

Dell OpenStack™-Powered Cloud Solution v1.6.1

A Dell Reference Architecture Guide

November 6, 2013

Contents

Contents	2
Tables	5
Figures	6
Trademarks	6
Terminology and Abbreviations	7
Overview	8
Approach.....	8
OpenStack Readiness	8
Taxonomy	8
Hardware Options	9
Networking and Network Services	9
Crowbar Management.....	10
OpenStack Architecture	11
Release Schedule.....	11
OpenStack Components.....	11
Server Infrastructure Options	14
PowerEdge C8000 Series	14
Hardware Configurations	15
Configuration Notes	16
PowerEdge R720(xd) Server	16
Hardware Configurations	17
Configuration Notes	18
PowerEdge C6220 Server	18
Hardware Configurations	19
Configuration Notes	19

Dell OpenStack™-Powered Cloud Solution Sizing Terms	19
Rack	19
Pod	20
Cluster	20
Cluster Sizing	20
Operational Notes	21
Backup/Recovery	21
Service Layout	21
Deployment	22
Dell OpenStack™-Powered Cloud Solution Hardware Monitoring and Alerting	23
Nagios	23
Ganglia	23
Dell OpenStack™-Powered Cloud Solution Network Architecture	24
Network Components	24
Server Nodes	24
Access Switch or Top of Rack (ToR)	24
Aggregation Switches	25
Core	25
Layer-2 and Layer-3	25
Single Rack Configuration	25
Multi-rack Configuration for 10GbE with Force10 S4810	28
Multi-rack Configuration for 10GbE Using Force10 Z9000 Switches	29
Management Network	29
Dell Open Switch Solution	29
IPv6 Capabilities	30
Network Connectivity	30
Appendix A : Physical Configuration – PowerEdge C8000 Series	32
Appendix B : Bill of Materials – PowerEdge C8000 Series	36
Appendix C : Physical Configuration – PowerEdge R720xd	43
Appendix D : Bill of Materials – PowerEdge R720xd Nodes	44

Appendix E :	Bill of Materials – PowerEdge R720 Nodes.....	45
Appendix F :	Bill of Materials – PowerEdge R720xd Data Node.....	46
Appendix G :	Physical Configuration – PowerEdge C6220	47
Appendix H :	Bill of Materials – PowerEdge C6220.....	49
Appendix I :	Bill of Materials – Force10 Network Equipment.....	51
Appendix J :	Bill of Materials – Dell 6248 Network Equipment	54
Appendix K :	Miscellaneous Equipment Notes.....	55
	Server Racks and Power.....	55
	Networking Equipment.....	55
Appendix L :	Bill of Materials – Software and Support	56
Appendix M :	JBOD versus single disk RAID 0 Configuration	57
Update History.....		58
	Changes in Version 1.6.1.....	58
Getting Help.....		59
	Contacting Dell.....	59
To Learn More.....		59

Tables

Table 1: Terminology and Abbreviations.....	7
Table 2: OpenStack Components.....	12
Table 3: Server Platform Attributes.....	14
Table 4: Hardware Configurations – PowerEdge C8000.....	15
Table 5: Hardware Configurations – PowerEdge C8000 Storage Sleds.....	16
Table 6: Hardware Configurations – PowerEdge R720xd.....	17
Table 7: Hardware Configurations – Mixed PowerEdge R720/R720xd.....	18
Table 8: Hardware Configurations—PowerEdge C6220 with PowerEdge R720.....	19
Table 9: Cluster Sizes by Server Model.....	20
Table 10: Node Type to Services.....	21
Table 11: Single Rack Network Equipment.....	26
Table 12: Multi Rack Network Equipment.....	28
Table 13: Chassis Configuration – PowerEdge C8000 Admin Chassis.....	32
Table 14: Chassis Configuration – PowerEdge C8000 Controller Chassis.....	32
Table 15: Chassis Configuration – PowerEdge C8000 Storage Chassis (Optional).....	32
Table 16: Chassis Configure - PowerEdge C8000 Compute Chassis.....	33
Table 17: Rack Configuration – PowerEdge C8000 (Maximum 60 Nodes).....	34
Table 18: Admin or Compute Chassis – PowerEdge C8000.....	36
Table 19: Controller Chassis – PowerEdge C8000.....	38
Table 20: Storage Node Chassis – PowerEdge C8000.....	41
Table 21: Rack Configuration – PowerEdge R720xd (or R720/R720xd).....	43
Table 22: Infrastructure and Nova-Compute Nodes – PowerEdge R720xd.....	44
Table 23: Infrastructure and Nova-Compute Nodes – PowerEdge R720.....	45
Table 24: Storage Node – PowerEdge R720xd.....	46
Table 25: Rack Configuration – PowerEdge C6220 (Maximum 60 Nodes).....	47
Table 26: Infrastructure and Nova-Compute Node – PowerEdge C6220.....	49
Table 27: Network Equipment – 1GbE and 10GbE – Dell Force10.....	51
Table 28: Network Equipment – 10GbE – Dell Force10.....	52
Table 29: Network Equipment – Dell 6248 (Optional).....	54
Table 30: Network Cables Required – 10GbE Configurations.....	55
Table 31: Network Cables Required – 1GbE Configurations.....	55

Figures

Figure 1: OpenStack Taxonomy	9
Figure 2: Crowbar Ops Management	10
Figure 3: OpenStack Architecture	11
Figure 4: PowerEdge C8000 Chassis	15
Figure 5: PowerEdge 720xd Server	17
Figure 6: PowerEdge C6220	18
Figure 7: Single Rack Networking Equipment	26
Figure 8: Multi-rack Networking Equipment	27
Figure 9: Multi-Rack View for 10G Servers Using Force10 s4810 Switches	28
Figure 10: Force10 s4810 Aggregation Cables	29
Figure 11: Multi-Rack View Using Force10 Z9000 Switches (Based on Layer-2)	29
Figure 12: Dell OpenStack™-Powered Cloud Compute PowerEdge R720xd Node Network Interconnects	30
Figure 13 Network Connections	31

Trademarks

This guide is for informational purposes only, and may contain typographical errors and technical inaccuracies. The content is provided as is, without express or implied warranties of any kind.

© 2011 – 2013 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell. Dell, the Dell logo, and the Dell badge, and PowerEdge are trademarks of Dell Inc.

Trademarks used in this text: Dell™, the DELL logo, Dell Precision™, OptiPlex™, Latitude™, PowerEdge™, PowerVault™, PowerConnect™, OpenManage™, EqualLogic™, KACE™, FlexAddress™, Force10™ and Vostro™ are trademarks of Dell Inc. Intel®, Pentium®, Xeon®, Core™ and Celeron® are registered trademarks of Intel Corporation in the U.S. and other countries. AMD® is a registered trademark and AMD Opteron™, AMD Phenom™, and AMD Sempron™ are trademarks of Advanced Micro Devices, Inc. Microsoft®, Windows®, Windows Server®, MS-DOS® and Windows Vista® are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Red Hat Enterprise Linux® and Enterprise Linux® are registered trademarks of Red Hat, Inc. in the United States and/or other countries. Novell® is a registered trademark and SUSE™ is a trademark of Novell Inc. in the United States and other countries. Oracle® is a registered trademark of Oracle Corporation and/or its affiliates. Citrix®, Xen®, XenServer® and XenMotion® are either registered trademarks or trademarks of Citrix Systems, Inc. in the United States and/or other countries. VMware®, Virtual SMP®, vMotion®, vCenter®, and vSphere® are registered trademarks or trademarks of VMware, Inc. in the United States or other countries.

Other trademarks and trade names may be used in this publication to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

Terminology and Abbreviations

Table 1: Terminology and Abbreviations

Term	Meaning	Substitute term
Admin	Initial server setup to manage/bootstrap other servers.	Crowbar
Barclamp	Software that evaluates the whole environment and proposes configurations, roles, and recipes that fit your infrastructure.	Crowbar Module
BMC	Baseboard management controller. An on-board microcontroller that monitors the system for critical events by communicating with various sensors on the system board and sends alerts and log events when certain parameters exceed their preset thresholds.	IPMI
CH1SL1 or CH2SL3	Notation used to indicate the chassis number from the bottom of the rack and the sled in the chassis. CH2SL3 is Chassis 2 Sled 3.	
Controller(s)	Infrastructure and management components installed in each chassis. Supports the following: Glance, Keystone, MySQL, Nova Multi Controller, Nova Volume, Swift Ring Compute, Swift Proxy.	
Crowbar	The code name for a bootstrap installer.	PXE Server
DevOps	An operational model for managing data centers using automated deployments.	Chef™, Puppet™
Glance	The OpenStack image cache.	
Hypervisor	Software that runs virtual machines (VMs).	KVM, Xen, VMware, HyperV
iDRAC	Baseboard management controller. An on-board microcontroller that monitors the system for critical events by communicating with various sensors on the system board and sends alerts and log events when certain parameters exceed their preset thresholds.	IPMI
LOM	LAN on motherboard.	
Node	One of the servers in the system. A single chassis (sometimes called a server) can have multiple nodes.	Host, Box, Unit
Nova	The OpenStack Compute module for virtual machine deployment.	EC2 API
Sled	A server that is part of a shared infrastructure chassis, such as a Dell™ PowerEdge™ C6220.	
Swift	A reference to OpenStack storage.	S3 API

Overview

Approach

This reference architecture focuses on helping organizations begin evaluating OpenStack™ software and/or create pilots. Dell can provide guidance for more sophisticated deployments; however, they are beyond the scope of this document. The expected focus for the Dell OpenStack™-Powered Cloud Solution encompasses software, hardware, operations, and integration.

This reference architecture advocates an operational approach based on highly automated solution deployments using the components of the Dell OpenStack-Powered Cloud Solution. We believe that this operational model (known as CloudOps and based on DevOps) is the best practice for both initial cloud evaluations and long-term maintenance of both moderate and hyperscale data centers.

The impact of CloudOps is that OpenStack solution deployments from the bare metal to the configuration of specific components can be completely scripted so that operators never configure individual servers. This highly automated methodology enables users to rapidly iterate through design and deployment options until the right model is determined. Once the architecture is finalized, the CloudOps model makes sure that the environment stays in constant compliance even as new hardware and software components are added.

OpenStack Readiness

The code base for OpenStack is evolving at a very rapid pace. The April 2013 OpenStack release is known as Grizzly. Crowbar ships with Grizzly Version 2013.1. See <https://wiki.openstack.org/wiki/ReleaseNotes/Grizzly>.

The seventh release of the OpenStack builds on previous releases, with more than 517 contributors and nearly 230 new features across the many shared services in the cloud platform. One of the most highly-anticipated features is OpenStack Networking (Neutron, previously known as Quantum). OpenStack Networking includes networking as a service, and enables advanced network automation. That automation enables users to control their networking technology of choice. <http://www.openstack.org/software/grizzly/press-release>

We designed this reference architecture to make it easy for Dell customers to use the current releases to build their own operational readiness and design their initial offerings.

Taxonomy

In the Grizzly design, the Dell OpenStack-Powered Cloud Solution contains the core components of a typical OpenStack solution (Compute (Nova), Object Storage (Swift), Block Storage (Cinder), Image Service (Glance), Authentication/Authorization (Keystone), Dashboard (Horizon), and Networking (Neutron), plus components that span the entire system (Crowbar, Chef, Nagios, DNS, NTP, etc.).

The taxonomy presented in Figure 1 reflects both included infrastructure components (shown in light green) and OpenStack-specific components that are under active development (shown in red) by the community, Dell, and Dell partners. The taxonomy reflects a CloudOps¹ perspective that there are two sides for cloud users: standards-based API (shown in pink) interactions and site-specific infrastructure. The standards-based APIs are the same between all OpenStack deployments and let customers and vendor ecosystems operate across multiple clouds. The site-specific infrastructure combines open and proprietary software, Dell hardware, and operational processes to deliver cloud resources as a service.

¹ For more information about CloudOps, please read the [CloudOps white paper](#) by Rob Hirschfeld.

The implementation choices for each cloud infrastructure are highly specific to the needs and requirements of each site. Many of these choices can be standardized and automated using the tools in this reference architecture (specifically Crowbar) and following the recommended CloudOps processes. Conforming to best practices helps reduce operational risk.

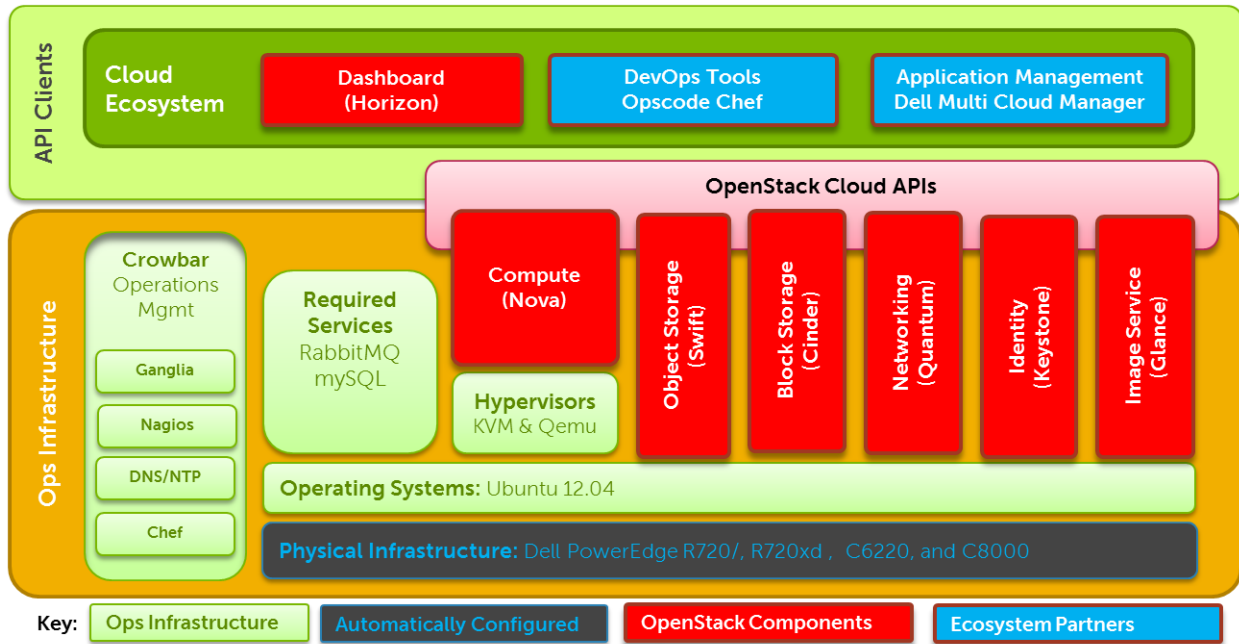


Figure 1: OpenStack Taxonomy

Hardware Options

To reduce time on hardware specification for a small system, this reference architecture offers specific choices for hardware and networking. For evaluations, the recommended hardware is general purpose and allows for a wide range of configuration options. For pilots, the recommended hardware has been optimized for infrastructure, compute, and storage roles. As noted throughout this reference architecture, we are constantly adding capabilities to expand this offering. We encourage you to discuss your plans with us to help us understand market drivers and expand the offering.

