JUNIPER NETWORKS | Engineering Simplicity

Getting Started with Contrail Cloud

Published
2023-07-14

RELEASE
Release 13

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Getting Started with Contrail Cloud
Release 13

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at https://support.juniper.net/support/eula/. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

CHAPTER 1

Getting Started with Contrail Cloud

About This Guide

Use this guide to understand Contrail Cloud basics. This guide provides an overview of Contrail Cloud, the Contrail Cloud architecture, Red Hat Openstack integration with Contrail Networking in Contrail Cloud, and summaries of the hardware, software, and networking components in a Contrail Cloud environment.

See Reference Architecture for Contrail Cloud 13.1 and 13.2 for an in-depth guide about Contrail Cloud that includes detailed descriptions of Contrail Cloud environments and configuration procedures.

See Contrail Cloud Deployment Guide for Contrail Cloud deployment instructions.

See the Contrail Cloud Product Information page and the Contrail Cloud TechLibrary page for additional Contrail Cloud information resources.

Contrail Cloud Overview

Juniper Networks Contrail Cloud provides cloud service providers with a bundled solution to build cloud platform infrastructures. The Contrail Cloud bundle simplifies your cloud building journey by integrating multiple software components—Red Hat Openstack orchestration, Contrail Networking, Contrail Command, Red Hat Ceph Storage, and Appformix—into a bundle with pre-configured files that is collectively installed on multiple devices using a single installation procedure.

Contrail Cloud:

- simplifies installations by bundling multiple software components into the same installation procedure.
- simplifies purchasing procedures by removing the guesswork associated with selecting your individual orchestration and SDN networking components for your cloud network.
- eases cloud network configuration and implementation by providing pre-configured YAML files that establish the initial networks at the server layer with minimal user intervention.
- avoids many of the unintended consequences associated with implementing components that have not been tested in the same network.
- provides a cloud infrastructure that is fully supported by the Juniper Networks technical support team.

The bundled Contrail Cloud package includes:

- Red Hat OpenStack
- Red Hat Ceph Storage
- Juniper Networks Contrail Networking
- Juniper Networks Contrail Command
- Juniper Networks Appformix
- Pre-configured YAML files to simplify the initial configuration procedure, and Ansible scripts to deploy configurations made in the YAML files.

All Contrail Cloud software components are downloaded from a Juniper Satellite device. You must send an email request to mailto:contrail_cloud_subscriptions@juniper.net to obtain the credentials to access the Juniper Satellite. See Deploying Contrail Cloud.

RELATED DOCUMENTATION

> Reference Architecture for Contrail Cloud 13.1 and 13.2
>
> Contrail Cloud Deployment Guide
>
> Contrail Cloud TechLibrary page

Contrail Cloud High-Level Architecture Overview

Contrail Cloud is a software package that arrives with detailed descriptions of how networking should be established on servers running in the topology and the options available for configuring the software running on the servers. The Contrail Cloud package establishes the storage, compute, and control nodes as well as the connections between these devices. These node functions and their interconnections are covered in greater detail later in this document. The Contrail Cloud package provides additional functionality beyond storage, compute, and control node setup, including Contrail Networking to provide an SDN Controller and Red Hat Openstack to provide an orchestrator.

Figure 1 on page 4 illustrates a representative Contrail Cloud high-level architecture.

Figure 1: Contrail Cloud High-Level Representative Architecture Overview

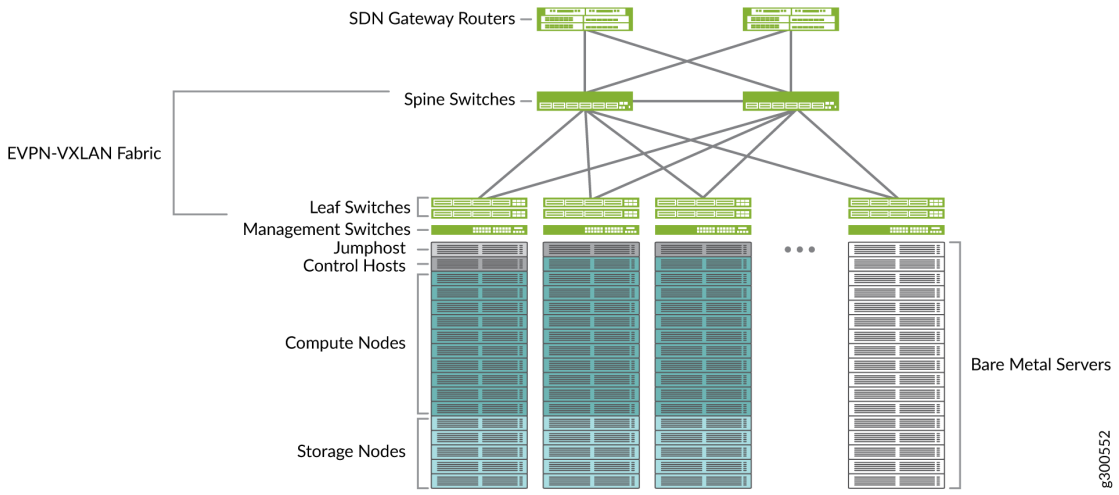

The server nodes—the storage, compute, and control nodes, as well as the bare metal servers (BMS) in environments supporting bare metal servers—are established using Contrail Cloud. These nodes also contain a variety of interconnections to one another and to the EVPN-VXLAN fabric to support multiple functions. See "Contrail Cloud Hardware Nodes" on page 7.

For more detailed information about connecting the storage, compute, and control nodes to the fabric, see Contrail Cloud Configuration in Reference Architecture for Contrail Cloud 13.1 and 13.2.

The fabric network in this representative architecture is a Clos fabric composed of spine and leaf switches. The switches run their own software—they are typically Juniper switches running Junos OS software—and the fabric itself is provisioned using Contrail Networking. See Data Center: Contrail Enterprise Multicloud for Fabric Management and Contrail Networking Fabric Lifecycle Management Guide for information on fabric provisioning and management.

The SDN Gateway routers in this architecture are layer 3 devices that connect to the spine devices in the EVPN-VXLAN fabric. The routers are responsible for all traffic exiting and entering the tenant overlay networks. The scope of configuring the SDN gateway router is beyond this guide. For a sample external gateway router configuration connecting to spine devices in EVPN-VXLAN fabrics, see the Data Center Interconnect Design and Implementation section of the Data Center Fabric Architecture Guide.

