ports, plus a single 1GbE NIC is used for the admin network from each resource node and is up-linked to both of the S4048-ON switches for redundancy. With each server having multiple NIC ports, forming a bonded link for all of the remaining cloud networks with at least one port on each switch in the stack provides both performance and redundancy across the multiple NICs and switch configuration as shown in Illustration 2. The appendices include the recommended configuration for each of these Dell Networking switches plus the types of cables to include in an order.

> *Note: Ensure that all similar switching devices are consistent and up-to-date with regard to firmware versions to reduce potential troubleshooting issues later.*

> *Tip: Meticulous care of the network wiring from the various resource nodes and switches makes troubleshooting much easier. Where possible, also label connections and stick to consistent patterns of port/placement of connections, as shown in the following illustration.*

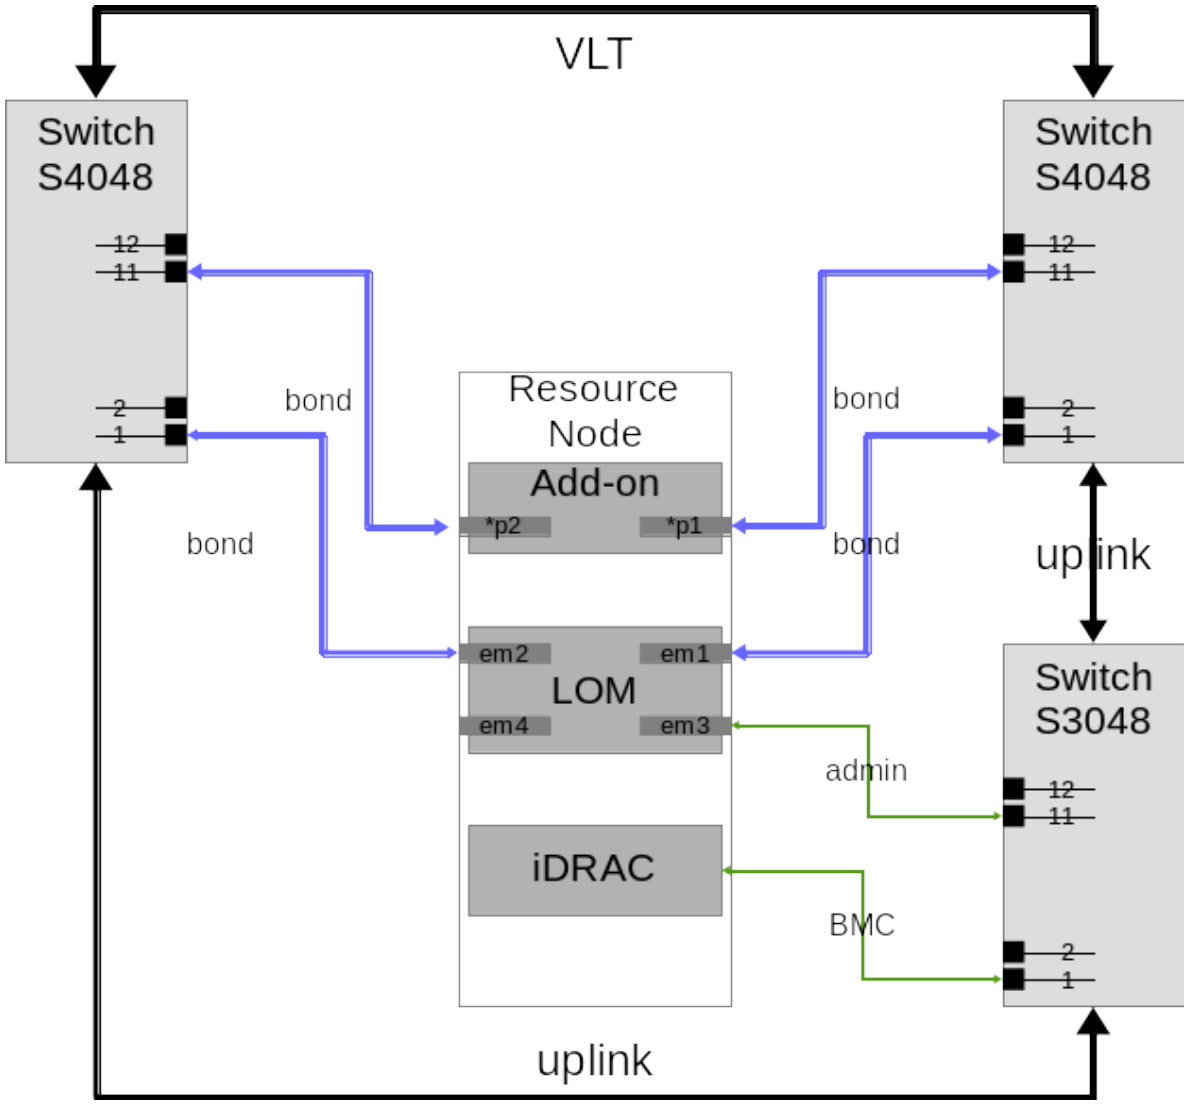

Illustration 2: System Network Interface and Switch Wiring Schematic

t

ive to introduce the necessary network name spaces for this private cloud instance. The name spaces and the functions they provide are fundamental to a private cloud implementation. The various network ranges are needed to manage, operate, access and provide isolation for all the components and services. The following table summarizes the various network name spaces used in this implementation:

Network Name Space	Network Description/Considerations
Admin	Private deployment and management 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. You have the following options for controlling access to this network: • Do not allow access from the outside and keep the admin network completely separate.
Storage VLAN:200	Private SUSE OpenStack 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 networks 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 public network. • The public network is 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 public and nova-floating functions are split to address security concerns, so that the OpenStack Dashboard and APIs are able to be assessed from one range, and the floating IPs assigned to the workload are in another network range.
Software Defined Network (os_sdn) VLAN:400	Private SUSE OpenStack Cloud internal virtual network. This network is used when OpenStack Networking (Neutron) is configured to use Open vSwitch with GRE tunneling for the virtual networks. This network is not accessed by users.
Private (nova-fixed) VLAN:500	Private SUSE OpenStack Cloud internal virtual network. This network is used for inter-instance communication and provides access to the outside world for the instances. The gateway required is also automatically provided by SUSE OpenStack Cloud.

Table 1: Network Name Space Descriptions

Computing and Storage Platform Considerations

One of the benefits of a private cloud implementation is that Dell PowerEdge servers can fulfill all the needs for various roles. 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 yet still meet the demands of the intended usage. This reference architecture is built on this plan for all of its components.

The Dell PowerEdge R630 and R730xd series of servers with their powerful and balanced performance, advanced I/O capabilities and flexible, scalable networking options are ideally suited for the various nodes in this private cloud installation reference configuration. The Dell PowerEdge R630 will serve as the control plane nodes and the compute nodes. For storage nodes, the expanded storage capabilities of the Dell PowerEdge R730xd platform make this an ideal choice. 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.

> *Note: Ensure that all similar system devices are consistent and up-to-date with regard to firmware versions to reduce potential troubleshooting issues later.*

> *Note: To ease later setup, follow the configuration tables included in the appendix for System Configuration and System Storage Volume settings on each system for its respective role.*

Software Component Considerations

The current, de facto open source standard of private cloud software is OpenStack. Because of its vibrant ecosystem of projects and governing body overseeing the inclusion and coordination of new offerings, each release provides an ever-expanding scope of functionality.

SUSE OpenStack Cloud provides an enterprise-grade implementation of OpenStack components and includes other tools necessary for a complete private cloud implementation. Augmenting the upstream OpenStack projects, SUSE OpenStack Cloud includes a complete deployment framework based upon the Crowbar and Chef projects in addition to providing core services such as the underlying database and message queue. With its multi-hypervisor support, SUSE OpenStack Cloud affords choice in the virtualization arena while leveraging existing expertise; however, this implementation focuses on a single hypervisor technology. To meet the needs of an inherently multi-user platform like a private cloud, both the control plane and the compute plane are made highly available to complement the inherent resiliency in the storage services.

There are four types of nodes in a full SUSE OpenStack Cloud deployment:

- An Administration Server provides all services needed to set up, deploy and maintain all of the other three types of resource nodes in the private cloud installation. The primary user of this node and its web interface is the cloud administrator.
- One or more Control Nodes host the core OpenStack services needed to orchestrate virtual machines deployed on the Compute Nodes and utilize volumes from the Storage Nodes in the private cloud installation. Multiple Control Nodes can be clustered together to provide highly available services.
- Two or more Storage Nodes act as a provider of persistent cloud-based storage elements and can offer OpenStack Object Storage (Swift) as well as Ceph-based options. Both object- and block-based storage types are supported, given the underlying technology. By design, these storage services are inherently resilient to protect against component failures. Users will typically interact with the Storage Nodes only through the use of volumes via the self-service dashboard.
- Any number of Compute Nodes, each of which is a physical, dedicated server running a hypervisor and serves as the host for the eventual deployment of workloads and services. Multiple Compute Nodes, using KVM, can be configured with a lightweight cluster technology to make failover of workload instances automatic. Users will typically interact with the Compute Nodes only through the self-service dashboard service.