Each of the Dell™ PowerEdge™ C6220, C8000, R720, and R720xd server configurations in this reference architecture is designed as a getting-started setup for OpenStack compute, OpenStack storage, or both simultaneously. We recommend starting with OpenStack software using components from this configuration because the hardware and operations processes are a flexible foundation to expand upon. By design, you can repurpose the reference architecture configuration as your cloud deployment grows so your investment is protected.

Networking and Network Services

As a starter configuration, no core or layered networking is included in this reference architecture. Nothing in this reference architecture prevents the addition of these components as the system grows. Their omission is to reduce the initial complexity during evaluation. For a production system, additional networking and redundancy configurations are required. This includes 10Gbe networking, NIC teaming, and redundantly trunking top-of-rack (ToR) switches into core routers. While not documented in this reference architecture, these designs are available to customers using Dell consulting services.

To further simplify and speed deployments, our installer includes all the needed components to operate *without* external connectivity. These services include PXE, DHCP, DNS, and NTP. DNS and NTP services can be integrated into customer environments that already offer them; however, our installation relies on PXE and DHCP to perform discovery and provisioning.

Crowbar Management

Crowbar is a software framework that provides the foundation for the CloudOps approach articulated in this reference architecture. Initially, Crowbar manages the OpenStack deployment from the initial server boot to the configuration of Nova, Swift, and other OpenStack components. Once the initial deployment is complete, use Crowbar to maintain, expand, and architect the complete solution.

Note: Crowbar is open-source software (Apache 2 license) built upon open-source components. The most significant part is Opscode™ Chef Server™ that provides the deployment orchestration. Chef is a widely used DevOps platform with a library of installation recipes.

Crowbar provides a user interface (UI) and command-line view into the state of the nodes as they join the environment. Once the nodes have joined, use the API-based interfaces to assign that node to a service to provide a specific function. Crowbar has preconfigured automation that deploys OpenStack and its required services.

Crowbar provides a modular extensibility feature that lets individual operational components be managed independently. Each module, known as a barclamp, contains the configuration logic and the Chef Deployment recipes needed to integrate a cloud component, such as Swift, Nova, or DNS.

The three main aspects of Crowbar barclamps are:

- A RESTful API that all barclamps provide. These provide programmatic ways to manage the life cycle of the barclamps as well as the nodes running the functions of the barclamp.
- A simple command line interface for each barclamp. The command line wraps the API calls into commands that manipulate the barclamp without having to write API calls.
- A UI interface for each barclamp. These provide a more directed and controlled experience for manipulating the barclamp and its configuration.

These three interfaces are used to control and configure the running services in the cloud (for example, Swift or Nova, in addition to the base Crowbar components).

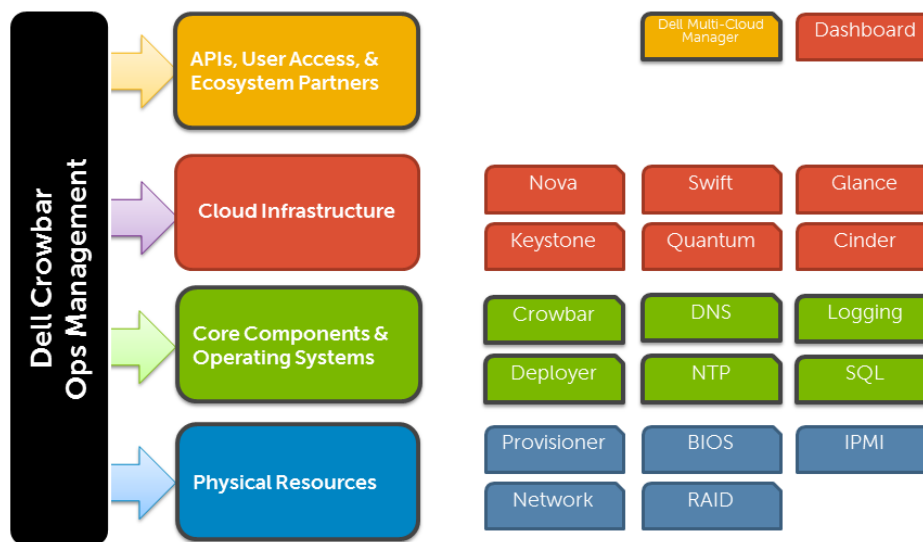


Figure 2: Crowbar Ops Management

Crowbar manages multiple layers of the operational environment. Barclamps can be applied at any layer. With the addition of importing barclamps, you can create your own and change the barclamp easily.

OpenStack Architecture

While OpenStack has many configurations and capabilities, Dell does not certify all configurations and options. This reference architecture is intended to specify which configurations are supported by Dell.

There are seven (7) primary components of OpenStack Grizzly:

- Compute (Nova)
- Object Storage (Swift)
- Block Storage (Cinder)
- Image Service (Glance)
- Authentication/Authorization (Keystone)
- Dashboard (Horizon)
- Networking (Neutron)

Note: For a complete overview of OpenStack software, visit www.OpenStack.org

Note: We highly recommend that you review the December 2011 update of the Dell white paper "Bootstrapping Open Source Clouds" as part of your preparation for deploying an OpenStack cloud infrastructure. This white paper is available at www.Dell.com/OpenStack

Release Schedule

Releases for OpenStack are named in alphabetical order on a six-month schedule. As of the publication of this reference architecture, Grizzly is the current stable OpenStack release. This release replaces the Folsom release.

The next release, Havana, will be delivered in Q4 2013. More information on OpenStack release schedules can be gathered at www.openstack.org.

OpenStack Components

Figure 3 shows Nova, Swift, and shared components for the OpenStack projects deployed by Crowbar. The yellow arrows indicate projects that provide an HTTP API or user interface (UI). This diagram shows the interconnections that are configured automatically. Users must adapt the placement of components to suit requirements and capabilities of each site.

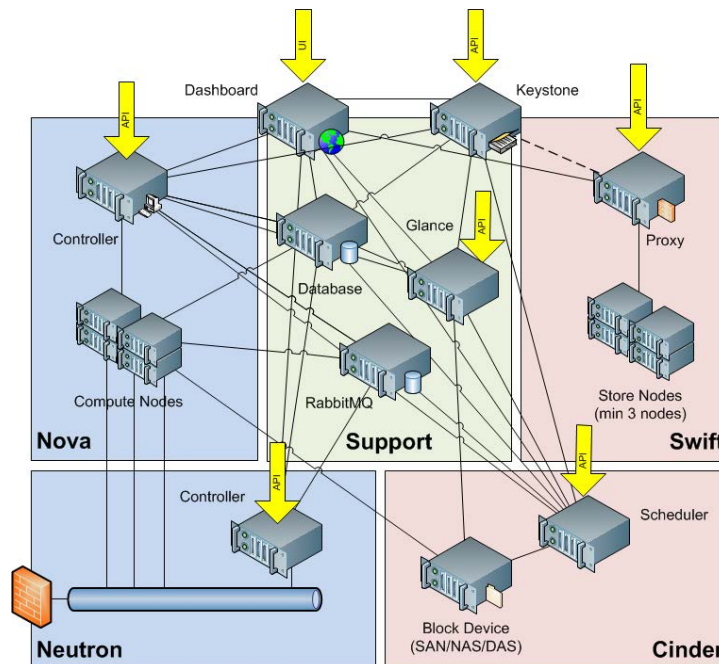


Figure 3: OpenStack Architecture

The following component descriptions are from the <http://OpenStack.org> site. Extensive documentation for the OpenStack components is available at <http://docs.openstack.org/>.

Table 2: OpenStack Components

Function	Code Name	Description
Authentication/Authorization	Keystone	http://www.openstack.org/software/openstack-shared-services/ Identity Service provides a central directory of users mapped to the OpenStack services they can access. It acts as a common authentication system across the cloud operating system and can integrate with existing backend directory services.
Dashboard/ Portal	Horizon	http://www.openstack.org/software/openstack-dashboard/ The OpenStack dashboard provides administrators and users a graphical interface to access, provision and automate cloud-based resources. The extensible design makes it easy to plug in and expose third party products and services.
Object Storage	Swift	http://www.openstack.org/software/openstack-storage/ OpenStack Object Storage (Swift) is open source software for creating redundant, scalable object storage using clusters of standardized servers to store petabytes of accessible data. It is not a file system or real-time data storage system, but rather a long-term storage system for a more permanent type of static data that can be retrieved, leveraged, and then updated if necessary. Primary examples of data that best fit this type of storage model are virtual machine images, photo storage, email storage, and backup archiving. Having no central "brain" or master point of control provides greater scalability, redundancy, and permanence. Objects are written to multiple hardware devices in the data center, with the OpenStack software responsible for ensuring data replication and integrity across the cluster. Storage clusters can scale horizontally by adding new nodes. Should a node fail, OpenStack works to replicate its content from other active nodes. Because OpenStack uses software logic to ensure data replication and distribution across different devices, inexpensive commodity hard drives and servers can be used in lieu of more expensive equipment.
Compute/IaaS	Nova	http://www.openstack.org/software/openstack-compute/ OpenStack Compute is open source software designed to provision and manage large networks of virtual machines, creating a redundant and scalable cloud computing platform. It gives you the software, control panels, and APIs required to orchestrate a cloud, including running instances, managing networks, and controlling access through users and projects. OpenStack Compute strives to be both hardware and hypervisor agnostic, currently supporting a variety of standard hardware configurations and seven major hypervisors.

Function	Code Name	Description
Virtual Images	Glance	http://www.openstack.org/software/openstack-shared-services/ OpenStack Image Service (Glance) provides discovery, registration, and delivery services for virtual disk images. The Image Service API server provides a standard REST interface for querying information about virtual disk images stored in a variety of back-end stores, including OpenStack Object Storage. Clients can register new virtual disk images with the Image Service, query for information on publicly available disk images, and use the Image Service's client library for streaming virtual disk images.
Block Storage	Cinder	http://www.openstack.org/software/openstack-storage/ OpenStack provides persistent block level storage devices for use with OpenStack compute instances. The block storage system manages the creation, attaching and detaching of the block devices to servers. Block storage volumes are fully integrated into OpenStack Compute and the Dashboard allowing for cloud users to manage their own storage needs. In addition to using simple Linux server storage, it has unified storage support for numerous storage devices. Block storage is appropriate for performance sensitive scenarios such as database storage, expandable file systems, or providing a server with access to raw block level storage. Snapshot management provides powerful functionality for backing up data stored on block storage volumes. Snapshots can be restored or used to create a new block storage volume.
Network	Neutron (formally known as Quantum)	http://www.openstack.org/software/openstack-networking/ OpenStack Networking is a pluggable, scalable and API-driven system for managing networks and IP addresses. Like other aspects of the cloud operating system, it can be used by administrators and users to increase the value of existing datacenter assets. OpenStack Networking ensures the network will not be the bottleneck or limiting factor in a cloud deployment and gives users real self-service, even over their network configurations.

Server Infrastructure Options

The Dell OpenStack™-Powered Cloud Solution includes three choices for server infrastructure:

- Dell™ PowerEdge™ C8000 series
- Dell™ PowerEdge™ R720/R720xd series
- Dell™ PowerEdge™ C6220 series

These alternatives provide density and capacity choices to match customer requirements. The appropriate choice of platform depends on the intended cluster usage and workload, cluster size, and the planned customer environment.

Table 3: Server Platform Attributes summarizes the high level attributed involved in a server platform choice.

Table 3: Server Platform Attributes		
Shared Infrastructure Platform	Customer Environment Attributes	Workload Attributes
R720/R720XD	<p>Choose if:</p> <ul style="list-style-type: none"> • Standardized on Monolithic • Rack Density 10-20 Server Per Rack • Power Per Rack < 10Kw • Standard Rack/Rear cabling 	<p>Choose if:</p> <ul style="list-style-type: none"> • Higher Frequency CPU • Require High Memory Density (768GB, 24 DIMMs) • Require High Spindle >12 x 2.5" Drives • Ideal for Small-Medium OpenStack Cluster
C8000 Series	<p>Choose if:</p> <ul style="list-style-type: none"> • Open to Shared Infrastructure • Rack Density 20+ Server Per Rack • Power Per Rack > 10Kw • Wide-Deep Rack/Front cabling 	<p>Choose if:</p> <ul style="list-style-type: none"> • Need High Spindle >12 x 3.5" Drives • Intend to run Multiple Server Types Per Chassis • Need Future Flexibility/Configuration • Ideal for Medium-Large OpenStack Cluster
C6220/R720xd	<p>Choose if:</p> <ul style="list-style-type: none"> • Open to Shared Infrastructure • Rack Density 20+ Server Per Rack • Power Per Rack < 7Kw • Standard Rack/Rear cabling 	<p>Choose if:</p> <ul style="list-style-type: none"> • Require Medium Memory Density (192GB, 12 DIMMs) • Require Medium Spindle <6 x 2.5" Drives Per Node • Ideal for Small-Medium OpenStack Cluster

The following sections describe the supported server models and configurations required. Detailed part lists and rack layouts are included in the appendices. The PowerEdge C8000 series and PowerEdge R720 series are recommended for new installations.

PowerEdge C8000 Series

The PowerEdge C8000 series is Dell's hyperscale-inspired 4U shared infrastructure server that allows the mixing and matching of compute, storage, and GPU sleds in one chassis. The PowerEdge C8000 chassis holds up to eight single-wide compute PowerEdge C8220 server sleds, up to four double-wide PowerEdge C8220X compute/GPU sleds, or PowerEdge C8000XD storage sleds, or a combination of these and two power sleds.

This design allows the right balance of CPU-to-memory-to-disk ratio and large-scale storage nodes requiring 24 or more hard drives to run big data applications faster. The flexible PowerEdge C8000 can run Nova-Compute, Nova-Controller and Storage nodes and multiple workloads from the same chassis or across racks, allowing for better use of IT resources, lower total cost of ownership over the lifecycle of the server, and more efficient use of space while increasing OpenStack POD compute/storage density and performance.



