

# Dell OpenStack- Powered Cloud Solution

A Dell Reference Architecture Guide



## Table of Contents

---

<b>Tables</b>	<b>3</b>
<b>Figures</b>	<b>4</b>
<b>Terminology and Abbreviations</b>	<b>5</b>
<b>Reference Architecture Overview</b>	<b>6</b>
Approach	6
OpenStack Readiness	6
Taxonomy	6
Hardware Options	7
Networking and Network Services	7
Crowbar Management	8
OpenStack Architecture	9
Operational Notes	11
<b>Hardware Components</b>	<b>14</b>
Multiple Solution Configuration Options	14
Solution Bundling and Delivery	14
Expanding the Solution Beyond the Six-node Solution	17
Site Preparation Needed for the Deployment	19
<b>Network Overview</b>	<b>19</b>
High-level Description	21
Network Cabling and Redundancy	21
Logical Network Configuration	21
Stacked Top-of-Rack Switches	22
Physical Configuration	23
Single Rack Expansion from Starter Configuration	23
Network Port Assignments	25
<b>Appendix A: PowerEdge C6220 Configuration—Hardware Bill of Materials</b>	<b>27</b>
<b>Appendix B: PowerEdge C6105 Configuration—Hardware Bill of Materials</b>	<b>29</b>
<b>Appendix C: PowerEdge R720 and R720xd Bill of Materials</b>	<b>32</b>
<b>Appendix D: Rack Bill of Materials</b>	<b>35</b>
<b>Appendix E: Network Equipment Bill of Materials</b>	<b>36</b>
<b>Appendix F: Solution and Installation</b>	<b>38</b>
<b>Appendix G: Solution Starter Configuration Expansion Options</b>	<b>39</b>
<b>Getting Help</b>	<b>40</b>
Contacting Dell	40
<b>To Learn More</b>	<b>40</b>

## Tables

Table 1: Terminology and Abbreviations.....	5
Table 2: OpenStack components.....	10
Table 3: Deployment Platforms.....	15
Table 4: Admin and Controller System .....	15
Table 5: PowerEdge C6220—Compute .....	16
Table 6: PowerEdge C6220/R720xd Compute/Object Store .....	16
Table 7: PowerEdge C6105 Hybrid—Compute/Object Store .....	17
Table 8: PowerEdge R720/R720xd Compute/Object Store .....	17
Table 9: Additional and Optional Hardware Recommendations .....	18
Table 10: Default Crowbar Networks .....	22
Table 11: Starter Solution Node Network Cabling.....	25
Table 12: Full Rack Node Network Cabling .....	26
Table 13: C6220 Config .....	27
Table 14: Admin/Controller PowerEdge C6105 Node.....	29
Table 15: 2 Sled Compute Node—PowerEdge C6105 Configuration.....	30
Table 16: R720 Admin/Controller node .....	32
Table 17: R720xd Storage Node.....	33
Table 18: PowerEdge 24U Rack.....	35
Table 19: PowerEdge 42U Rack.....	35
Table 20: Force 10 S60 .....	36
Table 21: PowerConnect PC6224 (Optional Only) .....	36
Table 22: Network Switch Add-Ons (For Information Only) .....	37
Table 23: Services.....	38
Table 24: Additional Servers for Design Solution .....	39

## Figures

---

Figure 1: OpenStack Taxonomy	7
Figure 2: Crowbar Ops Management	8
Figure 3: OpenStack Architecture	9
Figure 4: Crowbar Dashboard	12
Figure 5: Crowbar Overview	12
Figure 6: Nagios Provides Alerting and Node Status Monitoring	
Figure 7: Ganglia Provides Performance Monitoring Add-ins	
Figure 8: Network Wiring	20
Figure 9: Seven-node PowerEdge C6220/R720xd Compute/Object Store Starter Solution	23
Figure 10: Full Rack PowerEdge C6105	24
Figure 11: PowerEdge C6220 28-node Compute Deployment	24

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 – 2012 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
<b>Admin</b>	Initial server setup to manage/bootstrap other servers.	Crowbar
<b>Barclamp</b>	Software that evaluates the whole environment and proposes configurations, roles, and recipes that fit your infrastructure.	Crowbar Module
<b>BMC</b>	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
<b>CH1SL1 or CH2SL3</b>	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.	
<b>Controller(s)</b>	Infrastructure and management components installed in each chassis. Supports the following: Glance, Keystone, MySQL, Nova Multi Controller, Nova Volume, Swift Ring Compute, Swift Proxy.	
<b>Crowbar</b>	The code name for a bootstrap installer.	PXE Server
<b>DevOps</b>	An operational model for managing data centers using automated deployments.	Chef™, Puppet™
<b>Glance</b>	The OpenStack image cache.	
<b>Hypervisor</b>	Software that runs virtual machines (VMs).	KVM, Xen, VMware, HyperV
<b>iDRAC</b>	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
<b>LOM</b>	LAN on motherboard.	
<b>Node</b>	One of the servers in the system. A single chassis (sometimes called a server) can have multiple nodes.	Host, Box, Unit
<b>Nova</b>	The OpenStack Compute module for virtual machine deployment.	EC2 API
<b>Sled</b>	A server that is part of a shared infrastructure chassis, such as a Dell™ PowerEdge™ C6220 or C6105.	
<b>Swift</b>	A reference to OpenStack storage.	S3 API

### Reference Architecture 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 2012 OpenStack release is known as Essex. Crowbar ships with Essex Version 2012.1.2. <http://wiki.openstack.org/ReleaseNotes/2012.1.2>. This release builds on Diablo by focusing on quality, usability, and extensibility across enterprise, service provider, and high performance computing (HPC) deployments.

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. Planning for a migration is essential for the success of future releases. Upgrading to the latest stable software release is key to the CloudOps approach advocated by this reference architecture.

#### Taxonomy

In the Essex design, the Dell OpenStack-Powered Cloud Solution contains the core components of a typical OpenStack solution (Nova, Horizon, Swift, Glance, Keystone, Nova-Volume), 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 CloudOps<sup>1</sup> 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.

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.

---

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

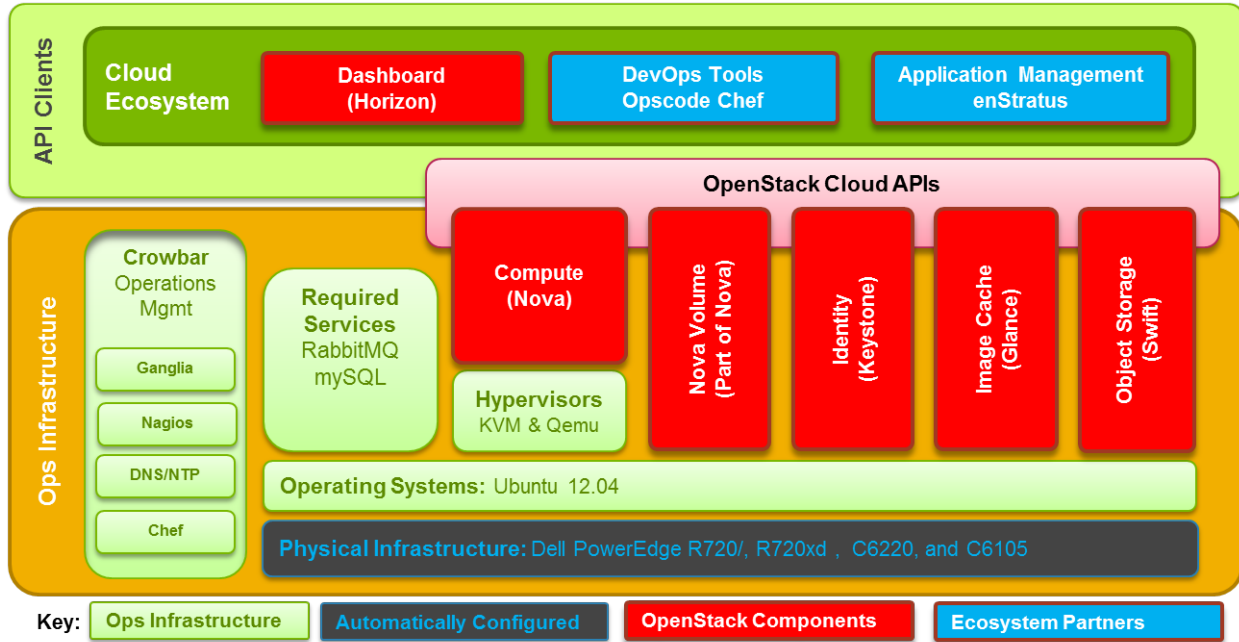


Figure 1: OpenStack Taxonomy

## 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, C6105, 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 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.*