The typical user and administrator interactions with the services and node types are shown in the following figure:

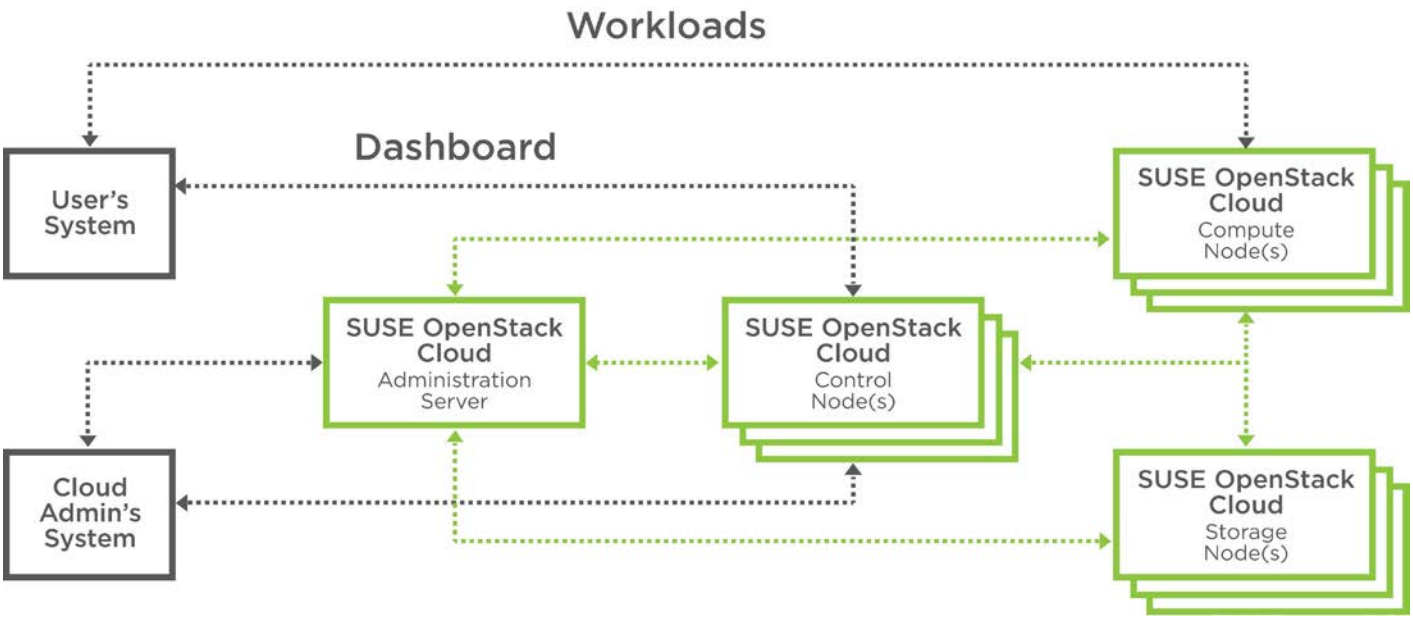

Once all of these basics are in hand, you can now build the private cloud instance. The following sections detail specific aspects of the solution components being deployed along with their related design decisions.

Solution Components

To coordinate all user-facing functionality and abstraction of infrastructure resources, a private cloud instance requires some dedicated administrative and control resources in addition to the actual storage and compute resources. As with most IT infrastructures, this covers the traditional disciplines of networking, systems and software, so the relevant aspects will be included in each section.

Here is a detailed account of deploying each of these components along with their respective design decisions and scalability recommendations. In general, this document is meant as a companion guide to the official network, system and software product deployment documentation, citing specific settings as needed to accomplish the respective design decisions.

Solution Admin Host

Because of the need for various administrative-like services, a convenient approach is to create a Solution Admin Host (SAH) that consolidates these services. Given a finite number of physical systems, this consolidation helps to preserve other system nodes for more resource-intensive use by deploying virtual machine guests for various administrative functions.

> *Design Decision: A simple hypervisor host, using KVM, provides the platform for the SAH and enables further grouping of administrative functions here as virtual machines.*

Process:

- Using one of the available Dell PowerEdge R630 systems, perform a bare-metal installation of the SUSE Linux Enterprise Server 12-SP1 operating system with either physical media or virtual media through iDRAC

  Note: The default partitioning scheme can be used, but remember to store any virtual machine images into the larger home directory partition.

- A minimal system can be installed, with at least the following patterns included:
  - base, minimal, kvm_server, kvm_tools
- Register the system in the SUSE Customer Center (SCC) during or after the installation to ensure all the latest software updates are present.
- After the installation completes, use YaST to:
  - Configure the desired networking including:
    - An external network interface for access from beyond the cloud environment (using one of the 1GbE NICs, e.g., em3)
    - A bond, mode 802.3ad to match the switch configuration, across all 10GbE NICs being used (e.g., em1, em2, p3p1, p3p2)
    - A bridge for virtualization on top of the previous bonded network interfaces, configured with an IP address in the admin network
    - VLANs for each of the desired administrative functions on top of the virtualization bridge (e.g., vlan200 for storage, vlan300 for public access)
  - Ensure an Administrative VNC server is set up and running to remotely access this system from other systems, which provides a graphical user interface

- Ensure the system is configured to have a valid Network Time Protocol (NTP) source, as this will become a reference for the cloud resource nodes as well

SUSE OpenStack Cloud Administration Server

The SUSE OpenStack Cloud Administration Server provides the deployment framework along with other tools and needed services for the cloud instance, such as 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, as 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.

Design Decision: *To preserve as many physical resource nodes for other functions of the cloud as possible, create a virtual machine on the SAH to become the Administration Server.*

Process:

- Using virt-manager (GUI) or virsh (CLI) on the SAH, create a virtual machine that meets or exceeds the minimum requirements from the deployment document.
- Allocate a virtual NIC for the admin network, tied to the virtualization bridge.
- Configure the following virtual CD drives:
  - SUSE Linux Enterprise Server 12-SP1 operating system ISO image (bootable)
  - SUSE OpenStack Cloud 6 ISO media (to be used during the install as an add-on)
- Complete the installation steps as described in the deployment guide, with the following guidance:
  - Install the Subscription Management Tool (SMT) on this system.
  - Include both the High Availability Extension and SUSE Enterprise Storage repositories.

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.

Note: Proceed only up to the "Crowbar Setup" section. Because this step is a critical point in the deployment, it is also convenient to snapshot the virtual machine's state, either via virsh or virt-manager, to preserve the effort up to this point.

Tip: Before proceeding beyond this critical juncture, the following network design decisions are necessary to lock down the configuration for the rest of this private cloud instance deployment. Other options are available, but beyond the scope of this document.

Design Decisions:

The network mode across the resource nodes will be "dual," making the admin network distinct from the remaining cloud networks, with the following settings:

Note: Refer to the Network Name Spaces (Table 1) and to appendices, for both the switch settings and for the system boot and NIC settings.

Admin and BMC networks will coexist as untagged VLANs on the 1GbE NIC (e.g., em3) and will provide the PXE booting interface. The four, available 10GbE NICs of all the Dell PowerEdge systems will be bonded together in an 802.3ad (LACP, mode 4)network configuration for all the remaining VLAN network name spaces.

For the remainder of this document, the following example IP addresses and subnets were used for the various network name spaces (substitute as appropriate for your local environment and for the remainder of this document):

Admin/BMC/BMC_VLAN : 192.168.110.0/24 (untagged),
Storage : 192.168.125.0/24 (VLAN 200),
Nova Floating : 10.148.45.0/24 (VLAN 300),
Public : 10.148.45.0/24 (VLAN 300),
OS_SDN : 192.168.130.0/24 (VLAN 400),
Nova Fixed : 192.168.123.0/24 (VLAN 500)

- At this point, carefully review the Crowbar Setup section of the deployment document:
  - Ensure the IP address of the Administration Server is as desired, resides in the Admin Network and has routing through the SAH to external resources.
  - Manually edit the "conduit map" section pertaining to the "dual" configuration section of /etc/crowbar/network.json file correctly references "1g1" for "intf0" and "10g1", 10g2", "10g3", 10g4" across the referenced "intf1" and "intf2" interface listings. See the Network Conduits section of the deployment guide for more details.
  - Furthermore, to match the switch settings, ensure that the "teaming mode" is set to "4" to match the LACP settings of the respective switch ports.
- To complete the setup, either utilize the "screen install-suse-cloud" setup script while directly logged into the Administration Server or start the Crowbar service and connect to the web server at http://<IPofAdministrationServer> to start the process

  Note: When this previous task completes, review the noted log file for any issues and resolve them before proceeding. More information can be found in the troubleshooting section of the deployment guide.

  Note: It is important to ensure that the Administration Server has all available package updates installed by registering with the SUSE Customer Center or by utilizing the local repositories of updates already available on this node.

- At this point, login to the web interface of the Administration Server to prepare for deployment of the remaining resource nodes and services:
  - Modify the IPMI barclamp, setting the "Enable BMC" to true, and update the login credentials to match the iDRAC settings; then apply.
  - Modify the NTP barclamp, adding the SAH as another NTP server; then apply.
- Now boot each of the remaining resource nodes via PXE. Each node should be displayed on the administrative web interface in a discovered state.

  Note: It is important to verify that each respective node has the expected disk drive volumes, network interfaces and other physical attributes as a means of correlating to the eventual role.

  Then each resource can be allocated, either in an iterative process by editing each node or in the bulk allocation mode, as per the following table:

Resource Node	Target Platform	Alias	Group	Intended Role
Dell PowerEdge R630	SUSE Linux Enterprise 12 SP1	cntl1	Control	Controller
Dell PowerEdge R630	SUSE Linux Enterprise 12 SP1	cntl2	Control	Controller
Dell PowerEdge R630	SUSE Linux Enterprise 12 SP1	cntl3	Control	Controller
Dell PowerEdge R630	SUSE Linux Enterprise 12 SP1	nova1	Compute	Compute
Dell PowerEdge R630	SUSE Linux Enterprise 12 SP1	nova2	Compute	Compute
Dell PowerEdge R630	SUSE Linux Enterprise 12 SP1	nova3	Compute	Compute
Dell PowerEdge R730xd	SUSE Linux Enterprise 12 SP1	ceph1	Storage	Storage
Dell PowerEdge R730xd	SUSE Linux Enterprise 12 SP1	ceph2	Storage	Storage
Dell PowerEdge R730xd	SUSE Linux Enterprise 12 SP1	ceph3	Storage	Storage

Table 2: Association of Resource Nodes to Roles and Functions

At this point, you are ready to begin assigning the services and roles to the various resource nodes as described in the following sections.

SUSE OpenStack Cloud Control Nodes

A SUSE OpenStack Cloud Control Node is managed and deployed through the SUSE OpenStack Cloud Administration Server web interface. To ensure the high availability of the control plan, you need to create a cluster of dedicated Control Nodes with at least one pair of nodes needed for a cluster. However, an odd number of cluster member nodes is preferred to establish and maintain a quorum. Multiple clusters are allowed, should services need to be isolated for performance or other considerations, such as whether the service is stateful or can be used only in an active-passive mode.