Figure 4: PowerEdge C8000 Chassis

- PowerEdge C8000 feature summary:
 - Up to eight independently serviceable PowerEdge C8220 compute sleds, four PowerEdge C8220X compute sleds, or four PowerEdge C8000XD storage sleds in a 4U rack chassis
 - Cold aisle service
 - Intel® E5-2600 series processors with up to eight cores and support for up to 130W TDP
 - Up to 256GB of memory with 16 DDR3 slots at 1600MHz per node (512GB RTS+)
- PowerEdge C8220 Single Width Compute (SWC)
 - Up to 2 x 2.5-inch non-hot-plug hard drives per PowerEdge C8220 compute sled
- PowerEdge C8220X Double Width Compute (DWC)
 - Up to 12 x 2.5-inch or 4 x 3.5-inch hot-plug hard drives per PowerEdge C8220X compute
 - Up to 2 x 2.5-inch non-hot-plug hard drives per PowerEdge C8220X compute
 - Up to 2 x 2.5-inch hot-plug hard drives per PowerEdge C8220X compute
- PowerEdge C8000XD Double Width Storage (DWS)
 - Up to 12 x 3.5-inch or 12 x 2.5-inch hot-plug hard drives or 24 x 2.5-inch SSDs per PowerEdge C8000XD storage sled

Hardware Configurations

Table 4: Hardware Configurations – PowerEdge C8000

Machine Function	Infrastructure Sled	Storage Sled
Sled 1	PowerEdge C8220X	PowerEdge C8220X
Processor	2 x E5-2670 (8-core)	2 x E5-2670 (8-core)
RAM (Minimum)	128 GB	64 GB
LOM	2 x 1GbE	2 x 1GbE
Network Controller	Intel X520 10GbE NIC, Dual Port, SFP+, Low Profile	Intel X520 10GbE NIC, Dual Port, SFP+, Low Profile
DISK (onboard)	None	None
DISK (hot-swap)	N/A	1 x 2.5-in. 1TB
DISK (side)	6 x 1 TB 2.5-in. SATA	4 x 3 TB 3.5-in. NL SAS

Machine Function	Infrastructure Sled	Storage Sled
DISK (expansion)	None	1 x C8000XD 36TB
Storage Controller	LSI 2008 (Mezzanine)	LSI 2008 (Mezzanine)
Storage Controller 2	None	LSI 9202 (PCI)
RAID	RAID 10	JBOD

Table 5: Hardware Configurations – PowerEdge C8000 Storage Sleds

Machine Function	Infrastructure Nodes	Storage Node
Sled 2	N/A	PowerEdge C8000XD
DISK	N/A	12 x 3 TB 3.5-in. Nearline SAS (NL-SAS)
Sled 3	N/A	PowerEdge C8000XD
DISK	N/A	12 x 3 TB 3.5-in. Nearline SAS (NL-SAS)

Configuration Notes

Appendix B : contains a complete bill of materials (BOM) listing for the C8000 server configurations.

Storage nodes are configured with the onboard chipset controller connected to the front hot-swap drives in the PowerEdge C8220X compute sled. The two 'rear' motherboard drives in the PowerEdge C8220X compute sled are not supported for any nodes.

Storage nodes require one PowerEdge C8220XD sled. Storage nodes can alternatively be configured with two PowerEdge C8220XD sleds, referred to as 'heavy' storage nodes.

Storage nodes use an LSI 9202 PCI HBA to connect to one or two PowerEdge C8220XD storage sleds. The connection requires one SAS extender cable per external sled.

The reference BOM's in the appendices are organized by chassis to simplify ordering. Some configurations may require sled blanks for empty slots; the reference BOMs in the appendices account for this.

A SAS extension cable is required for storage nodes, and connects from the compute sled to the storage sled. All required cables are included in the BOM listings.

The PowerEdge C8000 series is designed for cold-aisle service, with cabling in front of the chassis. Verify that rack configurations are compatible with this configuration. Be sure to consult your Dell account representative before changing the recommended disk sizes.

A minimum configuration can be implemented in three PowerEdge C8000 chassis, if one of the storage nodes is installed in the Controller chassis.

PowerEdge R720(xd) Server

The PowerEdge R720 and R720xd servers are Dell's 12G PowerEdge mainstream 2S 2U rack servers. They are designed to deliver the most competitive feature set, best performance, and best value. In this generation, Dell offers a large storage footprint, best-in-class I/O capabilities, and more advanced management features. The PowerEdge R720 and R720xd are technically similar except the R720xd has a backplane that can accommodate more drives (up to 24).



Figure 5: PowerEdge 720xd Server

PowerEdge R720xd feature summary:

- Intel® Romley platform and Intel® Xeon® E5-2600 processors
- 1600MHz DDR3
- Network daughter cards for customer choice of LOM speed, fabric, and brand at point of sale
- PCIe SSD in a front-accessible, hot-plug format
- Internal GPGPU support
- Intel® Node Manager power management technology
- Software RAID
- Platinum efficiency power supplies, common across 600 and 700 series platforms

Hardware Configurations

Table 6: Hardware Configurations – PowerEdge R720xd

Machine Function	Infrastructure Nodes	Storage Nodes
Platform	PowerEdge R720xd	PowerEdge R720xd
CPU	2 x E5-2640 (6-core)	2 x E5-2640 (6-core)
RAM (Minimum)	128 GB	64 GB
LOM	4 x 1GbE	4 x 1GbE
Add In Network	Intel X520 DP 10Gb DA/SFP+ (for 10GbE networking)	Intel X520 DP 10Gb DA/SFP+ (for 10GbE networking)
DISK	6 x 600-GB 10K SAS 2.5-inch	24 x 1-TB SATA 7.2K 2.5-inch
Storage Controller	PERC H710	PERC H710
RAID	RAID 10	Single Drive RAID 0

Notes:

- *Be sure to consult your Dell account representative before changing the recommended disk sizes.*

Table 7: Hardware Configurations – Mixed PowerEdge R720/R720xd

Machine Function	Infrastructure Nodes	Storage Nodes
Platform	PowerEdge R720	PowerEdge R720xd
CPU	2 x E5-2640 (6-core)	2 x E5-2640 (6-core)
RAM (Minimum)	128 GB	64 GB
LOM	4 x 1GbE	4 x 1GbE
DISK	6 x 600-GB 10K SAS 3.5-inch	24 x 1-TB SATA 7.2K 2.5-inch
Add In Network	Intel X520 DP 10Gb DA/SFP+ (for 10GbE networking)	Intel X520 DP 10Gb DA/SFP+ (for 10GbE networking)
Storage Controller	PERC H710	PERC H710
RAID	RAID 10	Single Drive RAID 0

Notes:

- Be sure to consult your Dell account representative before changing the recommended disk sizes.
- Refer to Appendix M : JBOD versus RAID 0 Configuration for more information.

Configuration Notes

Appendix D, Appendix E, and Appendix F contain the full bill of materials (BOM) listing for the PowerEdge R720 and R720Xd server configurations.

The R720 and R720xd configurations can be used with 10GbE networking. To use 10GbE networking support, an additional network card is required in each node. Refer to the BOM for the details on the supported card.

PowerEdge C6220 Server

The Dell PowerEdge C62220 is part of Dell's hyperscale inspired PowerEdge C server line designed to maximize compute power and minimize space and energy usage to lower operational costs. These servers have the right combination of what you need and mothering more.

These are purpose-built servers, designed for high performance computing, Web 2.0, data analytics, hosting and cloud building. They are best for rack deployments, large homogenous cloud/cluster application environments where the software stack provides primary platform availability and resiliency.



Figure 6: PowerEdge C6220

PowerEdge C6220 feature summary

- Up to 4 server nodes in 2U
- 2 x Intel Xeon E5-26000 Processor
- 16 x DDR3 RDIMM

Dell - Internal Use - Confidential - Customer Workproduct

- 24 x 2.5-inch or 12 x 3.5-inch HDD
- 2 x 1GbE Intel 82576

Hardware Configurations

Table 8: Hardware Configurations—PowerEdge C6220 with PowerEdge R720

Machine Function	Infrastructure Nodes	Storage Nodes
Platform	PowerEdge C6220	PowerEdge R720xd
CPU	2 x E5-2600 (6-core)	2 x E5-2640 (6-core)
RAM (Minimum)	128 GB	64 GB
LOM	4 x 1GbE	4 x 1GbE
Add In Network	Intel X520 DP 10Gb DA/SFP+ (for 10GbE networking)	Intel X520 DP 10Gb DA/SFP+ (for 10GbE networking)
DISK	6 x 600 SAS 10K 2.5-inch	24 x 1-TB SATA 7.2K 2.5-inch
Storage Controller	PERC H710	PERC H710
RAID	RAID 10	Single Drive RAID 0

Notes:

- *Be sure to consult your Dell account representative before changing the recommended disk sizes.*

Configuration Notes

Appendix H :contains the full bill of materials (BOM) for the PowerEdge C6220data infrastructure node configurations.

Appendix F :contains the BOM for the PowerEdge R720xd storage nodes that should be used with this configuration.

Configurations based on the C6220 use R720xd servers for the storage nodes, and C6220 for infrastructure nodes.

Dell OpenStack™-Powered Cloud Solution Sizing Terms

The Dell OpenStack™-Powered Cloud Solution Reference Architecture is organized into three components for sizing as the OpenStack environment grows. From smallest to largest they are:

- Rack
- Pod
- Cluster

Each has specific characteristics and sizing considerations documented in this reference architecture. The design goal for the OpenStack environment is to enable you to scale the environment by adding the additional capacity as needed, without the need to replace any existing components.

Rack

A rack is the smallest size designation for an OpenStack environment. A rack consists of all the necessary power, the network cabling, and the two Ethernet switches necessary to support up to 20 nodes. These nodes should utilize their own power connectivity and space within the data center, separate from other racks, and is treated as a fault zone.

Pod

A pod is an installation composed on three racks, based on server and network sizing. The three racks are capable of supporting enough OpenStack server nodes and network switches for a minimum commercial scale installation. In this reference architecture we discuss the administration and operational infrastructure to support three racks.

Cluster

A cluster is a set of racks dedicated to OpenStack that can be attached to a pair of distribution switches. It is a set of OpenStack nodes that share the same Controller Node and management tools for operating the OpenStack environment. The size of the cluster can vary depending on the capacity of the aggregation network. For example, a Dell™ Force10™ Z9000 aggregation switch can run a larger cluster than the Dell™ Force10™ s4810 switches.

Cluster Sizing

The minimum configuration supported is eight (8) nodes:

- Crowbar Administration Node
- OpenStack Controller Node – Supports Database, Neutron (Quantum), Nova-Controller, Keystone, Horizon, Glance, Cinder (optional), Swift-Proxy, and RabbitMQ
- Two (2) Nova-Compute Node – if no Cinder Node increase to three (3).
- Three (3) Swift Storage Nodes
- Cinder Node (optional)

Table 9: Cluster Sizes by Server Model

Server Model	Min Per Rack	Max Per Rack	Min per Pod	Max per Pod	Min per Cluster	Max per Cluster
R720 node	3	20	8	60	8	To be determined, based on sizing criteria
C8000 node	3	12	8	36	8	To be determined, based on sizing criteria
C6220 node	2 chassis (3 nodes)	10 chassis (20 nodes)	2 chassis (3 nodes)	30 chassis (60 nodes)	2 chassis (3 nodes)	To be determined, based on sizing criteria

Operational Notes

You can add new nodes at any time to expand the capacity of the cloud. The system is intended to use Crowbar to configure the services and the Nagios/Ganglia interfaces for monitoring.

Since this initial reference architecture is focused around exploration, Crowbar provides functions to reset or reinstall nodes to allow for trying various configurations or deployments.

Backup/Recovery

Since the system is designed for exploration that could later be extended to a production stage, backup and recovery have not been addressed in this configuration. The admin node, while not needed for normal operations of the services, is not redundant or backed-up.

Service Layout

During the deployment each service configured by the solution needs to be on a particular hardware type. For each server platform, two types of nodes have been designed: Infrastructure and Storage. The configurations in the Appendices have them separated that way. Crowbar is designed for flexibility, enabling you to try different configurations in order to find the optimal service placement for your workload. Table 10 is the recommended layout of each service.

Table 10: Node Type to Services

Hardware Type	Service	Node to Deploy Upon
Infrastructure	Crowbar/Ganglia/Nagios/DNS/NTP	Admin Node
Infrastructure	Database-server	Openstack Controller
Infrastructure	RabbitMQ-server	Openstack Controller
Infrastructure	Keystone-server	Openstack Controller
Infrastructure	Swift-ring-compute	Openstack Controller
Infrastructure	Swift-proxy	Openstack Controller
Infrastructure	Swift-dispersion	Openstack Controller
Storage	Swift-storage	1 or more Storage Nodes
Infrastructure	Cinder-scheduler	Openstack Controller
Storage	Cinder-volume (RAW mode)	1 or more Storage Node
Infrastructure	Cinder-volume (Local mode)	Openstack Controller or Unused Infrastructure nodes, configured with RAID10
EqualLogic Array	Cinder-Volume (EQLX mode)	See Sales Representative ¹
Infrastructure	Quantum-server (Neutron-server)	Openstack Controller
Infrastructure	Nova-multi-controller	Openstack Controller
Infrastructure	Nova-multi-compute	1 or more Compute Nodes
Infrastructure	Nova dashboard-server	Openstack Controller
Infrastructure	Tempest	Openstack Controller

¹ Cinder has multiple options for the volume type. The current version of the Dell OpenStack™-Powered Solution supports the Raw, Local and Eqlx (EqualLogic) types. Please contact your sales representative about the EqualLogic option.

The Cinder service can be deployed to one or more servers in Raw mode. The setup of Cinder in this mode creates a single Logical Volume across all unused drives in each individual server; and as such will have no data protection value. Alternatively, the service can be deployed to one or more server in Local mode; the service creates a logical volume on the existing local hard drive, and will inherit any data protection that drive already has. If data protection is required, Dell recommends that you use a device such as an EqualLogic array.

Deployment

Deployment consists of two phases:

- Admin node installation and configuration – a one-time process
- Cluster nodes installation and configuration – can be performed at will

The admin node installation and configuration phase installs the admin node with the components to run the Crowbar system. The initial installation is done either through a DVD installation or a network installation through a cross-connected laptop. Once installation is finished, the admin node must be configured and finalized via editing JSON files and running a script. Once this task is complete, the system is ready for the next phase.

- At this point, additional nodes may be added to the environment as needed. The general deployment model for the non-admin nodes:
 1. Unbox, rack, and cable the nodes.
 2. Turn on the nodes.
 3. Wait until the Crowbar UI reports that the nodes are in the Discovered state.
 4. Allocate the nodes.
 5. Wait until the Crowbar UI reports that the nodes are in the Allocated state.