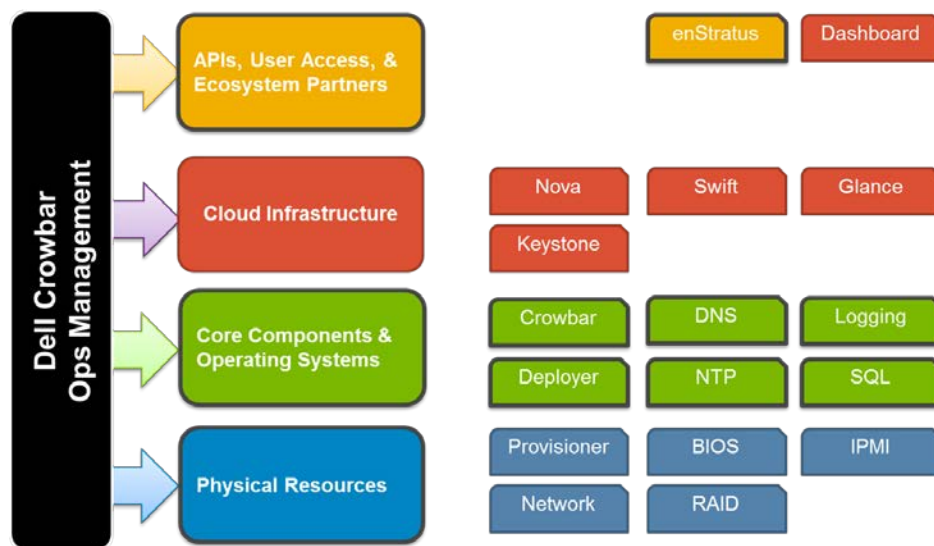
Crowbar provides a user interface (UI) and command-line view into the state of the nodes as they join the environment. Once the nodes have joined, use the API-based interfaces to assign that node to a service to provide a specific function. Crowbar has preconfigured automation that deploys OpenStack 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 six (6) primary components of OpenStack Essex: Compute (Nova), Object Storage (Swift), an Image Service (Glance), Keystone, Horizon, and Nova-Volume. Additional components will be included in the next release.

*Note: For a complete overview of OpenStack software, visit [www.OpenStack.org](http://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](http://www.Dell.com/OpenStack).*

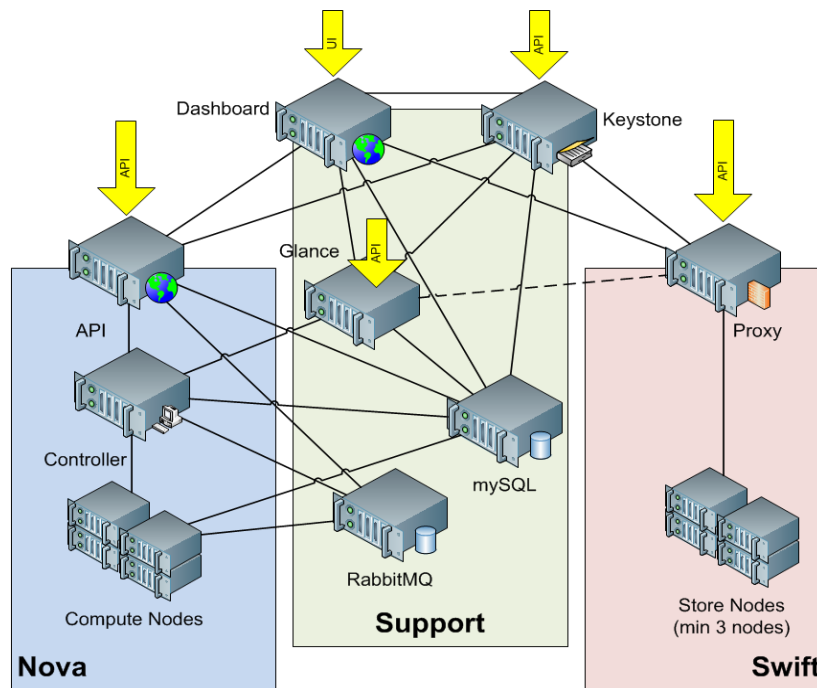
## Release Schedule

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

The next release, Folsom, will be delivered in Q4 2012, with Grizzly slated for Q2 2013. More information on OpenStack release schedules can be gathered at [www.openstack.org](http://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	URL
<b>Authentication</b>	Keystone	<p><a href="http://openstack.org/projects/">http://openstack.org/projects/</a></p> <p>Identity Service provides unified authentication across all OpenStack projects and integrates with existing authentication systems.</p> <p><a href="http://openstack.org/projects/">http://openstack.org/projects/</a></p>
<b>Dashboard/ Portal</b>	Horizon	<p><a href="http://openstack.org/projects/storage/">http://openstack.org/projects/storage/</a></p> <p>OpenStack Dashboard enables administrators and users to access and provision cloud-based resources through a self-service portal.</p>
<b>Object Storage</b>	Swift	<p>OpenStack Object Storage (code-named 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.</p> <p>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.</p> <p><a href="http://openstack.org/projects/compute/">http://openstack.org/projects/compute/</a></p>
<b>Compute/ IaaS</b>	Nova	<p>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.</p>
<b>Virtual Images</b>	Glance	<p><a href="http://openstack.org/projects/image-service">http://openstack.org/projects/image-service</a></p> <p>OpenStack Image Service (code-named 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.</p>

<http://docs.openstack.org/diablo/openstack-compute/starter/content/Nova-volume-d1e2000.html>

## Nova-Volume

Nova

Nova-Volume provides persistent block storage comparable with Amazon's Elastic Block Store. The storage on the instances is non-persistent in nature and hence any data you generate and store on the file system on the first disk of the instance gets lost when the instance is terminated. You will need to use persistent volumes provided by Nova-Volume if you want any data generated during the life of the instance to persist after the instance is terminated.

---

## 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.

### Deployment

Deployment consists of two phases. The first phase requires installing and configuring the admin node. The admin node controls the second phase for the rest of the machines in the deployment.

The first 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 will need to be configured and finalized. This will be done through editing some files and running a script. Once the one-time task of installing and configuring the admin node completes, 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 node on.
3. Wait until the Crowbar UI reports complete.

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.

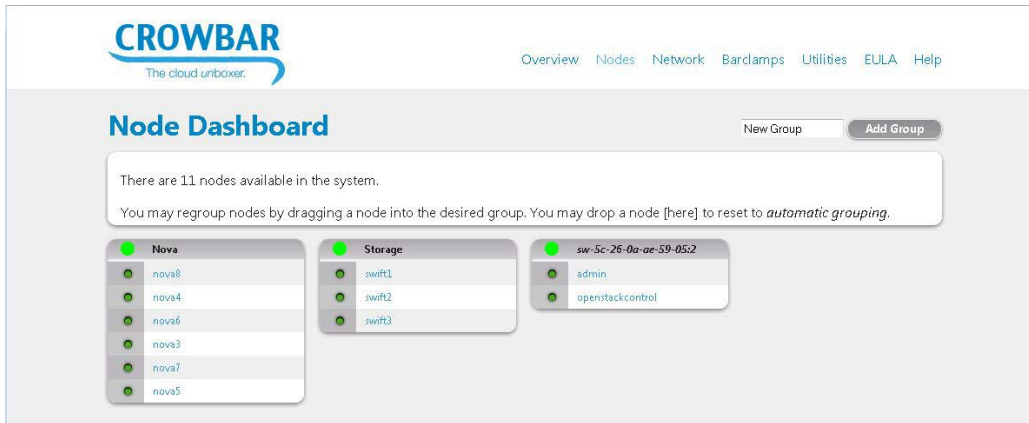


Figure 4: Crowbar Dashboard

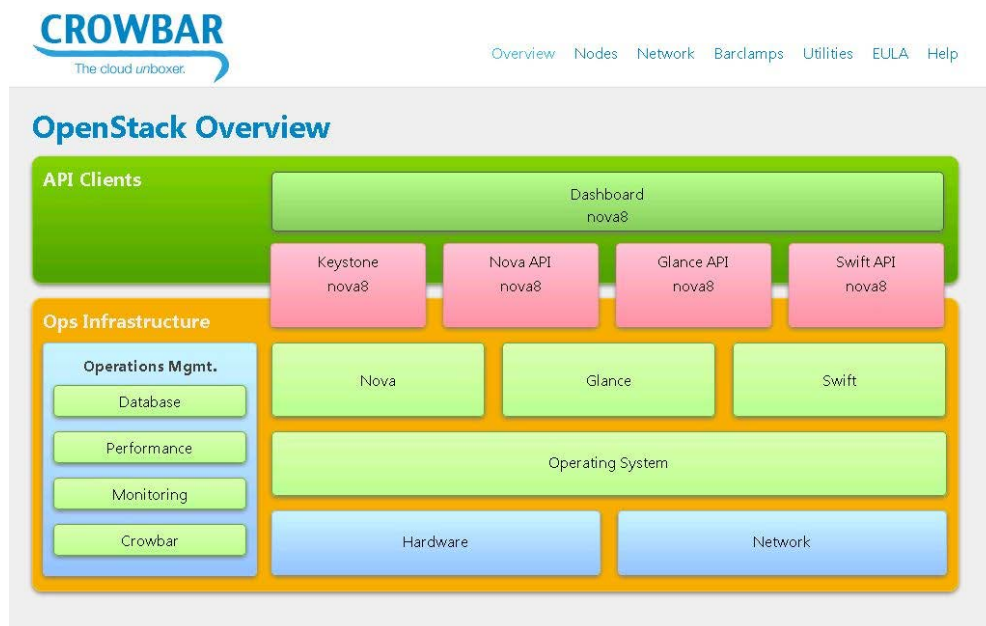


Figure 5: Crowbar Overview

## Nagios/Ganglia

Once the nodes have joined, the system employs Nagios and Ganglia to provide additional status monitoring, performance data gathering, and alerting. Nagios is the primary agent for alerting and node status monitoring. Ganglia has performance monitoring add-ins that directly tie in with the OpenStack integration. Configure Nagios to create alerts from Ganglia. Administrators can decide to turn off or replace these systems with their own.