> ***Design Decision****: Combine three resource nodes (cntl1, cntl2, cntl3) into an HA cluster configuration for the control plane via the Pacemaker barclamp. This cluster can then be assigned to many of the OpenStack services.*

Process:

- Apply the Pacemaker barclamp to the respective cluster nodes:
  - Create a recognizable proposal name (e.g., “control_ha”).
  - 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=lanplus
    - Include all three nodes into “pacemaker-cluster-member” and “hawk-server.”
- With a control-plane cluster in place, two of the foundational services needed by the private cloud instance are shown below and can be layered on to the control_ha cluster:

Foundational Services	OpenStack Project Name	Description
Database	N/A	Backend database for OpenStack Services that tracks current information about users and instances.
Message Queue	N/A	Enables OpenStack services to communicate with the other nodes via Advanced Message Queue Protocol (AMQP).

Table 3: Foundational Services for OpenStack

- For each of these active-passive, stateful services, the control_ha cluster can be assigned to the respective service role in the Barclamps → OpenStack proposal.
- A shared file system is required for each of these services and is configured on the SAH by creating respective NFS exports (i.e., /exports/db and /exports/mq) with the parameters noted in the deployment guide to use for these services.

> *Note: As you may notice when visiting the High Availability Web Console (HAWK) interface (via the link on any of the clustered Control Node pages), these services will tend to exist on different nodes, assuming two or more nodes are active in the cluster.*
>
> *Tip: Proceed to set up the SUSE OpenStack Cloud Storage Nodes as noted in that section; then return here to complete the setup of the control plane.*

Process (post Storage Node):

- Assign the following roles to the control_ha cluster for each of the following services, subject to the process notes following the table:

Core OpenStack Services	OpenStack Project Name	Description
Identity	Keystone	Provides an authentication and authorization service for other OpenStack services. Provides a catalog of endpoints for all OpenStack services.
Image	Glance	Stores and retrieves virtual machine disk images. OpenStack Compute makes use of this during instance provisioning.
Block Storage	Cinder	Provides persistent block storage to running instances. Its pluggable driver architecture facilitates the creation and management of block storage devices.
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.

Table 4: Core OpenStack Services

- For Glance, select "Rados" as the "Default Storage Store" to use the Storage Nodes from the Ceph setup.
- For Cinder, select "RBD" as the "Type of Volume" for the "Add new Cinder Backend" and delete the "Backend: default" Cinder Backend. Also assign the control_ha cluster to the "cinder-controller" role.

Tip: Proceed to set up the SUSE OpenStack Cloud Compute Nodes as noted in that section; then return here to complete the setup of the control plane.

Process (post Compute Node):

- Assign the following roles to the control_ha cluster for each of the following services, subject to the process notes following the table:

OpenStack Services	OpenStack Project Name	Description
Compute	Nova	Manages the lifecycle of compute instances in an OpenStack environment. Responsibilities include spawning, scheduling and decommissioning virtual machines on demand.
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.
Telemetry	Ceilometer	Monitors and meters the OpenStack cloud for billing, benchmarking, scalability and statistical purposes. Note the MongoDB backend is still in a technical preview status.
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.
Integration	Tempest	Set of integration tests to be run against a live OpenStack instance,

Test Suite (Optional)		include batteries of tests for OpenStack API validation, scenarios and other specific tests useful in validating an OpenStack deployment

Table 5: OpenStack Services

- For Nova, assign the control_ha cluster to the designated controller role(s), and assign control_ha (3 remote nodes) to the "nova- multi-compute-kvm" role.
- For Horizon, Ceilometer and Heat, assign the control_ha cluster to the respective roles.
- For Tempest, assign any one of the Control Nodes to the role.

Design Decision: *The following OpenStack services available in SUSE OpenStack Cloud are* ***not*** *deployed in this configuration: FileShare (Manilla), nor any of the Tech Preview services, for example, Database-as-a-Service (Trove), Elastic Map Reduce (Sahara), Bare Metal Install (Ironic), nova-docker.*

All SUSE OpenStack Cloud services should be deployed and ready for use, so proceed to the Operational Aspects section.

SUSE OpenStack Cloud Storage Nodes

A SUSE OpenStack Cloud Storage Node is managed and deployed through the SUSE OpenStack Cloud Administration Server web interface. These nodes act as a provider of persistent cloud-based storage elements and offer a management interface. This storage can be presented in both object or block format, furnish volumes for a workload instance or even act as the underlying storage for some of the core OpenStack services. Multiple Storage Nodes should be deployed to provide data redundancy and resiliency in case of component or system failures, as detailed in this document. Storage Nodes can be added over time to increase the storage capacity.

Design Decisions:

Three Resource Nodes (ceph1, ceph2, ceph3) and one virtual machine will be combined into a Ceph configuration to provide backend storage for the private cloud instance.

An additional resource node, in the form of a virtual machine, will provide the management interface, via Calamari. Since the specified storage resource nodes have SSD mechanisms, these will be configured as the journals for the OSDs to maximize performance.

Note: The MON processes also run on these same storage nodes. No other OpenStack services are permitted on these nodes. It is sufficient to use only these three initial MON nodes and not add additional ones even for very large storage clusters.

Process:

- Create an additional virtual machine on the SAH to, at least, meet the minimum requirements for a Calamari monitor and management node as described in the deployment guide.

  Note: Configure five virtual NICs (matching the number of active interfaces of other resource nodes) for the virtual machine guest, all tied to the virtualization bridge of the SAH. Ensure that the Storage network VLAN is present on the SAH itself and tied to the same virtualization bridge as well.

- PXE boot the newly created Calamari virtual machine to act as the Storage Nodes; then allocate it as a SUSE Linux Enterprise Server 12 SP1 target platform, an alias of Calamari, into the Storage group, and with a storage intended role.
- Assign the "ceph-calamari" role to the virtual machine.
- Assign the "ceph-osd" and "ceph-mon" roles to the ceph1, ceph2 and ceph3 resource nodes.
- Modify the Ceph barclamp, editing in raw mode to insert the respective SSD device designations (i.e., /dev/sdb, /dev/sdc, /dev/sdd) into the "journal devices" portion of the osd section.

- Apply the barclamp proposal.

Note: By default, all unused disk drives will be allocated to the OSD function, and partitions of each SSD will be assigned as the journal for one of the local OSDs.

Tip: With the Storage Nodes setup, return to the post storage node portion of the Control Node setup.

SUSE OpenStack Cloud Compute Nodes

A SUSE OpenStack Cloud Compute Node is managed and deployed through the SUSE OpenStack Cloud Administration Server web interface. 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 systems and then to allow selective spreading of workloads across these zones for availability concerns.

Design Decisions:

Three resource nodes (nova1, nova2, nova3) will be incorporated into the existing HA control plane cluster configuration to provide workload failover capabilities.

All Compute Nodes will provide the KVM hypervisor.

Process:

- Apply the pacemaker-remote role barclamp to the respective cluster nodes of the existing control_ha Pacemaker proposal

Tip: With the Compute Nodes prepared, return to the post Compute Node portion of the Control Node setup.

Operational Aspects

At this point, you should be able to utilize your private cloud instance, according to the SUSE OpenStack Cloud Admin, User and Supplement Guides, with the ability to upload images for deployment, launch workloads and manipulate data volumes. For reference, an export of the deployed instance is included in the appendices, showing the nodes and the roles deployed. The overall, logical view of the public cloud instance constructed per the steps of this document is shown below:

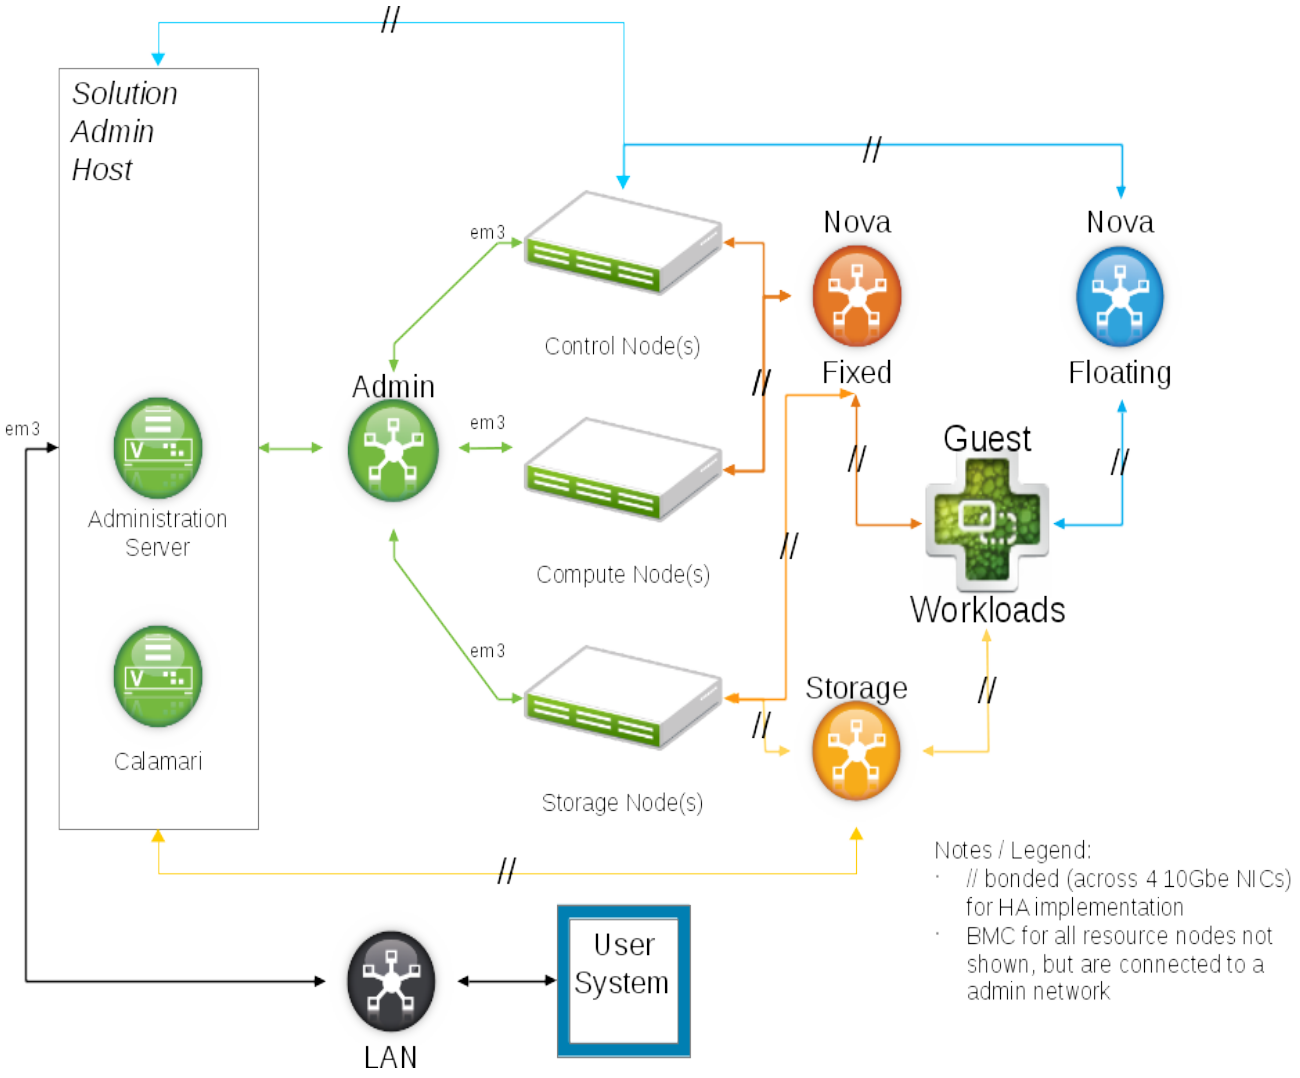

You can also optionally perform a basic functionality and API test of your overall installation, utilizing the OpenStack Integration Test Suite (Tempest) to exercise the "smoke" tests that are part of the OpenStack gate process as described in the deployment guide.

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. As you progress through the deployment outlined in the document, if any errors arise, address those first before proceeding onto the next step or service. If you cannot resolve the issue quickly, check if the troubleshooting section of the deployment guide can 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.

Appendices

Bill of Materials : Network

Dell Networking S4048-ON

Quantity	SKU	Description
2	210-ADUW	Dell Networking S4048-ON, 48x 10GbE SFP+ and 6x 40GbE QSFP+ ports, IO to PSU air, 1x AC PSUs, DNOS9
2	634-BCWZ	Dell Networking OS9, S4048-ON
40	470-AAGN	Dell Networking, Cable, SFP+to SFP+, 10GbE, Copper TwinaxDirect Attach Cable, 1 Meter
2	470-AAGP	Dell Networking, Cable, SFP+to SFP+, 10GbE, Copper TwinaxDirect Attach Cable, 3 Meter
2	470-AAIB	Dell Networking, Cable, QSFP+ to QSFP+, 40GbE Passive Copper Direct Attach Cable, 0.5 m
2	450-AAQB	Power Supply, AC, 460w, IO to PSU airflow, S4048-ON
4	450-AASX	Dell Networking, Jumper Cord, 250V, 12A, 2 Meters, C13/C14, US
2	634-BCWX	Dell Networking S4048-ON User Guide
2	332-1286	US Order
2	989-3439	Thank you choosing Dell ProSupport. For tech support, visit http://www.dell.com/support or call 1-800- 945-3355
2	997-6304	Dell Hardware Limited Warranty Extended Year(s)
2	997-6179	ProSupport: Next Business Day Onsite Service After ProblemDiagnosis, Initial Year
2	997-6306	Info 3rd Party Software Warranty provided by Vendor
2	997-6186	ProSupport: 7x24 HW / SW Tech Support and Assistance, 3 Years
2	997-6305	Dell Hardware Limited Warranty Initial Year
2	997-6174	ProSupport: Next Business Day Onsite Service After ProblemDiagnosis, 2 Year Extended
2	900-9997	On-Site Installation Declined

Dell Networking S3048-ON

Quantity	SKU	Description
1	210-AEDM	Dell Networking S3048-ON, 48x 1GbE, 4x SFP+ 10GbE ports, Stacking, IO to PSU air, 1x AC PSU, DNOS 9
1	634-BCXS	Dell Networking OS 9, S3048-ON
1	450-AEOP	Power Supply, 200w, IO to PSU airflow, S3048-ON
2	450-AASX	Dell Networking, Jumper Cord, 250V, 12A, 2 Meters, C13/C14, US
1	634-BCXR	Dell Networking S3048-ON User Guide
1	332-1286	US Order

Quantity	SKU	Description
1	802-7389	Dell Hardware Limited Warranty Initial Year
1	802-7403	ProSupport: Next Business Day Onsite Service After ProblemDiagnosis, 2 Year Extended
1	802-7394	ProSupport: Next Business Day Onsite Service After ProblemDiagnosis, Initial Year
1	802-7404	ProSupport: 7x24 HW / SW Tech Support and Assistance, 3 Years
1	997-6306	Info 3rd Party Software Warranty provided by Vendor
1	989-3439	Thank you choosing Dell ProSupport. For tech support, visit http://www.dell.com/support or call 1-800- 945-3355
1	802-7400	Dell Hardware Limited Warranty Extended Year(s)
1	900-9997	On-Site Installation Declined
1	973-2426	Declined Remote Consulting Service

Bill of Materials : System

Dell PowerEdge R630 Solution Admin Host

Quantity	SKU	Description
1	329-BCIY	PowerEdge R630 Motherboard
1	210-ACXS	PowerEdge R630 Server
1	321-BBKL	Chassis with up to 8, 2.5" Hard Drives, 3 PCIe Slots
1	340-AKPS	PowerEdge R630 Shipping- 8 Drive Chassis
1	338-BFFF	Intel Xeon E5-2650 v3 2.3GHz,25M Cache,9.60GT/s QPI,Turbo,HT,10C/20T (105W) Max Mem 2133MHz
1	374-BBGM	Upgrade to Two Intel Xeon E5-2650 v3 2.3GHz,25M Cache,9.60GT/s QPI,Turbo,HT,10C/20T (105W)
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, 2133 MT/s, Dual Rank, x4 Data Width
1	780-BBJC	No RAID for H330/H730/H730P (1-24 HDDs or SSDs)
1	405-AAEG	PERC H730 Integrated RAID Controller, 1GB Cache
4	400-AEEN	500GB 7.2K RPM NLSAS 6Gbps 2.5in Hot-plug Hard Drive,13G
1	540-BBHY	Intel X520 DP 10Gb DA/SFP+ Server Adapter, Low Profile

Quantity	SKU	Description
1	385-BBHO	iDRAC8 Enterprise, integrated Dell Remote Access Controller, Enterprise
1	429-AAQL	DVD+/-RW SATA Internal
1	350-BBBW	No Bezel
1	770-BBBC	ReadyRails Sliding Rails Without Cable Management Arm
1	384-BBBL	Performance BIOS Settings
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, North America
1	343-BBDK	Electronic System Documentation and OpenManage DVD Kit, PowerEdge R630
1	619-ABVR	No Operating System
1	421-5736	No Media Required
1	332-1286	US Order
1	540-BBBB	Intel X520 DP 10Gb DA/SFP+, + I350 DP 1Gb Ethernet, Network Daughter Card
1	951-2015	Thank you for choosing Dell ProSupport Plus. For tech support, visit http://www.dell.com/contactdell
1	976-7728	Dell Hardware Limited Warranty Plus On Site Service
1	976-7768	ProSupport Plus: 7x24 HW/SW Tech Support and Assistance,3 Year
1	976-7761	ProSupport Plus: 7x24 Next Business Day Onsite Service, 3 Year
2	900-9997	On-Site Installation Declined

Dell PowerEdge R630 Control and Compute Nodes

Quantity	SKU	Description
6	329-BCIY	PowerEdge R630 Motherboard
6	210-ACXS	PowerEdge R630 Server
6	321-BBKL	Chassis with up to 8, 2.5" Hard Drives, 3 PCIe Slots
6	340-AKPS	PowerEdge R630 Shipping- 8 Drive Chassis
6	338-BFFF	Intel Xeon E5-2650 v3 2.3GHz,25M Cache,9.60GT/s QPI,Turbo,HT,10C/20T (105W) Max Mem 2133MHz
6	374-BBGM	Upgrade to Two Intel Xeon E5-2650 v3 2.3GHz,25M Cache,9.60GT/s QPI,Turbo,HT,10C/20T (105W)
12	412-AAEE	120W Heatsink for PowerEdge R630
6	370-ABWE	DIMM Blanks for System with 2 Processors

Quantity	SKU	Description
6	370-ABUF	2133MT/s RDIMMs
6	370-AAIP	Performance Optimized
48	370-ABUG	16GB RDIMM, 2133 MT/s, Dual Rank, x4 Data Width
6	780-BBJC	No RAID for H330/H730/H730P (1-24 HDDs or SSDs)
6	405-AAEG	PERC H730 Integrated RAID Controller, 1GB Cache
12	400-AJRU	300GB 15K RPM SAS 12Gbps 2.5in Hot-plug Hard Drive
36	400-AEER	600GB 10K RPM SAS 6Gbps 2.5in Hot-plug Hard Drive,13G
6	540-BBHY	Intel X520 DP 10Gb DA/SFP+ Server Adapter, Low Profile
6	634-BBWU	OpenManage Essentials, Server Configuration Management
6	385-BBHO	iDRAC8 Enterprise, integrated Dell Remote Access Controller, Enterprise
6	429-AAQL	DVD+/-RW SATA Internal
6	350-BBBW	No Bezel
6	770-BBBC	ReadyRails Sliding Rails Without Cable Management Arm
6	384-BBBL	Performance BIOS Settings
6	450-ADWM	Dual, Hot-plug, Redundant Power Supply (1+1), 1100W
12	492-BBDH	C13 to C14, PDU Style, 12 AMP, 2 Feet (.6m) Power Cord, North America
6	343-BBDK	Electronic System Documentation and OpenManage DVD Kit, PowerEdge R630
6	619-ABVR	No Operating System
6	421-5736	No Media Required
6	632-BBDC	SanDisk DAS Cache, 90 Day Trial License
6	332-1286	US Order
6	540-BBBB	Intel X520 DP 10Gb DA/SFP+, + I350 DP 1Gb Ethernet, Network Daughter Card
6	976-7761	ProSupport Plus: 7x24 Next Business Day Onsite Service, 3 Year
6	976-7728	Dell Hardware Limited Warranty Plus On Site Service
6	976-7768	ProSupport Plus: 7x24 HW/SW Tech Support and Assistance,3 Year
6	951-2015	Thank you for choosing Dell ProSupport Plus. For tech support, visit http://www.dell.com/contactdell

