



CHAPTER 1

Architecture

OpenStack is one of the fastest growing open source projects today, with thousands of active developers and hundreds of actively supporting companies and individuals. Cisco has developed and maintains the Cisco OpenStack Installer to provide an automated installation of a packaged reference version of OpenStack. Customers can use the Cisco OpenStack Installer to easily and quickly stand up a production cloud.

This document looks to answer the next set of questions that arise after a cloud is up and running. These questions revolve around scalability. While both hardware and software scalability are of concern to a production cloud deployment, this document focuses on the scalability of the OpenStack control plane. This testing was conducted with the Cisco OpenStack Installer on a hardware architecture comprised of Cisco Nexus switches and Cisco UCS servers.

Hardware Architecture

The Cisco UCS servers chosen for this testing closely match the Mixed Workload Configuration from the [Cisco UCS Solution Accelerator Paks for OpenStack Cloud Infrastructure Deployments](#). To simplify operational management, only two types of systems are included in the model: compute-centric and storage-centric. Keep in mind, software performance is the focus of this testing, but minor variation from the hardware specification below should still result in similar data.

Mixed-Workload Server Configuration

- 6 Cisco UCS C220 M3 Rack Servers, each with:
 - 2 Intel Xeon processors E5-2665
 - 128 GB of memory
 - LSI MegaRAID 9266-CV 8i card
 - Cisco UCS VIC 1225
 - Redundant power supplies
 - 2 x 600-GB SAS hard disk drives
- 2 Cisco UCS C240 M3 Rack Servers, each with:
 - 2 Intel Xeon processors E5-2665
 - 256 GB of memory
 - LSI MegaRAID 9271-CV 8i card
 - Cisco UCS VIC 1225

- Redundant power supplies
- 12 x 900-GB SAS hard disk drives
- 2 Cisco UCS 6296UP 96-Port Fabric Interconnects
- 2 Cisco Nexus 2232PP 10GE Fabric Extenders

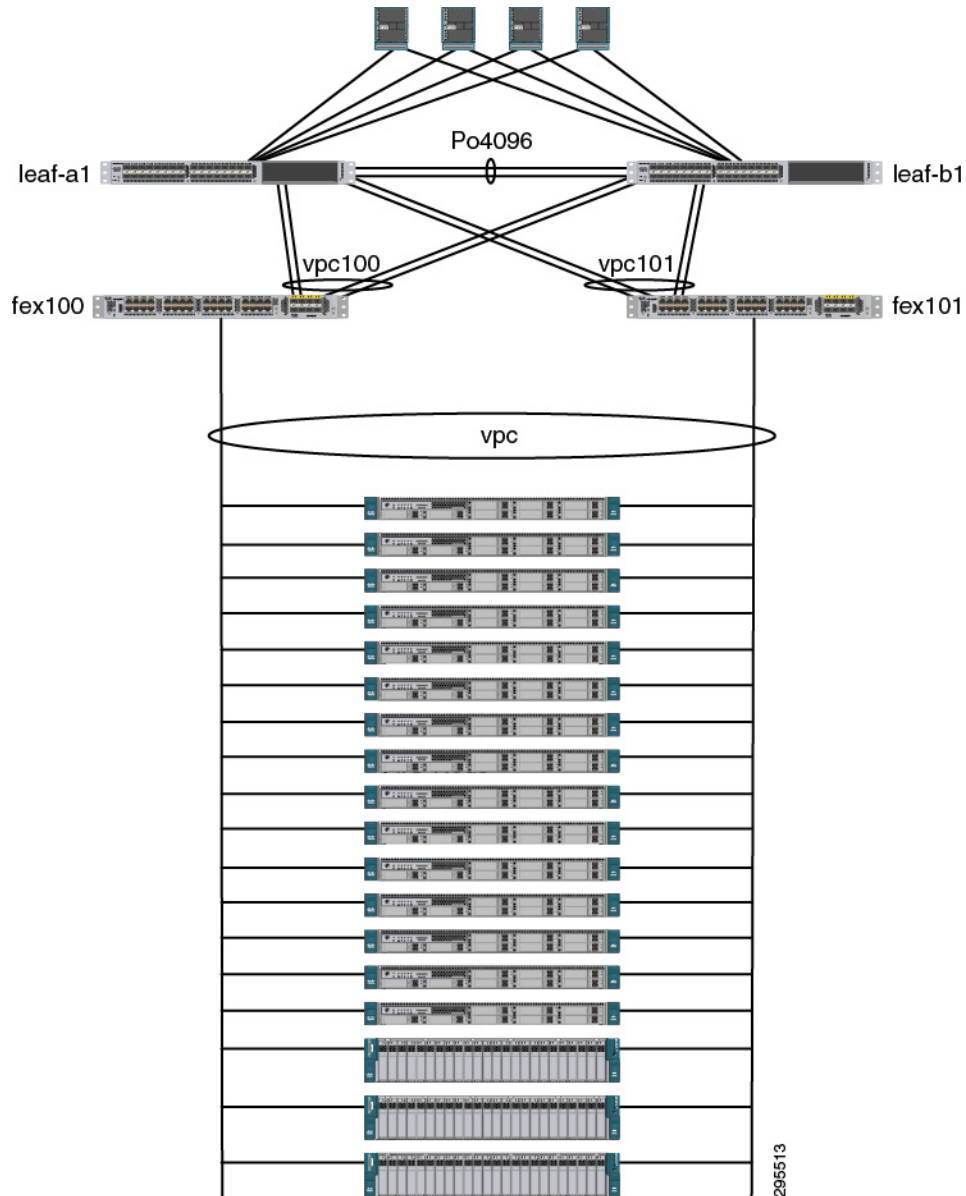
The compute system is based on the 1RU C220-M3 platform and leverages a low power 8 core CPU and 128GB of memory giving a memory-to-core ratio of 8:1. The storage subsystem is based on a high performance RAID controller and SAS disks for a flexible model for ephemeral, distributed Cinder, and/or Ceph storage. The network interface is based on the Cisco Virtual Interface Controller (VIC), providing dual 10Gbps network channels and enabling Hypervisor Bypass with Virtual Machine Fabric Extension (VM-FEX) functionality when combined with a Nexus 5500 series data center switch as the Top of Rack (ToR) device, Fibre Channel over Ethernet (FCOE) storage, and Network Interface Card (NIC) bonding for network path resiliency or increased network performance for video streaming, high performance data moves, or storage applications.

