
ORACLE®

Oracle Solaris OpenStack

Frankfurt Solaris User Group
Introduction, Experiences, Live Demonstration

Detlef Drewanz
Master Principal Sales Consultant
Northern Europe Systems Architects

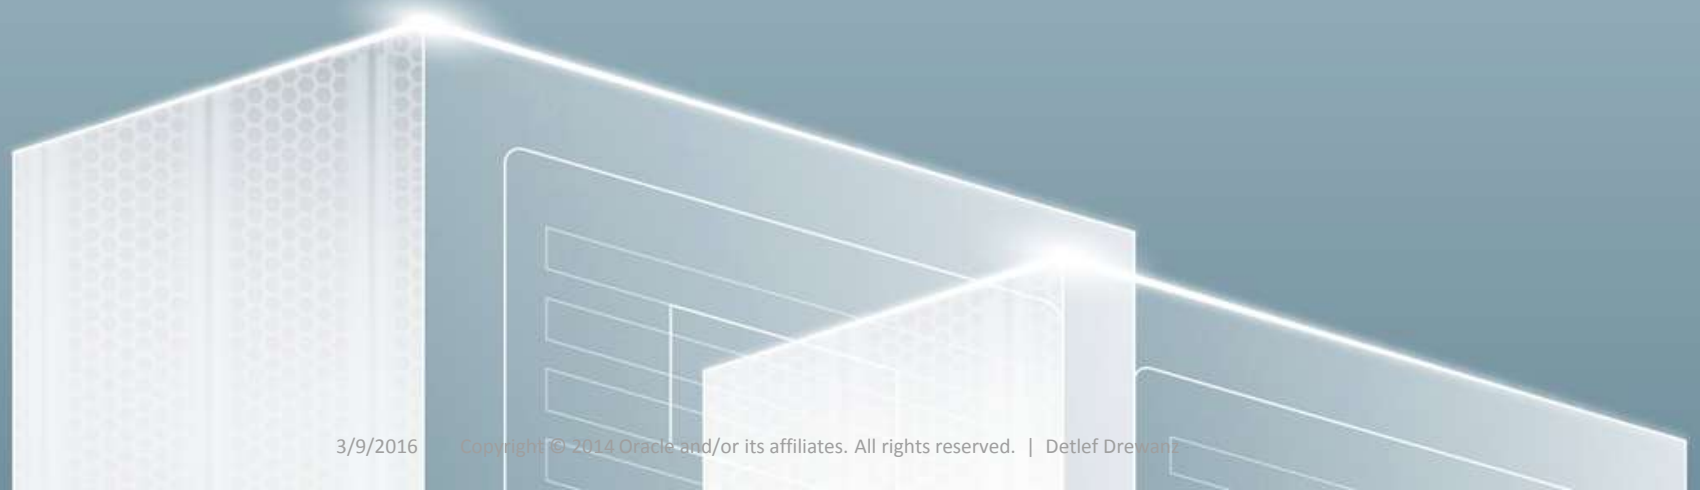

Safe Harbor Statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Agenda

1. The Road to private Cloud
2. The OpenStack Project
3. The OpenStack Architecture
4. OpenStack and Oracle
5. Cloud Use Cases
6. Building the Cloud

Journey to Private Cloud

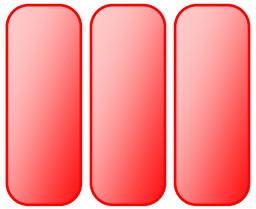

Traditional Silos

- Physical
- Dedicated
- Heterogeneous

Consolidated

- Virtual
- Shared
- Standardized

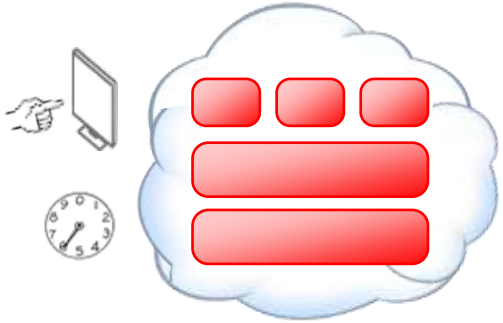

Private Cloud

- **Security**
- **Self-service**
- **Auto-scaling**
- **Metering and chargeback**

Start with consolidation • Extend to private cloud • Use public cloud where appropriate

Design Considerations – The planned Use Cases

- Targeting new software development only ?
- Moving existing services also ?
- Organization-wide, or just for certain parts ?
- Offer self-service ?

The OpenStack Project

OpenStack Overview

What is OpenStack?

- Open source cloud software
  - Generic solution for IaaS, PaaS and SaaS
  - Modular architecture
    - Web portal for cloud admins and self-service users
    - Cloud services exposed through APIs
    - CLI, Python libraries, ...
- Combines compute, network and storage resources
  - Self-service dashboard
  - Services exposed through REST APIs

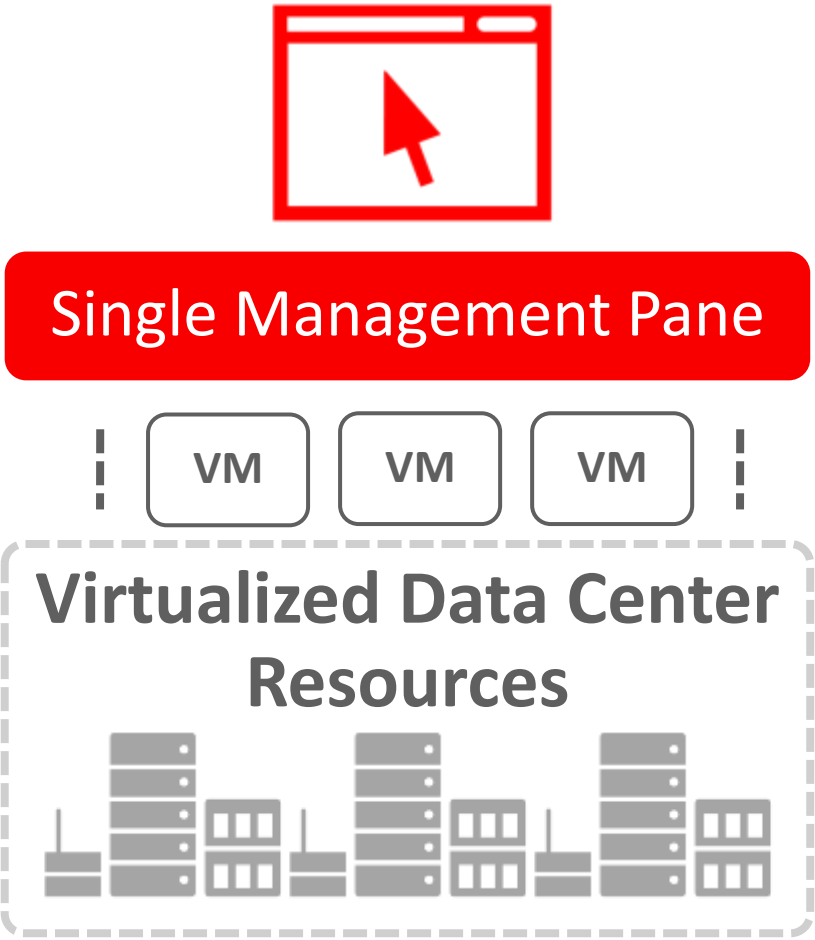

Drivers for OpenStack Cloud Infrastructure

What's exciting customers about OpenStack?

- Broad Industry Adoption and Participation in the Community
  - Over **400** companies involved!
- Rich technology Ecosystem
  - Provides choice / freedom to leverage underlying Compute, Storage, Networking Infrastructure
- "Standardized" Cloud API layer
  - Abstracts higher level Cloud Software & Services from *Undercloud* infrastructure specifics

OpenStack Releases

https://wiki.openstack.org/wiki/Releases

Release Name	Status	Latest Release Info
Austin, Bexar, Cactus, Diablo	Deprecated	2010.1, 2011.1, 2011.2, 2011.3
Essex, Folsom, Grizzly	EOL	2012.1, 2012.2, 2013.1
Havana	EOL	2013.2
Icehouse	EOL	2014.1
Juno	Security-supported	2014.2
Kilo	Security-supported	2015.1
Liberty	Current stable release, security supported	(10/15/2015)
Mitaka	Under development	

What OpenStack is not…

- Out-of-the-Box Ready to Use
  - Lots of work to configure underlying infrastructure
    - Overall Integration
    - Lots of moving parts
    - Dependent on expertise, experience, vendors, maturing technology

- Mature
  - Release cycle; short support lifespan
  - Update capability is limited, maybe “rip and replace”
  - However … it is improving
- Free
  - OpenSource does not mean free
  - Large investment to set up and support environment

The OpenStack Architecture

OpenStack Logical Layout

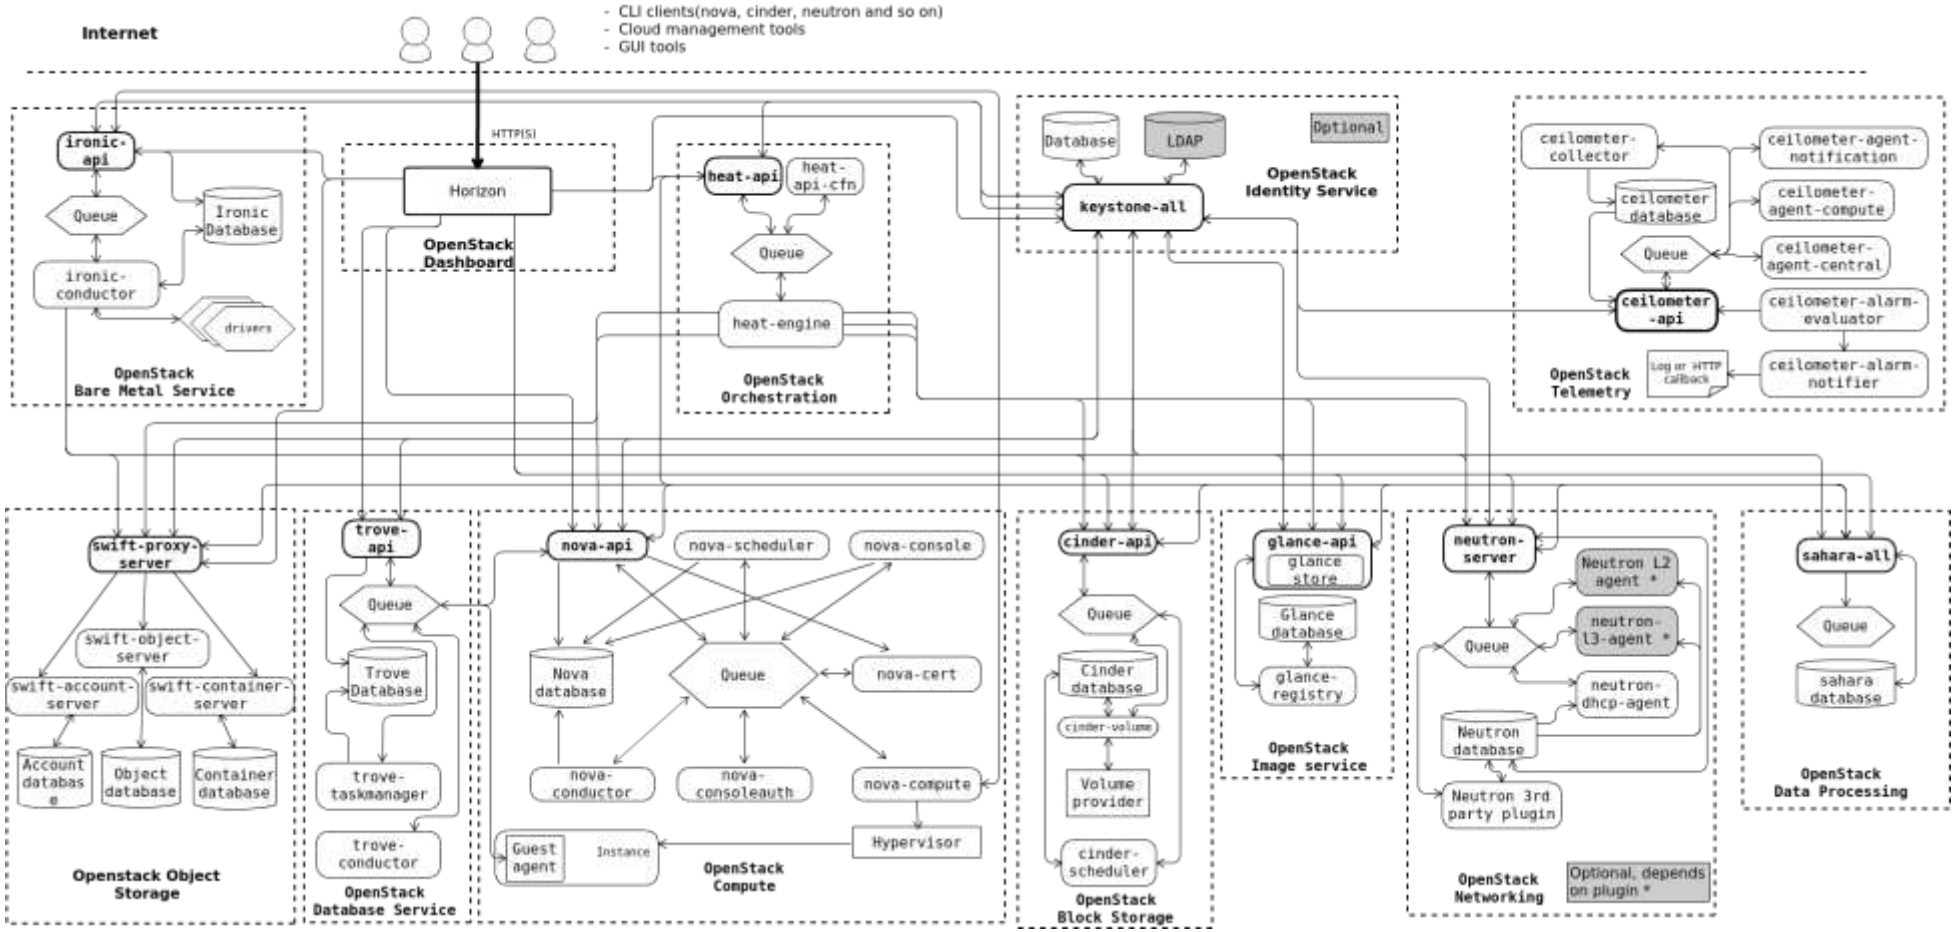

OpenStack Services

Overview of Core Components

Component	Description	Component	Description
Nova	Compute virtualization	Glance	Image management and deployment
Cinder	Block storage	Swift	Object storage
Neutron	Software defined networking	Heat	Application and VM orchestration
Keystone	Authentication between cloud services and simple Authorization	Murano	Application catalog
Horizon	Web based dashboard	Trove	Database as a Service

