

Cisco Service Provider Technology Webinar

Cisco Virtual Infrastructure Manager (CVIM)

Overview & Introduction

Tugrul Acikgoz – Technical Solution Architect



Agenda

- 5G and Openstack Chat
- CVIM Introduction
- CVIM POD Types
- CVIM Deployment
- Compute & Network Technology Overview
- CVIM Software Update
- 5G Transition
- Future

5G & Openstack Chat

What is 5G anyway?

ETSI →

The minimum requirements:

for peak data rate: Downlink: 20 Gbit/s, Uplink: 10 Gbit/s

for peak spectral efficiencies: Downlink: 30 bit/s/Hz, Uplink: 15 bit/s/Hz

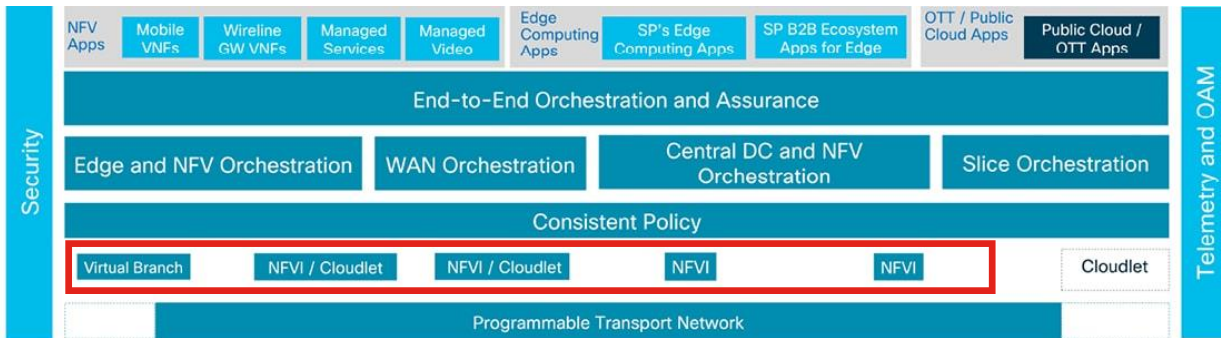
user plane latency (single user, small packets): 4 ms for eMBB, 1 ms for URLLC

control plane latency (idle => active): 10-20ms

Other requirements:

maximum aggregated system bandwidth: at least 100 MHz, up to 1GHz in higher frequency bands (above 6GHz)

mobility: up to 500km/h in rural eMBB



2G -> 3G -> 4G -> 5G

● Is it all about higher speed?

- ➔ Low Latency
- ➔ High Availability
- ➔ Automation
- ➔ Scale



OpenStack: Industry Choice today for Telco Workloads

But It isn't Simple

OpenStack can be Complex to Operate:



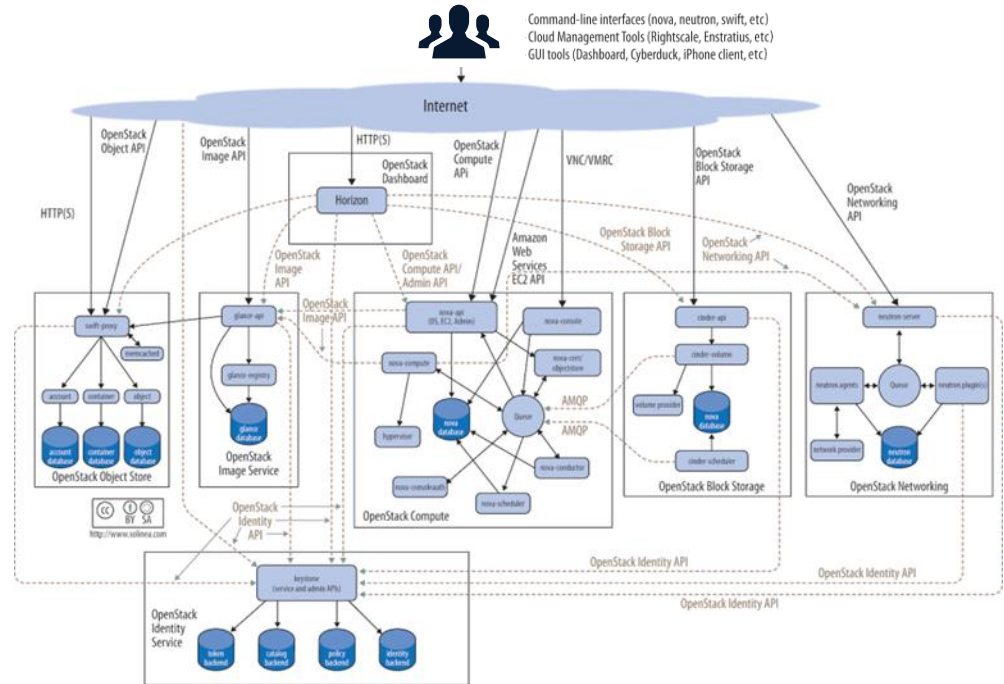
Complex interactions between services, databases, messaging queues etc.