The non-admin nodes are required to network boot. This is the default boot order configuration from the factory. Upon first network boot, the node PXE boots to the admin node, registers to the system, and receives a LiveCD image to make sure that the box is inventoried and able to run Linux. Once this is successfully executed, the node waits for the user to determine the use of the node in order to transition to the next state.

The node transitions into a hardware-installing state. At this point, the node will receive BIOS, BMC, and other hardware firmware updates as well as configuration for these components. Once the node has been successfully updated, the node reboots into an installing state.

During the installing state, the node receives a base image and prepares to have its configuration managed by the Chef server at the core of the Crowbar System. Upon rebooting, the node contacts the Chef server and finalizes its installation. Initially, a node receives minimal configuration but can be added to other applications as needed.

Once the node is ready, it can be consumed by other services in the cloud. For example, once a node is discovered, the system may decide that this node should be a Swift storage node. Once the node is installed, you can provide the additional configuration needed to make it part of the Swift system as well as other nodes that need to know about the new node. All of this process is controlled by the various barclamp-based applications in the cloud.

Dell OpenStack™-Powered Cloud Solution Hardware Monitoring and Alerting

To automate the alert and response to unexpected events and failures within the Dell OpenStack™-Powered Cloud Solution, the software stack includes Nagios and Ganglia. The Dell OpenStack™-Powered Cloud Solution includes capabilities for three primary components of the monitoring environment:

Monitoring of cluster activities — The Dell OpenStack™-Powered Cloud Solution utilizes Nagios to monitor the cluster, including hardware, software, and users. The Nagios deployment, as part of the Dell OpenStack™-Powered Cloud Solution, keeps historical information regarding system availability, maintenance, and failure events.

- Alerts on unexpected events — The Dell OpenStack™-Powered Cloud Solution utilizes Nagios to alert system operations staff to events that occur that deviate from normal operation, if the administrator has designated them for notification.
- Debugging of cluster runtime operations — The Dell OpenStack™-Powered Cloud Solution utilizes Ganglia to provide the users and administrators of the OpenStack environment with the necessary tools for tracking, debugging, and monitoring job performance and characteristics.

The Dell OpenStack™-Powered Cloud Solution is designed to include the necessary components to monitor and respond to events in your OpenStack environment. It is flexible enough to allow integration with existing operations management frameworks in your environment.

The monitoring components of the Dell OpenStack™-Powered Cloud Solution Reference Architecture are designed to be proactive in nature; they alert the IT operations team when a failure in the environment occurs, and they do so before the failure causes an outage that affects product workloads and users.

Nagios

Nagios is an open source solution for enterprise monitoring. Its pluggable architecture allows for consistent event handling. It supports a wide variety of sensors, plug-ins, applications, servers, and hardware platforms.

The Dell OpenStack™-Powered Cloud Solution includes Nagios as part of all default installations. The Dell OpenStack™-Powered Cloud Solution will automatically install the Nagios console and the necessary Nagios plug-ins for monitoring the OpenStack cluster, including processes, operating systems, and physical servers.

See Nagios: <http://www.nagios.org> for more information about Nagios.

Ganglia

Ganglia is a scalable, distributed monitoring system for high-performance computing systems, such as clusters and grids. It is based on a hierarchical design targeted at federations of clusters. It leverages widely used technologies, using carefully engineered data structures and algorithms to achieve very low per-node overheads and high concurrency. The implementation is robust. It has been ported to an extensive set of operating systems and processor architectures used to link clusters across university campuses and around the world. It can scale to handle clusters with 2,000 nodes.

The Dell OpenStack™-Powered Cloud Solution automates the installation and configuration of Ganglia within the OpenStack cluster, enabling IT operations staff to have detailed reporting on the status and utilization of all OpenStack nodes.

See <http://ganglia.sourceforge.net/> for more information about Ganglia.

Dell OpenStack™-Powered Cloud Solution Network Architecture

The Dell OpenStack™-Powered Cloud Solution uses Dell™ Force10™ S60 Gigabit or Dell Force10 S4810 Ten-Gigabit Ethernet switches as the top-of-rack connectivity to all OpenStack-related nodes. This reference architecture is used to support consistency in rapid deployments through the minimal differences in the network configuration.

This reference architecture implements at a minimum five (5) distinct, separate VLANs:

- **Admin Cluster Data LAN**—connects the compute node NICs into the fabric used for sharing data and distributing work tasks among compute nodes.
- **Cluster Management LAN**—connects all the iDRAC/BMCs in the cluster nodes.
- **Nova-fixed** – sets up the initial backend network for nova
- **OS_SDN** – Openstack Software Defined Network, used for creating GRE tunnels between the controller node and the nova-compute nodes.
- **OpenStack Cluster Public LAN**—connects the cluster to the outside world.
 - **Nova Floating subnet** of the public LAN

The network consists of three major network infrastructure layouts:

- **Data network infrastructure**—the data network consists of the server NICs, the top-of-rack (ToR) switches, and the aggregation switches.
- **Management network infrastructure**—the BMC management network, consisting of iDRAC ports and the out-of-band management ports of the switches, is aggregated into a 1-RU s55 switch in one of the three racks in the POD. This 1-RU switch in turn can connect to one of the Aggregation or Core switches to create a separate network with a separate VLAN.
- **Core network infrastructure**—the connectivity of aggregation switches to the core for external connectivity.

Network Components

The data network is primarily composed of the ToR and the aggregation switches. Configurations for 1GbE and 10GbE are included in this reference architecture. The following component blocks make up this network:

Server Nodes

Server connections to the network switches could be one of four possible configurations:

- Active-Active LAG in load-balance bond formation
- Active-Backup in failover/failback formation
- Active-Active round robin based on gratuitous ARP
- Single port

In the first case the connectivity on the switch side must be in a LAG (or port-channel). In cases B and C, we recommend that you do the configuration as a LAG but the ports should still be part of the same layer-2 domain. In some cases all members of the LAG connect to a single ToR switch. In others the LAG splits into two ToR switches. This is an optional setup as OpenStack has redundancy built into the application, and high-availability is not compromised by connecting into a single switch.

The teaming configuration that Dell recommends is transmit-lb (mode = 5). This configuration setting is explained in greater detail in the *Dell Openstack Solution Deployment Guide*. Please contact your sales representative for a copy of the deployment guide. The Dell Crowbar deployment software automatically configures this setting for OpenStack environments.

Access Switch or Top of Rack (ToR)

The servers connect to ToR switches. Typically there are two in each rack. The switches recommended by Dell are the Force10 S60 for 1GbE connectivity and S4810 for 10G servers. The PowerEdge C8000 requires the 10GbE option. PowerEdge R720 and C6220 configurations can use 1GbE or 10GbE.

The Force10 S60 ToR switches stack together in the same rack for 1GbE. This is useful in managing the two switches as a single unit and allowing the servers to connect into two different switches for redundancy. The ToR switches each have two expansion slots that can accept a two-port 10G module or a two-port stacking

module. This architecture recommends one of each type in the two slots. The 10GbE module would be used to connect into the pod-interconnect switches, one port to each switch, forming a LAG. The stacking module connects the switches together as a single unit. The uplinks to the aggregation pair would be a single LAG of four 10GbE ports, two from each switch. Each rack connects to the pod-interconnect independently; thereby scaling is easier.

For the 10GbE configuration, the ToR switches are Force10 S4810, and we recommend this pair of switches run a high availability feature called the Virtual Link Trunking (VLT). This feature allows the servers to terminate their LAG interfaces into two different switches instead of one. This allows HA as well as active-active bandwidth utilization. This feature gives redundancy within the rack if one switch fails or needs maintenance. The uplink to the aggregation pair is 80 Gb, using a LAG from each ToR switch. This is achieved using two 40G interfaces in a LAG connecting to the aggregation pair. Therefore, from each rack there is a collective bandwidth of 160G available.

Each rack is managed as a separate entity from a switching perspective, and ToR switches connect only to the aggregation switches.

Aggregation Switches

For a medium-scale deployment of one to three PODs of 1G server (12 racks max) the Dell Force10 S4810 is the recommended aggregation switch. It is both 10GbE and 40GbE capable. The 40GbE interfaces on the S4810 could be converted into four 10GbE, thereby converting this switch into 64 10GbE-capable ports. This potentially scales OpenStack deployments into tens of nodes. ToR switches connect to aggregate switches via uplinks of 10GbE interfaces from the ToR Force10 S60 to the Force10 S4810.

The recommended architecture uses Virtual Link Trunking (VLT) between the two Force10 S4810 switches in aggregation. This feature enables a multi-chassis LAG from the stacked ToR switches in each rack. The stacks in each rack would divide their links between this pair for switches to achieve the powerful capability of active-active forwarding while using full bandwidth capability, in absence of any requirement for spanning tree. Running 40GbE Ethernet switches like the Dell™ Force10™ Z9000 in aggregation can achieve a scale of up to hundreds of 1G deployed nodes.

For the 10G server deployment, depending on the scale at which the PODs are planned and depending on the how much future scale is required, we recommend the Force10 S4810 for aggregation for smaller scale and the Force10 Z9000 for larger deployment. The Force10 Z9000 is a 32-port, 40G high-capacity switch. It can aggregate up to 15 racks of high-density PowerEdge C8000 servers. The rack-to-rack bandwidth needed in OpenStack would be most suitably handled by a 40G-capable, non-blocking switch. The Force10 Z9000 can provide a cumulative bandwidth of 1.5TB of throughput at line-rate traffic from every port.

Core

The aggregation layer could itself be the network core in many cases but where it's not, it would connect to a larger core, which is represented by the cloud in Figure 8. Details on this topic are beyond the scope of this document.

Layer-2 and Layer-3

The layer-2 and layer-3 boundaries are separated at the aggregation layer. The colors blue and red in Figure 8 represent the layer-2 and layer-3 boundaries. The Reference Architecture uses layer-2 as the reference up to the aggregation layer. That is why VLT is used on the aggregation switches. Crowbar requires a layer-2 domain in order to provision servers.

Single Rack Configuration

Figure 7 shows the single rack equipment. Dell recommends using Force10 S60 ToR switches in the rack. Each rack could have a maximum of 20 servers in some configurations, while a dense packing of sleds in the C-series server chassis can hold even more. Each rack has two ToR Force10 S60 switches that are stacked, and this stack connects to the two Force10 S4810 aggregation switches. The Force10 S60 stack offers a single switch view to the servers. Each node can have up to four 1G NIC ports. It forms a LAG of two ports with one port on each switch in the stack. The LAG of 2GbE offers a switch redundancy within the rack and enables high availability.

This recommendation is for 1GbE only. For 10GbE installations use the Force 10 S4810 as the TOR, and single Force 10 S55 or S60 switch for the BMC management. See Multi-rack Configuration for 10GbE with Force10 S4810.

Table 11: Single Rack Network Equipment

Total Racks	1 (6-20 nodes)
Top-of-rack switch	2 Force10 S60 (2 per rack)
Aggregation switch	Not needed for a single rack
Server	2RU PowerEdge R720/R720xd/C6220/C8000
Over-subscription at ToR	1:1
Modules in each ToR	1x 12-2port Stacking, 1x 10G -2 port uplink




Figure 7: Single Rack Networking Equipment

Figure 8 shows the Force10 S4810 switch aggregating the pods to enable inter-rack traffic and the management network. There are multiple VLANs for data, VM, and management. All port-channels on the Force10 S4810 and ToR are tagged in these VLANs.




Figure 8: Multi-rack Networking Equipment

Table 12: Multi Rack Network Equipment shows the network inventory details in a cluster of three racks.

Table 12: Multi Rack Network Equipment	
Total racks	3 (15-20 nodes per rack)
Top-of-rack switch	6 Force10 S60 (2 per rack)
Pod-interconnect switch	2 Force10 S4810
Server	PowerEdge R720/R720xd/C6220/C8000C
Over-subscription at ToR	1:1
Modules in each ToR	1x 12-2port Stacking, 1x 10G -2 port uplink

Multi-rack Configuration for 10GbE with Force10 S4810

In this reference architecture we define a 10G solution with the PowerEdge C6220, C8000 and R720X servers. OpenStack applications are increasingly being deployed on 10GbE servers for the scale and price advantages they bring. That brings about an enormous economy of scale in the usage of hardware. That in turn needs 10GbE switches in the racks. This can be achieved using the Force10 S4810 as a top-of-rack switch and the option of using Force10 S4810 or Z9000, the 10G/40G high-density switch in the aggregation. The scale that is achieved by that configuration can grow into thousands of nodes using a CLOS architecture, which was used in the 1GbE solution above. Running 40GbE switches like the Force10 Z9000 in aggregation can achieve a scale of hundreds of nodes using high-density data center class switches.




Figure 9: Multi-Rack View for 10G Servers Using Force10 s4810 Switches

In Figure 9 we see that each rack with a pair of switches aggregates into the pair of Force10 S4810 switches. This connection should be based on layer-2, with the aggregation running on the VLT feature. All links utilize the full bandwidth on all links due to load-balancing. The VLT scenario creates a complete layer-2 domain between all racks in the cluster.

The Force10 s4810-based aggregation design is preferred for lower cost and medium scalability. This design can handle up to six racks or two PODs. In this design, the ToR Force10 s4810 uplinks using its 40G interface in quad-mode, where each 40G interface runs as 4x 10G. Figure 10: Force10 s4810 Aggregation Cables shows the cables needed for this design. Using a passive copper break-out cable eliminates the need for any QSFP+ or SFP+ since these are built into the copper twin-ax cable. This reduces the cost compared to a pure fiber option, where there is a need for a QSFP+ optic; four SFP+ optics; plus the break-out fiber. The benefit of the fiber option is the longer reach achieved with it compared to twin-ax, which is limited to 5 meters.




Figure 10: Force10 s4810 Aggregation Cables

Multi-rack Configuration for 10GbE Using Force10 Z9000 Switches

For a scale-out version of the deployment that is looking to expand its OpenStack environment into a larger setup or needs the OpenStack cluster co-located with other applications in different racks, the recommended option is the Force10 Z9000 core switch. The Force10 Z9000 does not need to connect into any other higher-tier core switches as the capacity is enough for a data center with hundreds of servers.

Each switch should form a Layer-2 LAG as shown in Figure 10. This assumes that the Z9000 pair in the aggregation forms a VLT pair for HA. Now we have 2 tiers of VLT, one forming at the ToR for servers and another at the aggregation for the top of rack switches.