OpenStack Logical Architecture with Service Names

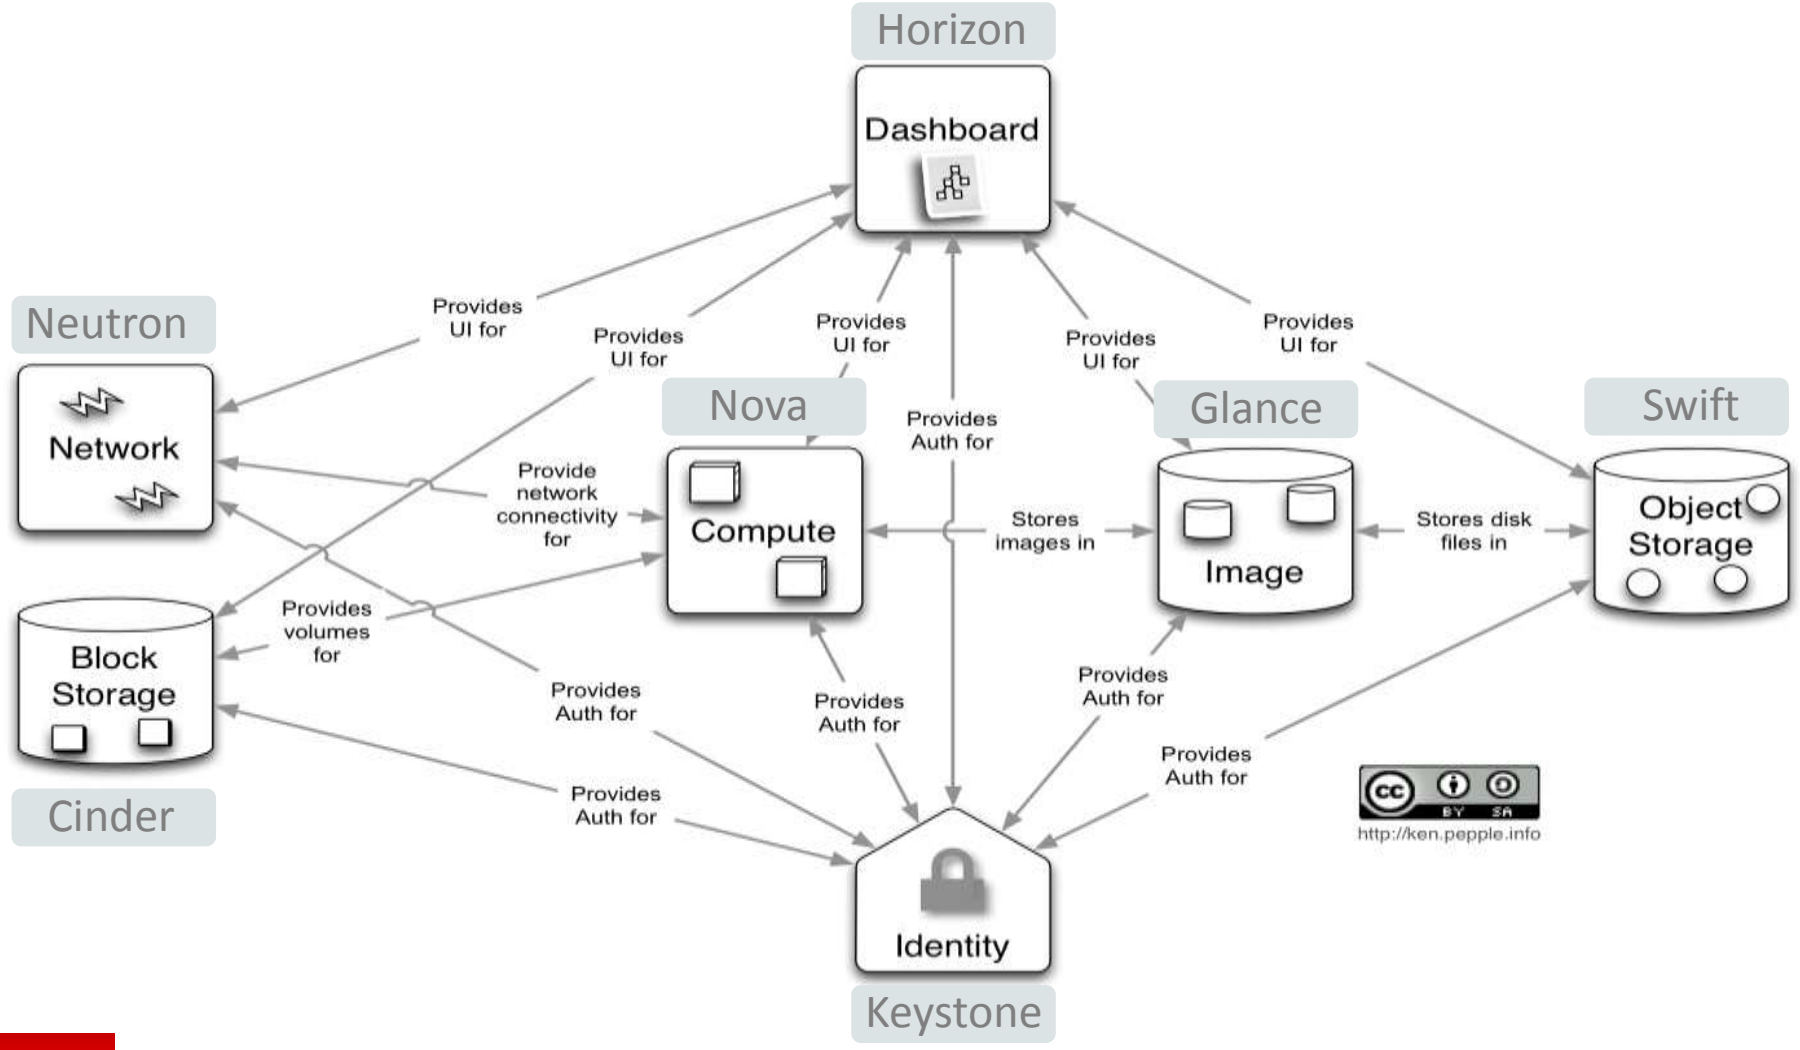

OpenStack - Modular Architecture

- Web portal / dashboard for cloud admins and self-service users
- Cloud services exposed through APIs
- CLI, Python libraries, ...
- Interoperating services with REST APIs

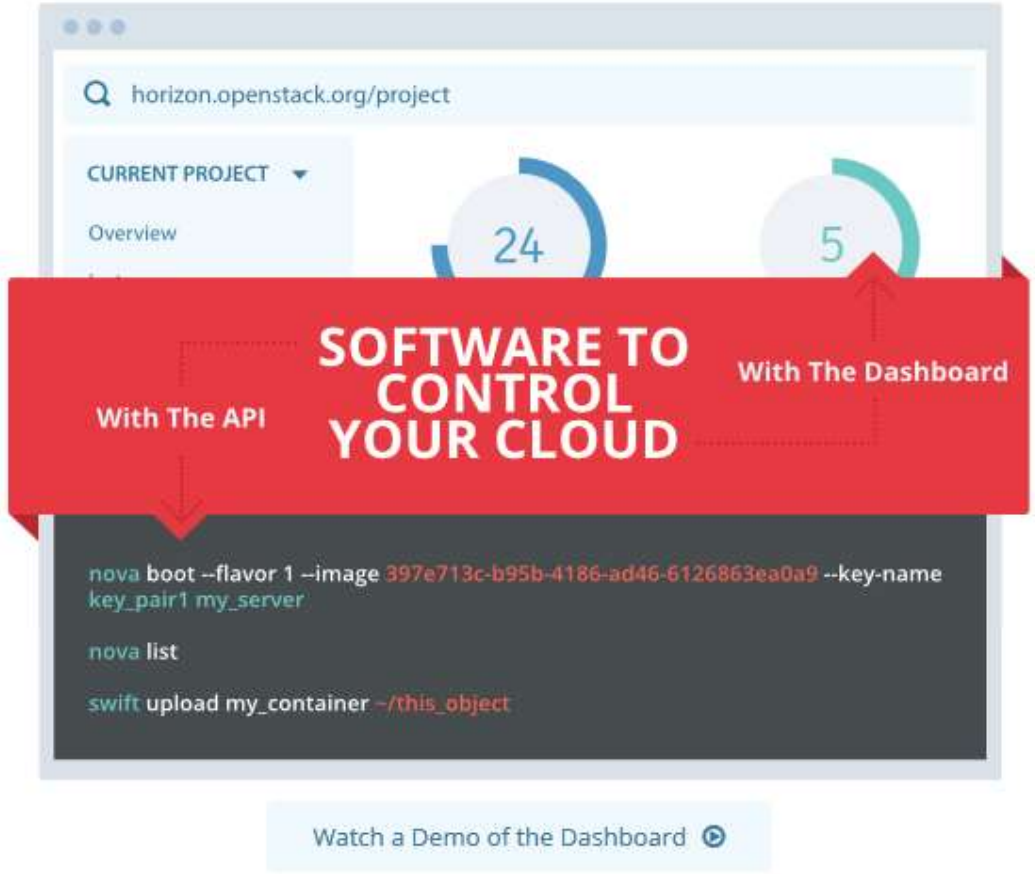

www.openstack.org

OpenStack CLI

Procedure 1.1. To create and attach a volume

1. You create a volume.

   For example, you might create a 30 GB volume called `vol1`, as follows:

```
$ cinder create --display-name vol1 30
```

   The command returns the `521752a6-acf6-4b2d-bc7a-119f9148cd8c` volume ID.

2. You attach that volume to a virtual machine (VM) with the `616fb98f-46ca-475e-917e-2563e5a8cd19` ID, as follows:

   For example:

```
$ nova volume-attach 616fb98f-46ca-475e-917e-2563e5a8cd19 521752a6-acf6-4b2d-bc7a-119f9148cd8c /dev/vdb
```

OpenStack and Oracle

OpenStack and Oracle Storage

Oracle ZFS Storage Appliance Simplifies OpenStack Deployments

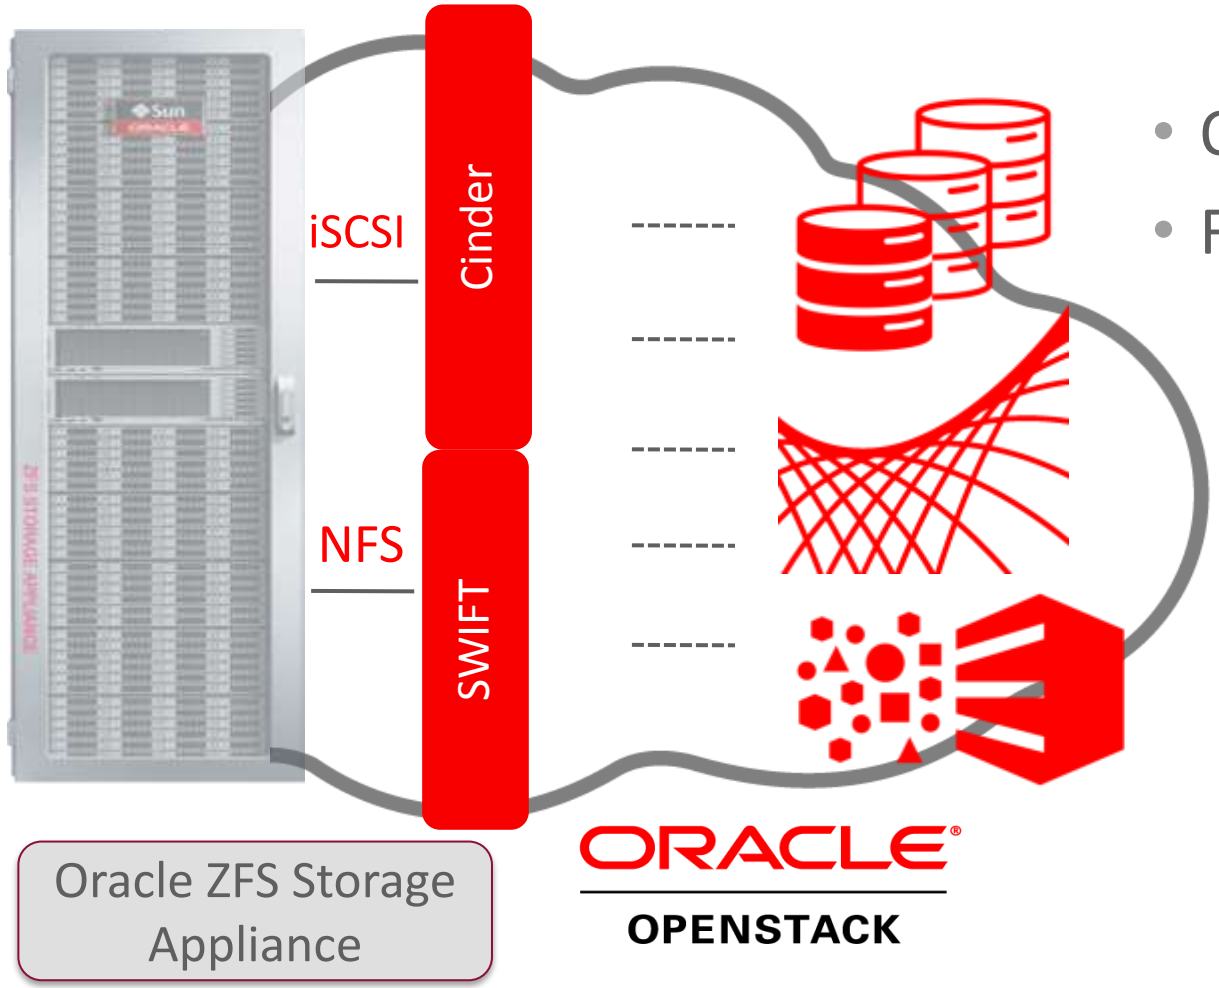

- OpenStack drivers
- Fully implemented RESTful APIs

Oracle FS1 Flash Storage System
OpenStack Ready

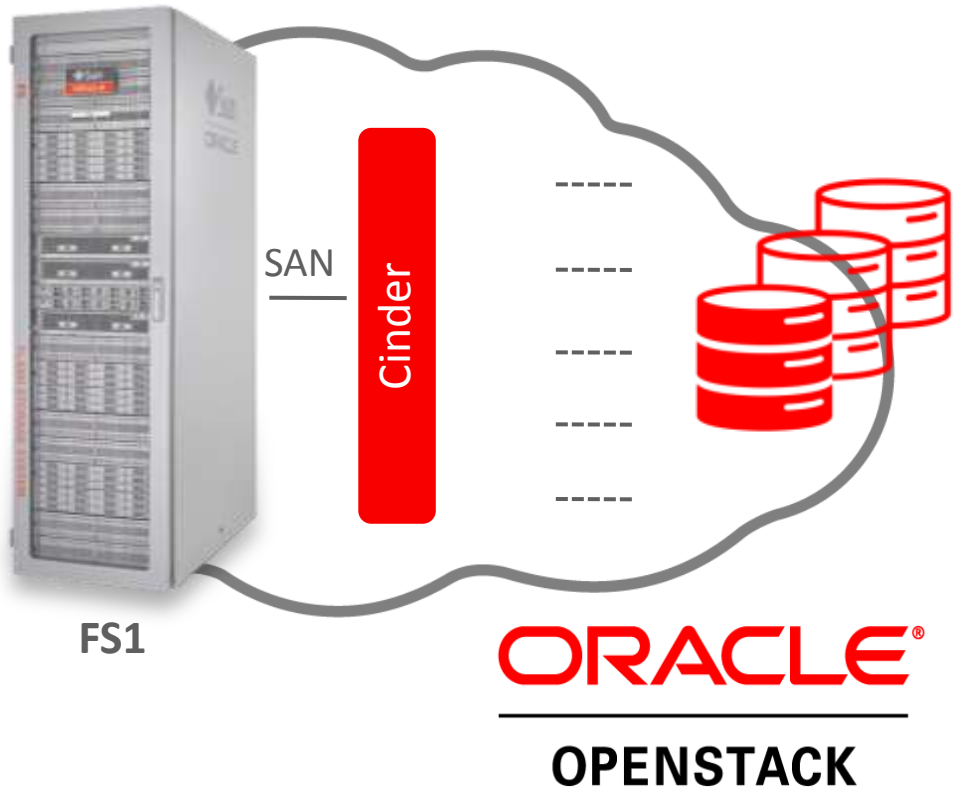

- Massively Scalable Flash Array
- Extreme Performance
- Unparalleled Flexibility
- QoS Plus Autonomous Tiering
- Perfect for Virtualized IT/OpenStack