Health and performance of a cloud is difficult to quantify, verify and monitor.



Updates/upgrades require extensive human effort and are prone to issues.



Telco cloud characteristics

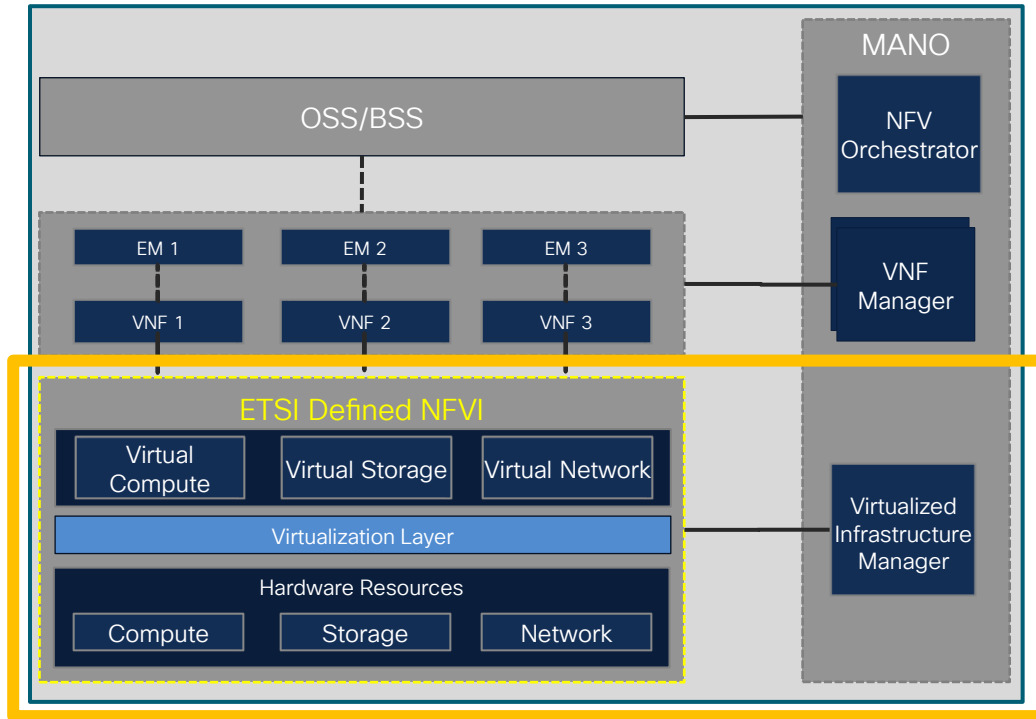


Distributed, Consistent, and Automated


- Abstracted infrastructure layer
- Automated lifecycle management
- Centralized, end-to-end monitoring
- End-to-end service orchestration
- Network functions deployed anywhere

CVIM Introduction

CVIM in ETSI NFV Reference Architecture and NFVI




- ETSI NFVI – Physical + Virtual Infrastructure resources (Compute, Storage and Network)
- NFVI Solutions – ETSI NFVI + VIM (*SDN Controller) + Management