Quantity	SKU	Description
6	900-9997	On-Site Installation Declined

Dell PowerEdge R730xd Storage Nodes

Quantity	SKU	Description
3	210-ADBC	PowerEdge R730xd Server
3	591-BBCH	PowerEdge R730/R730xd Motherboard
3	350-BBEX	Chassis with up to 12 + 4 Internal, 3.5" Hard Drives and 2, 2.5" Flex Bay Hard Drives
3	338-BFFG	Intel Xeon E5-2660 v3 2.6GHz,25M Cache,9.60GT/s QPI,Turbo,HT,10C/20T (105W) Max Mem 2133MHz
3	374-BBGN	Upgrade to Two Intel Xeon E5-2660 v3 2.6GHz,25M Cache,9.60GT/s QPI,Turbo,HT,10C/20T (105W)
3	330-BBCR	R730/xd PCIe Riser 1, Right
3	330-BBCO	R730/xd PCIe Riser 2, Center
3	370-ABUF	2133MT/s RDIMMs
3	370-AAIP	Performance Optimized
24	370-ABUG	16GB RDIMM, 2133 MT/s, Dual Rank, x4 Data Width
3	780-BBLR	RAID 1+Unconfigured RAID forH330/H730/H730P (2 + 1-20 HDDs or SSDs)
3	405-AAEG	PERC H730 Integrated RAID Controller, 1GB Cache
6	400-AEOE	300GB 15K RPM SAS 6Gbps 2.5in Flex Bay Hard Drive,13G
9	400-AEPY	200GB Solid State Drive SAS Mix Use MLC 12Gbps 2.5in Hot-plug Drive,3.5in HYB CARR,13G
27	400-AEGB	2TB 7.2K RPM NLSAS 6Gbps 3.5in Hot-plug Hard Drive,13G
12	400-AENO	2TB 7.2K RPM NLSAS 6Gbps 3.5in Internal Bay Hard Drive,13G
3	540-BBBB	Intel X520 DP 10Gb DA/SFP+, + I350 DP 1Gb Ethernet, Network Daughter Card
3	540-BBCT	Intel X520 DP 10Gb DA/SFP+ Server Adapter
3	350-BBEJ	Bezel
3	770-BBBQ	ReadyRails Sliding Rails Without Cable Management Arm
3	384-BBBL	Performance BIOS Settings
3	450-ADWM	Dual, Hot-plug, Redundant Power Supply (1+1), 1100W
6	492-BBDH	C13 to C14, PDU Style, 12 AMP, 2 Feet (.6m) Power Cord, North America
3	631-AAJG	Electronic System Documentation and OpenManage DVD Kit, PowerEdge R730/xd

Quantity	SKU	Description
3	619-ABVR	No Operating System
3	421-5736	No Media Required
3	387-BBHY	Energy Star, PowerEdge R720
3	332-1286	US Order
3	366-0193	Std Bios Setting Power Management* - Maximum Performance
3	340-AKPM	PowerEdge R730xd Shipping
3	370-ABWE	DIMM Blanks for System with 2 Processors
6	374-BBHR	Heatsink for 12 + 4 Chassis PowerEdge R730xd
3	634-BBWU	OpenManage Essentials, Server Configuration Management
3	385-BBHO	iDRAC8 Enterprise, integrated Dell Remote Access Controller, Enterprise
3	976-9030	ProSupport Plus: 7x24 HW/SW Tech Support and Assistance,3 Year
3	976-9007	Dell Hardware Limited Warranty Plus On Site Service
3	976-9029	ProSupport Plus: 7x24 Next Business Day Onsite Service, 3 Year
3	951-2015	Thank you for choosing Dell ProSupport Plus. For tech support, visit http://www.dell.com/contactdell
3	900-9997	On-Site Installation Declined

Bill of Materials : Software

Quantities shown cover the documented configuration (one Solution Admin Host with Administration Server and Calamari Node, three HA Control Nodes, three HA 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
4(c)	A8344781	SUSE Linux Enterprise Server, x86 & x86-64, 1-2 Sockets with Unlimited Virtual Machines, Priority Subscription, 3 Year
3(c)	A8619764	SUSE Linux High Availability Extension, x86 & x86-64, 1-2 Sockets with Inherited Virtualization, Inherited Subscription, 3 Year
1	A8703778	SUSE Enterprise Storage Base Configuration, x86-64, 4 OSD Nodes with 1-2 Sockets, Priority Subscription, 3 Year

(d)	A8703779	SUSE Enterprise Storage Expansion Node, x86-64, 1 OSD Node with 1-2 Sockets, 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) Increase quantity, if needed for each SUSE-based Compute Node running KVM or Xen, beyond the total of the Solution Admin host and three cited Compute Nodes