RELATED DOCUMENTATION

- Reference Architecture for Contrail Cloud 13.1 and 13.2
- Contrail Cloud Deployment Guide
- Contrail Cloud TechLibrary page

Data Center: Contrail Enterprise Multicloud for Fabric Management
Contrail Networking Fabric Lifecycle Management Guide
Data Center Fabric Architecture Guide

Red Hat OpenStack and Contrail Networking Integration in Contrail Cloud

Contrail Cloud uses Contrail Networking as the SDN controller and Red Hat OpenStack as the orchestration platform. This section covers how Red Hat Openstack and Contrail Networking integrate in a Contrail Cloud environment. This section assumes some existing knowledge of Contrail Cloud and Red Hat Openstack components. The Contrail Cloud components are largely discussed in the later sections of this documentation. The RedHat components used in Contrail Cloud are listed in the "Contrail Cloud Software Summary" on page 13. See Product Documentation for Red Hat OpenStack Platform 13 for additional information on the Red Hat Openstack components.

Figure 2 on page 6 illustrates the API-level interaction between OpenStack and Contrail Networking configurations. The Contrail Neutron plug-in enables the OpenStack Neutron service to make the necessary API calls into the Contrail configuration node to create, delete, and update network resources as defined in OpenStack. The Contrail configuration node has a northbound interface which exposes REST APIs to orchestration systems such as OpenStack and Contrail Web UI that can be used to make the network configurations.

Figure 2: Contrail Configuration and OpenStack—API Interaction

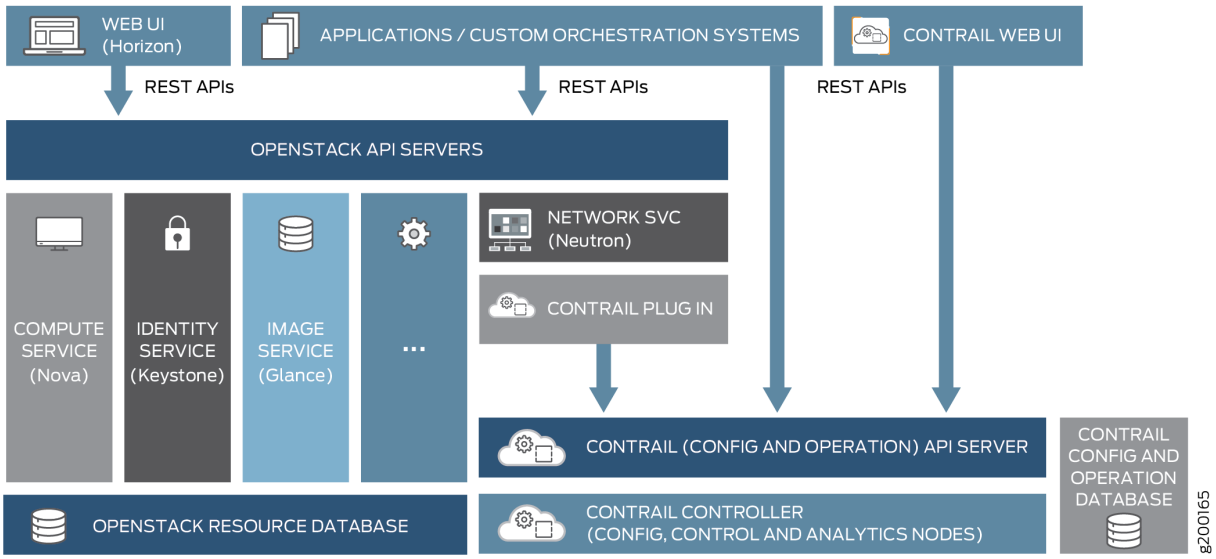

Figure 3 on page 7 illustrates a more detailed step-by-step interaction between components of OpenStack services, specifically Nova, Neutron, and Contrail services. At the control plane the Contrail plug-in for Neutron translates the configuration commands received from the northbound orchestration engines into corresponding Contrail-related configuration. On the data plane or forwarding plane, the Nova-agent running on each compute hypervisor interacts with the distributed forwarding engine in Contrail Networking, called the vRouter, to set up a virtual-interface and tap interfaces for VMs to obtain connectivity.