The open source pages are located at:

- Nagios: <http://www.nagios.org>
- Ganglia: <http://ganglia.sourceforge.net/>

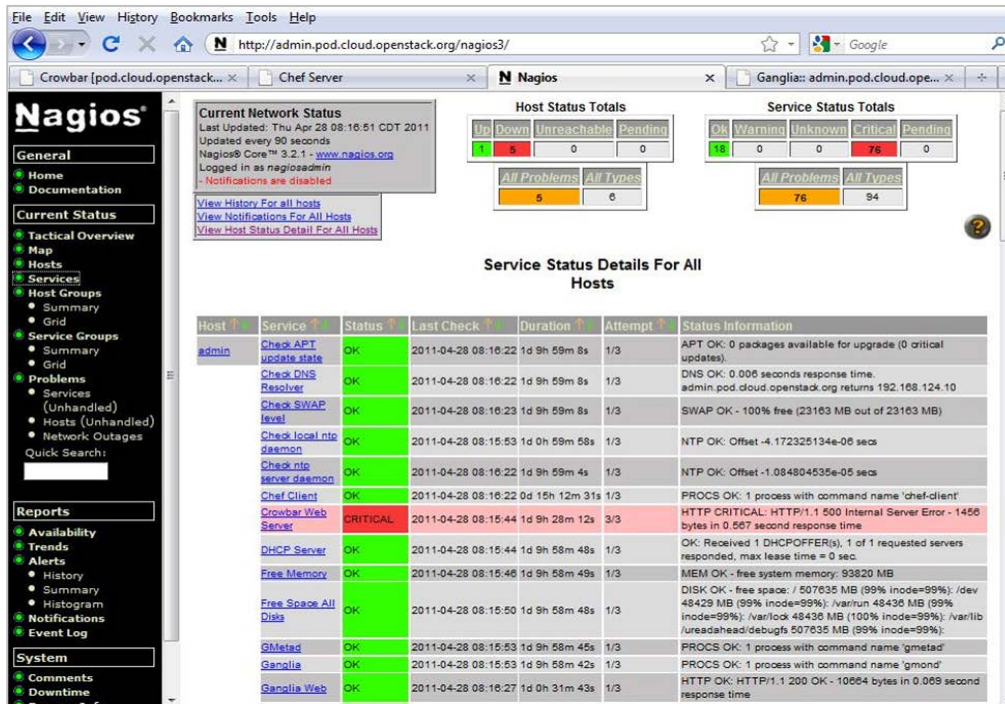


Figure 6: Nagios Provides Alerting and Node Status Monitoring

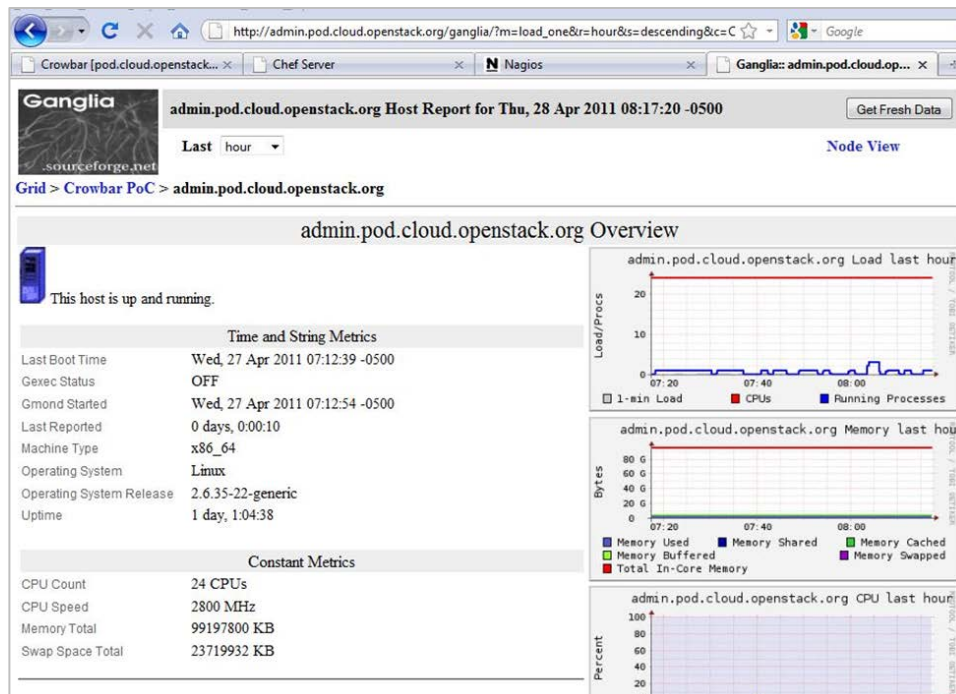


Figure 7: Ganglia Provides Performance Monitoring Add-ins

### Hardware Components

---

#### Multiple Solution Configuration Options

The Dell OpenStack-Powered Cloud Solution has been designed to grow with your needs. The building blocks of the evaluation configuration can be repurposed into a pilot system.

#### Evaluation / Proof of Concept Configuration

The minimum or starter configuration for the solution is dependent on the goal. For a Nova (Compute) only solution, you can use two (2) four-sled Dell™ PowerEdge C6220 servers, or three (3) two-sled Dell PowerEdge C6105 servers, or six (6) Dell PowerEdge R720 servers. For other configurations, please see Table 5: PowerEdge C6220—Compute, Table 6: PowerEdge C6220/R720xd Compute/Object Store, Table 7: PowerEdge C6105 Hybrid—Compute/Object Store, or Table 8: PowerEdge R720/R720xd Compute/Object Store. These configurations can be easily upgraded with little change to the network or power. Configurations beyond 15 chassis are easily accomplished by adding additional building blocks of servers and networking.

#### Solution Bundling and Delivery

In this release of the solution, hardware options are given to meet the customer needs. The three (3) platforms the Reference Architecture uses are the PowerEdge C6105, the PowerEdge C6220, or the PowerEdge R720/R720xd. Each of the platforms brings the best technology for high performance cloud environments.

The Dell PowerEdge C6220 server is designed to deliver the highest possible performance in a dense, scale-out package. This shared-infrastructure server offers a mix of performance, flexibility, and efficiency that makes it an exceptional building block for render farms, as well as high-performance computing, data analytics, cloud computing, and Web 2.0 environments. The PowerEdge C6220 chassis supports up to four independent server nodes in a 2U form factor, using the latest Intel® Xeon® processor E5-2600 product family, helping to boost performance when compared to the Intel® Xeon® processor 5600 series used in previous generations of the PowerEdge C6100 server.

The PowerEdge C6105 is a multi-node shared infrastructure platform delivered in a 2U chassis. There are two or four (in PowerEdge C6105) compute nodes (servers) that install in the chassis. The PowerEdge C6105 configuration for this release is the two-node configuration. The PowerEdge C6105 is a dual-socket AMD Opteron™ 4000 series processor server. The only difference is the form factor and number of PCIe slots. Each compute node in the PowerEdge C6105 chassis has access to 12 hot-swappable 2.5-inch drives.

The PowerEdge R720xd two sockets, 2U rack server combines greatly expandable storage, high density memory footprint, and up to 10-GB speeds. Best for large mail servers, big data analytics, Swift distributed computing where local storage is key, distributed storage, and hosting.