OpenStack and Oracle Solaris

Oracle OpenStack for Oracle Solaris

ORACLE
SOLARIS

Horizon
Cloud Management

Ironic
Bare Metal
SPARC/x86

Heat
Orchestration
Unified Archives

Murano
Application
Oracle Database and Fusion Middleware

Nova
Compute Virtualization
Zones & Kernel Zones

Neutron
Cloud Networking
Elastic Virtual Switch

Cinder/Swift
Cloud Storage
ZFS File System

Glance
Image Deployment
Unified Archives

OpenStack Juno in Oracle Solaris 11.3

Ready for Enterprise: OS. Virtualization. SDN. OpenStack.

ORACLE
SOLARIS

- Engineered for security and compliance
  - Minimal privileges for cloud services
  - Lock down infrastructure with immutability
- Assured reliability and scale
  - Automatic service restart and node dependencies
  - Guaranteed data integrity
- Seamless upgrade, instant roll-back

ORACLE
OPENSTACK

Oracle Solaris 11 Packaging

- Secure and seamless software delivery mechanism
- Impossible to install mismatched software, including firmware
- No modifications to running OS, complete safety with BEs
- Fully integrated with Zones

Image Packaging System

Easy and fast cloud update

- Seamless integration with IPS
  - Full dependency management
  - Fail proof updates with rollback
- Integrated with Oracle Solaris Zones and Unified Archives for seamless lifecycle management
- Foundation for cloud update strategy
  - Configuration and database schemas updated through SMF update services with full rollback

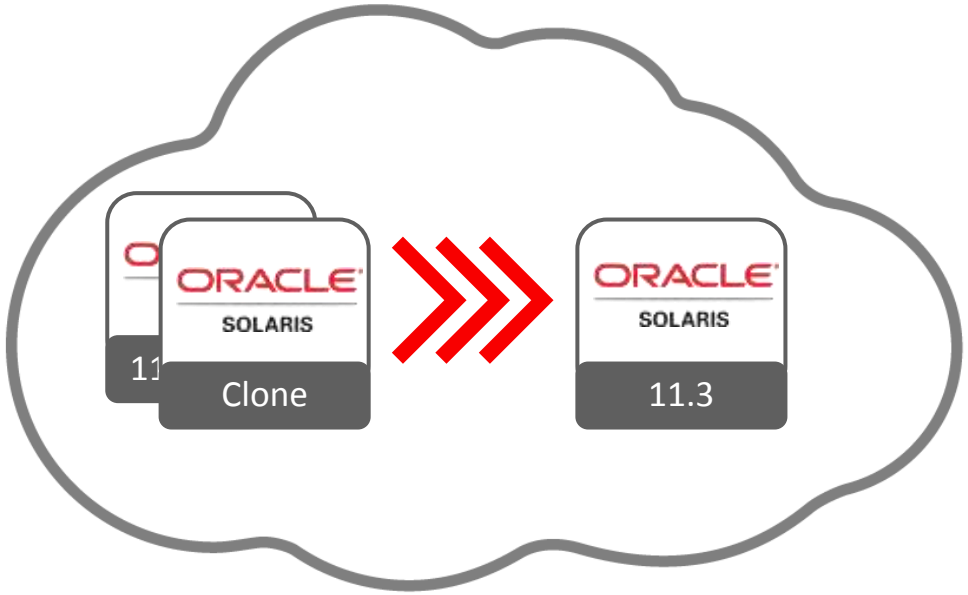

OpenStack and Oracle Solaris 11.3.3 - Package List

- + RabbitMQ as Messaging Broker

```
# pkg list -af | grep openstack
cloud/openstack                 0.2014.2.2-0.175.3.0.0.30.0     i--
cloud/openstack/cinder          0.2014.2.2-0.175.3.2.0.2.0      i--
cloud/openstack/glance          0.2014.2.2-0.175.3.0.0.30.0     i--
cloud/openstack/heat            0.2014.2.2-0.175.3.0.0.30.0     i--
cloud/openstack/horizon         0.2014.2.2-0.175.3.0.0.30.0     i--
cloud/openstack/ironic          0.2014.2.1-0.175.3.0.0.30.0     i--
cloud/openstack/keystone        0.2014.2.2-0.175.3.0.0.30.0     i--
cloud/openstack/neutron         0.2014.2.2-0.175.3.0.0.30.0     i--
cloud/openstack/nova            0.2014.2.2-0.175.3.0.0.30.0     i--
cloud/openstack/swift           2.2.2-0.175.3.0.0.30.0          i--
...
```

Openstack and Oracle Solaris 11.3 - Versioning

```
# pkg info openstack
Name:              cloud/openstack
Summary:           OpenStack
Description:       OpenStack is a cloud operating system that controls large pools
                   of compute, storage, and networking resources throughout a data
                   center, all managed through a dashboard that gives
                   administrators control while empowering their users to provision
                   resources through a web interface.
Category:          Meta Packages/Group Packages (org.opensolaris.category.2008)
                   System/Administration and Configuration (org.opensolaris.category.2008)
                   System/Enterprise Management (org.opensolaris.category.2008)
                   System/Virtualization (org.opensolaris.category.2008)
                   Web Services/Application and Web Servers (org.opensolaris.category.2008)
State:             Installed
Publisher:         solaris
Version:           0.2014.2.2 (Juno 2014.2.2)
Build Release:     5.11
Branch:            0.175.3.0.0.30.0
Packaging Date:    August 21, 2015 04:14:44 PM
Size:              5.46 kB
FMRI:              pkg://solaris/cloud/openstack@0.2014.2.2,5.11-0.175.3.0.0.30.0:20150821T161444Z
```

Service Management Facility

Secure and Highly available and reliable cloud services

- OpenStack services run with minimum privileges necessary, and don't run as root
  - Create users for different OpenStack services
  - Leverage RBAC to enable privileged actions
- Automatic service restart from failure
  - Integrated with Oracle Solaris fault management
  - Full dependency checking for precise and efficient cloud start up

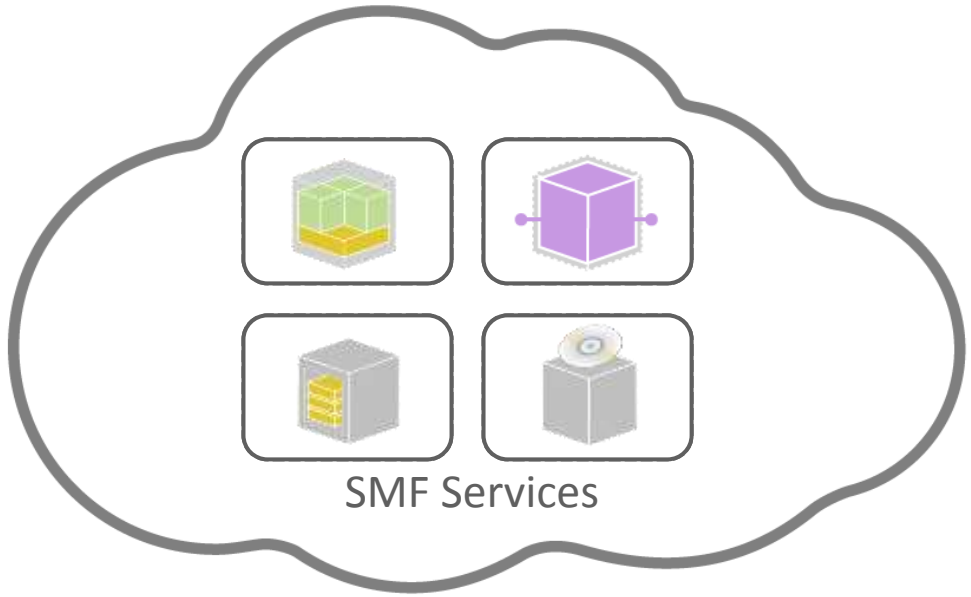

OpenStack Cinder/Swift Data Management – ZFS

Production ready data management, no compromises

- ZFS is foundation for Cinder and Swift
  - iSCSI or FC LUN provisioning
  - Leverage integrated data services including snapshots, compression and encryption
  - These data services are completely transparent to the guests
- Integrated OpenStack support for ZFSSA

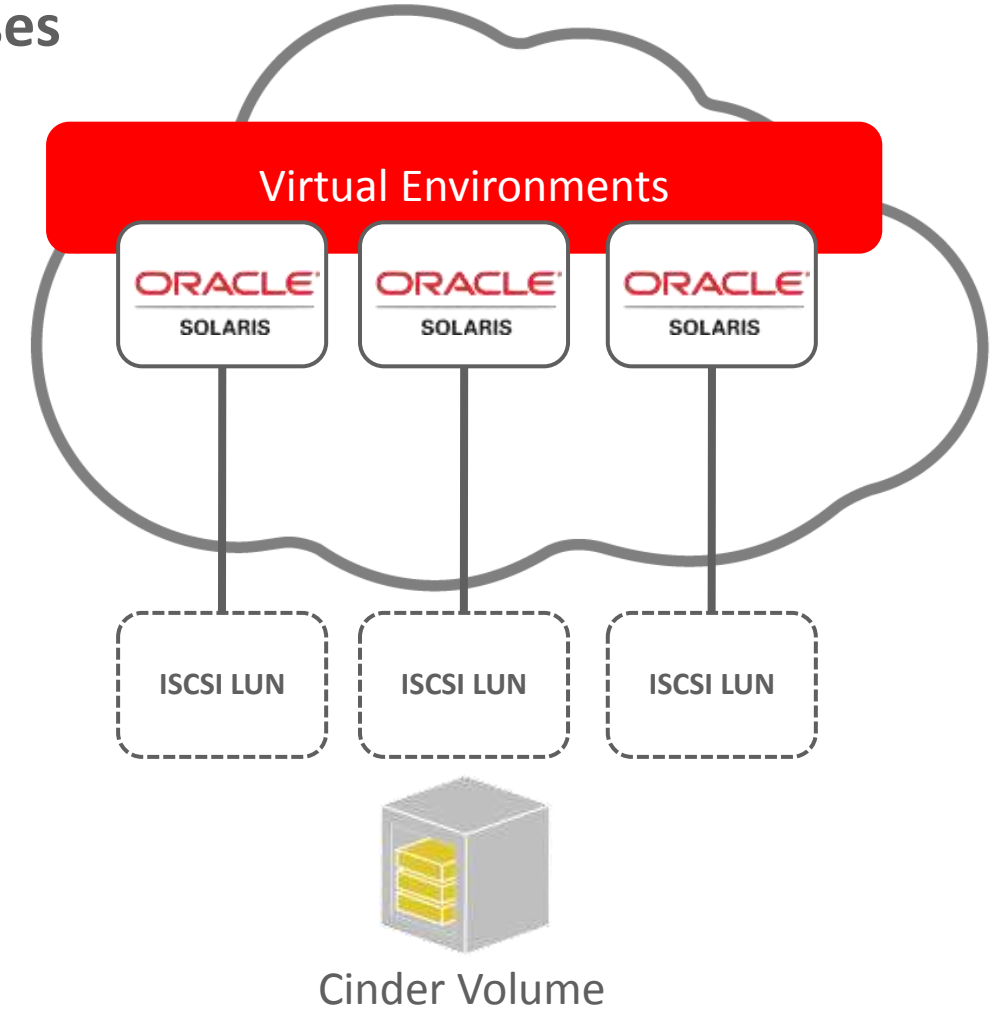

OpenStack Cinder Data Management – ZFS

Volume Driver

- Choose volume driver in /etc/cinder/cinder.conf
  - ZFSVolumeDriver
    - Supports creation of local volumes for use by Nova on the same node as the Cinder volume service.
  - ZFSISCSIDriver
    - Supports creation and export of iSCSI targets for use by remote Nova compute nodes.
  - ZFSFCDriver
    - Supports creation and export of Fibre Channel LUNs for use by remote Nova compute nodes.
  - ZFSSAISCSIDriver
    - Supports creation and export of iSCSI targets from a remote Oracle ZFS Storage Appliance for use by remote Nova compute nodes.