Figure 3: OpenStack Services Interaction

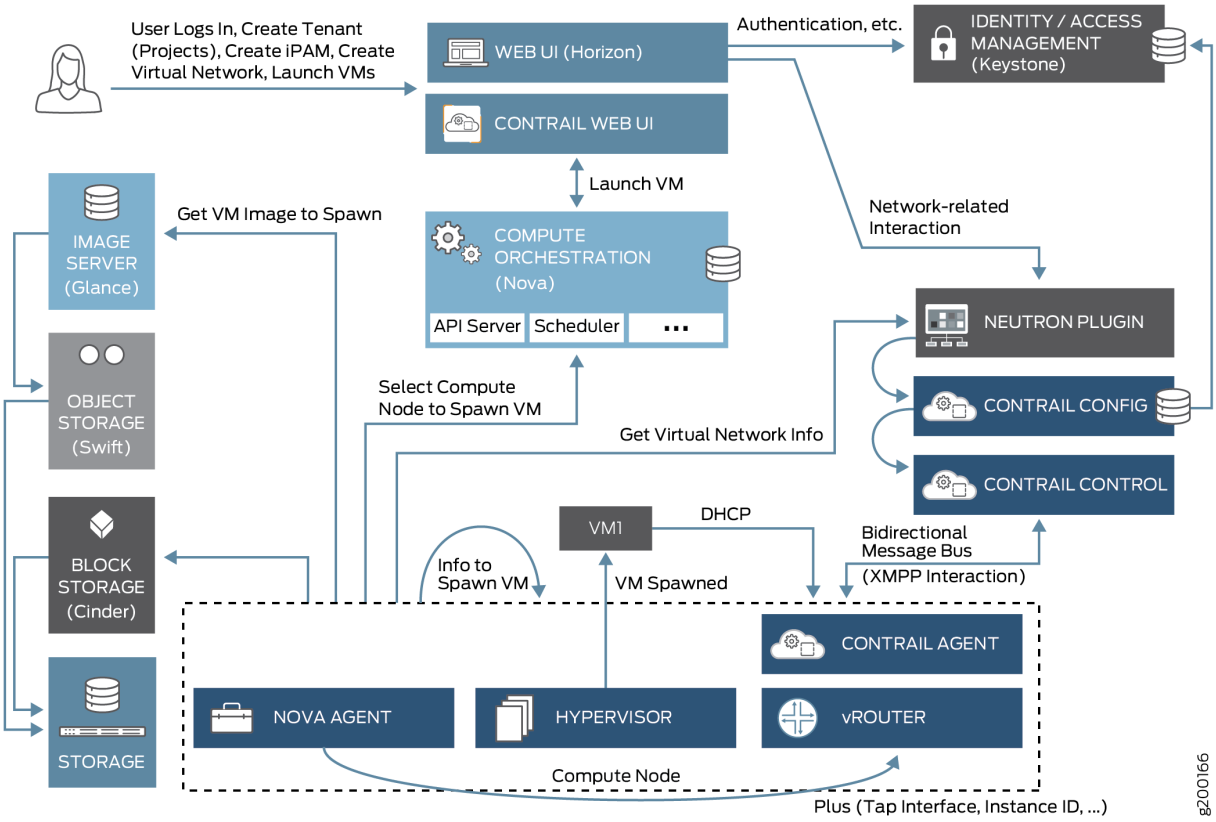

RELATED DOCUMENTATION

- Reference Architecture for Contrail Cloud 13.1 and 13.2
- Contrail Cloud Deployment Guide
- Contrail Cloud TechLibrary page

Contrail Cloud Hardware Nodes

IN THIS SECTION

A Contrail Cloud environment includes servers that function as compute nodes, storage nodes, and management nodes. The compute nodes provide compute services, the storage nodes provide storage services, and a variety of management nodes are used to manage components for the compute and storage nodes. These nodes are interconnected to one another and connect to the fabric to connect to the larger network.

This section provides an overview of the compute, storage, and management nodes in the Contrail Cloud infrastructure. Other system nodes, including the SDN controller managed using Contrail Networking and the orchestrator that is managed using Red Hat Openstack, are discussed in other sections of this document.

Hardware Nodes Overview

Figure 4 on page 8 illustrates the hardware nodes in a Contrail Cloud infrastructure.

Figure 4: Compute, Storage, and Management Nodes Overview

The compute, storage, and management nodes are interconnected endpoint servers. The servers at this layer connect directly to the fabric to access the larger network.

During the Contrail Cloud installation, the compute, storage, and control nodes are instantiated and the networks connecting these devices are established. The configuration of these nodes and the initial

network is provided by pre-configured YAML files that are downloaded onto the jump host. These YAML files can then be customized by users after the initial installation. See Contrail Cloud Configuration File Structure Overview in the Reference Architecture for Contrail Cloud.

The other devices that work at the endpoint layer with the compute, storage, and management devices include the jump host, the management switches, and the BMS nodes. Table 1 on page 9 provides a summary description of each hardware node.

Table 1: Hardware Node Summary

Node	Description
Bare metal server (BMS)	Physical server that is not provisioned to support virtualization but requires overlay network access to participate in the Contrail Cloud environment. BMS nodes are only present in Contrail Cloud environments that require the non virtualized functions provided by the BMS and are therefore not present in all Contrail Cloud environments. Contrail Cloud manages some services for the bare metal servers using the Contrail Services Node, which is also called the TOR Services Node (TSN) in Red Hat OpenStack.
Compute node	Server provisioned to support virtualization; hosts virtual machines.
Control host	Server that hosts controller nodes. Controller nodes are VMs responsible for controlling Contrail Cloud functions. See "Control Host" on page 11.

Jump host	A server that performs the following functions in Contrail Cloud: • hosts the undercloud. The undercloud runs as a VM on the jump host and is responsible for provisioning and managing all control hosts, compute nodes, and storage nodes. • hosts Contrail Cloud automation. All Contrail Cloud scripts and YAML configuration files are stored and instantiated on the jump host. • hosts the Contrail Command web user interface virtual machine. • serves as the entry point for access to the undercloud or overcloud. You can "jump" between the overcloud and the undercloud from the jump host.
Management switch	Switches that are responsible for managing the IPMI, Management, Provisioning, and Intranet (for the jump host only) networks in Contrail Cloud. The management switches are not managed by Contrail and are not included in all Contrail Cloud implementations.
Storage node	Server whose purpose is storing data. Storage nodes run Red Hat Ceph storage software in Contrail Cloud.

Jump Host

The jump host is a server that hosts the undercloud VM. Figure 5 on page 11 illustrates the high-level function of the jump host.

Figure 5: Jump Host Overview

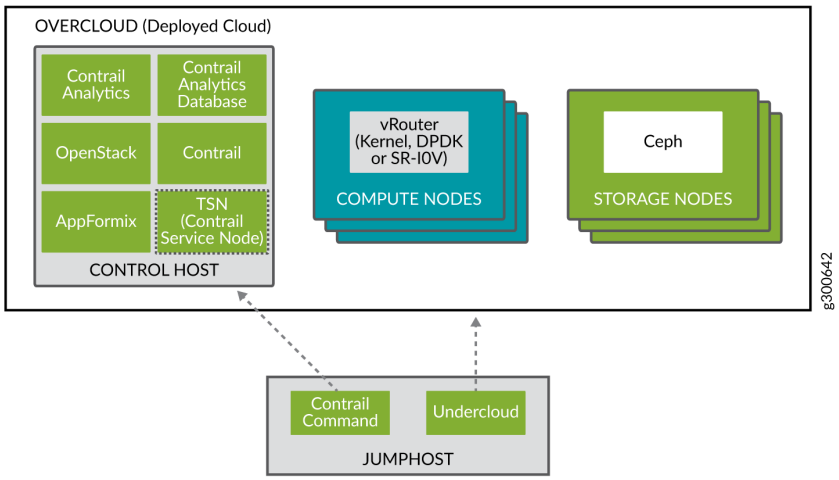

The jump host:

- hosts the undercloud. The undercloud is a VM responsible for provisioning and managing all control hosts, storage nodes, and compute nodes in a Contrail Cloud. All Contrail-related setup and configuration is performed through the undercloud in a Contrail Cloud.
- stores Contrail Cloud configuration-related files. The YAML files that configure Contrail Cloud are stored on the jump host. The Ansible scripts that apply the configurations made in the YAML files to the Contrail Cloud nodes are also stored on the jump host.
- hosts the Contrail Command web user interface virtual machine.
- runs Red Hat Enterprise Linux with only base packages installed.