**Table 3: Deployment Platforms**

<b>Compute Node</b>	PowerEdge C6220 4-sled	PowerEdge C6105 2-sled	PowerEdge R720xd
<b>Platform</b>	PowerEdge 6220 6 2.5-inch per node	PowerEdge C6105 12 2.5-inch drive, 2- node BP per node	PowerEdge R720xd
<b>CPU</b>	Intel E56-2670 (8-core)	AMD Opteron 4000 series processor	Intel Xeon E5-2640(6- core)
<b>RAM</b>	128 GB per node	96 GB (1333 MHz) per node	96 GB
<b>Additional Network Controller</b>	None	None	None
<b>RAID Controller</b>	LSI 9265 Card	LSI 2008	PERC710
<b>Disk</b>	6 x 1-TB SATA 7.2K 2.5- inch per node	12 x 1-TB 2.5-inch Near Line SAS 7.2K per node	24 x 1-TB SATA 7.2K 2.5- inch
<b>RAID</b>	(See Deployment Guide)	(See Deployment Guide)	(See Deployment Guide)
<b>Cluster Switch</b>	Force10 S60	Force10 S60	Force10 S60

**Table 4: Admin and Controller System**

<b>Compute Node</b>	PowerEdge C6105 2-sled	PowerEdge R720
<b>Platform</b>	PowerEdge C6105 6 2.5-inch drive, 2-node BP per node	PowerEdge R720
<b>CPU</b>	AMD Opteron 4000 series processor	Intel Xeon E5-2640(6- core)
<b>RAM</b>	96 GB (1333-MHz) per node	96 GB
<b>Additional Network Controller</b>	None	None
<b>RAID Controller</b>	LSI 2008	PERC710
<b>Disk</b>	6 x 1-TB 2.5-inch Near Line SAS 7.2K per node	6 x 600-GB 2.5-inch
<b>RAID</b>	(See Deployment Guide)	(See Deployment Guide)
<b>Cluster Switch</b>	Force10 S60	Force10 S60



### Note about the AMD Opteron™ 4000 series processor:

The AMD Opteron 4100 Series processor is the world's lowest power-per-core server processor, making the PowerEdge C6105 server optimized for performance per watt per dollar. Ideal for workloads such as web/cloud and IT infrastructure, the AMD Opteron 4000 Series platform was designed from the ground up to handle demanding server workloads at the lowest available energy draw. AMD Opteron 4000 series processors:

- Offer four-core performance at less than 6W per core ACP and TDP
- Use up to 24 percent less power than previous AMD processor generations
- Provide the right mix of power, price, and performance with a long product life cycle ideal for scale-out data centers

The tables below highlight the four configurations using the six-node configuration and how each node/sled is allocated in a configuration that allows you to test both Nova Compute and Swift Storage. Please discuss your use case with your Dell sales representative to determine the best solution for you.

**Table 5: PowerEdge C6220—Compute**

Role	Node Type	Raid Config	Bios Config
Admin	PowerEdge C6220	RAID 10 – 6 drives	Virtualization
Controller	PowerEdge C6220	RAID 10 – 6 drives	Virtualization
Nova-Compute	PowerEdge C6220	RAID 10 – 6 drives	Virtualization
Nova-Compute	PowerEdge C6220	RAID 10 – 6 drives	Virtualization
Nova-Compute	PowerEdge C6220	RAID 10 – 6 drives	Virtualization
Nova-Compute	PowerEdge C6220	RAID 10 – 6 drives	Virtualization
Nova-Compute	PowerEdge C6220	RAID 10 – 6 drives	Virtualization
Nova-Compute	PowerEdge C6220	RAID 10 – 6 drives	Virtualization

**Table 6: PowerEdge C6220/R720xd Compute/Object Store**

Role	Node Type	Raid Config	Bios Config
Admin	PowerEdge C6220	RAID 10 – 6 drives	Virtualization
Controller	PowerEdge C6220	RAID 10 – 6 drives	Virtualization
Nova-Compute	PowerEdge C6220	RAID 10 – 6 drives	Virtualization
Nova-Compute	PowerEdge C6220	RAID 10 – 6 drives	Virtualization
Swift-Storage	PowerEdge R720xd	JBOD – 24 drives	Storage
Swift-Storage	PowerEdge R720xd	JBOD – 24 drives	Storage
Swift-Storage	PowerEdge R720xd	JBOD – 24 drives	Storage



**Table 7: PowerEdge C6105 Hybrid—Compute/Object Store**

Role	Node Type	Raid Config	Bios Config
<b>Admin</b>	PowerEdge C6105	RAID 10 – 6 drives	Virtualization
<b>Controller</b>	PowerEdge C6105	RAID 10 – 6 drives	Virtualization
<b>Nova-Compute</b>	PowerEdge C6105	RAID 10 – 12 drives	Virtualization
<b>Nova-Compute</b>	PowerEdge C6105	RAID 10 – 12 drives	Virtualization
<b>Swift-Storage</b>	PowerEdge C6105	JBOD – 12 drives	Storage
<b>Swift-Storage</b>	PowerEdge C6105	JBOD – 12 drives	Storage

**Table 8: PowerEdge R720/R720xd Compute/Object Store**

Role	Node Type	Raid Config	Bios Config
<b>Admin</b>	PowerEdge R720	RAID 10 – 6 drives	Storage
<b>Controller</b>	PowerEdge R720	RAID 10 – 6 drives	Virtualization
<b>Nova-Compute</b>	PowerEdge R720	RAID 10 – 12 drives	Virtualization
<b>Nova-Compute</b>	PowerEdge R720	RAID 10 – 12 drives	Virtualization
<b>Swift-Storage</b>	PowerEdge R720xd	JBOD – 24 drives	Storage
<b>Swift-Storage</b>	PowerEdge R720xd	JBOD – 24 drives	Storage
<b>Swift-Storage</b>	PowerEdge R720xd	JBOD – 24 drives	Storage

### Expanding the Solution Beyond the Six-node Solution

Expanding the base configuration to full racks or larger is the logical growth path. Then by scaling in rack increments one can scale up to 60 PowerEdge C6220/C6105 or 45 PowerEdge R720/R720xd nodes. This scale system can be accommodated without major changes in the networking infrastructure.

To size your solution, please discuss your specific needs with your sales representative. These configurations are adjusted to meet individual needs.

**Table 9: Additional and Optional Hardware Recommendations**  
(See Appendix G: Solution Starter Configuration Expansion Options)

Compute Node	(Optional) PowerEdge C6100 (2- or 4-node)	(Optional) PowerEdge C6105 (4-node)	PowerEdge C2100
<b>Platform</b>	PowerEdge C6100	PowerEdge C6100 6 2.5-inch drive, 4-node BP per node	PowerEdge C2100 12 3.5-inch Drives
<b>CPU</b>	2x E5-2640 (6-core)	Intel Xeon E5620	Intel Xeon X5650
<b>RAM</b>	96 GB	96 GB (1333-MHz) per node	48 GB (1333-MHz) per node
<b>Additional Network Controller</b>	None	None	None
<b>RAID Controller</b>	LSI 2008 Mezz	LSI 2008	LSI 2008
<b>Disk</b>	6 x 1-TB SATA 7.2K 2.5- inch	6 x 600-GB 2.5-inch SAS 10K per node	12 X 1-TB 3.5-inch SATA 7.2K
<b>RAID</b>	(See Deployment Guide)	(See Deployment Guide)	(See Deployment Guide)
<b>Cluster Switch</b>	Force10 S60	Force10 S60	Force10 S60

### Site Preparation Needed for the Deployment

Solution deployment needs some preliminary preparation. The solution does not supply any firewalls or load-balancers. You may want to use firewalls or load-balancing to access and use portions of the solution. Indirectly, a bastion host, installed behind appropriate site-specific security systems, can be used to access the solution and the VMs remotely. This means that direct access to local/internal/external networks should not be done.

For the setup of the admin node, connect a keyboard, mouse, and monitor. In addition, all that is required is a laptop or another machine that can run VM-Player and connect to the admin node via a crossover network cable.

Estimate the electrical power and cooling usage using the Dell Energy Smart Solution Advisor: [http://www.dell.com/content/topics/topic.aspx/global/products/pedge/topics/en/config\\_calculator?c=us&cs=555&l=en&s=biz](http://www.dell.com/content/topics/topic.aspx/global/products/pedge/topics/en/config_calculator?c=us&cs=555&l=en&s=biz)

You can use this tool to plan the appropriate PDU and make sure the cooling is adequate.

### Network Overview

---

Due to the nature of the different software used, the network is set up as flat as possible using a dedicated BMC port and bonded LOMs. Crowbar manages all networks, and comes out of the box preconfigured to allow the initial configuration to come up quickly by predefining the storage, admin, public, and BMC networks.

The Crowbar network configuration can be customized to better map to site-specific networking needs and conventions. These changes include adding additional vLANs, changing vLAN mappings, and teaming NICs. Making these changes is beyond the scope of this document but available on the Crowbar open source site (<http://github.com/dellcloudedge/crowbar>).

The Crowbar deployed Openstack solution needs at least four (4) vLANs to support the cluster. **Error! Reference source not found.** reviews what each of the networks are and how they are used. The defaults listed are what a base install will configure. The Admin and Public/Nova Floating networks can be extended to the Corporate LAN or the internet as needed. It highly recommended that a firewall and other security devices be placed in line. Figure 8: Network Wiring gives an overview of some possible ports that would need to be accessible to manage the Crowbar Cluster and its components. Please read the document *Dell Openstack Powered Cloud Solution Deployment Guide* for further ports and configurations.




Figure 8: Network Wiring

## High-level Description

All servers in an OpenStack Cluster are tied together using TCP/IP networks. These networks form a data interconnect across which individual servers pass data back and forth, return query results, and load/unload data. These networks are also used for management.

The admin node manages all the cluster compute and storage nodes. It assigns the other nodes IP addresses, PXE boots them, configures them, and provides them the necessary software for their roles. To provide these services, the admin node runs Crowbar, Chef, DHCP, TFTP, NTP, and other services, and this must be the only DHCP server visible to the compute and storage nodes. Details follow:

- **Crowbar Server**—manages all nodes, supplying configuration of hardware and software.
- **DHCP server**—assigns and manages IPs for the compute and storage nodes.
- **NTP server** (Network Time Protocol server)—makes sure all nodes are keeping the same clock.
- **TFTP server**—PXE boots compute and storage nodes with a Linux kernel. The TFTP server services any PXE boot request it receives with its default options.
- **DNS server**—manages the name resolution for the nodes and can be configured to provide external name forwarding.