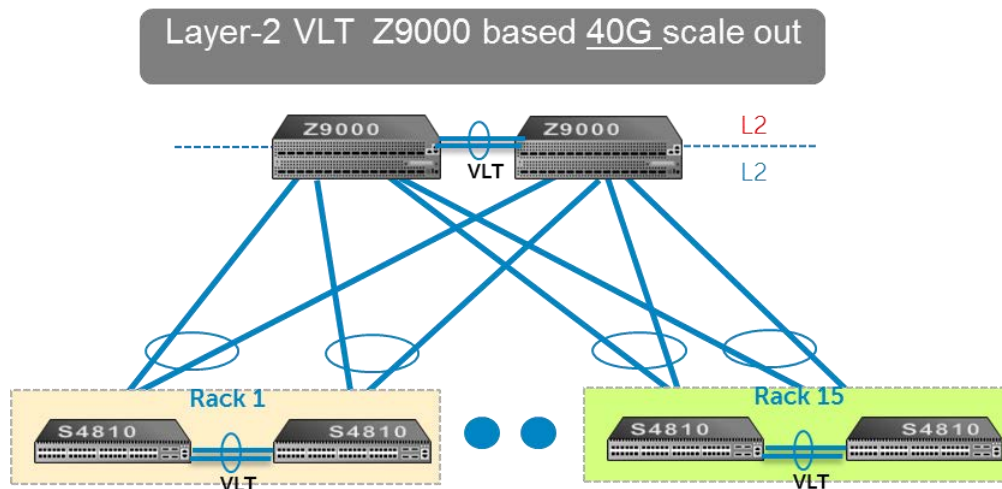


Figure 11: Multi-Rack View Using Force10 Z9000 Switches (Based on Layer-2)

In Figure 11 we see an example of a CLOS fabric that grows horizontally. This technique of network fabric deployment has been used by some of the largest Web 2.0 companies, whose businesses range from social media to public cloud in their data centers. Dell can help an enterprise solve its Cloud needs with a scalable end-to-end solution.

Management Network

The management network of all the servers and switches is aggregated into a Dell™ Force10™ S55 switch that is located in each rack of the POD. It uplinks on a 10G link to the S4810 switches.

Dell Open Switch Solution

In addition to the Dell switch-based reference architecture, Dell provides an open standard that allows you to choose other brands and configurations of switches for your OpenStack environment. The following list of requirements will enable other brands of switches to properly operate with the tools and configurations in the Dell OpenStack™-Powered Cloud Solution Reference Architecture:

- Support for IEEE 802.1Q VLAN traffic and port tagging

- Support using one untagged and multiple tagged VLANs on the same port
- Ability to provide a minimum of 170 Gigabit Ethernet ports in a non-blocking configuration within VLAN 100
 - Configuration can be a single switch or a combination of stacked switches to meet the additional requirements
- The ability to create link aggregation groups (LAGs) with a minimum of two physical links in each LAG
- If multiple switches are stacked:
 - The ability to create a LAG across stacked switches
 - Full-bisection bandwidth
 - Support for VLANs to be available across all switches in the stack
- The ability to provide a minimum 65 10/100 Ethernet ports on the untagged VLAN
- 250,000 packets-per-second capability per switch
- The ability to provide 12 10Gb ports for redundant uplinks contained in VLAN 10
- A managed switch that supports SSH and serial line configuration
- SNMP v3 support

IPv6 Capabilities

At this time, the Dell OpenStack™-Powered Cloud Solution does not support or allow for the use of IPv6 for network connectivity. All deployments are configured by Crowbar based on IPv4, with IPv6 explicitly disabled on all nodes within the OpenStack environment.

Network Connectivity

The network interconnects between various hardware components of the OpenStack solution are depicted in Figure 12 and Figure 13. For more information, please see the *Dell OpenStack™-Powered Cloud Solution Deployment Guide*.




Figure 12: Dell OpenStack™-Powered Cloud Compute PowerEdge R720xd Node Network Interconnects

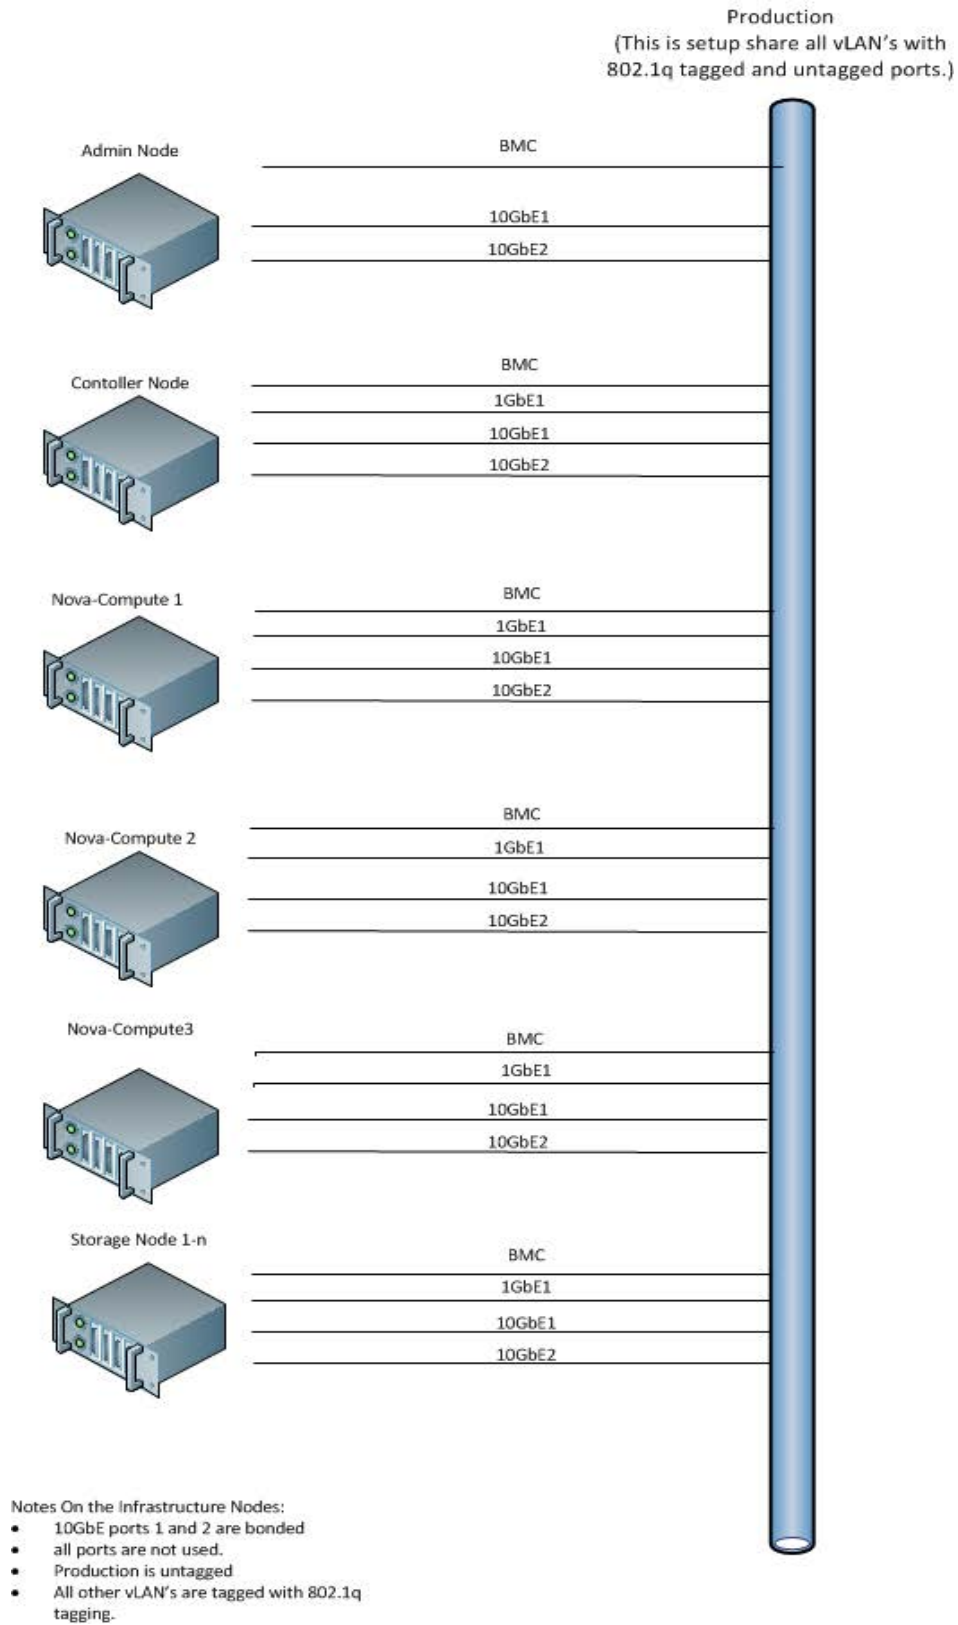


Figure 13 Network Connections

Appendix A : Physical Configuration – PowerEdge C8000 Series

C8000 Chassis Configuration

Table 13: Chassis Configuration – PowerEdge C8000 Admin Chassis

C8220X DWC (Admin)	C8220X DWC (Compute)	Power	Power	C8220X DWC (Cinder)	Empty	Empty
--------------------------	----------------------------	-------	-------	---------------------------	-------	-------

Refer to Table 18 in Appendix B : for the bill of materials for this chassis.

Table 14: Chassis Configuration – PowerEdge C8000 Controller Chassis

C8220X DWC (Controller)	C8220X DWC (Compute)	Power	Power	C8220X DWC (Storage)	C8220XD DWS
-------------------------------	----------------------------	-------	-------	----------------------------	----------------

←→

Refer to Table 19 in Appendix B : for the bill of materials for this chassis

Table 15: Chassis Configuration – PowerEdge C8000 Storage Chassis (Optional)

C8220X DWC (Storage)	C8220XD DWS	Power	Power	C8220X DWC (Storage)	C8220XD DWS
----------------------------	----------------	-------	-------	----------------------------	----------------

←→ ←→

Refer to Table 20 in Appendix B : for the bill of materials for this chassis

Table 16: Chassis Configure - PowerEdge C8000 Compute Chassis

C8220X DWC (Compute)	C8220X DWC (Compute)	Power	Power	C8220X DWC (Compute)	C8220X DWC (Compute)
----------------------------	----------------------------	-------	-------	----------------------------	----------------------------

Table 17: Rack Configuration – PowerEdge C8000 (Maximum 60 Nodes)

RU	RACK1	RACK2	RACK3
42	R1- Switch 2: 10Gb S4810	R2- Switch2: 10Gb S4810	R3- Switch2: 10Gb S4810
41	R1- Switch 1: 10Gb S4810	R2- Switch1: 10Gb S4810	R3- Switch1: 10Gb S4810
40	Cable Management	Cable Management	Cable Management
39	Cable Management	Cable Management	Cable Management
38	Admin Chassis	Controller Chassis	R3 - Switch 1: Force10 S4810 (1 RU) OR Force10 Z9000 (2 RU)
37			R3 - Switch 1: Force10 S4810 (1 RU) OR Force10 Z9000 (2 RU)
36			
35			
34	Cable Management	Cable Management	Cable Management
33	Cable Management	Cable Management	Cable Management
32	R1- Chassis06: Controller or Compute or Storage Chassis	Empty	Empty
31			
30			
29			
28	R1 - S55 iDRAC Mgmt switch	R2 - S55 iDRAC Mgmt switch	R3 - S55 iDRAC Mgmt switch
27- 21	Empty	Empty	Empty
20	R1- Chassis05: Storage Chassis	R2- Chassis05: Storage Chassis	R3- Chassis05: Storage Chassis
19			
18			
17			
16	R1- Chassis04: Compute or Storage Chassis	R2- Chassis04: Compute or Storage Chassis	R3- Chassis04: Compute or Storage Chassis
15			
14			
13			
12	R1- Chassis03: Compute or Storage Chassis	R2- Chassis03: Compute or Storage Chassis	R3- Chassis03: Compute or Storage Chassis
11			
10			
9			
8	R1- Chassis02 : Compute or Storage Chassis	R2- Chassis02: Compute or Storage Chassis	R3- Chassis02: Compute or Storage Chassis
7			
6			
5			
4	R1- Chassis01: Compute or Storage Chassis	R2- Chassis01: Compute or Storage Chassis	R3- Chassis01: Compute or Storage Chassis
3			
2			

1			
---	--	--	--

Appendix B : Bill of Materials – PowerEdge C8000 Series

For the PowerEdge C8000 series, the bill of materials is organized by chassis rather than node, to simplify ordering.

Table 18: Admin or Compute Chassis – PowerEdge C8000

The admin chassis includes the Administration Node and Nova-Compute Nodes. The compute chassis consists of Group 2 quantity 4 and can be used to have 4 Nova-Compute nodes.

SKU	Component
Group: 1	Quantity: 1
225-3550	PE C8000 Enclosure, Dual Power Supply
331-8341	PowerEdge C8000 Shipping
331-9573	SHIP,C8000,DAO
420-3323	No Factory Installed Operating System
318-2363	PowerEdge C8000 Sled Blank, Single Width - Quantity 2
330-7353	Power Cord, C13 to C14, PDU Style, 12 Amps, 2 meter, Qty 1 - Quantity 4
331-8218	PowerEdge C8000 Static Rails, Tool-less
936-3965	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, Initial Year
936-4695	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, 2 Year Extended
936-4705	ProSupport: 7x24 HW / SW Tech Support and Assistance, 3 Year
936-6035	Dell Hardware Limited Warranty Plus On Site Service Initial Year
936-6145	Dell Hardware Limited Warranty Plus On Site Service Extended Year
989-3439	Dell ProSupport. For tech support, visit http://support.dell.com/ProSupport or call 1-800-945-3355
900-9997	On-Site Installation Declined
331-3286	CLOUD COMPUTE NODE, DCS, INFO MOD ,OPENSTACK
Group: 2	Quantity: 3 (Quantity 4 for a Compute Chassis)
225-3555	PowerEdge C8220X, Double Width Compute Sled
331-4428	Performance Optimized
330-4118	System ordered as part of Multipack order
421-8663	No Factory Installed Operating System, v.2
430-3643	Intel DA 10GbE NIC, Dual Port, SFP+, Low Profile - Quantity 2
331-8223	C2 LSI 2008 Mezzanine Card supporting up to 8 Hard Drives
331-8996	Cable for 2.5in Rear Hard Drives, PE-C8220X
342-5079	LSI 2008 SAS Controller Card, 6G, PE C8XXX
317-4928	Dual Processor Option
317-9596	Intel Xeon E5-2670 2.60GHz, 20M Cache, 8.0GT/s QPI, Turbo, 8C, 115W, Max Mem 1600MHz
317-9610	Intel Xeon E5-2670 2.60GHz, 20M Cache, 8.0GT/s QPI, Turbo, 8C, 115W
318-2308	Thermal Heat Sink
318-2308	Thermal Heat Sink
317-8998	128GB Memory (16x8GB),1600Mhz, Dual Ranked RDIMMs for 2 Processors
468-7687	Info, Memory for Dual Processor selection

342-4986	2.5in HDD Enclosure, PE-C8220X
342-5057	2.5in HDD Blank, PE-C8220X - Quantity 2
342-4821	Hard Drive Carrier 2.5 C8000 - Quantity 6
342-4871	1TB 7.2K RPM SATA 3Gbps 2.5in Hard Drive - Quantity 6
342-4983	Hot Plug Hard Drive Carrier, PE-C8220X
934-0626	Dell Hardware Limited Warranty Plus On Site Service Extended Year
934-9845	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, Initial Year
935-0575	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, 2 Year Extended
935-0585	ProSupport: 7x24 HW / SW Tech Support and Assistance, 3 Year
989-3439	Dell ProSupport. For tech support, visit http://support.dell.com/ProSupport or call 1-800-945-3355
996-9927	Dell Hardware Limited Warranty Plus On Site Service Initial Year
900-9997	On-Site Installation Declined
331-3286	CLOUD COMPUTE NODE, DCS, INFO MOD ,OPENSTACK

Table 19: Controller Chassis – PowerEdge C8000

The Controller Chassis includes the OpenStack Controller node, a Nova-Compute node, and one storage node.