A jump host must be operational as a prerequisite for a Contrail Cloud installation. The jump host should not run any virtual machines besides the undercloud and the Contrail Command virtual machines. For a complete list of jump host requirements, see the Deploying Contrail Cloud section of the Contrail Cloud Deployment Guide.

Control Host

A control host is a hypervisor running on a server that hosts virtualized control functions as controller nodes. Controller nodes are VMs responsible for managing server functions.

Figure 6 on page 12 illustrates the controller nodes in a control host.

Figure 6: Control Host

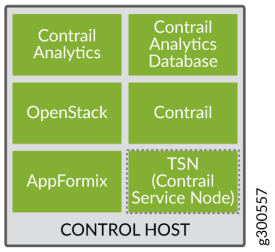

The controller nodes that can run on a control host:

- OpenStack Controller—manages Red Hat OpenStack & Ceph storage.
- Contrail Controller—manages the configuration and control functions of Contrail Networking.
- Contrail Analytics—manages the Analytics function of Contrail Networking.
- Contrail Analytics DB—manages the database used by Contrail Analytics.
- Appformix Controller—manages Appformix.
- Contrail Service Node—also called TOR Services Node in earlier Contrail releases. The Contrail Services node is an optional controller VM used by Contrail Networking to assist with some bare metal server (BMS)-related tasks.

Each control node in a Contrail Cloud environment has an OpenStack, Contrail Controller, Contrail Command, Contrail Analytics, Contrail Analytics DB, and Appformix controller node. The Contrail Service node runs only in environments where there is a need to support bare metal servers.

Compute Node

A compute node is a server that hosts virtual machines that provide services over the network. The services that run on compute nodes vary by network and the documentation of those services is beyond the scope of this guide.

Compute nodes use Ceph storage as a back-end OpenStack storage option for block, object, and file storage services. VM ephemeral disk storage is provided by Ceph. Ceph Storage is included as part of the standard Contrail Cloud bundle. If a compute node does not have access to Ceph storage, it must have locally attached disks to support it's storage requirements.

The compute nodes in Contrail Cloud use a vRouter to implement data plane functionality. The vRouter options are discussed in "Contrail Cloud Software Summary" on page 13.

Storage Node

A storage node is a server in the Contrail Cloud environment whose purpose is storing raw data. Storage nodes run Red Hat Ceph storage software in Contrail Cloud.

The storage nodes support Ceph OSD. For additional information on using Red Hat Ceph storage software, see Product Documentation for Red Hat Ceph Storage.

RELATED DOCUMENTATION

- Reference Architecture for Contrail Cloud 13.1 and 13.2
- Contrail Cloud Deployment Guide
- Contrail Cloud TechLibrary page

Contrail Cloud Software Summary

IN THIS SECTION

This section discusses overcloud and undercloud roles and the software running on the nodes in Contrail Cloud.

Overcloud Node and Jump Host Software

Figure 7 on page 14 illustrates the software running on the nodes in the overcloud and on the jump host.

Figure 7: Jump Host and Overcloud Node Component Summary

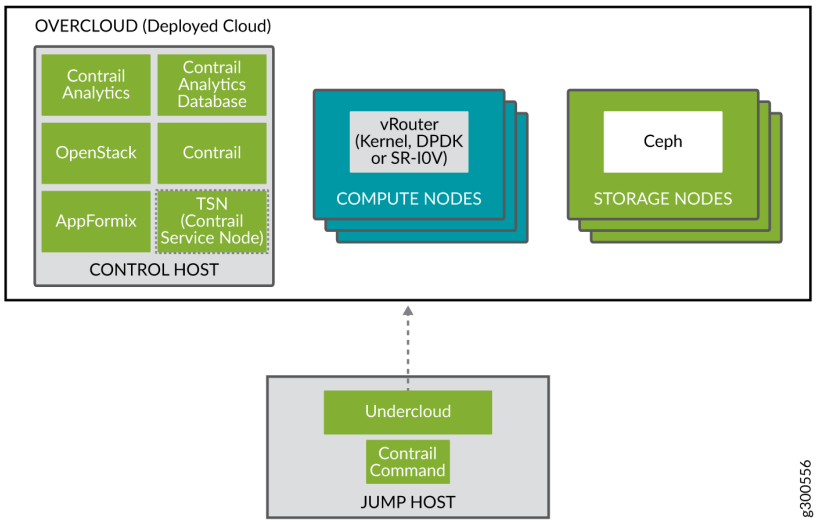

Table 2 on page 14 summarizes functions for the overcloud and undercloud nodes in Contrail Cloud and indicates whether the function is delivered as a VM in Contrail Cloud.

Table 2: Virtual Machine Summary

Function		Deployed as VM
Jump host		No
	Undercloud	Yes
	Contrail Command	Yes
Controller Hosts	OpenStack	Yes
	Contrail Controller	Yes
	Contrail Analytics	Yes
	Contrail Analytics DB	Yes
	Appformix	Yes

	Contrail Service Node (TOR Service Node)	Yes
Compute Nodes	vRouter in kernel mode	No
	vRouter using DPDK	No
	vRouter using SR-IOV	No
Storage Nodes (Ceph Software)		No

Undercloud VM

The undercloud is responsible for provisioning and managing all nodes in the overcloud, which are the controller, compute, and storage nodes. The undercloud runs as a VM on the jump host. The concept of an overcloud that is deployed from an undercloud is defined in the OpenStack TripleO (Openstack on Openstack) project—one of the OpenStack Life Cycle Managers—and implemented in Contrail Cloud with Red Hat OpenStack Director (RHOSPd).

The undercloud VM runs on top of the kernel-based virtual machine (KVM) hypervisor on the jump host. An initial Contrail Cloud deployment starts when the undercloud VM connects to a Juniper Satellite device to install Contrail Cloud software. See Deploying Contrail Cloud.