## Network Cabling and Redundancy

Figure 8 shows an example of network connectivity inside the cluster with 1GbE links. Please see your Dell Sales Representative about using 10GbE.

### Network Connectivity

All nodes have two 1Gb NICs. The admin node configures the BMC and the OS is configured to bond the two LOMs. Each NIC and BMC is cabled to the Dell Force10 S60 switch per the cabling setup found in Tables 10-13.

## Logical Network Configuration

The solution has been architected for minimal configuration tasks, but still maintains a logical segregation of traffic. There are many networks (admin/internal, storage, and external networks). Each is segmented into separate vLANs:

<b>Admin vLAN</b>	Used for administrative functions such as Crowbar node installation, TFTP booting, DHCP assignments, KVM, system logs, backups, and other monitoring. There is only one vLAN set up for this function and it is spanned across the entire network.  BMC vLAN. This is a subset of the Admin vLAN that is used for connecting to the BMC/iDRAC of each node.  These networks can have an external router, but the external/public vLAN is the preferred method.
<b>Storage vLAN</b>	Used by the Swift storage system for replication of data between machines, monitoring of data integrity, and other storage-specific functions. (802.1q tagged) This is used in the cluster only.
<b>Nova Fixed</b>	Nova internode communications, including guests talking to their host systems. This is used in the Cluster only.
<b>External/Public and Nova Floating vLANs</b>	Used for connections to devices that are external to the OpenStack cloud infrastructure. These include externally visible services, such as load balancers and web servers. Use one or many of these networks depending on the need to segregate traffic among groups of servers. (802.1q tagged)  Nova Floating vLAN. This is a subset of the Public network, and is used to give Nova guests publically accessible IP addresses.

*Note: Unlike the external and internal vLANs, the administrative vLAN does not use 802.1q vLAN tagging.*

**Table 10: Default Crowbar Networks**

Crowbar Network	Default VLAN	Default IP Range	802.1q tagging	Accessibility
Admin/MGMT (or Production)	100	192.168.124.0/24	No	Cluster Administrators
Storage	200	192.168.125.0/24	Yes	Cluster internal
Public	300	192.168.123.0/24	Yes	Public
Nova Floating	300	192.168.123.128/26	Yes	Yes (inherited from Public)
Nova Fixed	500	192.168.126.0/24	Yes	Cluster internal

## Stacked Top-of-Rack Switches

When deployed, the top-of-rack (ToR) switches are physically stacked together. Stacking the switches offers some significant benefits.


### Benefits

<b>Improved Manageability</b>	All switches in the stack are managed as a single switch.
<b>Efficient Spanning Tree</b>	The stack is viewed as a single switch by the Spanning Tree Protocol.
<b>Link Aggregation</b>	Stacking multiple switches in a chassis allows a link aggregation group (LAG) across ports on different switches in the stack.
<b>Reduced Network Traffic</b>	Traffic between the individual switches in a stack is passed across the stacking cable, reducing the amount of traffic passed upstream to network distribution switches.
<b>Higher Speed</b>	The stacking module supports a higher data rate than the 10GbE uplink module (supports 12Gb per stack port offering 24Gb between switches).
<b>Lower Cost</b>	Uplink ports are shared by all switches in the stack, reducing the number of distribution switch ports necessary to connect modular servers to the network.
<b>Simplified Updates</b>	The basic firmware management commands propagate new firmware versions and boot image settings to all switch stack members.

### Drawbacks

- Stacking cables are proprietary and come in only 1m and 3m lengths. This requires that the switches be in close proximity to each other.
- Stacking requires a ring topology for redundancy. This makes the distance limitation on the length of the stacking cables more of an issue.
- Errors in configuration propagate throughout the stack immediately.

## Physical Configuration



**Figure 9: Seven-node PowerEdge C6220/R720xd Compute/Object Store Starter Solution**

Figure 9 shows a starter solution with seven nodes installed. The cabling should be done as explained above in Figure 8 and per the Tables 10-13.

### Single Rack Expansion from Starter Configuration

You can build the four starter solutions to a full rack or multiple racks. For the PowerEdge C6220 Compute (**Table 5: PowerEdge C6220—Compute**). You can bring a single rack without adding any more networking to 28 nodes. This is done by adding 5(five) more PowerEdge C6220 Chassis to the solution. By adding a single network switch to handle all the management (BMC) traffic, you can increase this to 10 chassis for a total of 40 nodes in a rack. Beyond this will require additional networking and configuration.

You can build the PowerEdge C6220/R720XD (**Table 6: PowerEdge C6220/R720xd Compute/Object Store**), the PowerEdge C6105 (**Table 7: PowerEdge C6105 Hybrid—Compute/Object Store**), or the Power Edge R720/R720xd (**Table 8: PowerEdge R720/R720xd Compute/Object Store**) up to 30 total nodes with no network changes. Beyond 30 nodes will require network switch additions and configurations.




Figure 10: Full Rack PowerEdge C6105



Figure 11: PowerEdge C6220 28-node Compute Deployment

## Multi Rack Expansion

Expand the solution further by adding end-of-row (EoR) 10GbE switches and additional racks of equipment. The networking needs to use one of the various hyperscale networking deployments using multiple 10GB LAGs between each of the racks. This is not part of this reference architecture.






Figure 12: PowerEdge R720xd Swift Only Storage of 46 Nodes with C6105 Admin/Controller Node

## Network Port Assignments

Table 11: Starter Solution Node Network Cabling

Component	LOM0 <sup>4</sup>	LOM1 <sup>4</sup>	BMC
Node 1 <sup>1</sup>	SW1-1	SW2-1	SW1-31
Node 2 <sup>2</sup>	SW1-2	SW2-2	SW2-31
Node 3	SW1-3	SW2-3	SW1-32
Node 4	SW1-4	SW2-4	SW2-32
Node 5 <sup>3</sup>	SW1-5	SW2-5	SW1-33
Node 6	SW1-6	SW2-6	SW2-33
Node 7	SW1-7	SW2-7	SW1-34
Node 8	SW1-8	SW2-8	SW2-34

1. Node 1 is the Crowbar Admin Node.
2. Node 2 is the Openstack Controller when it is a R720/R720xd
3. Node 5 is the Openstack Controller when it is a C6220/C6105.
4. On Dell Force10 Switches you will need to start at Port 0 not Port

**Table 12: Full Rack Node Network Cabling**

Component	LOM0 <sup>6</sup>	LOM1 <sup>6</sup>	BMC
Node 1 <sup>1</sup>	SW1-1	SW2-1	SW1-31
Node 2 <sup>2</sup>	SW1-2	SW2-2	SW2-31
Node 3	SW1-3	SW2-3	SW1-32
Node 4	SW1-4	SW2-4	SW2-32
Node 5 <sup>3</sup>	SW1-5	SW2-5	SW1-33
Node 6	SW1-6	SW2-6	SW2-33
Node 7	SW1-7	SW2-7	SW1-34
Node 8	SW1-8	SW2-8	SW2-34
Node 9	SW1-9	SW2-9	SW1-35
Node 10	SW1-10	SW2-10	SW2-35
Node 11	SW1-11	SW2-11	SW1-36
Node 12	SW1-12	SW2-12	SW2-36
Node 13	SW1-13	SW2-13	SW1-38
Node 14	SW1-14	SW2-14	SW2-38
Node 15 <sup>4</sup>	SW1-15	SW2-15	SW1-39
Node 16	SW1-16	SW2-16	SW2-39
Node 17	SW1-17	SW2-17	SW1-40
Node 18	SW1-18	SW2-18	SW2-40
Node 19	SW1-19	SW2-19	SW1-41
Node 20	SW1-20	SW2-20	SW2-41
Node 21	SW1-21	SW2-21	SW1-42
Node 22	SW1-22	SW2-22	SW2-42
Node 23	SW1-23	SW2-23	SW1-43
Node 24	SW1-24	SW2-24	SW2-43
Node 25	SW1-25	SW1-44 – 1GBTSFP+	SW2-25
Node 26	SW1-26	SW2-44 – 1GBTSFP+	SW2-26
Node 27	SW1-27	SW1-45 – 1GBTSFP+	SW2-27
Node 28 <sup>5</sup>	SW1-28	SW2-45 – 1GBTSFP+	SW2-28
Node 29	SW1-29	SW1-46 – 1GBTSFP+	SW2-29
Node 30	SW1-30	SW2-46 – 1GBTSFP+	SW2-30

1. Node 1 is the Crowbar Admin Node.

2. Node 2 is the Openstack Controller when it is a PowerEdge R720/R720xd.

3. Node 5 is the Openstack Controller when it is a PowerEdge C6220/C6105.

4. Node 15 is the last node when using only the PowerEdge R720/R720xd platform.

5. Node 28 is the last node when using PowerEdge C6220 only.

6. On Dell Force10 switches you will need to start at Port 0 not Port 1.

## Appendix A: PowerEdge C6220 Configuration—Hardware Bill of Materials

Table 13: C6220 Config

(Please note there are two parts to a PowerEdge C6220 order the Chassis and the sleds)