Telco cloud composition and ownership

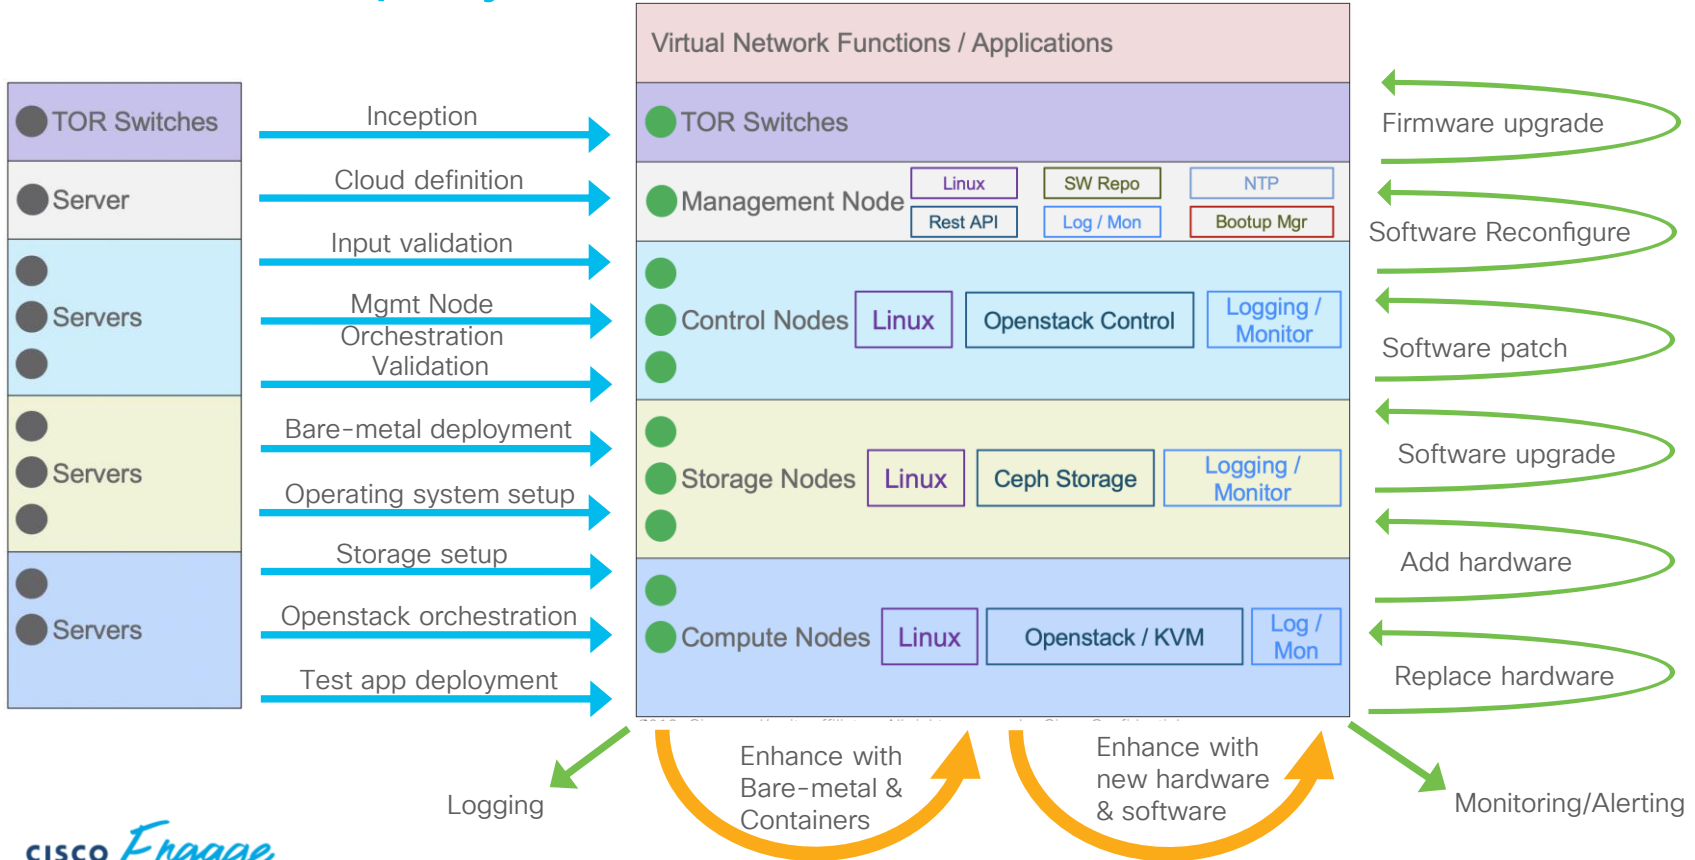
← Single Vendor ... Multi-Vendor →
← Pre-packaged ... Custom →