Contrail Cloud configuration updates are performed through the undercloud by updating YAML configuration files. The updates configured in YAML files are applied when an Ansible script is run. The YAML files template and Ansible scripts are downloaded to the jump host as part of the Contrail Cloud installation procedure. The Ansible scripts generate Heat templates and property files that are used by the Red Hat Openstack Director (RHOSPD) to deploy the configuration changes to the overcloud. The generation of the Heat template is automated; users only have to execute the Ansible scripts to apply a configuration change. See Contrail Cloud Configuration.

The OpenStack instance that constitutes the undercloud uses the following OpenStack services:

- Glance
- Heat
- Ironic
- Keystone

- Nova
- Neutron
- Swift

The OpenStack instance that constitutes the overcloud runs the following services in the control nodes.

- Cinder
- Glance
- Heat
- Horizon
- Keystone
- Nova
- Neutron
- Pacemaker
- Swift
- Galera (for HA services)

In Contrail Cloud, the following OpenStack services that are deployed in a basic Red Hat OpenStack overcloud instance are disabled:

- Ceilometer
- Gnochi

See "Red Hat OpenStack and Contrail Networking Integration in Contrail Cloud" on page 5 for additional information on Contrail Networking and Red Hat Openstack integration in Contrail Cloud.

Controller Host VMs

A control host is a server that hosts controller nodes. Controller nodes are a collection of services running inside VMs responsible for controlling Contrail Cloud functions.

The controller node VMs in Contrail Cloud are the Openstack Controller, Contrail Controller, Contrail Analytics, Contrail Analytics DB, Appformix, and the Contrail Service Node.

See "Contrail Cloud Hardware Nodes" on page 7.

Compute Node Data Plane Options

The compute nodes in Contrail networking—except compute nodes in SR-IOV mode—use a vRouter to implement data plane functionality.

This section discusses compute node vRouter options and includes the following topics:

Compute Nodes VRouter Options Summary

The compute nodes use a vRouter to implement data plane functionality.

Figure 8: vRouter Summary—Compute Nodes

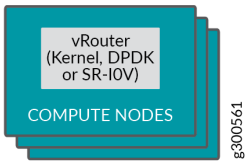

The data plane interfaces on compute nodes can be configured to support one of the following forwarding methods:

- Kernel mode—Linux kernel performs vRouter forwarding function.
- Data Plane Development Kit (DPDK)—a user space with a specified number of cores is defined to perform the vRouter forwarding function. See Configuring the Data Plane Development Kit (DPDK) Integrated with Contrail vRouter for additional information on DPDK in Contrail networking.
- Single root I/O virtualization (SR-IOV)—vRouter is bypassed. A VM or a container interface connects directly to the NIC.

Your traffic forwarding method choice depends on the traffic profile expectations for each individual compute node. A Contrail Cloud environment can have different compute nodes configured with different interface types, and workloads can be placed on the most appropriate compute node using various technologies, such as OpenStack availability zones.

Kernel vRouter

In kernel mode, vRouter is deployed as a Linux kernel module to perform the vRouter forwarding function.

Data Plane Development Kit (DPDK) Mode

A vRouter in DPDK mode runs in a user space on the compute node. Network traffic is handled by a special DPDK dedicated interface or interfaces that handle VLANs and bonds. A specified number of cores is assigned to perform the vRouter forwarding function.

DPDK vRouters provide higher throughput than kernel vRouters. See Configuring the Data Plane Development Kit (DPDK) Integrated with Contrail vRouter for additional information on DPDK in Contrail networking.

Single Root I/O Virtualization (SR-IOV) Mode

A compute node in SR-IOV mode provides direct access from the NIC to a VM. Because network traffic bypasses the vRouter in SR-IOV mode, no network policy or flow management is performed for traffic. See Configuring Single Root I/O Virtualization (SR-IOV) for additional information on SR-IOV in Contrail networking.

Storage Nodes Ceph Software

The storage nodes run Ceph software. For additional information on using Red Hat Ceph storage software, see Product Documentation for Red Hat Ceph Storage.

We recommend following these guidelines to optimize Ceph software performance in your Contrail Cloud deployment:

- Storage nodes must be separate from compute nodes in a Contrail Cloud environment. Hyperconverged nodes are not supported.
- Ceph requires a minimum of three storage nodes to operate.

RELATED DOCUMENTATION

- Reference Architecture for Contrail Cloud 13.1 and 13.2
- Contrail Cloud Deployment Guide
- Contrail Cloud TechLibrary page

Server Node Network Connections in Contrail Cloud

IN THIS SECTION

A variety of network connections run between servers to interconnect the storage, compute, and management nodes in a Contrail Cloud environment. Multiple network connections are also used to connect these nodes to the EVPN-VXLAN fabric layer to access the higher-layer networks.

These networks connections are covered in the following sections.

Server Node Network Connections in Contrail Cloud Overview

A variety of network connections run between servers to interconnect the storage, compute, and management nodes in a Contrail Cloud environment. Multiple network connections are also used to connect these nodes to the EVPN-VXLAN fabric layer to access the higher-layer networks.

The majority of the network connections—all networks except the IPMI network, which is established outside of the Contrail Cloud deployment, and the Intranet network—are created automatically using information in the YAML configuration files when Contrail Cloud is initially deployed. No user action is required on the server side devices to create these networks. VLANs, however, must be configured on the switches in the EVPN-VXLAN Fabric to ensure traffic can be passed between the server devices and the EVPN-VXLAN Fabric switches. See Contrail Cloud Deployment Guide.

Figure 9 on page 20 illustrates how the server nodes are connected to devices using the available networks in a representative architecture.