### Chassis:

<b>Base Unit:</b>	PowerEdge C6000, 2/4Motherboard, 2.5 in Hard Drives with Backplane (225-2582)
<b>Video Memory:</b>	SHIP,C6220,DAO (331-5531)
<b>Video Memory:</b>	PowerEdge C6000/C6220 Shipping (331-5532)
<b>Operating System:</b>	No Factory Installed Operating System (420-3323)
<b>TBU:</b>	PowerEdge C6220 Documentation (331-5527)
<b>CD-ROM or DVD-ROM Drive:</b>	PowerEdge C Static Rails, Toolless (331-1619)
<b>CD-ROM or DVD-ROM Drive:</b>	Power Supply, 1100W, Redundant Capable (331-5518)
<b>CD-ROM or DVD-ROM Drive:</b>	Power Supply, 1100W, Redundant Capable (331-5518)
<b>CD-ROM or DVD-ROM Drive:</b>	Label,Regulatory,1100W,C6220 (331-5522)
<b>CD-ROM or DVD-ROM Drive:</b>	Power Cord, C13 to C14, PDU Style, 12 Amps, 2 meter, Qty 1 (330-7353) - Quantity 2
<b>Service:</b>	GCP,CSTM,BASIC,PROGRAM,SPT,DCS,4 (929-5569)
<b>Service:</b>	Dell Hardware Limited Warranty Plus On Site Service Initial Year (937-8137)
<b>Service:</b>	Dell Hardware Limited Warranty Plus On Site Service Extended Year (937-8147)
<b>Service:</b>	Non-Mission Critical: 4-Hour 7x24 On-site Service After Problem Diagnosis, Initial Year (937-8427)
<b>Service:</b>	Non-Mission Critical: 4-Hour 7x24 On-site Service After Problem Diagnosis, 2 Year Extended (937-8457)
<b>Service:</b>	ProSupport: 7x24 HW / SW Tech Support and Assistance, 3 Year (937-8607)
<b>Service:</b>	Dell ProSupport. For tech support, visit <a href="http://support.dell.com/ProSupport">http://support.dell.com/ProSupport</a> or call 1-800-945-3355 (989-3439)
<b>Installation:</b>	On-Site Installation Declined (900-9997)

### Sled Order

<b>Base Unit:</b>	PowerEdge C6220,1U Motherboard with 2.5 SATA/SAS Hard Drives, Direct (225-3819)
<b>Memory:</b>	128GB Memory (16x8GB),1600Mhz, Dual Ranked RDIMMs for 2 Processors (317-8998)

## Dell OpenStack-Powered Cloud Solution: Reference Architecture Guide

---

<b>Memory:</b>	Info, Memory for Dual Processor selection (468-7687)
<b>Video Memory:</b>	System ordered as part of Multipack order (330-4118)
<b>Operating System:</b>	No Factory Installed Operating System (420-3323)
<b>TBU:</b>	C4, 1U, LSI 9265 Card supporting up to 3,3.5D, 6,2.5D/E Hard Drives (342-3931)
<b>TBU:</b>	LSI 9265 SAS Low Profile Card (342-5124)
<b>TBU:</b>	Dual Processor Option (317-4928)
<b>TBU:</b>	Thermal Heatsink (317-9107)
<b>TBU:</b>	Thermal Heatsink (317-9107)
<b>TBU:</b>	Intel Xeon E5-2670 2.60GHz, 20M Cache, 8.0GT/s QPI, Turbo, 8C, 115W, Max Mem 1600MHz (317-9596)
<b>TBU:</b>	Intel Xeon E5-2670 2.60GHz, 20M Cache, 8.0GT/s QPI, Turbo, 8C, 115W (317-9610)
<b>TBU:</b>	Carrier, Hard Drive,2.5,2LED,C6145 (342-2306) - Quantity 6
<b>TBU:</b>	1TB, 7.2K RPM,SATA, 3Gbps, 2.5in, Hot Plug Hard drive (342-3661) - Quantity 6
<b>Service:</b>	GCP,CSTM,BASIC,PROGRAM,SPT,DCS,4 (929-5569)
<b>Service:</b>	Dell Hardware Limited Warranty Plus On Site Service Initial Year (937-9827)
<b>Service:</b>	Dell Hardware Limited Warranty Plus On Site Service Extended Year (937-9837)
<b>Service:</b>	Non-Mission Critical: 4-Hour 7x24 On-site Service After Problem Diagnosis, Initial Year (938-0097)
<b>Service:</b>	Non-Mission Critical: 4-Hour 7x24 On-site Service After Problem Diagnosis, 2 Year Extended (938-0127)
<b>Service:</b>	ProSupport: 7x24 HW / SW Tech Support and Assistance, 3 Year (938-0277)
<b>Service:</b>	Dell ProSupport. For tech support, visit <a href="http://support.dell.com/ProSupport">http://support.dell.com/ProSupport</a> or call 1-800-945-3355 (989-3439)
<b>Installation:</b>	On-Site Installation Declined (900-9997)
<b>Support:</b>	Proactive Maintenance Service Declined (926-2979)

## Appendix B: PowerEdge C6105 Configuration—Hardware Bill of Materials

Table 14: Admin/Controller PowerEdge C6105 Node

Base Unit	PowerEdge C6105 Chassis w/ 2 System Boards and support for 2.5-inch Hard Drives (225-0024)
Processor	Dual Processor Option (317-4928)
Processor	4 THRM, HTSNK, CLE, 95W, PowerEdge C6105 (317-5758)
Processor	4 AMD Opteron 4184, 6C 2.8GHz, 3M L2/6M L3, 1333Mhz Max Mem (317-5767)
Memory	96GB Memory (12 x 8GB), 1333 MHz Dual Ranked RDIMMs, Low Volt (317-5559)
Memory	96GB Memory (12 x 8GB), 1333 MHz Dual Ranked RDIMMs, Low Volt (317-5559)
Memory	Info, Memory for Dual Processor selection (468-7687)
Video Memory	C6105 Shipping (331-2854)
Hard Drive	HD Multi-Select (341-4158)
Operating System	No Factory Installed Operating System (420-3323)
Operating System	No Factory Installed Operating System (420-3323)
Mouse	CARR, HD, 2.5, 2LED, PowerEdge C6100, MLK (342-1032) - Quantity 6
Mouse	CARR, HD, 2.5, 2LED, PowerEdge C6100, MLK (342-1032) - Quantity 6
Mouse	1TB, 7.2K RPM, Near Line SAS, 6Gbps, 2.5-inch, Hot Plug Hard drive (342-3160) - Quantity 6
Mouse	1TB, 7.2K RPM, Near Line SAS, 6Gbps, 2.5-inch, Hot Plug Hard drive (342-3160) - Quantity 6
Feature	LSI 2008 SATA/SATA Mezz Card, PowerEdge C6100 (342-1050)
Feature	LSI 2008 SATA/SATA Mezz Card, PowerEdge C6100 (342-1050)
Feature	Add-in LSI 2008 SAS/SATA Mezz Card supporting up to 12, 2.5-inch HDs SAS/SATA - No RAID (342-1883)
Feature	Add-in LSI 2008 SAS/SATA Mezz Card supporting up to 12, 2.5-inch HDs SAS/SATA - No RAID (342-1883)
Feature	PowerEdge C6100/C6105 Static Rails, Tool-less (330-8483)
Service	Dell Hardware Limited Warranty Initial Year (925-0527)
Service	DECLINED CRITICAL BUSINESS SERVER OR STORAGE SOFTWARE SUPPORT PACKAGE-CALL YOUR DELL SALES REP IF UPGRADE NEED (928-3229)
Service	GRBO, CSTM, BASIC, PROGRAM, SPT, DCS, 4 (929-5569)
Service	Basic Hardware Services: Business Hours (5X10) Next Business Day On Site Hardware Warranty Repair 2 Year Exten (929-5672)
Service	Dell Hardware Limited Warranty Extended Year(s) (931-6268)

<b>Service</b>	Basic Hardware Services: Business Hours (5X10) Next Business Day On Site Hardware Warranty Repair Initial Year (935-0160)
<b>Service</b>	Basic support covers SATA Hard Drive for 1 year only regardless of support duration on the system (994-4019)
<b>Installation</b>	On-Site Installation Declined (900-9997)
<b>Misc</b>	Power Supply, 1100W, Redundant Capable (330-8537)
<b>Misc</b>	Power Supply, 1100W, Redundant Capable (330-8537)
<b>Misc</b>	Label, Regulatory, DAO, 750/1100, PowerEdge 6105 (331-0588)
<b>Misc</b>	Power Cord, C13 to C14, PDU Style, 12 Amps, 2 meter, Qty 1 (330-7353)
<b>Misc</b>	Power Cord, C13 to C14, PDU Style, 12 Amps, 2 meter, Qty 1 (330-7353)