SKU	Component
Group: 1	Quantity: 1
225-3550	PE C8000 Enclosure, Dual Power Supply
331-8341	PowerEdge C8000 Shipping
331-9573	SHIP,C8000,DAO
420-3323	No Factory Installed Operating System
330-7353	Power Cord, C13 to C14, PDU Style, 12 Amps, 2 meter, Qty 1 - Quantity 4
331-8218	PowerEdge C8000 Static Rails, Tool-less
936-3965	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, Initial Year
936-4695	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, 2 Year Extended
936-4705	ProSupport: 7x24 HW / SW Tech Support and Assistance, 3 Year
936-6035	Dell Hardware Limited Warranty Plus On Site Service Initial Year
936-6145	Dell Hardware Limited Warranty Plus On Site Service Extended Year
989-3439	Dell ProSupport. For tech support, visit http://support.dell.com/ProSupport or call 1-800-945-3355
900-9997	On-Site Installation Declined
331-3286	CLOUD COMPUTE NODE, DCS, INFO MOD ,OPENSTACK
Group: 2	Quantity: 2
225-3555	PowerEdge C8220X, Double Width Compute Sled
331-4428	Performance Optimized
330-4118	System ordered as part of Multipack order
421-8663	No Factory Installed Operating System, v.2
430-3643	Intel DA 10GbE NIC, Dual Port, SFP+ ,Low Profile - Quantity 2
331-8223	C2 LSI 2008 Mezzanine Card supporting up to 8 Hard Drives
331-8996	Cable for 2.5in Rear Hard Drives, PE-C8220X
342-5079	LSI 2008 SAS Controller Card, 6G, PE C8XXX
317-4928	Dual Processor Option
317-9596	Intel Xeon E5-2670 2.60GHz, 20M Cache, 8.0GT/s QPI, Turbo, 8C, 115W, Max Mem 1600MHz
317-9610	Intel Xeon E5-2670 2.60GHz, 20M Cache, 8.0GT/s QPI, Turbo, 8C, 115W
318-2308	Thermal Heat Sink
318-2308	Thermal Heat Sink
317-8998	128GB Memory (16x8GB),1600Mhz, Dual Ranked RDIMMs for 2 Processors
468-7687	Info, Memory for Dual Processor selection
342-4986	2.5in HDD Enclosure, PE-C8220X
342-4821	Hard Drive Carrier 2.5 C8000 - Quantity 6

Dell OpenStack™-Powered Cloud Solution Reference Architecture Guide v1.6.1

342-4871	1TB 7.2K RPM SATA 3Gbps 2.5in Hard Drive - Quantity 6
342-5057	2.5in HDD Blank, PE-C8220X - Quantity 2
342-4983	Hot Plug Hard Drive Carrier, PE-C8220X
934-0626	Dell Hardware Limited Warranty Plus On Site Service Extended Year
934-9845	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, Initial Year
935-0575	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, 2 Year Extended
935-0585	ProSupport: 7x24 HW / SW Tech Support and Assistance, 3 Year
989-3439	Dell ProSupport. For tech support, visit http://support.dell.com/ProSupport or call 1-800-945-3355
996-9927	Dell Hardware Limited Warranty Plus On Site Service Initial Year
900-9997	On-Site Installation Declined
331-3286	CLOUD COMPUTE NODE, DCS, INFO MOD ,OPENSTACK
Group: 3	Quantity:1
225-3555	PowerEdge C8220X, Double Width Compute Sled
331-4428	Performance Optimized
330-4118	System ordered as part of Multipack order
421-8663	No Factory Installed Operating System, v.2
342-4851	LSI 9202-16E, LP, Controller, CE
430-3643	Intel DA 10GbE NIC, Dual Port, SFP+, Low Profile
331-8224	C2B LSI 2008 Mezzanine Card plus Onboard Controller supporting up to 12 Hard Drives
331-8999	SAS Controller Cable, PE-C8220X
342-5079	LSI 2008 SAS Controller Card, 6G, PE C8XXX
317-4928	Dual Processor Option
317-9596	Intel Xeon E5-2670 2.60GHz, 20M Cache, 8.0GT/s QPI, Turbo, 8C, 115W, Max Mem 1600MHz
317-9610	Intel Xeon E5-2670 2.60GHz, 20M Cache, 8.0GT/s QPI, Turbo, 8C, 115W
318-2308	Thermal Heat Sink
318-2308	Thermal Heat Sink
317-8810	Memory Filler Blank DIMM Quantity 8
317-8994	64GB Memory (8x8GB), 1600Mhz, Dual Ranked RDIMMs
468-7687	Info, Memory for Dual Processor selection
342-4987	3.5in HDD Enclosure, PE-C8220X
342-4820	Hard Drive Carrier 3.5 C8000 - Quantity 4
342-4874	3TB,7.2K RPM, Near Line SAS,6Gps,3.5in, Hard Drive - Quantity 4
342-5057	2.5in HDD Blank, PE-C8220X - Quantity 1
342-4841	Hard Drive,2.5 Rear Carrier,C8220
342-4861	1TB,7.2K RPM,SATA,3Gbps,2.5in, Hard Drive
342-4983	Hot Plug Hard Drive Carrier, PE-C8220X
934-0626	Dell Hardware Limited Warranty Plus On Site Service Extended Year
934-9845	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, Initial Year

Dell OpenStack™-Powered Cloud Solution Reference Architecture Guide v1.6.1

935-0575	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, 2 Year Extended
935-0585	ProSupport: 7x24 HW / SW Tech Support and Assistance, 3 Year
989-3439	Dell ProSupport. For tech support, visit http://support.dell.com/ProSupport or call 1-800-945-3355
996-9927	Dell Hardware Limited Warranty Plus On Site Service Initial Year
900-9997	On-Site Installation Declined
331-3286	CLOUD COMPUTE NODE, DCS, INFO MOD ,OPENSTACK
Group: 4	Quantity: 1
225-3558	PowerEdge C8220XD Storage Sled, Single, 12 Hard Drives
330-4118	System ordered as part of Multipack order
420-3323	No Factory Installed Operating System
342-3923	3TB, Near Line SAS 6Gps, 7.2K RPM, 3.5 in Hard Drive - Quantity 12
342-4824	Hard Drive Carrier,3.5,Expanded,Double Wide Storage,C8000 - Quantity 12
934-3976	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, Initial Year
934-4706	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, 2 Year Extended
934-4716	ProSupport: 7x24 HW / SW Tech Support and Assistance, 3 Year
934-6046	Dell Hardware Limited Warranty Plus On Site Service Initial Year
934-6156	Dell Hardware Limited Warranty Plus On Site Service Extended Year
989-3439	Dell ProSupport. For tech support, visit http://support.dell.com/ProSupport or call 1-800-945-3355
900-9997	On-Site Installation Declined
331-3286	CLOUD COMPUTE NODE, DCS, INFO MOD ,OPENSTACK
SOFTWARE & ACCESSORIES	
332-0727	External Cable for LSI9202, Customer Install C8xxx – Quantity: 1

Table 20: Storage Node Chassis – PowerEdge C8000

The Storage node chassis includes two storage nodes.

SKU	Component
Group: 1	Quantity: 1
225-3550	PE C8000 Enclosure, Dual Power Supply
331-8341	PowerEdge C8000 Shipping
331-9573	SHIP,C8000,DAO
420-3323	No Factory Installed Operating System
330-7353	Power Cord, C13 to C14, PDU Style, 12 Amps, 2 meter, Quantity 1 - Quantity 4
331-8218	PowerEdge C8000 Static Rails, Tool-less
936-3965	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, Initial Year
936-4695	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, 2 Year Extended
936-4705	ProSupport: 7x24 HW / SW Tech Support and Assistance, 3 Year
936-6035	Dell Hardware Limited Warranty Plus On Site Service Initial Year
936-6145	Dell Hardware Limited Warranty Plus On Site Service Extended Year
989-3439	Dell ProSupport. For tech support, visit http://support.dell.com/ProSupport or call 1-800-945-3355
900-9997	On-Site Installation Declined
331-3286	CLOUD COMPUTE NODE, DCS, INFO MOD ,OPENSTACK
Group: 2	Quantity: 2
225-3555	PowerEdge C8220X, Double Width Compute Sled
331-4428	Performance Optimized
330-4118	System ordered as part of Multipack order
421-8663	No Factory Installed Operating System, v.2
342-4851	LSI 9202-16E, LP, Controller, CE
430-3643	Intel DA 10GbE NIC, Dual Port, SFP+, Low Profile
331-8224	C2B LSI 2008 Mezzanine Card plus Onboard Controller supporting up to 12 Hard Drives
331-8999	SAS Controller Cable, PE-C8220X
342-5079	LSI 2008 SAS Controller Card, 6G, PE C8XXX
317-4928	Dual Processor Option
317-9596	Intel Xeon E5-2670 2.60GHz, 20M Cache, 8.0GT/s QPI, Turbo, 8C, 115W, Max Mem 1600MHz
317-9610	Intel Xeon E5-2670 2.60GHz, 20M Cache, 8.0GT/s QPI, Turbo, 8C, 115W
318-2308	Thermal Heat Sink
318-2308	Thermal Heat Sink
317-8810	Memory Filler Blank DIMM Quantity 8
317-8994	64GB Memory (8x8GB), 1600Mhz, Dual Ranked RDIMMs
468-7687	Info, Memory for Dual Processor selection

Dell OpenStack™-Powered Cloud Solution Reference Architecture Guide v1.6.1

342-4987	3.5in HDD Enclosure, PE-C8220X
342-4820	Hard Drive Carrier 3.5 C8000 - Quantity 4
342-4874	3TB,7.2K RPM, Near Line SAS,6Gps,3.5in, Hard Drive - Quantity 4
342-5057	2.5in HDD Blank, PE-C8220X - Quantity 1
342-4841	Hard Drive,2.5 Rear Carrier,C8220
342-4861	1TB,7.2K RPM,SATA,3Gbps,2.5in, Hard Drive
342-4983	Hot Plug Hard Drive Carrier, PE-C8220X
934-0626	Dell Hardware Limited Warranty Plus On Site Service Extended Year
934-9845	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, Initial Year
935-0575	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, 2 Year Extended
935-0585	ProSupport: 7x24 HW / SW Tech Support and Assistance, 3 Year
989-3439	Dell ProSupport. For tech support, visit http://support.dell.com/ProSupport or call 1-800-945-3355
996-9927	Dell Hardware Limited Warranty Plus On Site Service Initial Year
900-9997	On-Site Installation Declined
331-3286	CLOUD COMPUTE NODE, DCS, INFO MOD ,OPENSTACK
Group: 3	Quantity: 2
225-3558	PowerEdge C8220XD Storage Sled, Single, 12 Hard Drives
330-4118	System ordered as part of Multipack order
420-3323	No Factory Installed Operating System
342-3923	3TB, Near Line SAS 6Gps, 7.2K RPM, 3.5 in Hard Drive - Quantity 12
342-4824	Hard Drive Carrier,3.5,Expanded,Double Wide Storage,C8000 - Quantity 12
934-3976	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, Initial Year
934-4706	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, 2 Year Extended
934-4716	ProSupport: 7x24 HW / SW Tech Support and Assistance, 3 Year
934-6046	Dell Hardware Limited Warranty Plus On Site Service Initial Year
934-6156	Dell Hardware Limited Warranty Plus On Site Service Extended Year
989-3439	Dell ProSupport. For tech support, visit http://support.dell.com/ProSupport or call 1-800-945-3355
900-9997	On-Site Installation Declined
331-3286	CLOUD COMPUTE NODE, DCS, INFO MOD ,OPENSTACK
SOFTWARE & ACCESSORIES	
332-0727	External Cable for LSI9202, Customer Install C8xxx – Quantity: 2

Appendix C : Physical Configuration — PowerEdge R720xd

Table 21: Rack Configuration – PowerEdge R720xd (or R720/R720xd)

RU	RACK1	RACK2	RACK3
42	R1- Switch 2: Force10 S60	R2- Switch2: Force10 S60	R3- Switch2: Force10 S60
41	R1- Switch 1: Force10 S60	R2- Switch1: Force10 S60	R3- Switch1: Force10 S60
40	Cable Management	Cable Management	Cable Management
39	Cable Management	Cable Management	Cable Management
38	Admin Node:R720xd or R720	Controller Node : R720xd or R720	R3 - Switch 1: Force10 S4810 (1 RU) OR Force10 Z9000 (2 RU)
37			R3 - Switch 1: Force10 S4810 (1 RU) OR Force10 Z9000 (2 RU)
36	Cable Management	Cable Management	Cable Management
35	Cable Management	Cable Management	Cable Management
34	Cinder Node R720xd or R720	Compute Node R720xd or R720	Compute: R720xd or R720
33			
32	R1 - S55 iDRAC Mgmt switch	R2 - S55 iDRAC Mgmt switch	R3 - S55 iDRAC Mgmt switch
31	Empty	Empty	Empty
21-30	Empty	Empty	Empty
20	R1- Chassis10: R720/R720xd	R2- Chassis10: R720/R720xd	R3- Chassis10: R720/R720xd
19			
18	R1- Chassis09: R720/R720xd	R2- Chassis09: R720/R720xd	R3- Chassis09: R720/R720xd
17			
16	R1- Chassis08: R720/R720xd	R2- Chassis08: R720/R720xd	R3- Chassis08: R720/R720xd
15			
14	R1- Chassis07: R720/R720xd	R2- Chassis07: R720/R720xd	R3- Chassis07: R720/R720xd
13			
12	R1- Chassis06: R720/R720xd	R2- Chassis06: R720/R720xd	R3- Chassis06: R720/R720xd
11			
10	R1- Chassis05: R720/R720xd	R2- Chassis05: R720/R720xd	R3- Chassis05: R720/R720xd
9			
8	R1- Chassis04: R720/R720xd	R2- Chassis04: R720/R720xd	R3- Chassis04: R720/R720xd
7			
6	R1- Chassis03: R720/R720xd	R2- Chassis03: R720/R720xd	R3- Chassis03: R720/R720xd
5			
4	R1- Chassis02: R720/R720xd	R2- Chassis02: R720/R720xd	R3- Chassis02: R720/R720xd
3			
2	R1- Chassis01: R720/R720xd	R2- Chassis01: R720/R720xd	R3- Chassis01: R720/R720xd
1			