Figure 9: Server Network Connections

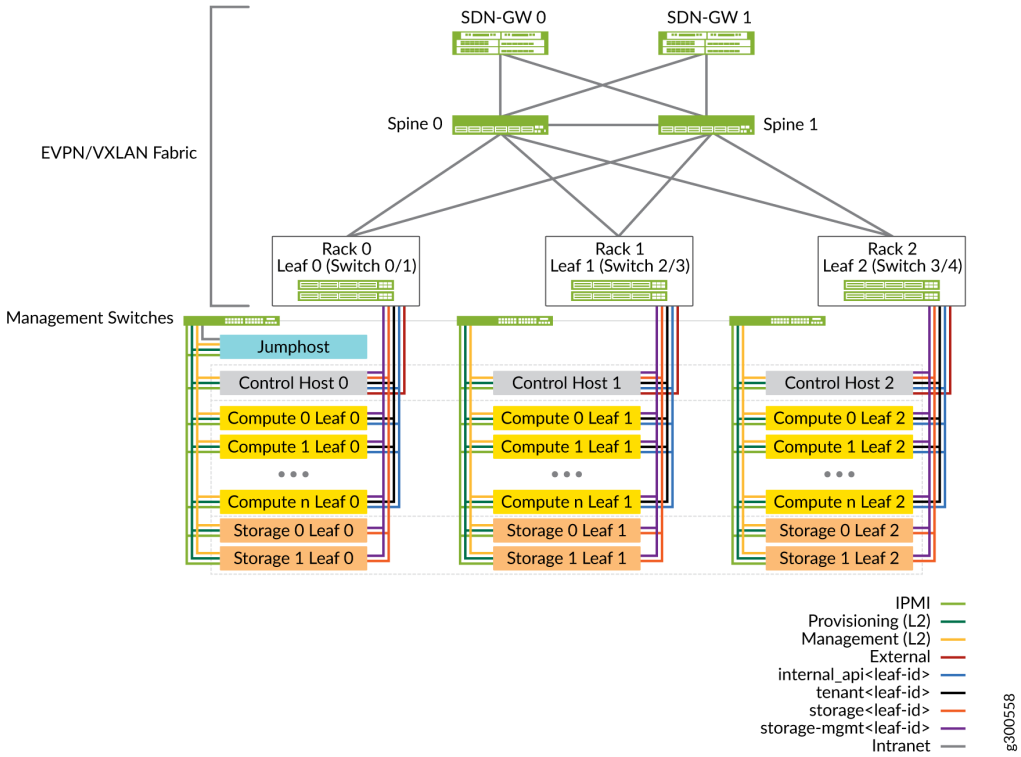

Table 3 on page 20 summarizes the purpose of each network connection.

Table 3: Networking in Contrail Cloud Summary

Network	Purpose
IPMI	Provides hardware management of servers. IPMI services are mostly used by the Openstack Ironic service and are established outside of the Contrail Cloud installation. IPMI services can be used by the server nodes in Contrail Cloud.

Intranet	Provides user access to the jump host. Provides access to satellite repositories for the jump host. Provides access to the Contrail Command web user interface. Provides external network access via SNAT from the control plane network.
Provisioning or Control Plane	Deploys new nodes using PXE booting and to install software packages on overcloud bare metal servers. The provisioning/ control plane network is used by Red Hat Director and is predefined before the creation of the undercloud.
Internal API	Provides communication with the OpenStack and Contrail Networking services, including Keystone, Nova, and Neutron using API communication, RPC messages, and database communication.
External	Provides tenant administrators access to the OpenStack Dashboard (Horizon) graphic interface, the public APIs for Openstack services, public Contrail APIs, and the Appformix WebUI. Commonly used as a path to the intranet.
Tenant	Supports overlay data-plane traffic—VXLAN and MPLS over UDP—and Contrail Controller to Contrail vRouter control plane traffic.
Storage	Supports storage data traffic for Ceph, block storage, NFS, iSCSI, and any other storage types.
Storage Management	Provides Ceph control, management, and replication traffic.
Management	(Optional) Provides direct access to compute nodes without having to send traffic through the jump host. Can be used for DNS and NTP services.

The IPMI, Management, and Provisioning networks are connected via management switches and the networks are stretched between switches using trunk connections for VLANs or by using a separate VXLAN tunnel for each network. The other networks connect to the leaf switches in the IP Fabric and use VXLAN with an EVPN control plane to connect between racks.

The External network contains the externally routable API and UI IP addresses for various controller functions. These functions are generally in the same subnet. The VLAN for the External network is configured to be in a VXLAN, which terminates in a VRF instance that is configured to route between the network where the tenant user packets arrive from the External network.

The other networks connected directly to the IP Fabric in this representative architecture will each have a separate subnet for each rack, and routing is used to connect the subnets within a VXLAN. The routing occurs by placing an IRB interface on either a leaf device (edge routed) or a spine device (centrally routed).

Table 4 on page 22 describes our recommendations regarding how the networks should be configured.

Table 4: Recommended Network Subnets and Implementations

Network	Subnet Description	Implementation
IPMI	IPMI is typically an external service that can be used by devices that can route to it. IPMI is not established by Contrail Cloud.	IPMI is typically not spanned over racks. It can be spanned over racks using VXLAN.
Intranet	Administrative access to jump host.	Specific port configuration
Provisioning	Layer 2 single subnet	Span over racks using VLAN or VXLAN
Internal API	Layer 2 single subnet	Span over racks using VLAN or VXLAN
Internal API (leaf)	One subnet per rack	Layer 3 routing between racks
External	Layer 2 single subnet	Span over racks using VLAN or VXLAN
Tenant	Layer 2 single subnet	Span over racks using VLAN or VXLAN
Tenant (leaf)	One subnet per rack	Layer 3 routing between racks
Storage	Layer 2 single subnet	Span over racks using VLAN or VXLAN
Storage (leaf)	One subnet per rack	Layer 3 routing between racks

Storage Management	Layer 2 single subnet	Span over racks using VLAN or VXLAN
Storage Management (leaf)	One subnet per rack	Layer 3 routing between racks
Management	Layer 2 single subnet	Span over racks using VLAN or VXLAN

When a network is composed of multiple subnets—the reference architecture illustrated earlier in this document has one subnet per rack—a subnet that contains all the rack subnets is created. This subnet is called the supernet. The subnet address of the supernet is used to configure static routes on servers to ensure proper traffic flow between racks.

Figure 10 on page 23 illustrates how VLANs, VTEPs, IRBs, and VRFs are configured in the networks that connect end systems to gateway routers, spine switches, and leaf switches when centralized routing—traffic is transferred between Layer 2 and Layer 3 using IRB interfaces on spine devices—is enabled.