**Table 15: 2 Sled Compute Node—PowerEdge C6105 Configuration**

<b>Base Unit</b>	<b>PowerEdge C6105 Chassis with 2 System Boards and support for 2.5-inch Hard Drives (225-0024)</b>
<b>Processor</b>	Dual Processor Option (317-4928)
<b>Processor</b>	4 THRM, HTSNK, CLE, 95W, PowerEdge C6105 (317-5758)
<b>Processor</b>	4 AMD Opteron 4184, 6C 2.8GHz, 3M L2/6M L3, 1333Mhz Max Mem (317-5767)
<b>Memory</b>	96GB Memory (12 x 8GB), 1333 MHz Dual Ranked RDIMMs, Low Volt (317-5559)
<b>Memory</b>	96GB Memory (12 x 8GB), 1333 MHz Dual Ranked RDIMMs, Low Volt (317-5559)
<b>Memory</b>	Info, Memory for Dual Processor selection (468-7687)
<b>Video Memory</b>	PowerEdge C6105 Shipping (331-2854)
<b>Hard Drive</b>	HD Multi-Select (341-4158)
<b>Operating System</b>	No Factory Installed Operating System (420-3323)
<b>Operating System</b>	No Factory Installed Operating System (420-3323)
<b>Mouse</b>	CARR, HD, 2.5, 2LED, PowerEdge C6100, MLK (342-1032) - Quantity 12
<b>Mouse</b>	CARR, HD, 2.5, 2LED, PowerEdge C6100, MLK (342-1032) - Quantity 12
<b>Mouse</b>	1TB, 7.2K RPM, Near Line SAS, 6Gbps, 2.5-inch, Hot Plug Hard drive (342-3160) - Quantity 12
<b>Mouse</b>	1TB, 7.2K RPM, Near Line SAS, 6Gbps, 2.5-inch, Hot Plug Hard drive (342-3160) - Quantity 12
<b>Feature</b>	LSI 2008 SATA/SATA Mezz Card, C6100 (342-1050)
<b>Feature</b>	LSI 2008 SATA/SATA Mezz Card, C6100 (342-1050)
<b>Feature</b>	Add-in LSI 2008 SAS/SATA Mezz Card supporting up to 12, 2.5-inch HDs SAS/SATA - No RAID (342-1883)
<b>Feature</b>	Add-in LSI 2008 SAS/SATA Mezz Card supporting up to 12, 2.5-inch HDs SAS/SATA - No RAID (342-188)

## Dell OpenStack-Powered Cloud Solution: Reference Architecture Guide

---

<b>Feature</b>	PowerEdge C6100/C6105 Static Rails, Tool-less (330-8483)
<b>Service:</b>	Dell Hardware Limited Warranty Initial Year (925-0527)
<b>Service</b>	DECLINED CRITICAL BUSINESS SERVER OR STORAGE SOFTWARE SUPPORT PACKAGE-CALL YOUR DELL SALES REP IF UPGRADE NEED (928-3229)
<b>Service</b>	GRBO,CSTM,BASIC,PROGRAM,SPT,DCS,4 (929-5569)
<b>Service</b>	Basic Hardware Services: Business Hours (5X10) Next Business Day On Site Hardware Warranty Repair 2 Year Exten (929-5672)
<b>Service:</b>	Dell Hardware Limited Warranty Extended Year(s) (931-6268)
<b>Service</b>	Basic Hardware Services: Business Hours (5X10) Next Business Day On Site Hardware Warranty Repair Initial Year (935-0160)
<b>Service</b>	Basic support covers SATA Hard Drive for 1 year only regardless of support duration on the system (994-4019)
<b>Installation</b>	On-Site Installation Declined (900-9997)
<b>Misc</b>	Power Supply, 1100W, Redundant Capable (330-8537)
<b>Misc</b>	Power Supply, 1100W, Redundant Capable (330-8537)
<b>Misc</b>	Label, Regulatory, DAO, 750/1100, PowerEdge 6105 (331-0588)
<b>Misc</b>	Power Cord, C13 to C14, PDU Style, 12 Amps, 2 meter, Qty 1 (330-7353)
<b>Misc</b>	Power Cord, C13 to C14, PDU Style, 12 Amps, 2 meter, Qty 1 (330-7353)

---

## Appendix C: PowerEdge R720 and R720xd Bill of Materials

**Table 16: R720 Admin/Controller node**

Base Unit:	PowerEdge R720 (225-2133)
TBU:	PowerEdge R720 Shipping (331-4437)
TBU:	Risers with up to 4, x8 PCIe Slots + 2, x16 PCIe Slot (331-4439)
TBU:	iDRAC7 Enterprise (421-5339)
TBU:	Broadcom 5720 QP 1Gb Network Daughter Card (430-4418)
TBU:	3.5" Chassis with up to 8 Hard Drives (342-3587)
TBU:	Bezel (318-1375)
TBU:	Power Saving Dell Active Power Controller (330-5116)
TBU:	No RAID for H310 (1-16 HDDs) (331-4379)
TBU:	PERC H710 Integrated RAID Controller, 512MB NV Cache [342-3529] Quantity:1
TBU:	Intel Xeon E5-2640 2.50GHz, 15M Cache, 7.2GT/s QPI, Turbo, 6C, 95W [317-9595] Quantity:1
TBU:	Heat Sink for PowerEdge R720 and R720xd [331-4508] Quantity:1
TBU:	DIMM Blanks for Systems with 2 Processors (317-8688)
TBU:	Intel Xeon E5-2640 2.50GHz, 15M Cache, 7.2GT/s QPI, Turbo, 6C, 95W [317-9595] Quantity:1
TBU:	Heat Sink for PowerEdge R720 and R720xd [331-4508] Quantity:1
TBU:	8GB RDIMM, 1600 MHz, Standard Volt, Dual Rank, x4 (317-9648) - Quantity 12
TBU:	1600 MHz RDIMMS (331-4424)
TBU:	Performance Optimized (331-4428)
TBU:	1TB 7.2K RPM Near-Line SAS 6Gbps 3.5in Hot-plug Hard Drive (342-2098) - Quantity 6
TBU:	No System Documentation, No OpenManage DVD Kit (310-5171)
CD-ROM or DVD-ROM Drive:	DVD+/-RW, SATA, INTERNAL (313-9090)
CD-ROM or DVD-ROM Drive:	ReadyRails Sliding Rails Without Cable Management Arm (331-4434)
CD-ROM or DVD-ROM Drive:	Dual, Hot-plug, Redundant Power Supply (1+1), 1100W (331-4607)
CD-ROM or DVD-ROM Drive:	Power Cord, C13 to C14, PDU Style, 12 Amps, 2 foot, Qty 1 (330-3150) - Quantity 2
CD-ROM or DVD-ROM Drive:	No Operating System (420-6320)
CD-ROM or DVD-ROM Drive:	No Media Required (421-5736)
Service:	Basic Hardware Services: Business Hours (5X10) Next Business Day On Site Hardware Warranty Repair 2 Year Exten (936-4543)
Service:	Dell Hardware Limited Warranty Plus On Site Service Extended Year (939-2678)



Service:	Dell Hardware Limited Warranty Plus On Site Service Initial Year (939-2768)
Service:	Basic Hardware Services: Business Hours (5X10) Next Business Day On Site Hardware Warranty Repair Initial Year (988-9131)
Service:	Basic support covers SATA Hard Drive for 1 year only regardless of support duration on the system (994-4019)
Service:	DECLINED CRITICAL BUSINESS SERVER OR STORAGE SOFTWARE SUPPORT PACKAGE-CALL YOUR DELL SALES REP IF UPGRADE NEED (996-8029)
Installation:	On-Site Installation Declined (900-9997)
Support:	Proactive Maintenance Service Declined (926-2979)