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

Configuration : Network Switches

Dell Networking S4048-ON

Note: Adjust highlighted values for each of the respective switches.

```
! Version 9.9(0.0)
! Last configuration change at Wed Apr  6 02:06:32 2016 by default
! Startup-config last updated at Wed Apr  6 02:06:36 2016 by default
!
boot system stack-unit 1 primary system://A
boot system stack-unit 1 secondary system://B
!
logging console informational
!
hostname S4048-1
!
protocol lldp
 advertise dot1-tlv port-protocol-vlan-id
 advertise dot3-tlv max-frame-size
 advertise management-tlv system-description system-name
 advertise med
!
redundancy auto-synchronize full
!
username crowbar password 7 b273ab54c2a8d82d privilege 15
!
protocol spanning-tree rstp
 no disable
 bridge-priority 20480
!
vlt domain 1
 peer-link port-channel 1
 back-up destination 10.9.244.39
 primary-priority 1
 unit-id 0
!
stack-unit 1 provision S4048-ON
!
interface TenGigabitEthernet 1/1
```

```
 no ip address
!
 port-channel-protocol LACP
  port-channel 11 mode active
 no shutdown
!
interface TenGigabitEthernet 1/2
 no ip address
!
 port-channel-protocol LACP
  port-channel 12 mode active
 no shutdown
!
interface TenGigabitEthernet 1/3
 no ip address
!
 port-channel-protocol LACP
  port-channel 13 mode active
 no shutdown
!
interface TenGigabitEthernet 1/4
 no ip address
!
 port-channel-protocol LACP
  port-channel 14 mode active
 no shutdown
!
interface TenGigabitEthernet 1/5
 no ip address
!
 port-channel-protocol LACP
  port-channel 15 mode active
 no shutdown
!
interface TenGigabitEthernet 1/6
 no ip address
!
 port-channel-protocol LACP
  port-channel 16 mode active
 no shutdown
!
interface TenGigabitEthernet 1/7
 no ip address
!
 port-channel-protocol LACP
  port-channel 17 mode active
 no shutdown
!
interface TenGigabitEthernet 1/8
 no ip address
!
 port-channel-protocol LACP
  port-channel 18 mode active
 no shutdown
!
interface TenGigabitEthernet 1/9
 no ip address
!
 port-channel-protocol LACP
```

```
  port-channel 19 mode active
 no shutdown
!
interface TenGigabitEthernet 1/10
 no ip address
!
 port-channel-protocol LACP
  port-channel 20 mode active
 no shutdown
!
interface TenGigabitEthernet 1/11
 no ip address
 no shutdown
!
interface TenGigabitEthernet 1/12
 no ip address
 no shutdown
!
interface TenGigabitEthernet 1/13
 no ip address
!
 port-channel-protocol LACP
  port-channel 11 mode active
 no shutdown
!
interface TenGigabitEthernet 1/14
 no ip address
!
 port-channel-protocol LACP
  port-channel 12 mode active
 no shutdown
!
interface TenGigabitEthernet 1/15
 no ip address
!
 port-channel-protocol LACP
  port-channel 13 mode active
 no shutdown
!
interface TenGigabitEthernet 1/16
 no ip address
!
 port-channel-protocol LACP
  port-channel 14 mode active
 no shutdown
!
interface TenGigabitEthernet 1/17
 no ip address
!
 port-channel-protocol LACP
  port-channel 15 mode active
 no shutdown
!
interface TenGigabitEthernet 1/18
 no ip address
!
 port-channel-protocol LACP
  port-channel 16 mode active
 no shutdown
```

```
!
interface TenGigabitEthernet 1/19
 no ip address
!
 port-channel-protocol LACP
  port-channel 17 mode active
 no shutdown
!
interface TenGigabitEthernet 1/20
 no ip address
!
 port-channel-protocol LACP
  port-channel 18 mode active
 no shutdown
!
interface TenGigabitEthernet 1/21
 no ip address
!
 port-channel-protocol LACP
  port-channel 19 mode active
 no shutdown
!
interface TenGigabitEthernet 1/22
 no ip address
!
 port-channel-protocol LACP
  port-channel 20 mode active
 no shutdown
!
interface TenGigabitEthernet 1/23
 no ip address
 shutdown
!
interface TenGigabitEthernet 1/24
 no ip address
 shutdown
!
interface TenGigabitEthernet 1/25
 no ip address
!
 port-channel-protocol LACP
  port-channel 6 mode active
 no shutdown
!
interface TenGigabitEthernet 1/26
 no ip address
 shutdown
!
interface TenGigabitEthernet 1/27
 no ip address
 shutdown
!
interface TenGigabitEthernet 1/28
 no ip address
 shutdown
!
interface TenGigabitEthernet 1/29
 no ip address
 no shutdown
```

```
!
interface TenGigabitEthernet 1/30
 no ip address
 no shutdown
!
interface TenGigabitEthernet 1/31
 no ip address
 no shutdown
!
interface TenGigabitEthernet 1/32
 no ip address
 no shutdown
!
interface TenGigabitEthernet 1/33
 no ip address
 no shutdown
!
interface TenGigabitEthernet 1/34
 no ip address
 no shutdown
!
interface TenGigabitEthernet 1/35
 no ip address
 no shutdown
!
interface TenGigabitEthernet 1/36
 no ip address
 no shutdown
!
interface TenGigabitEthernet 1/37
 no ip address
 no shutdown
!
interface TenGigabitEthernet 1/38
 no ip address
 no shutdown
!
interface TenGigabitEthernet 1/39
 no ip address
 no shutdown
!
interface TenGigabitEthernet 1/40
 no ip address
 no shutdown
!
interface TenGigabitEthernet 1/41
 no ip address
 no shutdown
!
interface TenGigabitEthernet 1/42
 no ip address
 no shutdown
!
interface TenGigabitEthernet 1/43
 no ip address
 no shutdown
!
interface TenGigabitEthernet 1/44
 no ip address
```

```
 no shutdown
!
interface TenGigabitEthernet 1/45
 no ip address
 no shutdown
!
interface TenGigabitEthernet 1/46
 no ip address
 portmode hybrid
 switchport
 no shutdown
!
interface TenGigabitEthernet 1/47
 no ip address
 no shutdown
!
interface TenGigabitEthernet 1/48
 no ip address
 no shutdown
!
interface fortyGigE 1/49
 no ip address
 shutdown
!
interface fortyGigE 1/50
 no ip address
 shutdown
!
interface fortyGigE 1/51
 no ip address
 shutdown
!
interface fortyGigE 1/52
 no ip address
 shutdown
!
interface fortyGigE 1/53
 no ip address
 no shutdown
!
interface fortyGigE 1/54
 no ip address
 no shutdown
!
interface ManagementEthernet 1/1
 ip address 10.9.244.38/23
 no shutdown
!
interface ManagementEthernet 2/1
 no shutdown
!
interface ManagementEthernet 3/1
 no shutdown
!
interface ManagementEthernet 4/1
 no shutdown
!
interface ManagementEthernet 5/1
 no shutdown
```

```
!
interface ManagementEthernet 6/1
 no shutdown
!
interface Port-channel 1
 description VLT_domain 1
 no ip address
 channel-member fortyGigE 1/53,1/54
 no shutdown
!
interface Port-channel 2
 no ip address
 portmode hybrid
 switchport
 channel-member TenGigabitEthernet 1/48
 vlt-peer-lag port-channel 2
 no shutdown
!
interface Port-channel 6
 no ip address
 portmode hybrid
 switchport
 spanning-tree rstp edge-port
 vlt-peer-lag port-channel 6
 no shutdown
!
interface Port-channel 11
 no ip address
 portmode hybrid
 switchport
 spanning-tree rstp edge-port
 vlt-peer-lag port-channel 11
 no shutdown
!
interface Port-channel 12
 no ip address
 portmode hybrid
 switchport
 spanning-tree rstp edge-port
 vlt-peer-lag port-channel 12
 no shutdown
!
interface Port-channel 13
 no ip address
 portmode hybrid
 switchport
 spanning-tree rstp edge-port
 vlt-peer-lag port-channel 13
 no shutdown
!
interface Port-channel 14
 no ip address
 portmode hybrid
 switchport
 spanning-tree rstp edge-port
 vlt-peer-lag port-channel 14
 no shutdown
!
interface Port-channel 15
```

```
 no ip address
 portmode hybrid
 switchport
 spanning-tree rstp edge-port
 vlt-peer-lag port-channel 15
 no shutdown
!
interface Port-channel 16
 no ip address
 portmode hybrid
 switchport
 spanning-tree rstp edge-port
 vlt-peer-lag port-channel 16
 no shutdown
!
interface Port-channel 17
 no ip address
 portmode hybrid
 switchport
 spanning-tree rstp edge-port
 vlt-peer-lag port-channel 17
 no shutdown
!
interface Port-channel 18
 no ip address
 portmode hybrid
 switchport
 spanning-tree rstp edge-port
 vlt-peer-lag port-channel 18
 no shutdown
!
interface Port-channel 19
 no ip address
 portmode hybrid
 switchport
 spanning-tree rstp edge-port
 vlt-peer-lag port-channel 19
 no shutdown
!
interface Port-channel 20
 no ip address
 portmode hybrid
 switchport
 spanning-tree rstp edge-port
 vlt-peer-lag port-channel 20