OpenStack Nova Compute – Oracle Solaris Zones

High density virtual environments – ideal for multi-tenant cloud

- Integrated with Oracle Solaris Zones
  - Zero overhead virtualization
  - Native non-global zone, Kernel zones
- Fully portable with Unified Archives
  - Physical-to-virtual and virtual-to-physical transforms
- VM lockdown with Immutable Zones

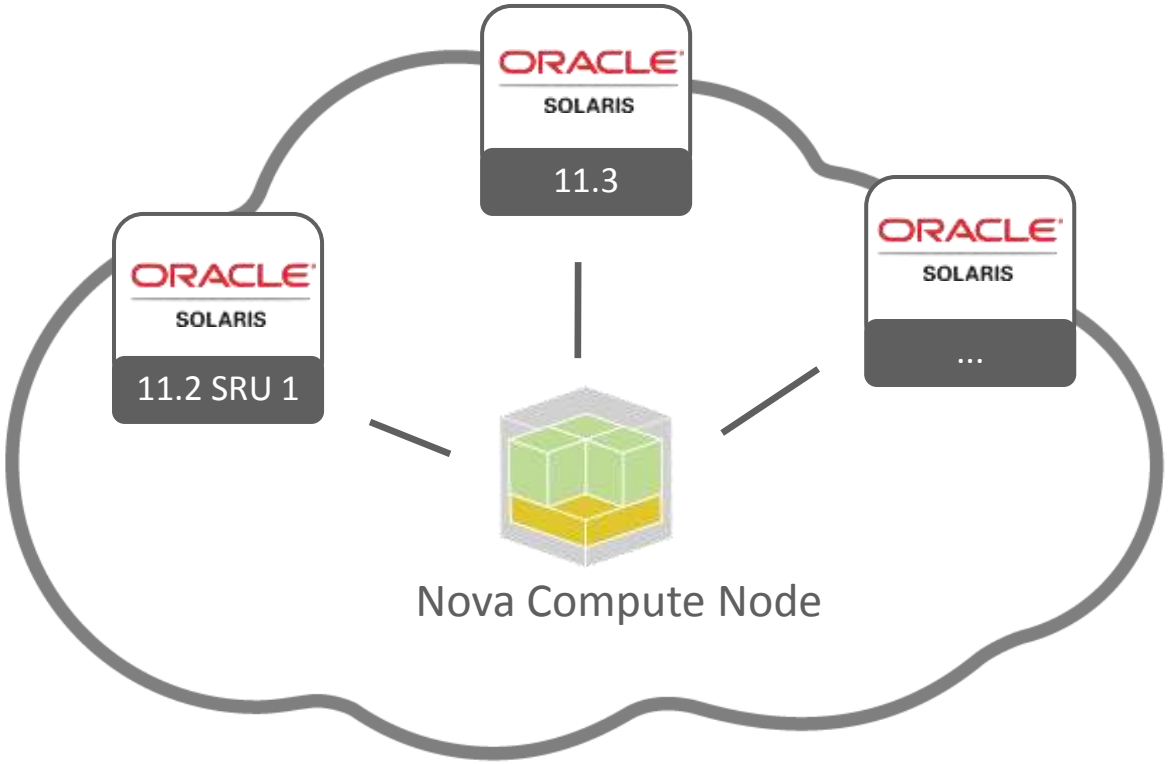

OpenStack Glance Image Management – Unified Archives

Rapid deployment through VM templates

- Golden image deployment
  - Single archive for virtual and bare metal
- Live VM snapshots in Glance with ZFS
- Capture live production systems
  - Clone archives for cloud-like images
  - Recovery archives for bare metal backup

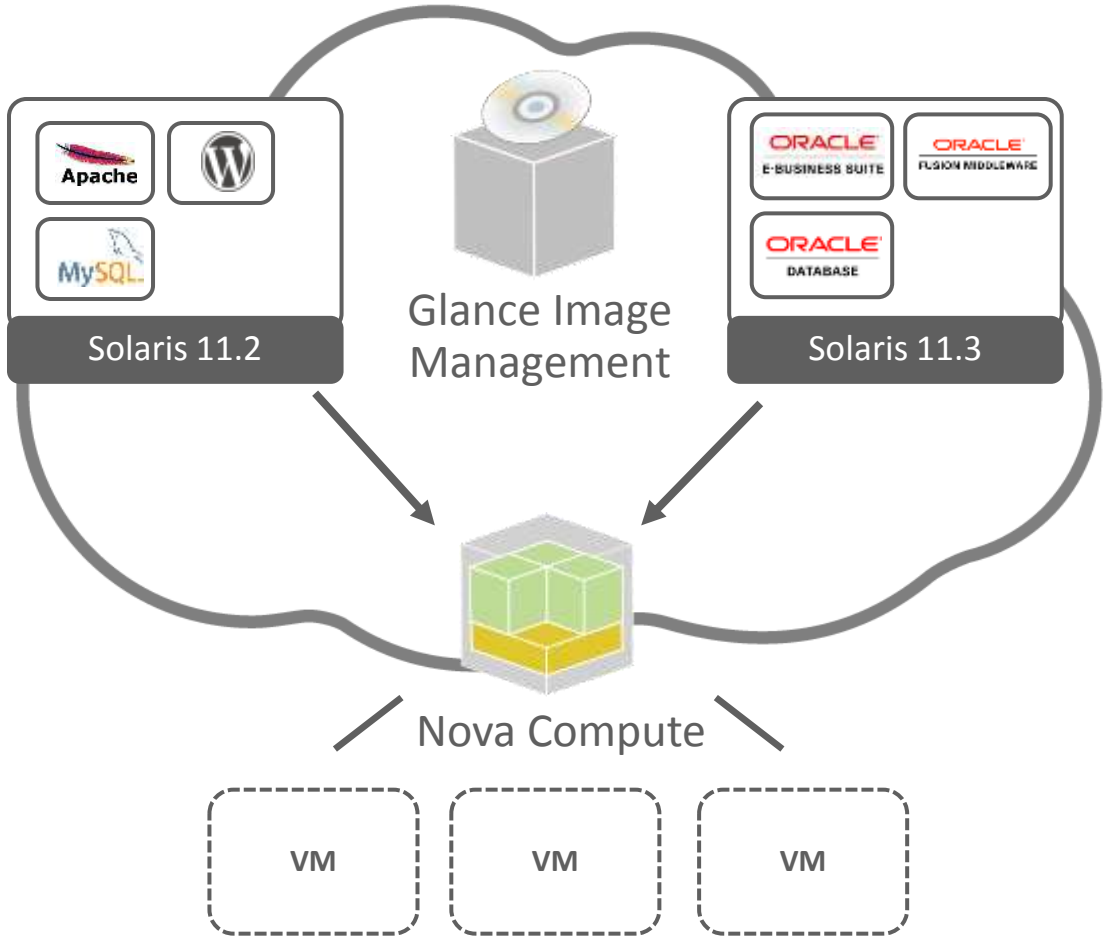

Creating and Importing an Unified Archive into Glance

```
global# zonecfg -z myzone create
global# zoneadm -z myzone install
global# ...
global# archiveadm create -z myzone /var/tmp/myzone.uar
global# glance image-create --container-format bare --disk-format raw \
--is-public true --name "Oracle Solaris 11.2 x86 NGZ" \
--property architecture=x86_64 \
--property hypervisor_type=solariszones \
--property vm_mode=solariszones < /var/tmp/myzone.uar
```

OpenStack Neutron Networking – EVS

SDN for servers and switches

- Elastic Virtual Switch spans compute nodes
  - Isolated VLAN or VXLAN networks, or plumbed into an existing fixed network
  - Integrated with Solaris virtual networking
- Ensure network critical SLAs
  - Flexible resource management
  - Application driven SDN

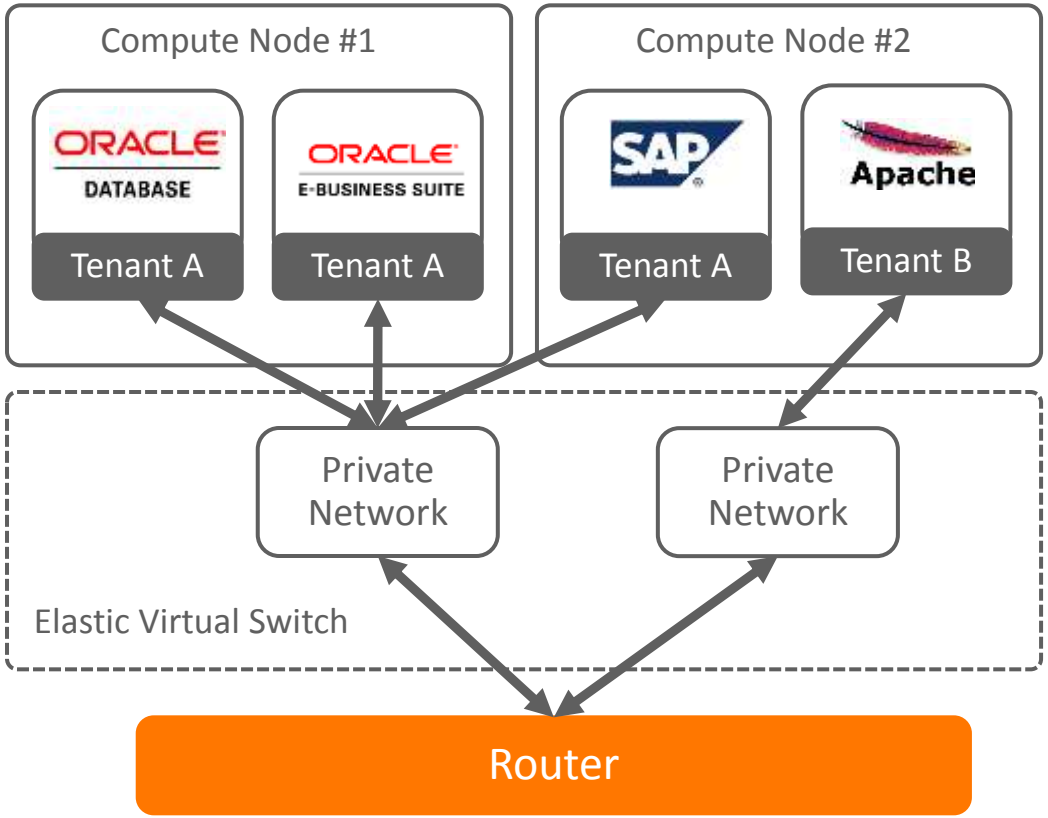

General Guidelines Concerning OpenStack on Solaris

OpenStack packaging on Solaris

- Each of the eight major components delivered by a single package
  - pkg:/cloud/openstack/<component>
  - The corresponding client APIs and CLI are in pkg:/library/python/<component>client
- The packages typically deliver their configuration files under /etc/<component>
- Horizon uses /etc/openstack_dashboard
- Runtime information is stored under /var/lib/<component>
- Group package, pkg:/cloud/openstack, installs all components.

General Guidelines Concerning OpenStack on Solaris

OpenStack services on Solaris

- Each component is represent by one or more smf(5) services
  - Some are meant to be run on only one node
  - Others can be replicated for reliability or to meet expected demand
    - Common case is deploying additional nodes running Nova compute or Cinder volume services
  - Horizon is enabled through the use of Apache and a configuration file
- Services names are of the form
  - svc:/application/openstack/<component>/<component>-<service>
- Corresponding smf(5) logs contain a wealth of debugging information
  - Components support both a debug and a verbose mode settable in the configuration

General Guidelines Concerning OpenStack on Solaris

OpenStack and RBAC on Solaris

- Each package delivers a RBAC profile for administering the component
  - Provide authorizations for managing corresponding smf(5) services and properties
  - Allows the modification of corresponding configuration files via pfedit(1M)
  - Allows the reading of the service log files which are normally not world-readable
- Each package delivers a unique user and group for the component
  - The smf(5) services run a method context with this user and group
  - Services only include minimum necessary privileges in order to operate