The storage system is based on the 2RU C240-M3 platform, which is similar at the baseboard level to the C220-M3, but provides up to 24 2.5” drive slots. With 24 spindles, this platform is geared more toward storage. The reference configuration makes use of low power 8 core CPUs, and a larger memory space at 256GB total. This specific configuration of these nodes is designed for Swift, Cinder, or Ceph-focused usage. This platform also includes the Cisco VIC for up to 20Gbps of storage forwarding with link resiliency when combined with the dual ToR model.

Network Configuration

The OpenStack control-plane traffic is carried on the server LOM management port connected to an out-of-band management switch. Server installation and management is also performed via this network. The OpenStack data-plane (VM) traffic is carried on the 10Gb Cisco VIC connected to an upstream network based on Nexus series switches, enabling the use of a number of advanced scale-out network services in the Layer 2 (Link Local) and Layer 3 (Routed Network) services. The TOR switches are configured as a virtual Port Channel (vPC) pair, with a set of 10Gb connections between them as the VPC peer link, and a vPC connecting to each FEX. Enhanced vPC is then used to bond the links to the servers into a port-channel as well.

Figure 1-1 OpenStack Data-Plane Network



Logically, the network is segregated via VLANs. VLANs are used for tenant segmentation as well as storage networks. To provide resiliency and high performance, provider networks are used. DHCP is provided by Neutron, but the physical network provides all other L3 services. Security will be provided by the IPtables security functionality driven by Neutron.


OpenStack Architecture

The Havana version of Cisco OpenStack installer provides four architectural scenarios: 2 Node, Full HA, Compressed HA, and All in One (AIO). While AIO provides the simplest setup, it is not intended for large-scale production deployment. However, it does provide a simplified model to evaluate scale points of single controller architecture, and a basis for comparison to other architectures.

AIO (All-in-One)

The AIO system is ideal for a first foray into OpenStack. Previous versions of Cisco OpenStack Installer required a separate build node to configure a server to be an AIO node. However, with the [COI Havana](#) release, the install script will turn a single node on which it is run into an AIO, so no separate build node is required. An AIO node puts all the OpenStack services (Control, Compute, Storage, Network, etc) onto a single node. RabbitMQ is used for message passing between services. MySQL database on the local drive is used to backend services.

Figure 1-2 All-in-One Services Layout



295514