Figure 10: VLANs, VTEPs, IRBs, and VRFs

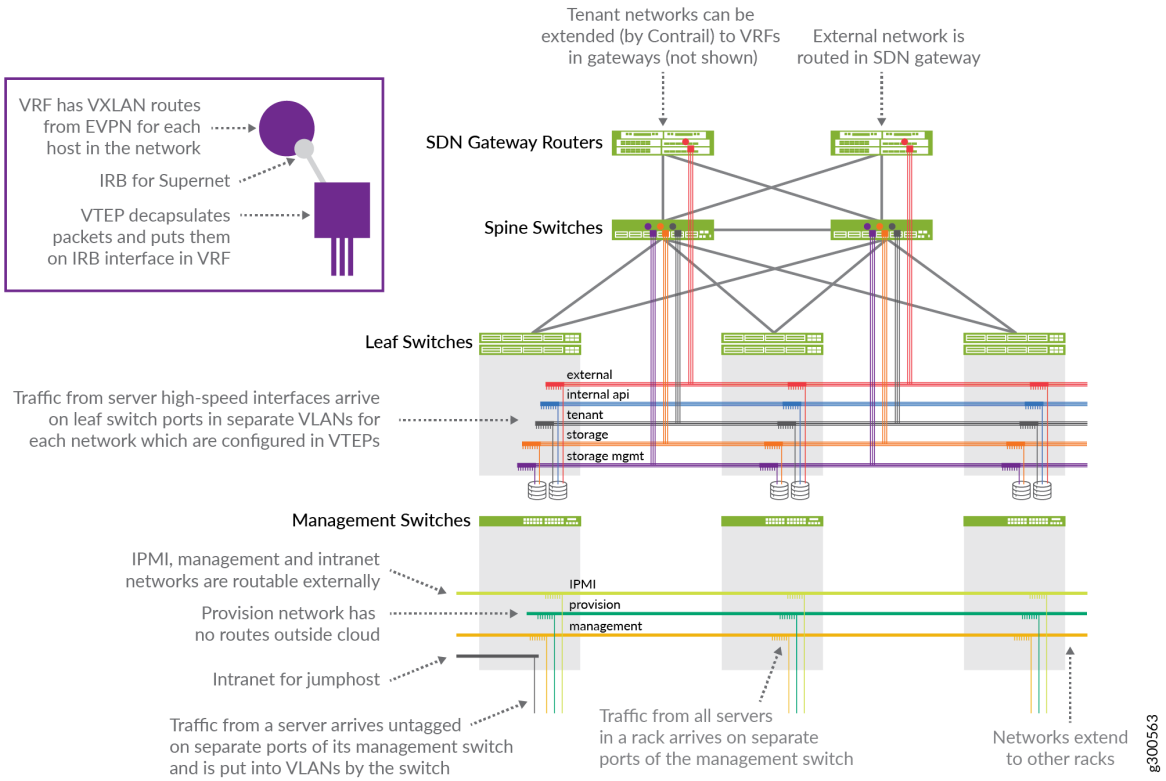

Table 5 on page 24 summarizes how VLANs are deployed per device.

Table 5: VLAN Summaries

Device	VLAN Summary
Management Switch	The IPMI, Management, Provisioning, and Intranet networks all connect the management switch to a server on different ports. The port traffic for each network arrives on the management switch untagged and is configured into the same VLAN. The VLAN is extended between the management switches on different racks.
Leaf Switch (EVPN-VXLAN Fabric)	Traffic from the external, internal API, tenant, storage, and storage management networks arrives on the leaf switch's high-speed interfaces from the servers in different VLANs. The VLANs are configured in VTEPs. Leaf ports are configured with logical interfaces that have VLANs for the networks that will be attached by servers to that port. Each VLAN on each switch is configured into a VXLAN VTEP and EVPN advertises host routes for each connected server to the spines. The VLANs used for each network are specified in the file overcloud-nics.yml.
Spine Switch (EVPN-VXLAN Fabric)	The spine switches are configured with a VTEP for each of the Internal API, Tenant, Storage and Storage-Mgt networks, and each of these are connected to an IRB interface whose IP address is that of the supernet. Each spine switch has a VRF that receives routes to each host from the leaf switches.
SDN Gateway Router	The SDN gateways are configured with a VTEP and VRF for the External network. Each SDN gateway will advertise a route for the External network to peers outside the Layer 2 network.

Contrail Cloud Network Types

Table 6 on page 25 summarizes the networks used by devices or VMs at the server layer in Contrail Cloud.

Table 6: Device & VM Network Summary

Role	Network							
	Intranet	Provisioning	Internal API	External	Tenant	Storage	Storage Management	Management
Jump host (Undercloud VM)	✓	✓						
Openstack Controller		✓	✓	✓	Optional	✓	✓	Optional
Contrail Controller		✓	✓	Optional	✓			Optional
Contrail Analytics		✓	✓	Optional	✓			Optional
Contrail Analytics DB		✓	✓	Optional	✓			Optional
Contrail Service Node (ToR Service Node)		✓	✓	Optional	✓			Optional
Appformix Controller		✓	✓	✓	Optional		✓	Optional

Control Host (Running all Controller VMs)		✓						Optional
Compute Node		✓	✓		✓	✓		Optional
Storage Node		✓				✓	✓	Optional

The use of a management network to enable direct access to nodes without going through the jump host is optional, but often recommended.

Contrail Cloud Roles and Networks

The networks in Contrail Cloud connect nodes to management switches or to devices that provide network access. In this reference architecture, the devices that connect the server nodes in the Contrail Cloud to the higher-layer network are the leaf devices in the EVPN-VXLAN IP Fabric.

Table 7 on page 26 summarizes the network connections for the nodes in Contrail Cloud.

Table 7: Server Layer Node Connections

Node	Management Switch Connections	Networking Device Connections
Jump host	IPMI Provisioning Management Intranet	None
Controller Host	IPMI Provisioning Management	External Internal API Tenant Storage Storage Management

Compute Node	IPMI Provisioning Management	Internal API Tenant Storage
Storage Node	IPMI Provisioning Management	Storage Storage Management

Updating Configurations and Software in Contrail Cloud

To ensure consistent configuration of your network nodes, the following changes to nodes in the overcloud—compute nodes, storage nodes, and control hosts—should always be made through Contrail Cloud:

- Software updates.
- Configuration changes, including configuration changes to reflect the removal or addition of a new node.

Bypassing Contrail Cloud to apply a change in a lower layer often leads to configurations that are later overwritten.