Appendix D : **Bill of Materials – PowerEdge R720xd Nodes**

Table 22: Infrastructure and Nova-Compute Nodes – PowerEdge R720xd

SKU	Component
225-2110	PowerEdge R720xd
331-4437	PowerEdge R720 Shipping
430-4445	Intel X520 DP 10Gb DA/SFP+ Server Adapter - Quantity 1
421-5339	iDRAC7 Enterprise
430-4418	Broadcom 5720 QP 1Gb Network Daughter Card
342-3566	Chassis with up to 24, 2.5" Hard Drives
318-1375	Bezel
330-5116	Power Saving Dell Active Power Controller
331-4537	RAID 10 for H710P/H710/H310 (4-24 HDDs in pairs)
342-3529	PERC H710 Integrated RAID Controller, 512MB NV Cache
317-9595	Intel Xeon E5-2640 2.50GHz, 15M Cache, 7.2GT/s QPI, Turbo, 6C, 95W, Max Mem 1333MHz
331-4508	Heat Sink for PowerEdge R720 and R720xd
317-8688	DIMM Blanks for Systems with 2 Processors
317-9609	Intel Xeon E5-2640 2.50GHz, 15M Cache, 7.2GT/s QPI, Turbo, 6C, 95W
331-4508	Heat Sink for PowerEdge R720 and R720xd
317-9644	8GB RDIMM, 1333 MT/s, Low Volt, Dual Rank, x4 Data Width - Quantity 16
331-4422	1333 MHz RDIMMs
331-4428	Performance Optimized
342-0847	600GB 10K RPM SAS 6Gbps 2.5in Hot-plug Hard Drive - Quantity 6
331-5914	Electronic System Documentation and OpenManage DVD Kit for R720 and R720xd
331-4433	ReadyRails Sliding Rails With Cable Management Arm
331-4605	Dual, Hot-plug, Redundant Power Supply (1+1), 750W
310-8509	Power Cord, NEMA 5-15P to C13, 15 amp, wall plug, 10 feet / 3 meter - Quantity 2
420-6320	No Operating System
421-5736	No Media Required
936-0967	Dell Hardware Limited Warranty Plus On Site Service Initial Year
936-7243	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, 2 Year Extended
936-7263	ProSupport: 7x24 HW / SW Tech Support and Assistance, 3 Year
939-3398	Dell Hardware Limited Warranty Plus On Site Service Extended Year
989-2701	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, Initial Year
989-3439	Dell ProSupport. For tech support, visit http://support.dell.com/ProSupport or call 1-800-945-3355
900-9997	On-Site Installation Declined
926-2979	Proactive Maintenance Service Declined
331-3310	CLOUD COMPUTE NODE,PEC,CROWBAR
331-3286	CLOUD COMPUTE NODE, DCS, INFO MOD ,OPENSTACK

Appendix E : **Bill of Materials – PowerEdge R720 Nodes**

Table 23: Infrastructure and Nova-Compute Nodes – PowerEdge R720

SKU	Component
225-2133	PowerEdge R720
331-4437	PowerEdge R720 Shipping
331-4440	Risers with up to 6, x8 PCIe Slots + 1, x16 PCIe Slot
430-4445	Intel X520 DP 10Gb DA/SFP+ Server Adapter - Quantity 1
421-5339	iDRAC7 Enterprise
430-4418	Broadcom 5720 QP 1Gb Network Daughter Card
317-8474	2.5" Chassis with up to 16 Hard Drives
318-1375	Bezel
330-5116	Power Saving Dell Active Power Controller
331-4383	RAID 10 for H710P/H710/H310 (4-16 HDDs in pairs)
342-3529	PERC H710 Integrated RAID Controller, 512MB NV Cache
317-9595	Intel Xeon E5-2640 2.50GHz, 15M Cache, 7.2GT/s QPI, Turbo, 6C, 95W, Max Mem 1333MHz
331-4508	Heat Sink for PowerEdge R720 and R720xd
317-8688	DIMM Blanks for Systems with 2 Processors
317-9609	Intel Xeon E5-2640 2.50GHz, 15M Cache, 7.2GT/s QPI, Turbo, 6C, 95W
331-4508	Heat Sink for PowerEdge R720 and R720xd
317-9644	8GB RDIMM, 1333 MT/s, Low Volt, Dual Rank, x4 Data Width - Quantity 16
331-4422	1333 MHz RDIMMs
331-4428	Performance Optimized
342-0847	600GB 10K RPM SAS 6Gbps 2.5in Hot-plug Hard Drive - Quantity 6
331-5914	Electronic System Documentation and OpenManage DVD Kit for R720 and R720xd
331-4612	No Internal Optical Drive
331-4433	ReadyRails Sliding Rails With Cable Management Arm
331-4605	Dual, Hot-plug, Redundant Power Supply (1+1), 750W
310-8509	Power Cord, NEMA 5-15P to C13, 15 amp, wall plug, 10 feet / 3 meter - Quantity 2
420-6320	No Operating System
421-5736	No Media Required
936-4593	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, 2 Year Extended
936-4603	ProSupport: 7x24 HW / SW Tech Support and Assistance, 3 Year
939-2678	Dell Hardware Limited Warranty Plus On Site Service Extended Year
939-2768	Dell Hardware Limited Warranty Plus On Site Service Initial Year
988-9281	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, Initial Year
989-3439	Dell ProSupport. For tech support, visit http://support.dell.com/ProSupport
900-9997	On-Site Installation Declined
926-2979	Proactive Maintenance Service Declined
331-3310	CLOUD COMPUTE NODE,PEC,CROWBAR
331-3286	CLOUD COMPUTE NODE, DCS, INFO MOD ,OPENSTACK

Appendix F : **Bill of Materials – PowerEdge R720xd Data Node**

Table 24: Storage Node – PowerEdge R720xd

SKU	Component
225-2110	PowerEdge R720xd
331-4437	PowerEdge R720 Shipping
430-4445	Intel X520 DP 10Gb DA/SFP+ Server Adapter
421-5339	iDRAC7 Enterprise
430-4418	Broadcom 5720 QP 1Gb Network Daughter Card
342-3566	Chassis with up to 24, 2.5" Hard Drives
318-1375	Bezel
330-5116	Power Saving Dell Active Power Controller
331-3765	UEFI BIOS Setting
331-4557	Non-configured RAID for H710P/H710/H310 (1-24 HDDs)
342-3529	PERC H710 Integrated RAID Controller, 512MB NV Cache
317-9595	Intel Xeon E5-2640 2.50GHz, 15M Cache, 7.2GT/s QPI, Turbo, 6C, 95W, Max Mem 1333MHz
331-4508	Heat Sink for PowerEdge R720 and R720xd
317-8688	DIMM Blanks for Systems with 2 Processors
317-9609	Intel Xeon E5-2640 2.50GHz, 15M Cache, 7.2GT/s QPI, Turbo, 6C, 95W
331-4508	Heat Sink for PowerEdge R720 and R720xd
317-9644	8GB RDIMM, 1333 MT/s, Low Volt, Dual Rank, x4 Data Width - Quantity 8
331-4422	1333 MHz RDIMMs
331-4428	Performance Optimized
342-1998	1TB 7.2K RPM SATA 3Gbps 2.5in Hot-plug Hard Drive - Quantity 24
331-5914	Electronic System Documentation and OpenManage DVD Kit for R720 and R720xd
331-4433	ReadyRails Sliding Rails With Cable Management Arm
331-4605	Dual, Hot-plug, Redundant Power Supply (1+1), 750W
310-8509	Power Cord, NEMA 5-15P to C13, 15 amp, wall plug, 10 feet / 3 meter - Quantity 2
420-6320	No Operating System
421-5736	No Media Required
936-0967	Dell Hardware Limited Warranty Plus On Site Service Initial Year
936-7243	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, 2 Year Exte
936-7263	ProSupport: 7x24 HW / SW Tech Support and Assistance, 3 Year
939-3398	Dell Hardware Limited Warranty Plus On Site Service Extended Year
989-2701	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, Initial Year
989-3439	Dell ProSupport. For tech support, visit http://support.dell.com/ProSupport
900-9997	On-Site Installation Declined
926-2979	Proactive Maintenance Service Declined
331-3310	CLOUD COMPUTE NODE,PEC,CROWBAR
331-3286	CLOUD COMPUTE NODE, DCS, INFO MOD ,OPENSTACK

Appendix G : Physical Configuration — PowerEdge C6220

Table 25: Rack Configuration – PowerEdge C6220 (Maximum 60 Nodes)

RU	RACK1	RACK2 ²	RACK3
42	R1- Switch 1: Force10 S4810	R2- Switch 1: Force10 S4810	R3- Switch 1: Force10 S4810
41	R1- Switch 2: Force10 S4810	R2- Switch 2: Force10 S4810	R3- Switch 2: Force10 S4810
40	Cable Management	Cable Management	Cable Management
39	Cable Management	Cable Management	Cable Management
38	Admin/ Compute x3 Node: C6220(4 node)	Controller Node/Computex3: C6220	R3 - Switch 1: Force10 S4810 (1 RU) OR Force10 Z9000 (2 RU)
37			R3 - Switch 2: Force10 S4810 (1 RU) OR Force10 Z9000 (2 RU)
36	Cable Management	Cable Management	Cable Management
35	Cable Management	Cable Management	Cable Management
34	Controller/Compute x 3 Node: C6220 ¹	Nova-Compute Node: C6220	Storage R720xd or Compute C6220(4-Node)
33			
32	R1 - S55 iDRAC Mgmt switch	R2 - S55 iDRAC Mgmt switch	R3 - S55 iDRAC Mgmt switch
31	Empty	Empty	Empty
21- 30	Empty	Empty	Empty
20	R1- Chassis10: Storage R720xd or Empty	Empty	Empty
19			
18	R1- Chassis09: Storage R720xd or Empty	R2- Chassis09: Storage R720xd) or Empty	R3- Chassis09: Storage R720xd or Empty
17			
16	R1- Chassis08: Storage R720xd or Empty	R2- Chassis08: Storage R720xd or Empty	R3- Chassis08: Storage R720xd or Empty
15			
14	R1- Chassis07: Storage R720xd or Empty	R2- Chassis07: Storage R720xd or Empty	R3- Chassis07: Storage R720xd or Empty
13			
12	R1- Chassis06: Storage R720xd or Empty	R2- Chassis06: Storage R720xd or Empty	R3- Chassis06: Storage R720xd or Empty
11			
10	R1- Chassis05: Storage R720xd or Empty	R2- Chassis05: Storage R720xd or Empty	R3- Chassis05: Storage R720xd or Empty
9			
8	R1- Chassis04: Storage R720xd or Empty	R2- Chassis04: Storage R720xd or Empty	R3- Chassis04: Storage R720xd or Empty
7			
6	R1- Chassis03: Storage R720xd or	R2- Chassis03: Storage R720xd or	R3- Chassis03: Storage R720xd or

² Controller node should be in in Rack 2 if using a multi-rack configuration.

Dell OpenStack™-Powered Cloud Solution Reference Architecture Guide v1.6.1

RU	RACK1	RACK2 ²	RACK3
5	Compute C6220(4-node)	Empty	Compute C6220(4-node)
4	R1- Chassis02: Storage R720xd or Compute C6220(4-node)	R2- Chassis02: Storage R720xd or Compute C6220(4-node)	R3- Chassis02: Storage R720xd or Compute C6220(4-node)
3			
2	R1- Chassis01: Storage R720xd or Compute C6220(4-node)	R2- Chassis01: Storage R720xd or Compute C6220(4-node)	R3- Chassis01: Storage R720xd or Compute C6220(4-node)
1			

Appendix H : **Bill of Materials – PowerEdge C6220**

Table 26: Infrastructure and Nova-Compute Node – PowerEdge C6220

SKU	Description
225-2582	PowerEdge C6000, 2/4 Motherboard, 2.5 in Hard Drives with Backplane
331-5531	SHIP,C6220,DAO
331-5532	PowerEdge C6000/C6220 Shipping
420-3323	No Factory Installed Operating System
331-5527	PowerEdge C6220 Documentation
331-1619	PowerEdge C Static Rails, Tool-less
331-5518	Power Supply, 1100W, Redundant Capable
331-5518	Power Supply, 1100W, Redundant Capable
331-5522	Label,Regulatory,1100W,C6220
330-7353	Power Cord, C13 to C14, PDU Style, 12 Amps, 2 meter, Qty 2
937-8137	Dell Hardware Limited Warranty Plus On Site Service Initial Year
937-8147	Dell Hardware Limited Warranty Plus On Site Service Extended Year
937-8427	Non-Mission Critical: 4-Hour 7x24 On-site Service After Problem Diagnosis, Initial Year
937-8457	Non-Mission Critical: 4-Hour 7x24 On-site Service After Problem Diagnosis, 2 Year Extended
937-8607	ProSupport: 7x24 HW / SW Tech Support and Assistance, 3 Year
989-3439	Dell ProSupport. For tech support, visit http://support.dell.com/ProSupport or call 1-800-945-3355
900-9997	On-Site Installation Declined
Sled Order	Quantity 4
225-3819	PowerEdge C6220,1U Motherboard with 2.5 SATA/SAS Hard Drives, Direct
317-9095	Memory Filler Blank DIMM Quantity 6
317-8810	Memory Filler Blank Dimm Quantity 8
319-1811	8GB RDIMM, 1600MT/s, Low Volt, Dual Rank, x4 Data Width, Qty 16