Oracle Solaris OpenStack Clouds

General Architecture and Use Cases

Cloud Base Architecture

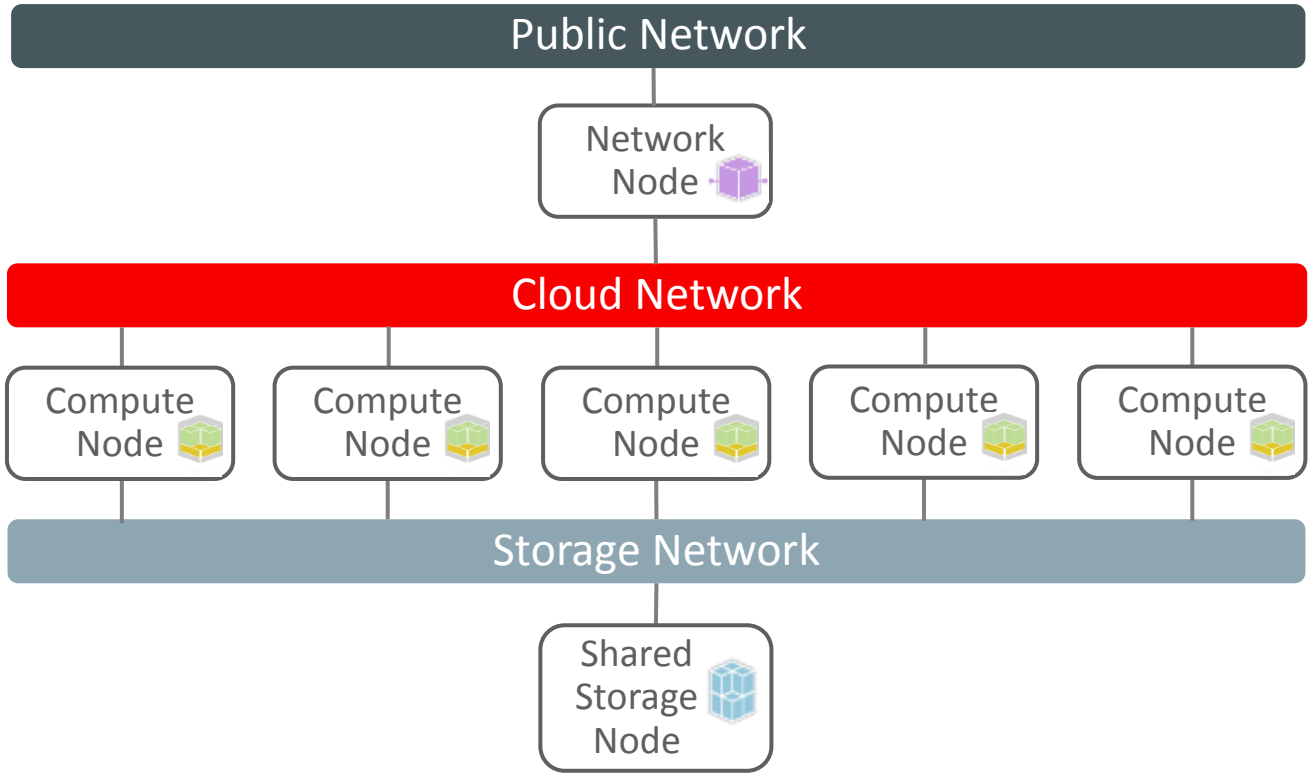

Cloud with Shared Storage and Controller

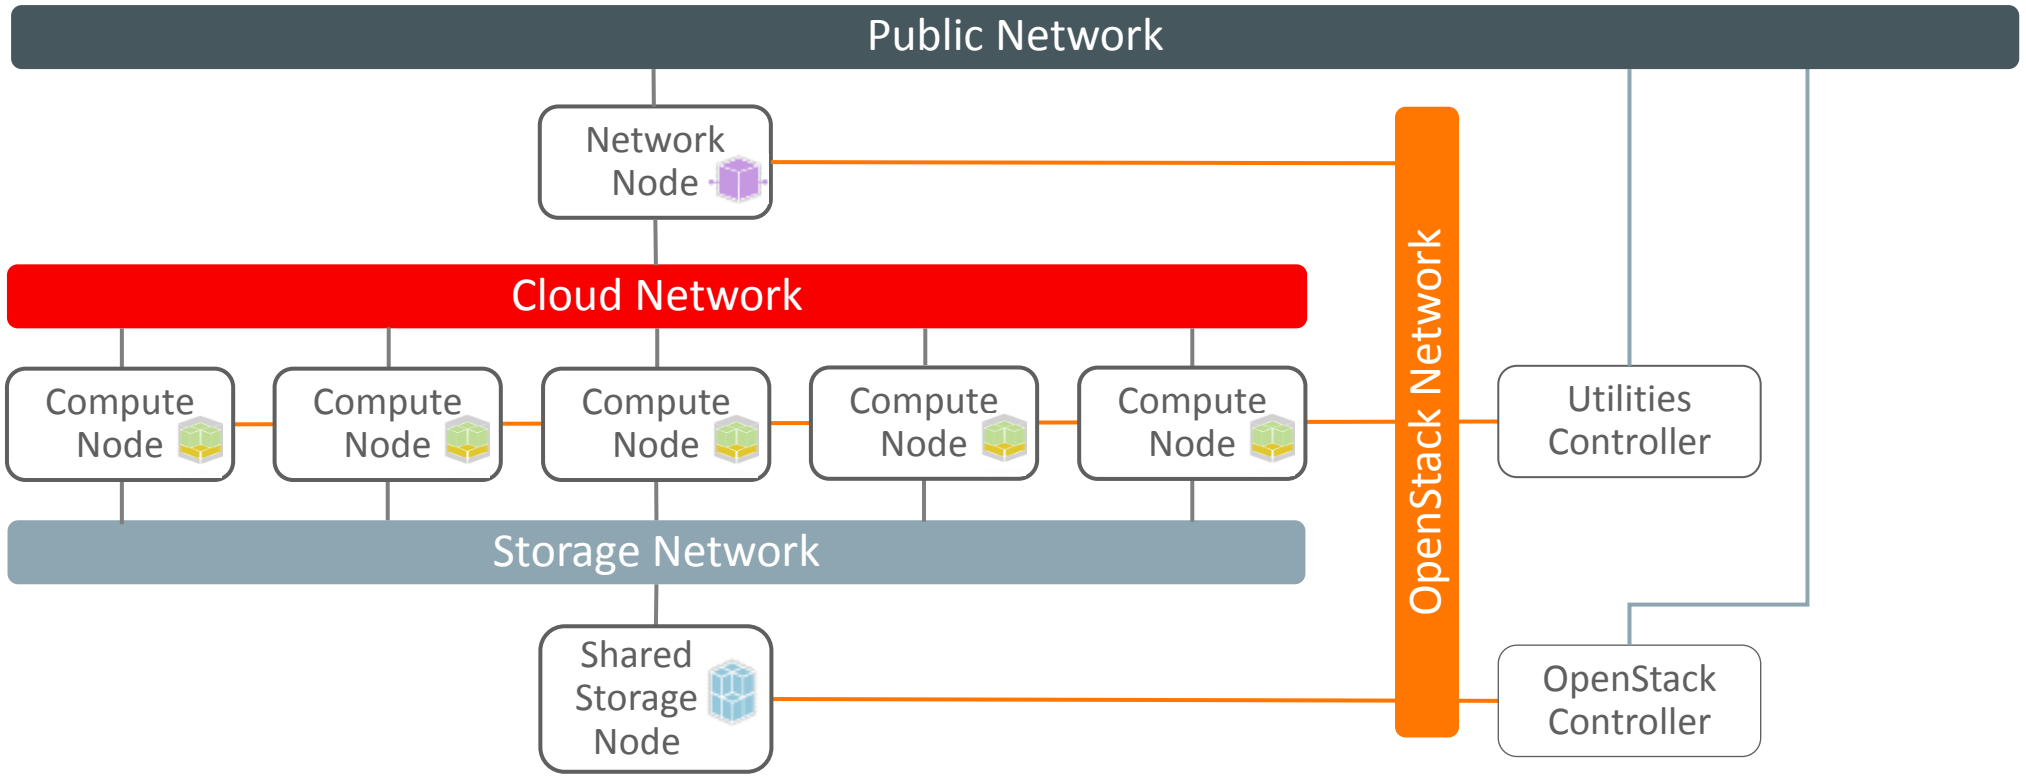

Use Case 1: Converged Cloud Infrastructure (CCI)

- OpenStack Setup for comprehensive, isolated Clouds
  - Compute Nodes are "invisible" to the corporate network
  - The Cloud Instances are connected via a L3 Node to the Corporate network

- Shared Storage enables OpenStack Instance Migration and Node Evacuation
  - Shared Storage with ZFSSA (**CCIZ**)
  - Shared Storage with a Solaris Node as iSCSI target server (**CCIN**)
  - Shared Storage with a FC SAN Storage System (**CCIS**)

- Local storage for best Storage Performance
  - Local Storage through the Compute Nodes internal disks (**CCIL**)

Use Case 2: Flat Cloud Infrastructure (FCI)

- OpenStack Setup for simple Clouds with best Network Performance
  - Compute Nodes are directly connected to the corporate network
  - The Cloud Instances are physical connected to the Corporate network and share VLAN/VXLAN networks
- Shared Storage enables OpenStack Instance Migration and Node Evacuation
  - Shared Storage with ZFSSA (**FCIZ**)
  - Shared Storage with a Solaris Node as iSCSI target server (**FCIN**)
  - Shared Storage with a FC SAN Storage System (**FCIS**)
- Local storage for best Storage Performance
  - Local Storage through the Compute Nodes internal disks (**FCIL**)