**Table 17: R720xd Storage Node**

<b>Base Unit:</b>	<b>PowerEdge R720xd [225-2110] Quantity:1</b>
<b>TBU:</b>	PowerEdge R720 Shipping [331-4437] Quantity:1
<b>TBU:</b>	Broadcom 5720 DP 1Gb Network Interface Card [430-4423] Quantity:1
<b>TBU:</b>	iDRAC7 Enterprise [421-5339] Quantity:1
<b>TBU:</b>	Broadcom 5720 QP 1Gb Network Daughter Card [430-4418] Quantity:1
<b>TBU:</b>	Chassis with up to 24, 2.5" Hard Drives [342-3566] Quantity:1
<b>TBU:</b>	Bezel [318-1375] Quantity:1
<b>TBU:</b>	Power Saving Dell Active Power Controller [330-5116] Quantity:1
<b>TBU:</b>	Unconfigured RAID for H710P/H710/H310 (1-24 HDDs) [331-4557] Quantity:1
<b>TBU:</b>	PERC H710 Integrated RAID Controller, 512MB NV Cache [342-3529] Quantity:1
<b>TBU:</b>	Intel Xeon E5-2640 2.50GHz, 15M Cache, 7.2GT/s QPI, Turbo, 6C, 95W [317-9595] Quantity:1
<b>TBU:</b>	Heat Sink for PowerEdge R720 and R720xd [331-4508] Quantity:1
<b>TBU:</b>	DIMM Blanks for Systems with 2 Processors [317-8688] Quantity:1
<b>TBU:</b>	Intel Xeon E5-2640 2.50GHz, 15M Cache, 7.2GT/s QPI, Turbo, 6C, 95W [317-9609] Quantity:1
<b>TBU:</b>	Heat Sink for PowerEdge R720 and R720xd [331-4508] Quantity:1
<b>TBU:</b>	8GB RDIMM, 1600 MHz, Standard Volt, Dual Rank, x4 [317-9648] Quantity:4
<b>TBU:</b>	4GB RDIMM,1600MHz,SV,DR,x8 [317-9647] Quantity:4
<b>TBU:</b>	1600 MHz RDIMMS [331-4424] Quantity:1
<b>TBU:</b>	Performance Optimized [331-4428] Quantity:1
<b>TBU:</b>	HDD, 1TB 7.2K SATA,3G,2.5,HP [342-1998] Quantity:24
<b>TBU:</b>	Electronic System Documentation and OpenManage DVD Kit for R720 and R720xd [331-5914] Quantity:1
<b>CD-ROM or DVD-ROM Drive;</b>	ReadyRails Sliding Rails With Cable Management Arm [331-4433] Quantity:1
<b>CD-ROM or DVD-ROM Drive;</b>	Dual, Hot-plug, Redundant Power Supply [331-4605] Quantity:1

## Dell OpenStack-Powered Cloud Solution: Reference Architecture Guide

---

<b>CD-ROM or DVD-ROM Drive;</b>	Power Cord, C13 to C14, PDU Style, 12 Amps, 2 meter, Qty 1 [330-3151] Quantity:2
<b>CD-ROM or DVD-ROM Drive;</b>	No Operating System [420-6320] Quantity:1
<b>CD-ROM or DVD-ROM Drive;</b>	No Media Required [421-5736] Quantity:1
<b>Service:</b>	Dell Hardware Limited Warranty Plus On Site Service Initial Year [936-0967] Quantity:1
<b>Service:</b>	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, 2 Year Extended [936-7243] Quantity:1
<b>Service:</b>	ProSupport: 7x24 HW / SW Tech Support and Assistance, 3 Year [936-7263] Quantity:1
<b>Service:</b>	Dell Hardware Limited Warranty Plus On Site Service Extended Year [939- 3398] Quantity:1
<b>Service:</b>	ProSupport: Next Business Day Onsite Service After Problem Diagnosis, Initial Year [989-2701] Quantity:1
<b>Service:</b>	Thank you choosing Dell ProSupport [989-3439] Quantity:1
<b>Installation:</b>	On-Site Installation Declined [900-9997] Quantity:1
<b>Support:</b>	Proactive Maintenance Service Declined [926-2979] Quantity:1

## Appendix D: Rack Bill of Materials

Table 18: PowerEdge 24U Rack

Description/Catalog Number	Product Code	Qty	SKU
<b>PowerEdge Rack 24U:</b>			
Dell 2420 24U Rack with Doors and Side Panels, Ground Ship, NOT for AK / HI	24FDSG	1	[224-4950]
<b>Hardware Support Services:</b>			
3Yr Basic Hardware Warranty Repair: 5x10 HW-Only, 5x10 NBD Parts	3PD	1	[992-1692] [992-4910] [993-4078] [993-4087]

Table 19: PowerEdge 42U Rack

Description/Catalog Number	Product Code	Qty	SKU
<b>PowerEdge Rack 4220:</b>			
Dell 4220 42U Rack with Doors and Side Panels, Ground Ship, NOT for AK / HI	42GFDS	1	[224-4934]
<b>Interconnect Kits:</b>			
PE4220 42U Rack Interconnect Kit, PS to PS	42UPSPS	1	[330-3601]
<b>Dell Rack Accessories:</b>			
120 Volt Powered Fan Kit for Dell Racks	42UF120	1	[310-1285]
<b>Hardware Support Services:</b>			
3Yr Basic Hardware Warranty Repair: 5x10 HW-Only, 5x10 NBD Parts	3PD	1	[992-1802] [992-5080] [993-4108] [993-4117]

## Appendix E: Network Equipment Bill of Materials

Table 20: Force 10 S60  
(Quantity is 2 for the Solution Starter Configuration.)

Description/Catalog Number	Product Code	Qty	SKU
<b>Force 10 S60</b>			
Force 10 S60, 48 GbE Ports, Managed Switch, 10GbE and Stacking Capable	Force 10 S60	1	[222-6714]
<b>Modular Upgrade Bay 1</b>			
Modules: Stacking Module, 48Gbps, Includes 1m Stacking Cable	48GSTCK	1	[320-5171]
Cables (optional)			[980-5492]
Stacking Cable, 3m	3MSTKCL	1	[320-5168]
<b>Hardware Support Services:</b>			
3Yr Basic Hardware Warranty Repair: 5x10 HW-Only, 5x10 NBD Onsite	U3OS	1	[981-1260] [985-6027] [985-6038] [991-8459]

Table 21: PowerConnect PC6224 (Optional Only)

Description/Catalog Number	Product Code	Qty	SKU
PCT6224,MG,24P,10GBE CAPABLE			[222-6710]
PESS BASIC NBD PARTS,4YR EXT,PC6224			[960-1044]
PESS BASIC NBD PARTS,INIT,PC6224			[981-0890]
HW WRTY,PC6224,INIT			[985-5977]
HW WRTY,PC6224,EXT			[985-5988]
NO INSTL,PCT			[950-8997]
STACKING MODULE,1M CABLE,CUST			[320-5171]
10GBE CX-4 OR STACK CABLE,3M,CUST			[320-5168]
<b>Front-end SFP Fiber Transceivers:</b>			
Four SFP Optical Transceivers, 1000BASE-SX, LC Connector			[320-2881]
PowerConnect 6xxx SFP+ Module supports up to two SFPs			[330-2467]

**Table 22: Network Switch Add-Ons (For Information Only)**

Description/Catalog Number	Product Code	Qty	SKU
<b>Modular Upgrade Bay 1:</b>			
Modules: Stacking Module, 48Gbps, Includes 1m Stacking Cable	48GSTCK	1	[320-5171]
<b>Modular Upgrade Bay 2:</b>			
Modules: PowerConnect 6xxx SFP+ Module supports up to two SFPs	C107D	1	[330-2467]
<b>Front-end SFP Fiber Transceivers:</b>			
Four SFP Optical Transceivers, 1000BASE-SX, LC Connector			[320-2881]
PowerConnect 6xxx SFP+ Module supports up to two SFPs			[330-2467]
<b>Front-end SFP Fiber Transceivers:</b>			
Four SFP Optical Transceivers, 1000BASE-SX, LC Connector	4SFPSX	1	[330-2881]
<b>Cables:</b>			
Stacking Cable, 3m	3MSTKCL	1	[320-5168]

## Appendix F: Solution and Installation

Table 23: Services

Description	SKU
<b>Required</b>	
Cloud Compute Node, Dell PowerEdge C Server, Crowbar	331-3310
DCS OpenStack Info SKU	331-3286
EDT Installation and Implementation of The Dell OpenStack Cloud Solution, with 6 Nodes	961-4959
EDT Installation and Implementation of The Dell OpenStack Cloud Solution, with up to 8 Nodes—Add-on	961-4969
<b>Optional *</b>	
Cloud Computing Workshop	Custom Quote (GICS)
Cloud Computing Assessment	Custom Quote (GICS)
OpenStack Design Services	Custom Quote (GICS)
OpenStack Integration & Customization Services	Custom Quote (GICS)

\* Dell recommends the Cloud Computing Workshop and Assessment for customers who are just beginning to define a cloud computing roadmap, and have not yet settled on a platform such as OpenStack.

\*OpenStack Design, Integration & Customization Services are required for customers who are interested in deploying an OpenStack configuration, which is outside of the Dell Reference Architecture.

## Appendix G: Solution Starter Configuration Expansion Options

The solution is not intended to end at just 6 or 10 nodes or even one rack. The solution is designed to let you bring up a known configuration and start using it quickly. By adding additional components, compute nodes, storage nodes, load balancer ... an ever-evolving system is created.

Below are just some of the additional servers you can add to this solution.

**Table 24: Additional Servers for Design Solution**

Model	Focus	Base Config/Unit	Comments
<b>PowerEdge C6100 2-sled</b>	General Purpose	12 2.5-inch HDDs/sled	Available
	Pilots, Administration	96 GB RAM	
		16 cores	
<b>PowerEdge C6105 2-sled</b>	Alternate to the C6100-2-sled		Lower power, more cores
<b>PowerEdge C6100 4-sled</b>	Compute (Nova)	Same as 2 sled except	Available
		6 2.5-inch HDDs	
		32 cores	
<b>PowerEdge C6105 4-sled</b>	Alternate to C6100 4-sled		Lower power, more cores
<b>PowerEdge C2100</b>	Storage Node	12 3.5-inch HDDs	Available
		48 GB RAM	

### 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 **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](http://www.dell.com/openstack)

©2011–2012 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.