331-4424	1600 MHz RDIMMS
330-4118	System ordered as part of Multipack order
421-8663	No Factory Installed Operating System , v.2
342-3931	C4, 1U, LSI 9265 Card supporting up to 3,3.5D, 6,2.5D/E Hard Drives
342-5124	LSI 9265 SAS Low Profile Card
317-4928	Dual Processor Option
317-9107	Thermal Heat Sink
317-9107	Thermal Heat Sink
317-9596	Intel Xeon E5-2670 2.60GHz, 20M Cache, 8.0GT/s QPI, Turbo, 8C, 115W, Max Mem 1600MHz
317-9610	Intel Xeon E5-2670 2.60GHz, 20M Cache, 8.0GT/s QPI, Turbo, 8C, 115W
342-2306	Carrier, Hard Drive,2.5,2LED,C6145 -Quantity 6
342-3661	1TB, 7.2K RPM,SATA, 3Gbps, 2.5in, Hot Plug Hard drive - Quantity 6
342-4082	10G Controller Card, Mezz, Dual Port, PE C6220
937-9827	Dell Hardware Limited Warranty Plus On Site Service Initial Year
937-9837	Dell Hardware Limited Warranty Plus On Site Service Extended Year
938-0097	Non-Mission Critical: 4-Hour 7x24 On-site Service After Problem Diagnosis, Initial Year
938-0127	Non-Mission Critical: 4-Hour 7x24 On-site Service After Problem Diagnosis, 2 Year Extended
938-0277	ProSupport: 7x24 HW / SW Tech Support and Assistance, 3 Year
989-3439	Dell ProSupport. For tech support, visit http://support.dell.com/ProSupport or call 1-800-945-3355
900-9997	On-Site Installation Declined
926-2979	Proactive Maintenance Service Declined

Appendix I : **Bill of Materials – Force10 Network Equipment**

Table 27: Network Equipment – 1GbE and 10GbE – Dell Force10

SKU	Description
225-2446	Force10, Z9000, 2U, 32 x 40Gbe QSFP+ Ports, 1 AC Power Supply, Fan w/IO Panel to PSU (Normal) Airflow (Non-Redundant Power)
331-5996	Force10, Power Cord, 125V, 15A, 10 Feet, NEMA 5-15/C13, S-Series
331-5343	Force10, Z9000, AC Power Supply for Chassis with IO Panel to PSU (Normal) Airflow
430-4543	Force10, Transceiver, 40GE QSFP+ Short Reach Optics, 850nmWavelength, 100-150m Reach onOM3/OM4
331-7279	Force10, Z9000 Cable Management Kit
225-2477	Force10, S4810P, 48 x 10GbE SFP+, 4 x QSFP 40GbE, 1 x AC PSU, 2 x Fans, IO Panel to PSU Airflow
225-2479	Force10, S4810P, 48 x 10GbE SF P+, 4 x QSFP 40GbE, 1 x AC PSU , 2 x Fans, PSU to IO Panel Airflow
331-5103	Force10, S4810, AC Power Supply, IO Panel to PSU Airflow
331-5105	Force10, S4810, AC Power Supply, PSU to IO Panel Airflow
331-5258	Force10, Cable, SFP+ to SFP+, 10GbE, Copper Twinax Direct Attach Cable, 2 Meters
331-5996	Force10, Power Cord, 125V, 15A, 10 Feet, NEMA 5-15/C13, S-Series
421-6981	Force10, Software, L3 Latest Version, S4810
430-4543	Force10, Transceiver, 40GE QSFP+ Short Reach Optics, 850nmWavelength, 100-150m Reach onOM3/OM4
331-5274	Force10, Transceiver, SFP+, 10GbE, SR, 850nm Wavelength, 300m Reach
430-4543	Force10, Transceiver, 40GE QSFP+ Short Reach Optics, 850nmWavelength, 100-150m Reach onOM3/OM4
331-5393	Force10, Rear Rack Mounting Bracket, 4 Post, S4810
225-2450	Force10, S60, 44 x 10/100/1000 BASE-T, 4 x SFP, 2 Expansion Slots, 1 x AC PSU, 2 x fans, P SU to IO Panel Airflow
331-5233	Force10, SFP+ Expansion Module , 2 x 10 GbE Ports, S60 Series (SFP+ optics required)
331-5996	Force10, Power Cord, 125V, 15A , 10 Feet, NEMA 5-15/C13, S-Series
331-5226	Force10, S60, AC Power Supply, PSU to IO Panel Airflow
331-5398	Force10, Rear Rack Mounting Bracket, Metal, 4 Post, S60
	Force10 S60 2 port, 12G, Stacking module
	Force10 S60 12 Gig 60cms stacking cable
331-3286	CLOUD COMPUTE NODE, DCS, INFO MOD ,OPENSTACK

Table 28: Network Equipment – 10GbE – Dell Force10

Quantity	SKU	Description
Cluster Network		
50	331-5274	Dell Networking, Transceiver, SFP+, 10GbE, SR, 850nm Wavelength, 300m Reach
36	330-8723	SFP+, Short Range, Optical Transceiver, LC Connector, 10Gb and 1Gb compatible(Intel 10G SFP+)
2	225-2477	Force10, S4810P, 48 x 10GbE SFP+, 4 x QSFP 40GbE, 1 x AC PSU, 2 x Fans, IO Panel to PSU Airflow
2	331-5996	Force10, Power Cord, 125V, 15A, 10 Feet, NEMA 5-15/C13, S-Series
2	331-5272	Dell Networking, Transceiver, SFP, 1000BASE-LX, 1310nm Wavelength, 10km Reach
2	331-5393	Force10, Rear Rack Mounting Bracket, 4 Post, S4810
2	331-6279	Force10, User Documentation for S4810, DAO/BCC
2	935-0103	SW Support,Force10 Software ,3 Years
2	935-0143	ProSupport: 7x24 HW / SW Tech Support and Assistance, 3 Years
2	931-3856	ProSupport: 4-Hour 7x24 Parts Only After Problem Diagnosis, Initial Year
2	989-3439	Dell ProSupport. For tech support, visit http://support.dell.com/ProSupport or call 1-800-945-3355
2	996-2760	Dell Hardware Limited Warranty Extended Year(s)
2	935-0123	ProSupport: 4-Hour 7x24 Parts Only After Problem Diagnosis, 2 Year Extended
2	996-2670	Dell Hardware Limited Warranty Initial Year
2	900-9997	On-Site Installation Declined
2	996-3080	ProSupport for, Force10,Layer 3 Enablement, 1 Year
2	331-9460	Force10, Software, iSCSI-Optimized Configuration, S4810
1	331-5217	Customer Kit, Dell Networking, Cable, QSFP+, 40GbE SFP+ Passive Copper Direct Attach Cable, 1 Meter
1	331-3286	CLOUD COMPUTE NODE, DCS, INFO MOD ,OPENSTACK
Administration Network		
1	225-2503	Force10, S55, 44 x 10/100/1000 BASE-T, 4 x SFP, 2 Expansion Slots, 1 x AC PSU, 2 x Fans, IO Panel to PSU Airfl (225-2503)
1	331-5233	Forcd10 SFP+ Expansion Module 2x10 GbE Ports
1	331-5243	Force10, S55, AC Power Supply, IO Panel to PSU Airflow (331-5243)
1	331-5996	Force10, Power Cord, 125V, 15A, 10 Feet, NEMA 5-15/C13, S-Series (331-5996)
1	331-5252	Force10, Rear Rack Mounting Bracket, 4 Post, S55 (331-5252)
1	331-9233	No Returns Allowed on Dell Force10 Switches (331-9233)
1	331-6271	Force10, User Documentation for S55/S60, DAO/BCC (331-6271)
1	935-1367	Dell Hardware Limited Warranty Initial Year (935-1367)

Dell OpenStack™-Powered Cloud Solution Reference Architecture Guide v1.6.1

1	938-7578	Dell Hardware Limited Warranty Extended Year(s) (938-7578)
1	989-3439	Dell ProSupport. For tech support, visit http://support.dell.com/ProSupport or call 1-800-945-3355 (989-3439)
1	995-0592	ProSupport: Next Business Day Parts Delivery, 2 Year Extended (995-0592)
1	995-0622	ProSupport: 7x24 HW / SW Tech Support and Assistance, 3 Years (995-0622)
1	995-9649	SW Support,Force10 Software ,5 Years (995-9649)
1	996-0530	ProSupport: Next Business Day Parts Delivery, Initial Year (996-0530)
1	996-0540	Force10, 5 Year Return To Depot Service, Base Warranty (996-0540)
1	990-9997	On-Site Installation Declined (900-9997)
1	331-3286	CLOUD COMPUTE NODE, DCS, INFO MOD ,OPENSTACK

Appendix J : **Bill of Materials – Dell 6248 Network Equipment**

Table 29: Network Equipment – Dell 6248 (Optional)

Component	Description	SKU
PowerConnect 6248P	PowerConnect 6248, 48 GbE Ports, Managed Switch, 10GbE and Stacking Capable	[222-6714]
Front-end SFP Fiber Transceivers	None	-
Modular Upgrade Bay 1: Modules	Stacking Module, 48Gbps, Includes 1m Stacking Cable	[320-5171]
Modular Upgrade Bay 1: Optics	None	-
Modular Upgrade Bay 2: Modules	None	-
Modular Upgrade Bay 2: Optics	None	-
Cables	Stacking Cable, 3m	[320-5171]
External Redundant Power Supply	None	-
Hardware Support Services	3 Year ProSupport and NBD On-site Service	[980-5492]
		[981-1260]
		[985-6027]
		[985-6038]
		[991-8459]
Installation Services	No Installation Services Selected	
Asset Recovery Services	None	
Cables (optional)	Stacking Cable, 3m	[320-5168]

Appendix K : Miscellaneous Equipment Notes

Server Racks and Power

The above list of SKU's for the servers include many items. However, they do not include racks or power distribution units, as they are generally site specific. The C8000 server line requires 240V power and other servers are dual voltage (110/240). The physical dimensions and power requirements need to be reviewed, as the C8000 requires extra space for front side cable management and rear power distribution, in addition to extra depth. The R720 and C6220 lines both require rear cable management and power distribution.

Networking Equipment

The above list of SKUs includes switches that have specific air flow options. There are both I/O to PSU SKU numbers and PSU to I/O side options available for reverse air flow. Redundant FANs (other than the minimum supplied with chassis) should also be same direction as the base switch. The airflow cannot be reversed in the field at this time.

The above list shows the AC power supplies only. All switch models are available in DC as well.

The above list includes the necessary cables for the connections between the switches for uplinks and interconnects.

The BOM's do not include the cables required for connecting the individual servers into the cluster, since the exact cables required depend on the final chosen rack layout, and choice of cable is often based on customer preference. The following tables summarize the cable counts required, including BMC connections:

Table 30: Network Cables Required – 10GbE Configurations

Description	1GbE Cables Required	10GbE Cables Required
All Nodes	2 x number of nodes	2 x number of nodes
All Nodes	2 x number of nodes	2 x number of nodes

Table 31: Network Cables Required – 1GbE Configurations

Description	1GbE Cables Required	10GbE Cables Required
All Nodes	3 x number of nodes	N/A
All Nodes	3 x number of nodes	N/A

Appendix L : **Bill of Materials – Software and Support**

Software, training, and support SKUs change regularly, and are related to specific global regions. Please see <http://salesedge.dell.com/cmsdoc?id=0901bc828071678d&ll=md>, or contact your Dell account representative for the latest information.

The Sample Bill of Materials appendices include service and support SKUs for the United States. These SKUs need to be changed for other regions.

Appendix M : JBOD versus single disk RAID 0 Configuration

The OpenStack™ community's strong advocacy for the non-RAID drives configuration known as "Just a Bunch of Disks," or JBOD, has caused some confusion for readers of our reference architecture. We fully endorse this approach but feel a need for clarification because there are multiple valid ways to achieve this configuration.

Normally, the optimum disk configuration for OpenStack data storage nodes is considered to be JBOD mode rather than RAID. This is because HDFS swift provides its own data replication, eliminating the need for the redundancy provided by various RAID levels 1-6. levels HDFS Swift also implements efficient round robin parallel I/O across multiple drives, eliminating the need for the parallelism provided by the striping capabilities of RAID 0.

Some drive controllers support only RAID mode, and so can't be used in a plain host bus adapter (HBA) mode for JBOD. For these situations, configuring the controllers as multiple RAID 0 "arrays" allows HDFS Swift to own them as a single drive. In this configuration, the controller is effectively operating just like a standard HBA in JBOD mode, and the RAID 0 and JBOD performance characteristics are comparable. While having a RAID controller adds a minor latency, it is offset by adaptive read-ahead caching.

Update History

Changes in Version 1.6.1

The following changes have been made to this guide since the 1.6 release:

- Updated reference to Dell Multi-Cloud Manager in Figure 2.
- Corrected the Server Infrastructure Options bulleted list of servers.
- Added the Service Layout section.
- Updated Figure 8: Multi-rack Networking Equipment to display VLT/VRRP between switches.
- Corrected the Add-in Network and Disk information in Table 8: Hardware Configurations—PowerEdge C6220 with PowerEdge R720.
- Added S55 iDRAC Management switches to each rack in Table 17: Rack Configuration – PowerEdge C8000 (Maximum 60 Nodes).
- Updated the description of Table 18: Admin or Compute Chassis – PowerEdge C8000.
- Updated the description of Table 19: Controller Chassis – PowerEdge C8000.
- Added S55 iDRAC Management switches to each rack in Table 21: Rack Configuration – PowerEdge R720xd (or R720/R720xd).
- Added S55 iDRAC Management switches to each rack in Table 25: Rack Configuration – PowerEdge C6220 (Maximum 60 Nodes).
- Updated Table 26: Infrastructure and Nova-Compute Node – PowerEdge C6220.

Getting Help

Contacting Dell

For customers in the United States, call 800-WWW-DELL (800-999-3355).

Note: If you do not have an active Internet connection, you can find contact information on your purchase invoice, packing slip, bill, or Dell product catalog.

Dell provides several online and telephone-based support and service options. Availability varies by country and product, and some services may not be available in your area. To contact Dell for sales, technical support, or customer service issues:

- Visit support.dell.com.
- Click your country/region at the bottom of the page. For a full listing of country/region click **All**.
- Click **All Support** from the **Support** menu.
- Select the appropriate service or support link based on your need.

Choose the method of contacting Dell that is convenient for you.

To Learn More

For more information on the Dell OpenStack™-Powered Cloud Solution, visit:
www.dell.com/openstack

©2011–2013 Dell Inc. All rights reserved. Trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Specifications are correct at date of publication but are subject to availability or change without notice at any time. Dell and its affiliates cannot be responsible for errors or omissions in typography or photography. Dell's Terms and Conditions of Sales and Service apply and are available on request. Dell service offerings do not affect consumer's statutory rights.

Dell, the DELL logo, and the DELL badge, PowerConnect, and PowerVault are trademarks of Dell Inc.