Cloud Networking

Oracle Solaris and SDN

Converged Infrastructure (CI) with OpenStack

Tenants

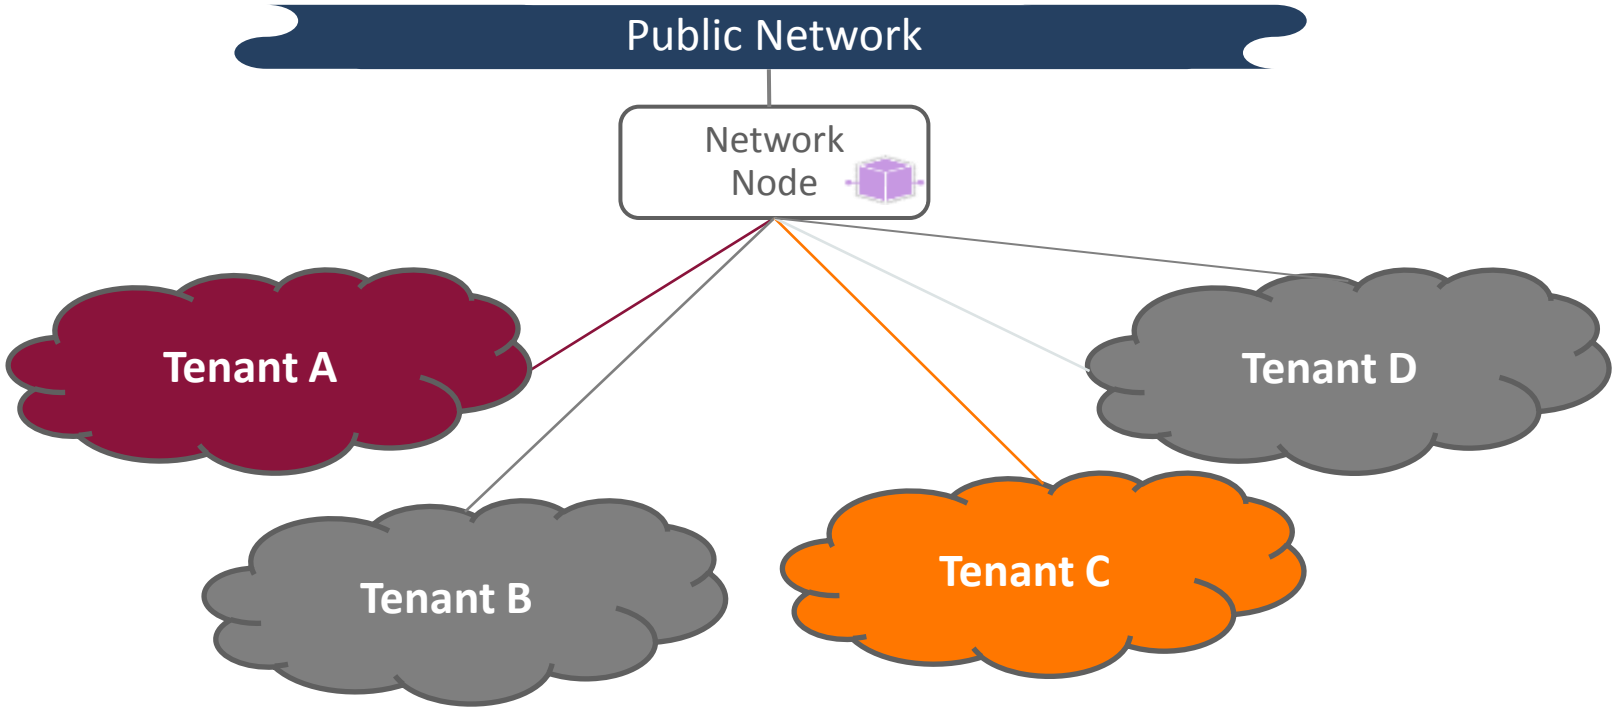

Cloud with Shared Storage and Controller

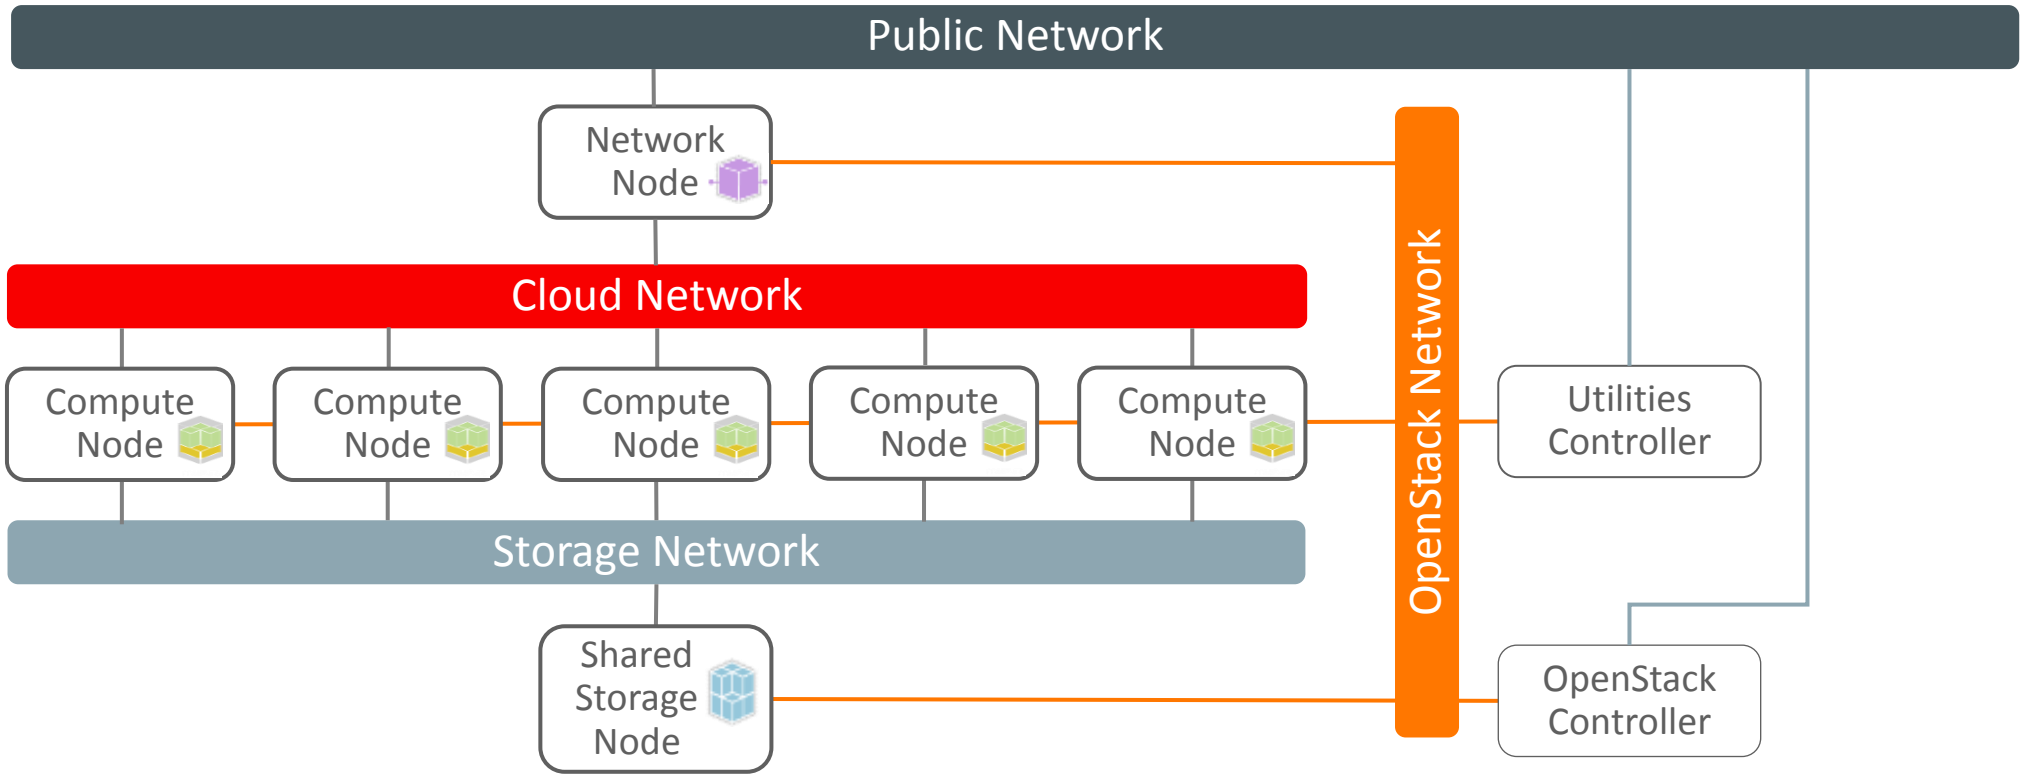

Converged Infrastructure (CI) with OpenStack

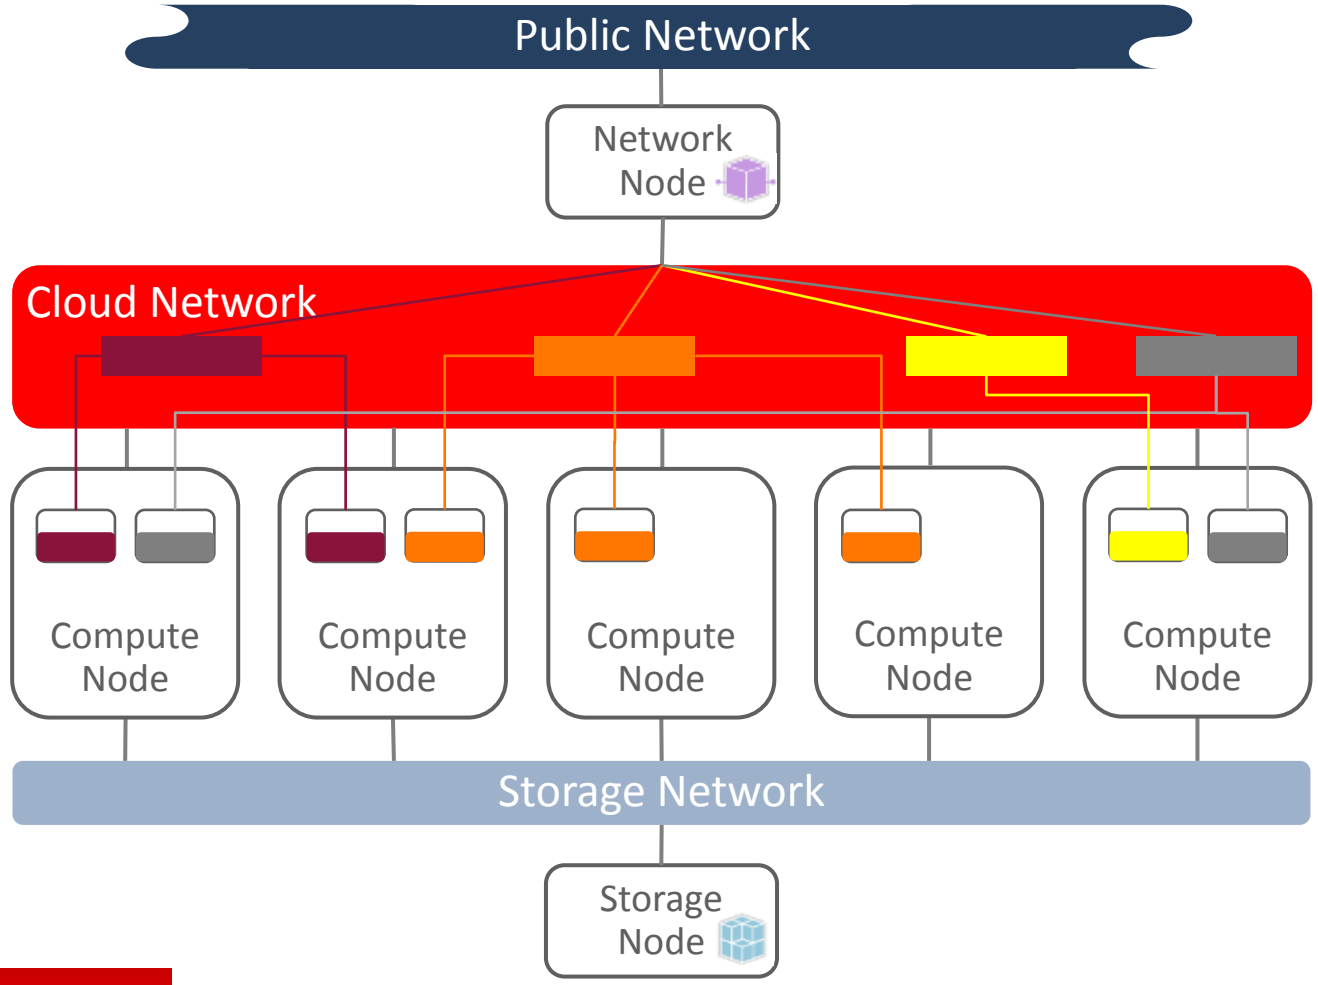

VLAN vs. VXLAN

- It's all about creating logical Layer-2 Networks
- VXLAN (Virtual Extensible LAN)
  - Encapsulation-Protocol for Layer-2 Overlay-Networks above Layer-3
  - 16 Mio Network-IDs
  - No Switch Support required
  - Need Multicast Forwarding to span multiple Subnets
- VLAN (Virtual LAN)
  - Encapsulation-Protocol for Layer-2 Overlay-Networks above Layer-2
  - 4096 Network-IDs
  - Switch Support required