 no shutdown
!
interface Vlan 1
!untagged TenGigabitEthernet 1/46
!untagged Port-channel 1-2,6
!
interface Vlan 100
 ip address 192.168.110.2/24
 tagged TenGigabitEthernet 1/46
 tagged Port-channel 2,6
 untagged Port-channel 11-20
 no shutdown
!
interface Vlan 200
```

```
 ip address 192.168.125.1/24
 tagged Port-channel 6,11-20
 no shutdown
!
interface Vlan 300
 no ip address
 tagged TenGigabitEthernet 1/46
 tagged Port-channel 2,6,11-20
 no shutdown
!
interface Vlan 400
 no ip address
 tagged Port-channel 6,11-20
 no shutdown
!
interface Vlan 500
 no ip address
 tagged Port-channel 6,11-20
 no shutdown
!
management route 0.0.0.0/0 10.9.244.1
!
ip ssh server enable
!
line console 0
line vty 0
line vty 1
line vty 2
line vty 3
line vty 4
line vty 5
line vty 6
line vty 7
line vty 8
line vty 9
!
reload-type
 boot-type normal-reload
 config-scr-download enable
!
end
```

Dell Networking S3048-ON

```
! Version 9.9(0.0)
! Last configuration change at Tue Mar 29 12:00:22 2016 by default
! Startup-config last updated at Tue Mar 29 12:00:24 2016 by default
!
boot system stack-unit 1 primary system://A
boot system stack-unit 1 secondary system://B
!
hostname S3048
!
protocol lldp
 advertise dot1-tlv port-protocol-vlan-id
 advertise dot3-tlv max-frame-size
 advertise management-tlv system-description system-name
```

```
 advertise med
!
redundancy auto-synchronize full
!
username crowbar password 7 b273ab54c2a8d82d privilege 15
!
protocol spanning-tree rstp
 no disable
 bridge-priority 16384
!
stack-unit 1 provision S3048-ON
!
interface GigabitEthernet 1/1
 no ip address
 switchport
 spanning-tree rstp edge-port
 no shutdown
!
interface GigabitEthernet 1/2
 no ip address
 switchport
 spanning-tree rstp edge-port
 no shutdown
!
interface GigabitEthernet 1/3
 no ip address
 switchport
 spanning-tree rstp edge-port
 no shutdown
!
interface GigabitEthernet 1/4
 no ip address
 switchport
 spanning-tree rstp edge-port
 no shutdown
!
interface GigabitEthernet 1/5
 no ip address
 switchport
 spanning-tree rstp edge-port
 no shutdown
!
interface GigabitEthernet 1/6
 no ip address
 switchport
 spanning-tree rstp edge-port
 no shutdown
!
interface GigabitEthernet 1/7
 no ip address
 switchport
 spanning-tree rstp edge-port
 no shutdown
!
interface GigabitEthernet 1/8
 no ip address
 switchport
 spanning-tree rstp edge-port
 no shutdown
```

```
!
interface GigabitEthernet 1/9
 no ip address
 switchport
 spanning-tree rstp edge-port
 no shutdown
!
interface GigabitEthernet 1/10
 no ip address
 switchport
 spanning-tree rstp edge-port
 no shutdown
!
interface GigabitEthernet 1/11
 no ip address
 switchport
 spanning-tree rstp edge-port
 shutdown
!
interface GigabitEthernet 1/12
 no ip address
 switchport
 spanning-tree rstp edge-port
 shutdown
!
interface GigabitEthernet 1/13
 no ip address
 switchport
 spanning-tree rstp edge-port
 shutdown
!
interface GigabitEthernet 1/14
 no ip address
 switchport
 spanning-tree rstp edge-port
 shutdown
!
interface GigabitEthernet 1/15
 no ip address
 switchport
 spanning-tree rstp edge-port
 shutdown
!
interface GigabitEthernet 1/16
 no ip address
 switchport
 spanning-tree rstp edge-port
 shutdown
!
interface GigabitEthernet 1/17
 no ip address
 switchport
 spanning-tree rstp edge-port
 no shutdown
!
interface GigabitEthernet 1/18
 no ip address
 switchport
 spanning-tree rstp edge-port
```

```
 no shutdown
!
interface GigabitEthernet 1/19
 no ip address
 switchport
 spanning-tree rstp edge-port
 no shutdown
!
interface GigabitEthernet 1/20
 no ip address
 switchport
 spanning-tree rstp edge-port
 no shutdown
!
interface GigabitEthernet 1/21
 no ip address
 switchport
 spanning-tree rstp edge-port
 no shutdown
!
interface GigabitEthernet 1/22
 no ip address
 switchport
 spanning-tree rstp edge-port
 no shutdown
!
interface GigabitEthernet 1/23
 no ip address
 switchport
 spanning-tree rstp edge-port
 no shutdown
!
interface GigabitEthernet 1/24
 no ip address
 switchport
 spanning-tree rstp edge-port
 no shutdown
!
interface GigabitEthernet 1/25
 no ip address
 switchport
 spanning-tree rstp edge-port
 no shutdown
!
interface GigabitEthernet 1/26
 no ip address
 switchport
 spanning-tree rstp edge-port
 no shutdown
!
interface GigabitEthernet 1/27
 no ip address
 switchport
 spanning-tree rstp edge-port
 shutdown
!
interface GigabitEthernet 1/28
 no ip address
 switchport
```

```
 spanning-tree rstp edge-port
 shutdown
!
interface GigabitEthernet 1/29
 no ip address
 switchport
 spanning-tree rstp edge-port
 shutdown
!
interface GigabitEthernet 1/30
 no ip address
 switchport
 spanning-tree rstp edge-port
 shutdown
!
interface GigabitEthernet 1/31
 no ip address
 switchport
 spanning-tree rstp edge-port
 shutdown
!
interface GigabitEthernet 1/32
 no ip address
 switchport
 spanning-tree rstp edge-port
 shutdown
!
interface GigabitEthernet 1/33
 no ip address
 switchport
 spanning-tree rstp edge-port
 shutdown
!
interface GigabitEthernet 1/34
 no ip address
 switchport
 spanning-tree rstp edge-port
 no shutdown
!
interface GigabitEthernet 1/35
 no ip address
 switchport
 spanning-tree rstp edge-port
 shutdown
!
interface GigabitEthernet 1/36
 no ip address
 switchport
 spanning-tree rstp edge-port
 shutdown
!
interface GigabitEthernet 1/37
 no ip address
 switchport
 spanning-tree rstp edge-port
 shutdown
!
interface GigabitEthernet 1/38
 no ip address
```

```
 switchport
 spanning-tree rstp edge-port
 shutdown
!
interface GigabitEthernet 1/39
 no ip address
 switchport
 spanning-tree rstp edge-port
 shutdown
!
interface GigabitEthernet 1/40
 no ip address
 switchport
 spanning-tree rstp edge-port
 shutdown
!
interface GigabitEthernet 1/41
 no ip address
 switchport
 spanning-tree rstp edge-port
 shutdown
!
interface GigabitEthernet 1/42
 no ip address
 switchport
 spanning-tree rstp edge-port
 no shutdown
!
interface GigabitEthernet 1/43
 no ip address
 portmode hybrid
 switchport
 no shutdown
!
interface GigabitEthernet 1/44
 no ip address
 portmode hybrid
 switchport
 no shutdown
!
interface GigabitEthernet 1/45
 no ip address
 portmode hybrid
 switchport
 no shutdown
!
interface GigabitEthernet 1/46
 no ip address
 portmode hybrid
 switchport
 no shutdown
!
interface GigabitEthernet 1/47
 no ip address
 portmode hybrid
 switchport
 no shutdown
!
interface GigabitEthernet 1/48
```

```
 no ip address
 portmode hybrid
 switchport
 no shutdown
!
interface TenGigabitEthernet 1/49
 no ip address
 shutdown
!
interface TenGigabitEthernet 1/50
 no ip address
 portmode hybrid
 switchport
 no shutdown
!
interface TenGigabitEthernet 1/51
 no ip address
 no shutdown
!
interface TenGigabitEthernet 1/52
 no ip address
 no shutdown
!
interface ManagementEthernet 1/1
 ip address 10.9.244.40/23
 no shutdown
!
interface ManagementEthernet 2/1
 shutdown
!
interface ManagementEthernet 3/1
 shutdown
!
interface ManagementEthernet 4/1
 shutdown
!
interface ManagementEthernet 5/1
 shutdown
!
interface ManagementEthernet 6/1
 shutdown
!
interface Port-channel 2
 no ip address
 portmode hybrid
 switchport
 channel-member TenGigabitEthernet 1/51-1/52
 no shutdown
!
interface Vlan 1
!untagged GigabitEthernet 1/11-1/16,1/27-1/41
!untagged Port-channel 2
!
interface Vlan 100
 ip address 192.168.110.1/24
 tagged Port-channel 2
 untagged GigabitEthernet 1/1-1/10,1/17-1/26
 no shutdown
!
```

```
interface Vlan 300
 no ip address
 tagged Port-channel 2
 untagged GigabitEthernet 1/42-1/48
 untagged TenGigabitEthernet 1/50
 no shutdown
!
management route 0.0.0.0/0 10.9.244.1
!
ntp server 10.148.44.11 version 1
ntp source Vlan 300
ntp master 8
!
clock timezone central -6
!
ip ssh server enable
!
line console 0
line vty 0
line vty 1
line vty 2
line vty 3
line vty 4
line vty 5
line vty 6
line vty 7
line vty 8
line vty 9
!
reload-type
 boot-type normal-reload
 config-scr-download enable
!
end
```