Configuration Files in Contrail Cloud

A series of pre-configured YAML file templates are provided as part of a Contrail Cloud installation. These user-configurable YAML files are downloaded onto the jump host server during the initial phase of the Contrail Cloud installation. The YAML files can be accessed and edited by users from within the jump host and the updated configurations can be deployed using Ansible playbook scripts. See Deploying Contrail Cloud for additional information on YAML file locations and configuration updating procedures.

Table 8 on page 27 lists commonly-used YAML file parameters in Contrail Cloud and provides a summary of the purpose of the parameter.

Table 8: YAML File Parameters

YAML File Parameter	*Purpose*

site.yml	
global:	DNS, NTP, domain name, time zone, satellite URL, and proxy configuration for the deployment environment.
jumphost:	Provision NIC name definition and PXE boot interface for the jump host.
control_hosts:	Control host parameters. Includes disk mappings for bare metal servers and control plane VM sizing per role for functions like analytics.
compute_hosts:	Parameters for SR-IOV, DPDK, and TSN in compute nodes. Root disk configuration per hardware profile.
storage_hosts:	Ceph and block storage profiles definition for storage nodes.
undercloud:	Nova flavors for roles. Applicable when using additional hardware profiles.
overcloud:	Hardware profile and leaf number-based: • disk mappings • network definitions—names, subnets, VLANs, DHCP pools, and roles for network. Other network definitions like TLS cert, keystone LDAP backend enablement, post deployment extra actions, tripleO extra configurations
ceph:	Ceph enablement and disk assignments (pools, OSDs) on storage nodes.
ceph_external:	Externally deployed Ceph integration parameters.
appformix:	Enable HA, VIP IPs, and network devices monitoring for Appformix.
inventory.yml	
inventory_nodes:	Name, IPMI IP, Ironic driver used for LCM, root disk, and other related functions for all Contrail cluster nodes.

control-host-nodes.yml	
control_host_nodes:	Internal IP and DNS (per control node) for control hosts and the control plane.. Statically added IPs for controllers need to be outside of DHCP pools for networks that use them.
control_host_nodes_network_ config:	Bridges, bonds, DHCP/IP, and MTU for control hosts.
control_hosts:	VM interface to bridge on control-host mapping.
overcloud-nics.yml	
contrail_network_config: controller_network_config: appformixController_network_config: computeKernel_network_config: compute_dpdk_network_config: cephStorage_network_config:	Interface to network mapping, routes, DHCP-IP allocation, bonds, VLAN to interface maps, and bond options for control, storage, and compute nodes.
compute-nodes.yml	
compute_nodes_kernel: compute_nodes_dpdk: compute_nodes_sriov:	Mapping hosts from inventory to compute roles and profiles for compute nodes.
storage-nodes.yml	
storage_nodes:	Names of storage nodes.
vault-data.yml	
global:	satellite key and contrail user password for the Red Hat Open Stack Vault function.
undercloud: overcloud: control_hosts:	VM & Bare metal server (BMS) passwords for Contrail Cluster nodes and the undercloud when using the Red Hat Open Stack Vault function.

appformix:	MySQL and RabbitMQ passwords for Appformix when using the Red Hat Open Stack Vault function.
ceph_external:	Client key used by Ceph External with the Red Hat Open Stack Vault function..
inventory_nodes:	IPMI credentials for Contrail cluster nodes when using the Red Hat Open Stack Vault function.

RELATED DOCUMENTATION

- Reference Architecture for Contrail Cloud 13.1 and 13.2
- Contrail Cloud Deployment Guide
- Contrail Cloud TechLibrary page

Controller VM Requirements and Physical Server Specifications

The controller VM requirements vary by Contrail Cloud environment.

Table 9 on page 30 lists the recommended CPU, memory, and disk requirements for the controller VMs.

Table 9: Controller VMs Requirements

Role	vCPU (Threads)	Memory (GB)	Disk (GB)
Undercloud VM	28	128	500
OpenStack Controller VM	8	48	500

Table 9: Controller VMs Requirements *(Continued)*

Role	vCPU (Threads)	Memory (GB)	Disk (GB)
Contrail Analytics DB VM	12	48	500 & 1000
Contrail Analytics VM	12	48	250
Contrail Controller VM	16	64	250
AppFormix VM	16	32	500
TSN (Contrail Service Node)	4	8	100
Control Host OS	4	8	100

In order to address the controller VM requirements listed in Table 9 on page 30, the typical physical server specifications are listed in Table 10 on page 31. Depending on the actual workloads and the data plane traffic requirements, the network interfaces could be 25G/40G versus 10G.

Table 10: Physical Server Specifications

Role	Cores	Memory	Disk	Network
Jumphost	20 (2x10)	128 GB	2 x 1 TB HDD (SW RAID) — Host OS **NOTE**: HW Raid recommended if RAID controller is available in your environment. 3 x 1 TB HDD — Data	1G — IPMI 1G/10G — Provisioning 1G/10G — Intranet

Table 10: Physical Server Specifications *(Continued)*

Role	Cores	Memory	Disk	Network
Control Host	36 (2x18)	256 GB	1 x 1 TB SSD (Cassandra Journaling) **NOTE**: 1x1TB is the recommended disk memory option for control hosts. 2 x 1 TB HDD (HW RAID) —Host OS 3 x 2 TB HDD (HW RAID) —Data	1 x 1G — IPMI 1x1G/10G — Provisioning 1x1G/10G — Management (optional) 4 x 10G/25G/40G ports on two Intel Fortville (XL710/X710/XXV710) family NICs. Both high-speed NICs must be on NUMA 0.
Compute Node	28 (2x14)	256 GB	2 x 1 TB HDD (HW RAID) — Host OS 4 x 1 TB HDD — Data (only needed if Ceph storage not used to provide storage)	1 x 1G — IPMI 1x1G/10G — Provisioning 1x1G/10G — Management (optional) 4 x 10G/25G/40G ports on two Intel Fortville (XL710/X710/XXV710) family NICs. Both high-speed NICs must be on NUMA 0.
Storage Node	16 (2x8)	128 GB	2 x 1 TB HDD (HW RAID) — Host OS 2 x 480 GB SSDs — Ceph Journaling 10 x 1 TB HDDs — Ceph OSDs	1 x 1G — IPMI 1x1G/10G — Provisioning 1x1G/10G — Management (optional) 4 x 10G/25G/40G ports on two Intel Fortville (XL710/X710/XXV710) family NICs. Both high-speed NICs must be on NUMA 0.