Converged Infrastructure (CI) with OpenStack

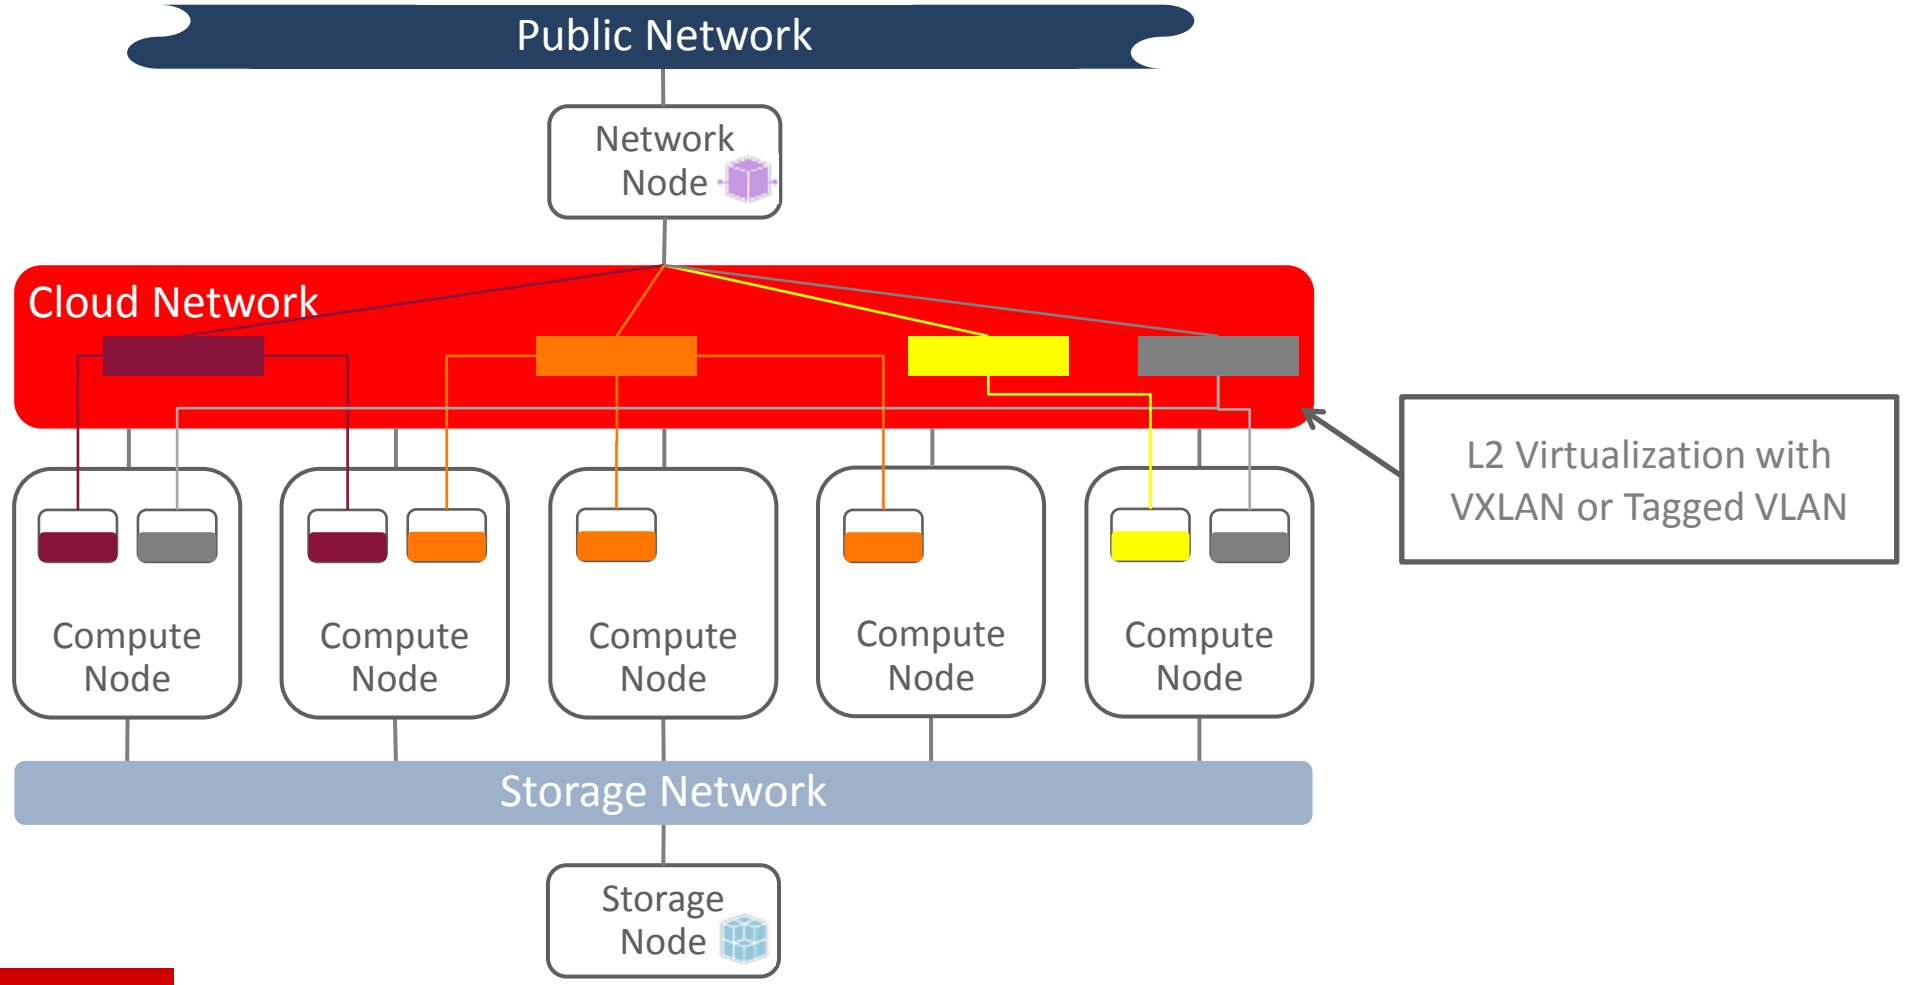

Converged Infrastructure (CI) with OpenStack

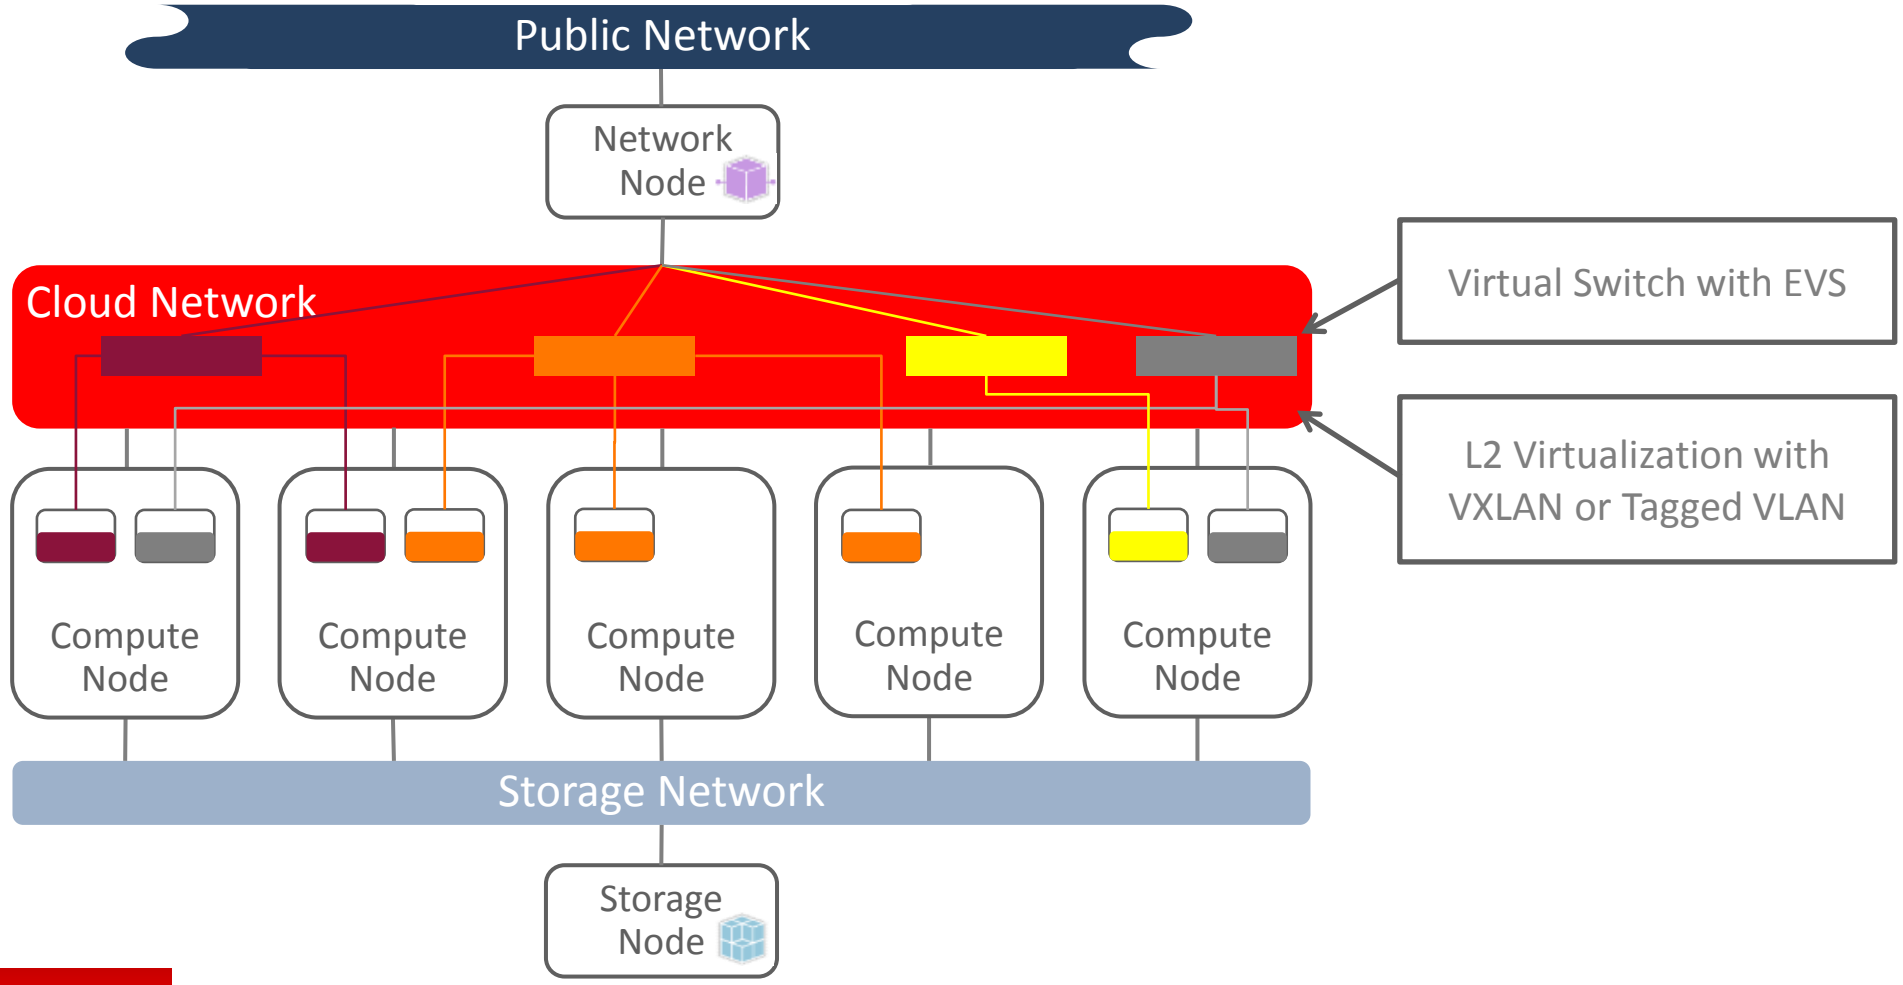

Converged Infrastructure (CI) with OpenStack

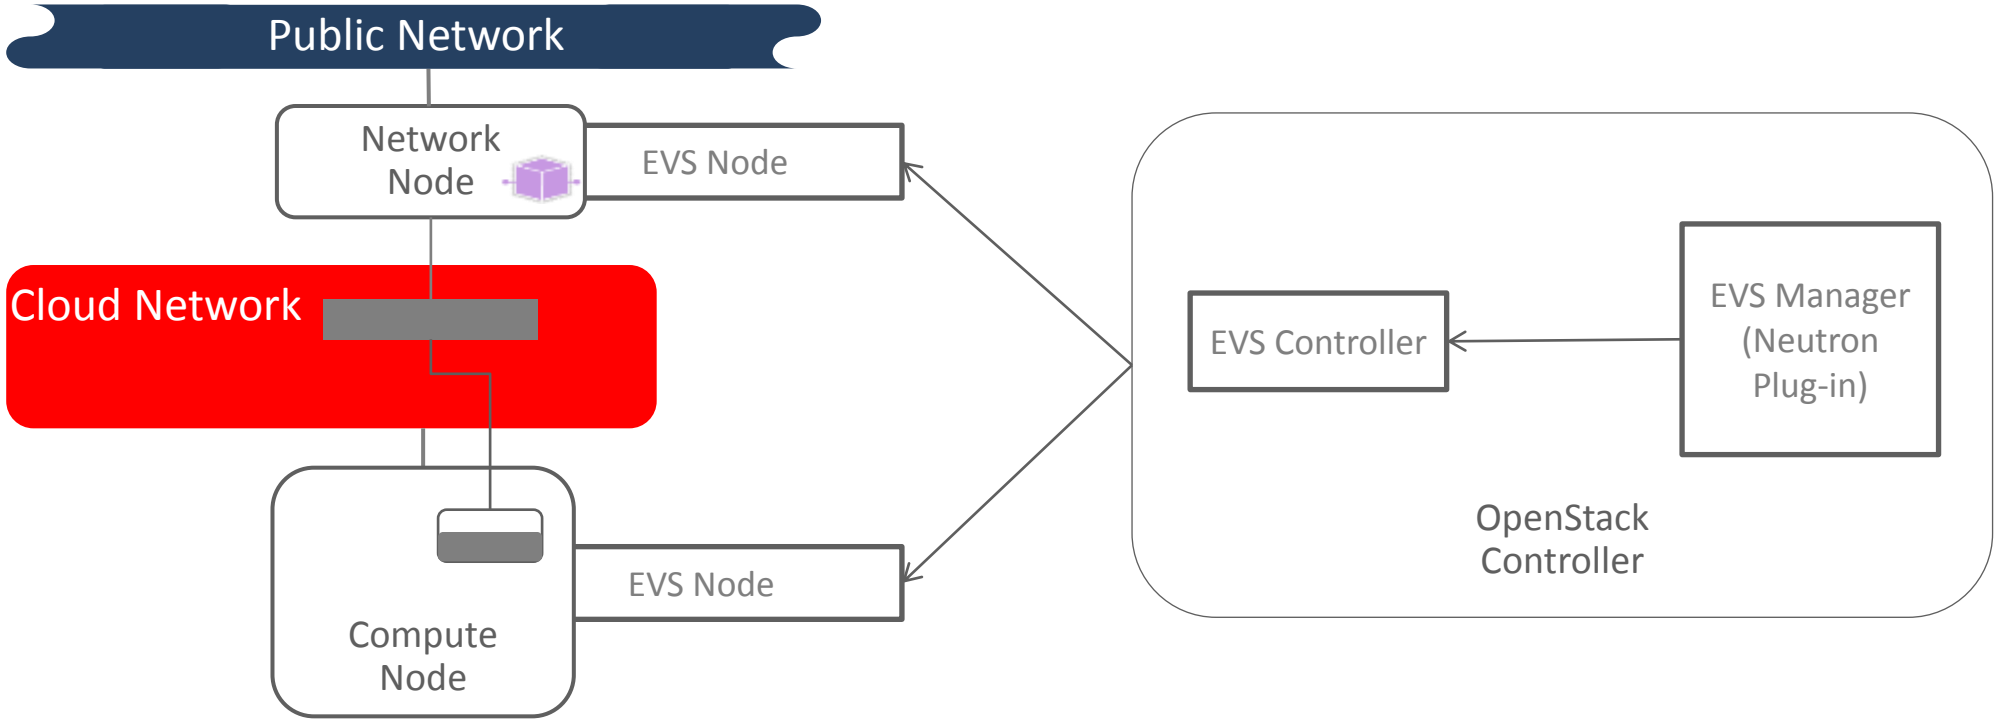

Converged Infrastructure (CI) with OpenStack

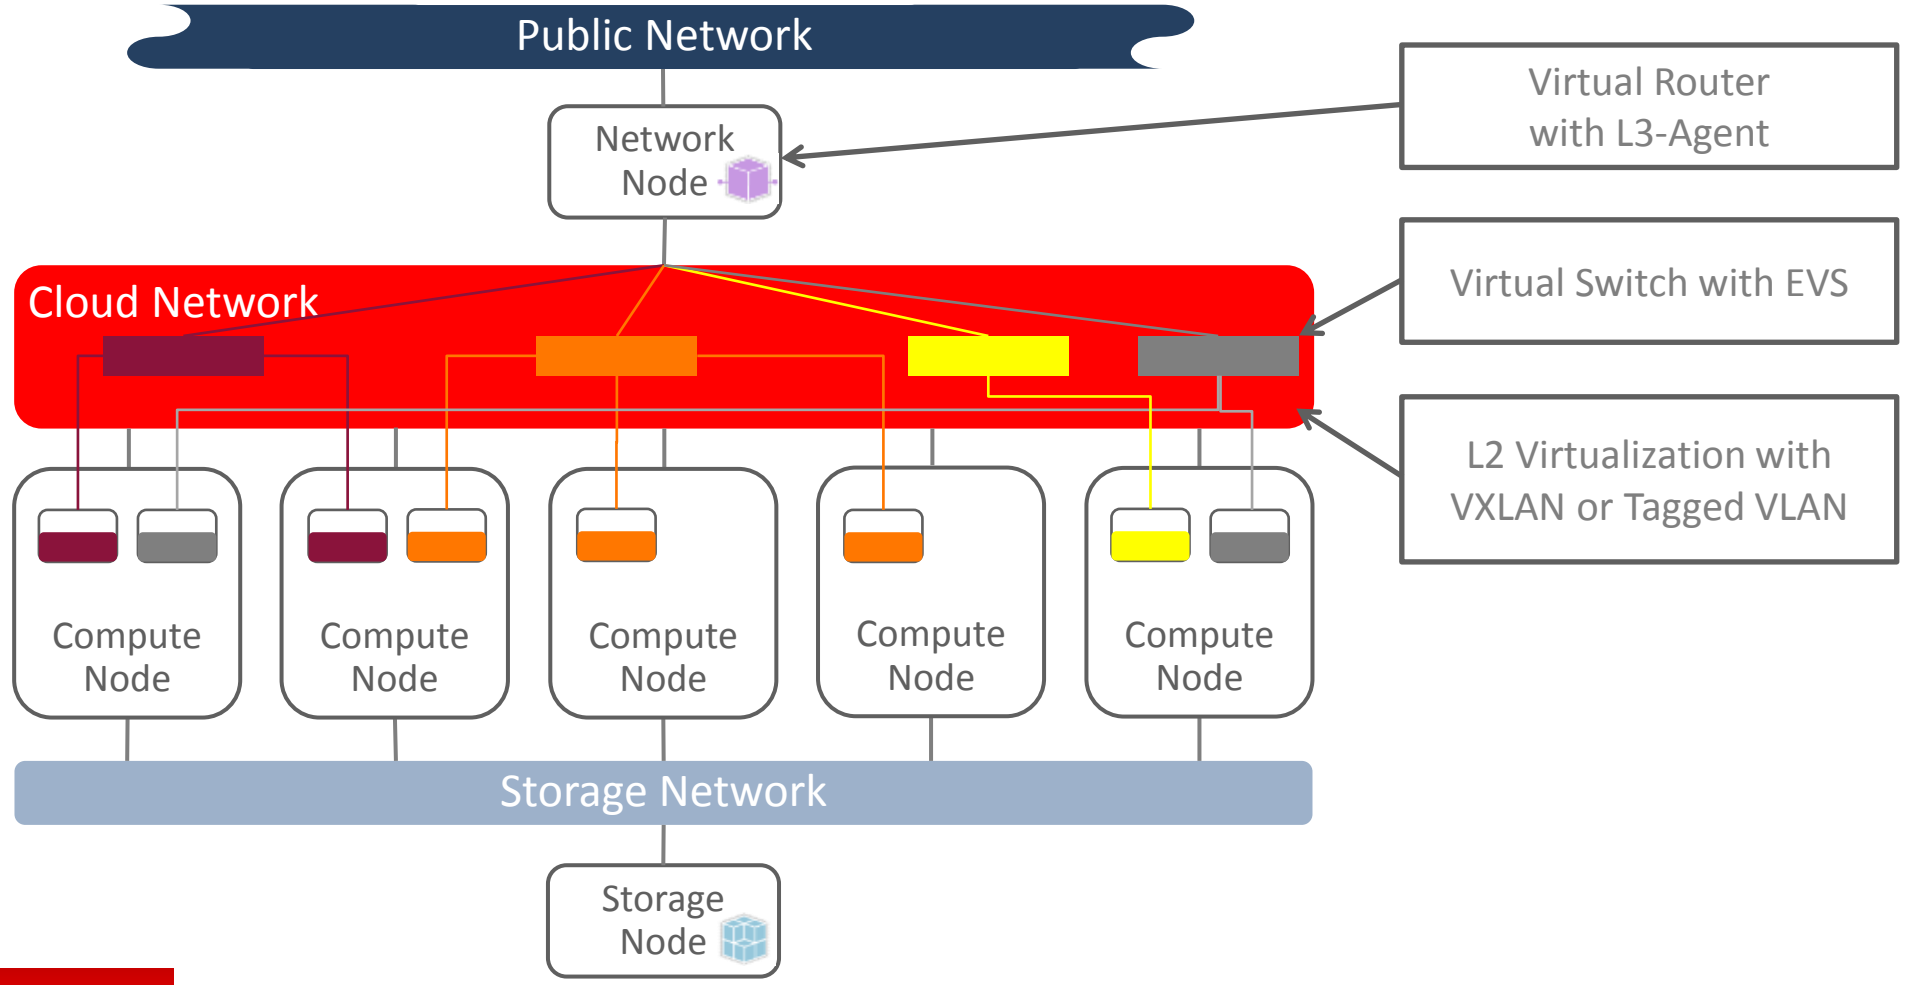

Converged Infrastructure (CI) with OpenStack

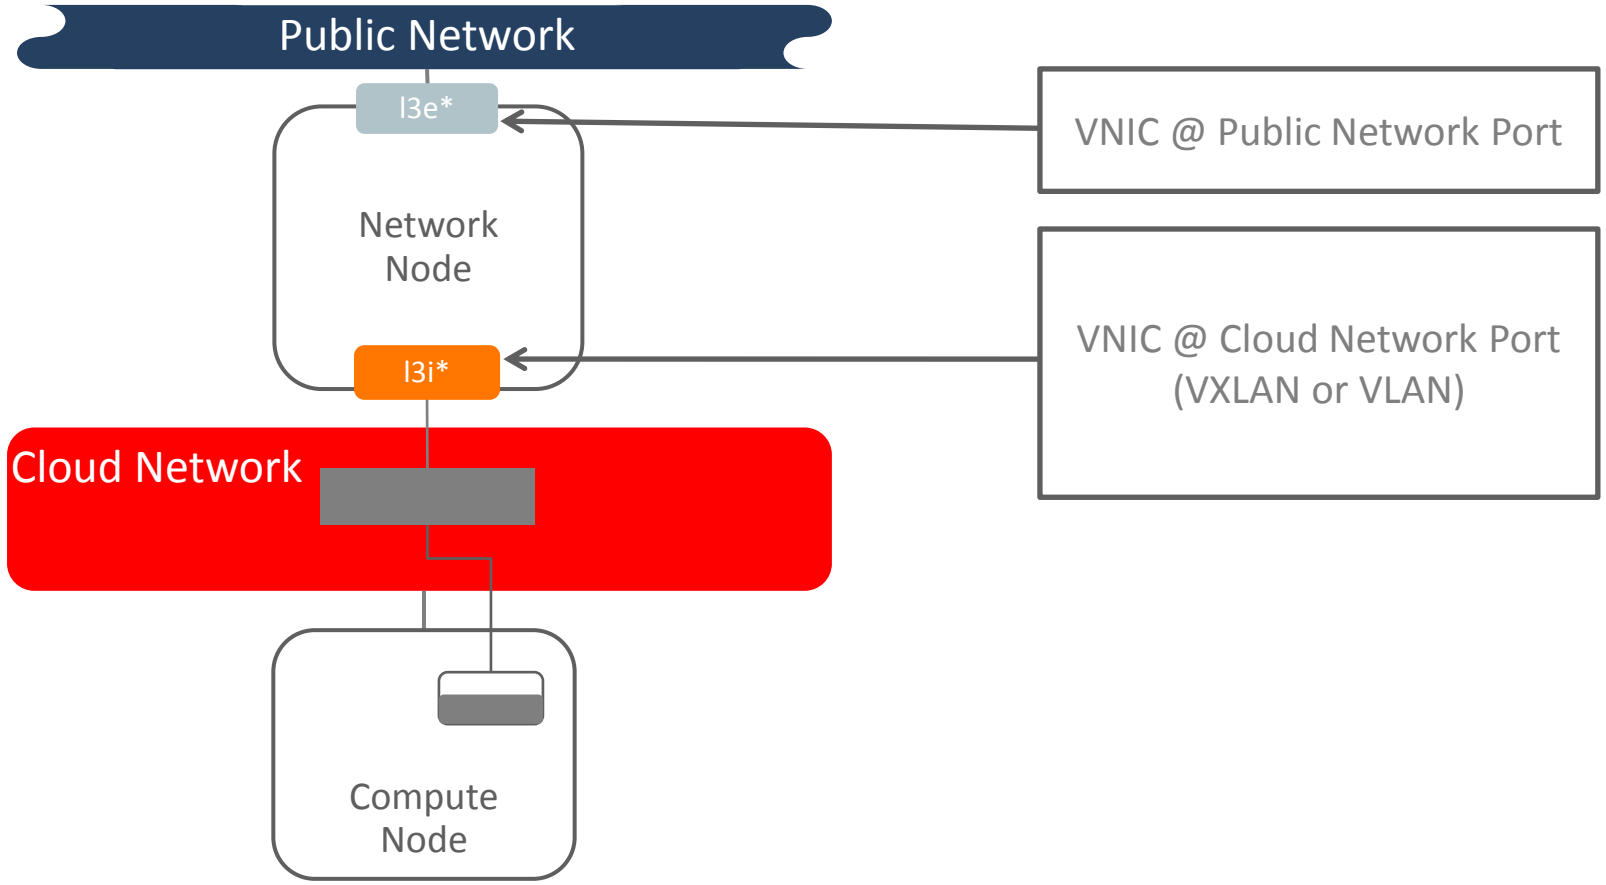

Cloud with Shared Storage and Controller

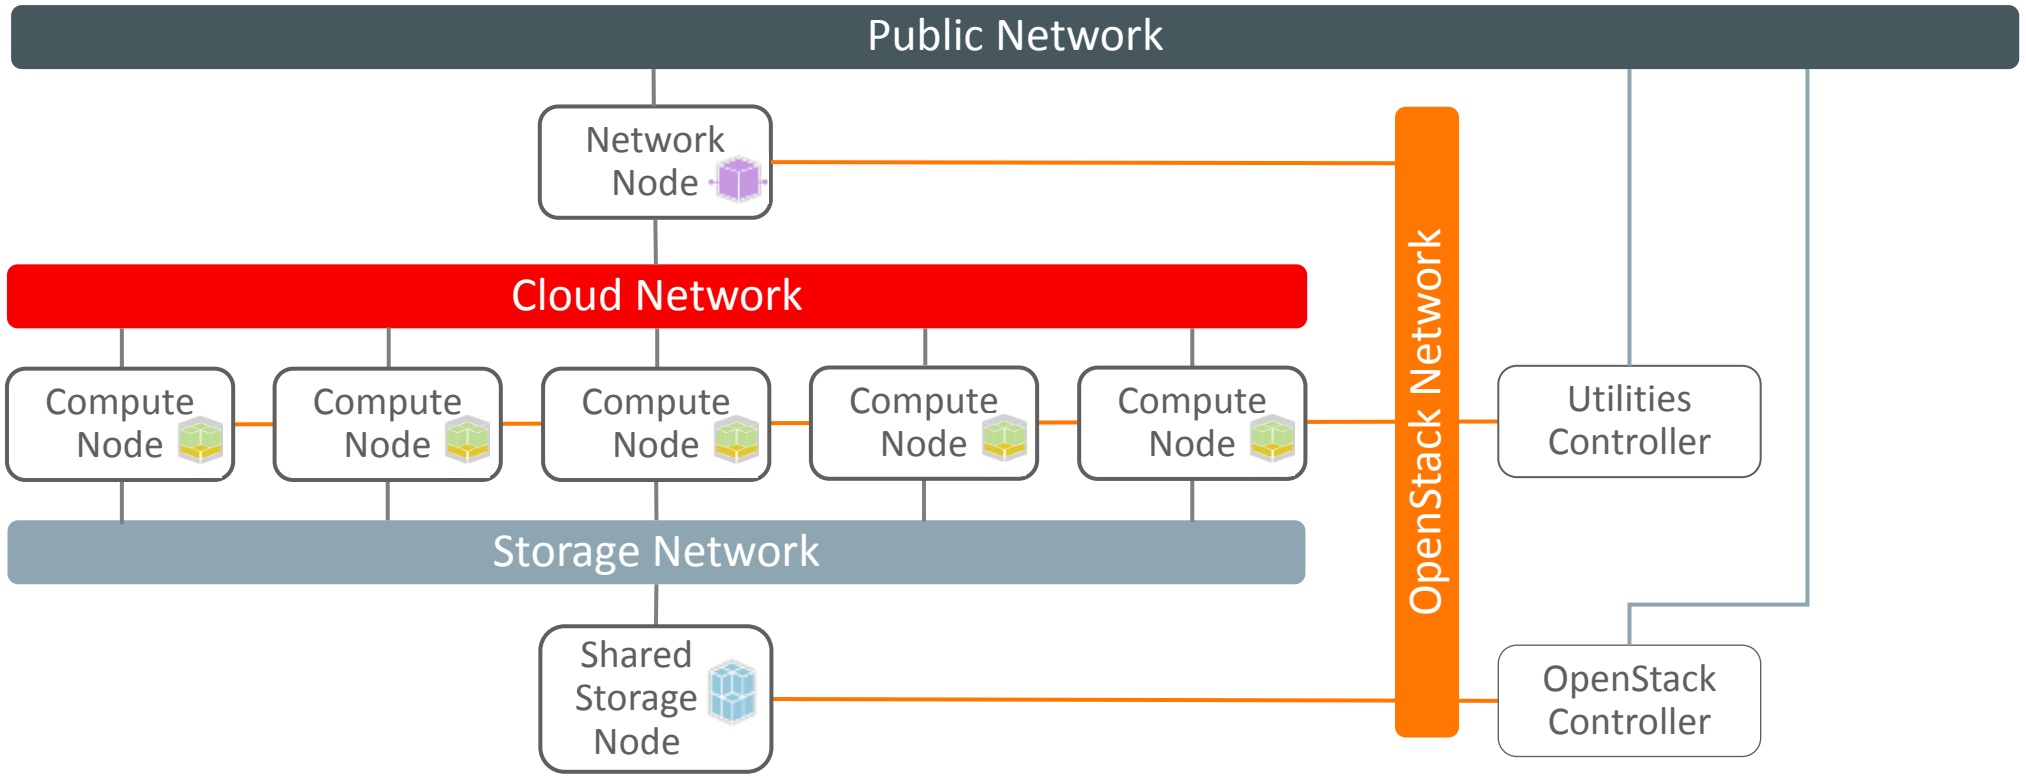

Neutron CLI - List Network and Subnet

```
# neutron net-list
+--------------------------------------+-----------------+-------------------------------------------------------+
| id                                   | name            | subnets                                               |
+--------------------------------------+-----------------+-------------------------------------------------------+
| 1a4fbee4-bf32-4d5c-8983-f3f94ffa7f43 | net1            |9611e2c2-33f3-4fb0-9bda-9efb4a2bbe4b 192.168.201.0/24|
|+-------------------------------------+-----------------+-------------------------------------------------------+

# neutron subnet-list
+------------------------------+------+----------------+-------------------------------------------------------+
| id                                   | name            | cidr             | allocation_pools                   |
+--------------------------------------+-----------------+------------------+------------------------------------+
| 9611e2c2-33f3-4fb0-9bda-9efb4a2bbe4b | subnet1         | 192.168.201.0/24 | {"start": "192.168.201.10", "end":
"192.168.201.250"}   |
+--------------------------------------+-----------------+------------------+------------------------------------+
```

EVS CLI - List Network, Ports and Properties

```
# evsadm
NAME          TENANT        STATUS VNIC         IP                HOST
net1 9844b7295b0b41a1a7d3d65bb2f9f5f2 busy -- subnet1             l3node
     aa936a6d7ce342d2b601b529483475b8 -- used l3iaa936a6d_7_0 192.168.201.1/24  l3node
     b427086b976c4a12a92bdba3bb9258f0 -- used dhb427086b_97_0 192.168.201.10/24 l3node

# evsadm show-prop
PROPERTY            PERM VALUE                          DEFAULT
controller          rw   ssh://evsuser@evscontroller    --

# evsadm show-controlprop
PROPERTY          PERM VALUE       DEFAULT   FLAT VLAN_RANGE  VXLAN_RANGE HOST
l2-type           rw   vxlan       vlan      --   --          --          --
uplink-port       rw   net1        --        no   --          201-300     --
uplink-port       rw   net2        --        yes  --          --          --
uri-template      rw   ssh://      ssh://    --   --          --          --
uuid              r-   f1caafa6-7c6b-11e5-8960-89bfff1027aa -- -- -- --   --
vlan-range        rw   --          --        --   --          --          --
vlan-range-avail  r-   --          --        --   --          --          --
vxlan-addr        rw   0.0.0.0     0.0.0.0   --   --          --          --
vxlan-ipvers      rw   v4          v4        --   --          --          --
vxlan-mgroup      rw   0.0.0.0     0.0.0.0   --   --          --          --
vxlan-range       rw   201-300     --        --   --          --          --
vxlan-range-avail r-   202-300     --        --   --          --          --
```

Solaris Data Link Status - Links and VNICs

```
# dladm show-link
LINK                CLASS     MTU    STATE    OVER
net1                phys      1500   up       --
net2                phys      1500   up       --
net0                phys      1500   up       --
l3edb9e210a_7_0     vnic      1500   up       net2