Configuration : System

Dell PowerEdge System Firmware Settings

Setting		Node Role			
Name	Default Value	Solution Admin Host	Control Node	Compute	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 / Boot					
BootMode	UEFI	BIOS	BIOS	BIOS	BIOS
BootSeqRetry	Enabled				
DcuIpPrefetcher.Enabled					
DcuStreamerPrefetcher	Enabled				
DynamicCoreAllocation	Disabled				
IntegratedRaid	Enabled				
InternalSdCard	Off				
loatEngine	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

Dell PowerEdge Storage Volume Settings

Volume Purpose	Solution Admin Host	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	N/A	RAID 0 (3 @ 200GB SSD)
Data : Storage	N/A	N/A	N/A	RAID 0 (13 @ 4TB)

Configuration : Software

SUSE OpenStack Cloud Crowbar Batch Export

```
---
proposals:
- barclamp: crowbar
  attributes:
    instances:
      ganglia: []
      nagios: []
      provisioner:
      - "/var/lib/crowbar/config/provisioner.json"
      network:
      - "/var/lib/crowbar/config/network.json"
      ntp:
      - "/var/lib/crowbar/config/ntp.json"
      dns:
      - "/var/lib/crowbar/config/dns.json"
    realm: SUSE OpenStack Cloud Admin Server
    users:
      machine-install:
        password:
63132649d209a7ac1698bda451c3b076fd594b13d39532b9c6e6cfc309423919999a5d17027dda2f7e61c734d1b
51675030d0822c1fa01b775487327fd2ff077
      crowbar:
        password: cr0wBar!
  deployment:
    elements:
      crowbar:
```

```
      - admin.cloud.mr.rcbd.lab
  - barclamp: deployer
    attributes:
    deployment:
      elements:
        deployer-client:
        - admin.cloud.mr.rcbd.lab
        - "@@nova3@@"
        - "@@nova1@@"
        - "@@ceph2@@"
        - "@@ceph3@@"
        - "@@cntl1@@"
        - "@@ceph1@@"
        - "@@calamari@@"
        - "@@nova2@@"
        - "@@cntl2@@"
        - "@@cntl3@@"
        - d52-54-00-16-58-0c.cloud.mr.rcbd.lab
  - barclamp: provisioner
    attributes:
      timezone: America/Chicago
    deployment:
      elements:
        provisioner-server:
        - admin.cloud.mr.rcbd.lab
        provisioner-base:
        - admin.cloud.mr.rcbd.lab
        - "@@nova3@@"
        - "@@nova1@@"
        - "@@ceph3@@"
        - "@@cntl1@@"
        - "@@ceph2@@"
        - "@@ceph1@@"
        - "@@calamari@@"
        - "@@nova2@@"
        - "@@cntl2@@"
        - "@@cntl3@@"
        - d52-54-00-16-58-0c.cloud.mr.rcbd.lab
        provisioner-bootdisk-finder:
        - admin.cloud.mr.rcbd.lab
        - "@@calamari@@"
        - "@@cntl1@@"
        - "@@cntl2@@"
        - "@@cntl3@@"
        - "@@nova1@@"
        - "@@nova3@@"
        - "@@nova2@@"
        - "@@ceph2@@"
        - "@@ceph3@@"
        - "@@ceph1@@"
        - d52-54-00-16-58-0c.cloud.mr.rcbd.lab
  - barclamp: ipmi
    attributes:
      bmc_enable: true
      ignore_address_suggestions: true
    deployment:
      elements:
        ipmi-discover:
        - admin.cloud.mr.rcbd.lab
```

```
      - "@@nova3@@"
      - "@@nova1@@"
      - "@@ceph2@@"
      - "@@ceph3@@"
      - "@@cntl1@@"
      - "@@ceph1@@"
      - "@@calamari@@"
      - "@@nova2@@"
      - "@@cntl2@@"
      - "@@cntl3@@"
      - d52-54-00-16-58-0c.cloud.mr.rcbd.lab
      ipmi-configure:
      - admin.cloud.mr.rcbd.lab
      - "@@nova3@@"
      - "@@nova1@@"
      - "@@ceph3@@"
      - "@@cntl1@@"
      - "@@ceph2@@"
      - "@@ceph1@@"
      - "@@calamari@@"
      - "@@nova2@@"
      - "@@cntl2@@"
      - "@@cntl3@@"
      - d52-54-00-16-58-0c.cloud.mr.rcbd.lab
      bmc-nat-router:
      - admin.cloud.mr.rcbd.lab
      bmc-nat-client:
      - "@@nova3@@"
      - "@@nova1@@"
      - "@@ceph3@@"
      - "@@cntl1@@"
      - "@@ceph2@@"
      - "@@ceph1@@"
      - "@@calamari@@"
      - "@@nova2@@"
      - "@@cntl2@@"
      - "@@cntl3@@"
      - d52-54-00-16-58-0c.cloud.mr.rcbd.lab
  - barclamp: network
    attributes:
      mode: dual
      teaming:
        mode: 4
      conduit_map:
      - conduit_list:
          intf0:
            if_list:
            - 1g1
          intf1:
            if_list:
            - 10g1
            - 10g2
            - 10g3
            - 10g4
          intf2:
            if_list:
            - "?1g1"
            - 10g2
            - 10g3
```

```
      - 10g4
    pattern: dual/.*/.*
  networks:
    public:
      subnet: 10.148.45.0
      broadcast: 10.148.45.255
      router: 10.148.45.254
      ranges:
        host:
          start: 10.148.45.117
          end: 10.148.45.129
    nova_floating:
      subnet: 10.148.45.0
      netmask: 255.255.255.0
      broadcast: 10.148.45.255
      ranges:
        host:
          start: 10.148.45.231
          end: 10.148.45.250
    bmc:
      subnet: 192.168.110.0
      broadcast: 192.168.110.255
      ranges:
        host:
          start: 192.168.110.117
          end: 192.168.110.130
    bmc_vlan:
      subnet: 192.168.110.0
      broadcast: 192.168.110.255
      ranges:
        host:
          start: 192.168.110.161
          end: 192.168.110.161
    admin:
      subnet: 192.168.110.0
      broadcast: 192.168.110.255
      router: 192.168.110.9
      ranges:
        admin:
          start: 192.168.110.10
          end: 192.168.110.11
        dhcp:
          start: 192.168.110.21
          end: 192.168.110.80
        host:
          start: 192.168.110.81
          end: 192.168.110.109
        switch:
          start: 192.168.110.241
          end: 192.168.110.250
deployment:
  elements:
    switch_config:
    - admin.cloud.mr.rcbd.lab
    network:
    - admin.cloud.mr.rcbd.lab
    - "@@nova3@@"
    - "@@nova1@@"
    - "@@ceph3@@"
```

```
      - "@@cntl1@@"
      - "@@ceph2@@"
      - "@@ceph1@@"
      - "@@calamari@@"
      - "@@nova2@@"
      - "@@cntl2@@"
      - "@@cntl3@@"
      - d52-54-00-16-58-0c.cloud.mr.rcbd.lab
- barclamp: dns
  attributes:
    domain: cloud.mr.rcbd.lab
    forwarders:
    - 192.168.110.10
    - 10.148.44.11
    - 10.148.44.12
  deployment:
    elements:
      dns-server:
      - admin.cloud.mr.rcbd.lab
      dns-client:
      - admin.cloud.mr.rcbd.lab
      - "@@nova3@@"
      - "@@nova1@@"
      - "@@ceph3@@"
      - "@@cntl1@@"
      - "@@ceph2@@"
      - "@@ceph1@@"
      - "@@calamari@@"
      - "@@nova2@@"
      - "@@cntl2@@"
      - "@@cntl3@@"
      - d52-54-00-16-58-0c.cloud.mr.rcbd.lab
- barclamp: logging
  attributes:
  deployment:
    elements:
      logging-server:
      - admin.cloud.mr.rcbd.lab
      logging-client:
      - "@@nova3@@"
      - "@@nova1@@"
      - "@@ceph3@@"
      - "@@cntl1@@"
      - "@@ceph2@@"
      - "@@ceph1@@"
      - "@@calamari@@"
      - "@@nova2@@"
      - "@@cntl2@@"
      - "@@cntl3@@"
      - d52-54-00-16-58-0c.cloud.mr.rcbd.lab
- barclamp: ntp
  attributes:
    external_servers:
    - 192.168.110.10
    - 192.168.110.9
  deployment:
    elements:
      ntp-server:
      - admin.cloud.mr.rcbd.lab
```

```
    ntp-client:
    - "@@ceph1@@"
    - "@@calamari@@"
    - "@@nova2@@"
    - "@@cntl1@@"
    - "@@cntl2@@"
    - "@@cntl3@@"
    - "@@nova1@@"
    - "@@nova3@@"
    - "@@ceph2@@"
    - "@@ceph3@@"
    - d52-54-00-16-58-0c.cloud.mr.rcbd.lab
- barclamp: pacemaker
  name: control_ha
  attributes:
    corosync:
      password: cr0wBar!
    stonith:
      mode: ipmi_barclamp
      sbd:
        nodes:
          d24-6e-96-04-4d-d4.cloud.mr.rcbd.lab:
            devices:
            - ''
          d24-6e-96-04-3e-dc.cloud.mr.rcbd.lab:
            devices:
            - ''
          d24-6e-96-04-35-7c.cloud.mr.rcbd.lab:
            devices:
            - ''
          d24-6e-96-04-4e-7c.cloud.mr.rcbd.lab:
            devices:
            - ''
          d24-6e-96-04-53-34.cloud.mr.rcbd.lab:
            devices:
            - ''
          d24-6e-96-04-4d-ec.cloud.mr.rcbd.lab:
            devices:
            - ''
      per_node:
        nodes:
          d24-6e-96-04-4d-d4.cloud.mr.rcbd.lab:
            params: ''
          d24-6e-96-04-3e-dc.cloud.mr.rcbd.lab:
            params: ''
          d24-6e-96-04-35-7c.cloud.mr.rcbd.lab:
            params: ''
          d24-6e-96-04-4e-7c.cloud.mr.rcbd.lab:
            params: ''
          d24-6e-96-04-53-34.cloud.mr.rcbd.lab:
            params: ''
          d24-6e-96-04-4d-ec.cloud.mr.rcbd.lab:
            params: ''
    drbd:
      shared_secret: VRMIP8FzUmZl
  deployment:
    elements:
      pacemaker-cluster-member:
      - "@@cntl1@@"
```

```
      - "@@cntl2@@"
      - "@@cntl3@@"
      hawk-server:
      - "@@cntl1@@"
      - "@@cntl2@@"
      - "@@cntl3@@"
      pacemaker-remote:
      - "@@nova1@@"
      - "@@nova2@@"
      - "@@nova3@@"
# WARNING: no proposals exist for nfs_client barclamp
# WARNING: no proposals exist for suse_manager_client barclamp
- barclamp: database
  attributes:
    ha:
      storage:
        shared:
          device: 192.168.110.9:/exports/db
          fstype: nfs
          options: rw,async
  deployment:
    elements:
      database-server:
      - cluster:control_ha
- barclamp: rabbitmq
  attributes:
    password: u1ZXx53LaGG0
    ha:
      storage:
        shared:
          device: 192.168.110.9:/exports/mq
          fstype: nfs
          options: rw,async
    trove:
      password: 4KYDHLLsV759
  deployment:
    elements:
      rabbitmq-server:
      - cluster:control_ha
- barclamp: keystone
  attributes:
    database_instance: default
    rabbitmq_instance: default
    db:
      password: qFsbohVjzNLQ
    admin:
      password: cr0wBar!
    service:
      token: fq2XGefELwKm
    default:
      password: cr0wBar!
  deployment:
    elements:
      keystone-server:
      - cluster:control_ha
- barclamp: ceph
  attributes:
    config:
      fsid: 1a0d41d2-8da5-4707-8358-8e7590c27359
```

```
    keystone_instance: default
    osd:
      journal_devices:
      - "/dev/sdb"
      - "/dev/sdc"
      - "/dev/sdd"
  deployment:
    elements:
      ceph-calamari:
      - "@@calamari@@"
      ceph-mon:
      - "@@ceph1@@"
      - "@@ceph2@@"
      - "@@ceph3@@"
      ceph-osd:
      - "@@ceph1@@"
      - "@@ceph3@@"
      - "@@ceph2@@"
      ceph-radosgw: []
# WARNING: no proposals exist for swift barclamp
- barclamp: glance
  attributes:
    service_password: YVX90FCDwHQ5
    db:
      password: ZxUcOMHYBLO1
    default_store: rbd
    keystone_instance: default
    database_instance: default
    rabbitmq_instance: default
  deployment:
    elements:
      glance-server:
      - cluster:control_ha
- barclamp: cinder
  attributes:
    rabbitmq_instance: default
    keystone_instance: default
    glance_instance: default
    database_instance: default
    service_password: hFZnC9nZjPvn
    volumes:
    - backend_driver: rbd
      backend_name: rbd
      rbd:
        use_crowbar: true
        config_file: "/etc/ceph/ceph.conf"
        admin_keyring: "/etc/ceph/ceph.client.admin.keyring"
        pool: volumes
        user: cinder
        secret_uuid: 5bb235f0-3f1c-4ef7-93b3-e097327a03fb
    db:
      password: sjbP7obkYMDC
  deployment:
    elements:
      cinder-controller:
      - cluster:control_ha
      cinder-volume:
      - "@@ceph1@@"
      - "@@ceph3@@"
```

```
      - "@@ceph2@@"
  - barclamp: neutron
    attributes:
      service_password: 1aebrmMonkNr
      rabbitmq_instance: default
      keystone_instance: default
      database_instance: default
      db:
        password: rAyUclt15DPp
    deployment:
      elements:
        neutron-server:
        - cluster:control_ha
        neutron-network:
        - cluster:control_ha
  - barclamp: nova
    attributes:
      service_password: 8I0Z8vtia1Zo
      neutron_metadata_proxy_shared_secret: oQD0RFNvt0py
      database_instance: default
      rabbitmq_instance: default
      keystone_instance: default
      glance_instance: default
      cinder_instance: default
      neutron_instance: default
      itxt_instance: ''
      use_migration: true
      db:
        password: qJnYcwQEx2mC
    deployment:
      elements:
        nova-controller:
        - cluster:control_ha
        nova-compute-hyperv: []
        nova-compute-kvm:
        - remotes:control_ha
        nova-compute-qemu: []
        nova-compute-xen: []
  - barclamp: horizon
    attributes:
      nova_instance: default
      keystone_instance: default
      database_instance: default
      db:
        password: LV7EFYJ0VAH3
    deployment:
      elements:
        horizon-server:
        - cluster:control_ha
  - barclamp: heat
    attributes:
      rabbitmq_instance: default
      database_instance: default
      stack_domain_admin_password: UzlcWoXYMJ6H
      keystone_instance: default
      service_password: 5nxuxMvPhlgc
      auth_encryption_key: hVtFXlXxWNeDKGVKqa59DBCb8byA7BGPhnVO
      db:
        password: hWl11CBi6x02
```

```
  deployment:
    elements:
      heat-server:
      - cluster:control_ha
- barclamp: ceilometer
  attributes:
    metering_secret: fvwpJlsz1FyF
    rabbitmq_instance: default
    database_instance: default
    keystone_instance: default
    service_password: 5GAybxwyjH7J
    db:
      password: CD9HTcVgD3xl
  deployment:
    elements:
      ceilometer-agent:
      - "@@nova3@@"
      - "@@nova2@@"
      - "@@nova1@@"
      ceilometer-agent-hyperv: []
      ceilometer-polling:
      - cluster:control_ha
      ceilometer-server:
      - cluster:control_ha
      ceilometer-swift-proxy-middleware: []
# WARNING: no proposals exist for manila barclamp
# WARNING: no proposals exist for updater barclamp
# WARNING: no proposals exist for trove barclamp
- barclamp: tempest
  attributes:
    tempest_adm_password: mkc8OGrCj0IK
    tempest_adm_username: tempest-adm-rlUoWMnkyR1c
    tempest_user_password: pgCWCpnUM20Q
    tempest_user_username: tempest-user-xbG66FjsdvDF
    tempest_user_tenant: tempest-tenant-qXZOIXyABZs9
    nova_instance: default
  deployment:
    elements:
      tempest:
       - "@@cntl1@@"
```