evs-vxlan200        vxlan     1440   up       --
l3iaa936a6d_7_0     vnic      1440   up       evs-vxlan201
dhb427086b_97_0     vnic      1440   up       evs-vxlan201
# dladm show-phys
LINK              MEDIA                STATE      SPEED  DUPLEX    DEVICE
net1              Ethernet             up         1000   full      e1000g1
net2              Ethernet             up         1000   full      e1000g2
net0              Ethernet             up         1000   full      e1000g0
# dladm show-vnic
LINK            OVER           SPEED  MACADDRESS        MACADDRTYPE IDS
l3edb9e210a_7_0 net2           1000   fa:16:3e:c:39:20  fixed       VID:0
l3iaa936a6d_7_0 evs-vxlan201   1000   fa:16:3e:30:cc:2b fixed       VID:0
dhb427086b_97_0 evs-vxlan201   1000   fa:16:3e:f2:a9:32 fixed       VID:0
# dladm show-vxlan
LINK                ADDR                    VNI   MGROUP
evs-vxlan201        192.168.106.30          201   224.0.0.1
```

Solaris IP Status - Links and VNICs

```
# ipadm
NAME                 CLASS/TYPE STATE    UNDER                ADDR
dhb427086b_97_0      ip         ok         --                --
   dhb427086b_97_0/v4 static ok           --                192.168.201.10/24
l3edb9e210a_7_0   ip         ok           --                --
   l3edb9e210a_7_0/v4 static ok           --                192.168.175.240/24
l3iaa936a6d_7_0   ip         ok           --                --
   l3iaa936a6d_7_0/v4 static ok           --                192.168.201.1/24
lo0               loopback   ok           --                --
   lo0/v4         static     ok           --                127.0.0.1/8
   lo0/v6         static     ok           --                ::1/128
net0              ip         ok           --                --
   net0/onv4      static     ok           --                192.168.101.30/24
net1              ip         ok           --                --
   net1/cnv4      static     ok           --                192.168.106.30/24
net2              ip         down         --                --
```

Building the Cloud

Converged Infrastructure

Source: https://en.wikipedia.org/wiki/Converged_infrastructure

- Grouping multiple information technology components
- One single, optimized computing package
- Components may include
  - Networking Equipment
  - Data-Storage Devices
  - Servers
  - Software for IT infrastructure Management
  - Software for Automation
  - Software for Orchestration

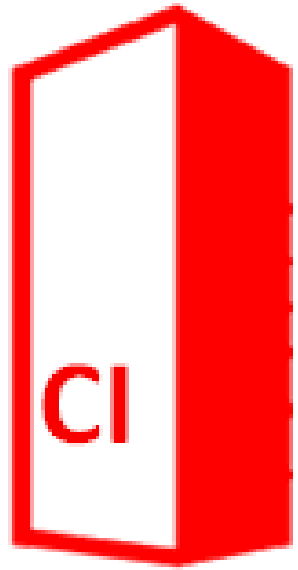

Converged Cloud Infrastructure with Oracle

Design Choices

Compute

Reserved for Expansion Storage or Compute

Network

Storage

Cloud Mgnt Package

Empty

Reserved for Expansion

Empty

Server
- SPARC Server
- x86 Server

Network
- 10 GbE
- Infiniband

Storage
- DAS
- NAS
- SAN

Operating System
- Oracle Solaris
- Oracle Linux

ORACLE OPENSTACK

ORACLE SOLARIS

ORACLE LINUX

CI

ORACLE SPARC SERVERS

ORACLE x86 SERVERS

ORACLE ZFS STORAGE APPLIANCE

ORACLE NETWORKING

Converged Cloud Infrastructure with local FS1 Storage

- Power Compute Node (x5-4)
- Dual Node Storage Controller
- Dual Node OpenStack Cloud Controller

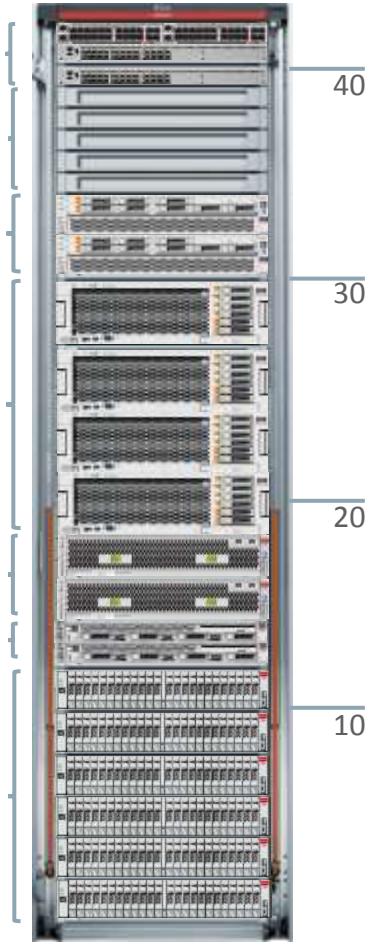

Getting Started

OpenStack on Oracle Solaris resources

- OpenStack on Oracle Solaris Discussion

  mailto:solaris_openstack_interest@openstack.java.net

- Oracle Solaris on Oracle Technology Network

  http://www.oracle.com/technetwork/server-storage/solaris11/technologies/openstack-2135773.html

- Source Code

  https://java.net/projects/solaris-userland/sources/gate/show/components/openstack

Oracle Solaris OpenStack Live Demonstration

Q & A

Detlef.Drewanz@oracle.com

Integrated Cloud

Applications & Platform Services


ORACLE®