↔ Pre-integrated
↔ Custom integration

← Infrastructure ... Services →
← Predictable Ops ... Agile Dev →

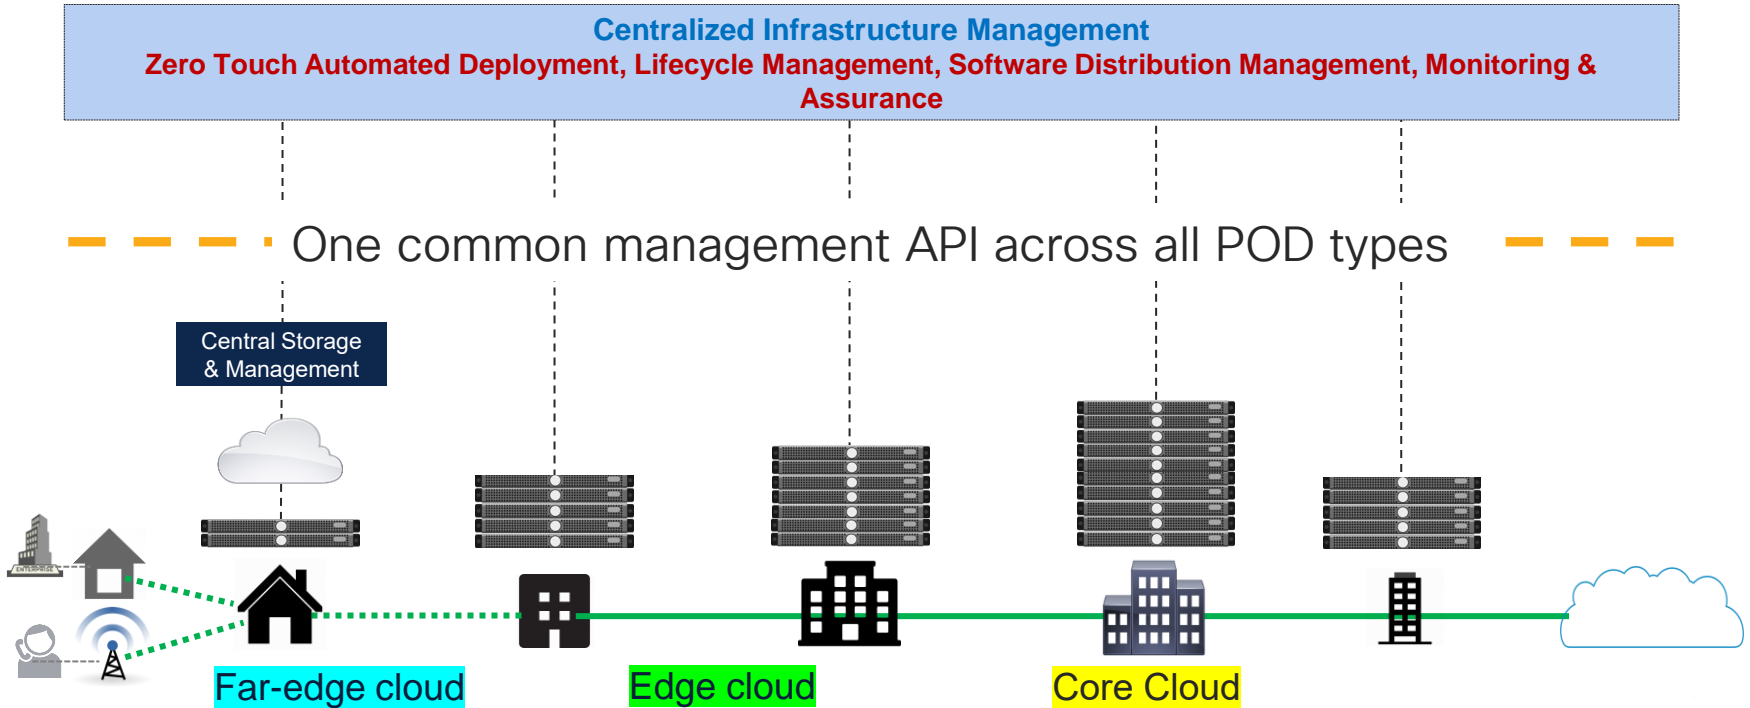
CVIM : Deploy → Operate → Evolve



CVIM : Deploy → Operate → Evolve



Network-wide Deployment with Cisco VIM



CVIM at-a-glance

Telco cloud platform
that can be used to ...




Grow Revenue



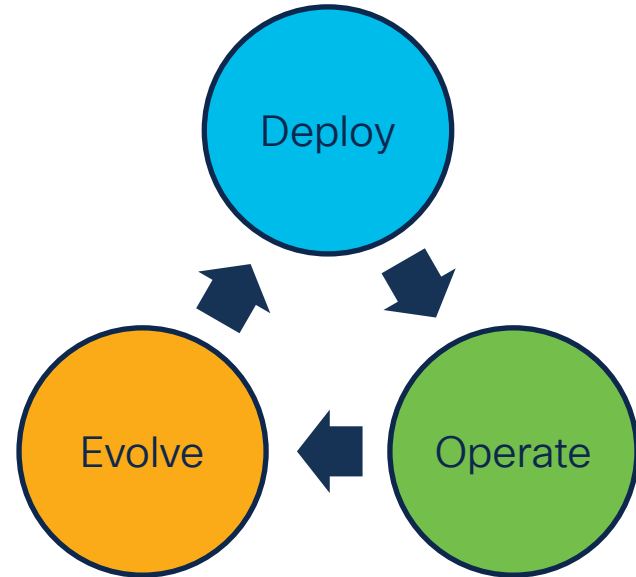
Reduce Costs



Mitigate Risk



... through automation
and focus on lifecycle



Our Contributions to Open Stack / Opensource

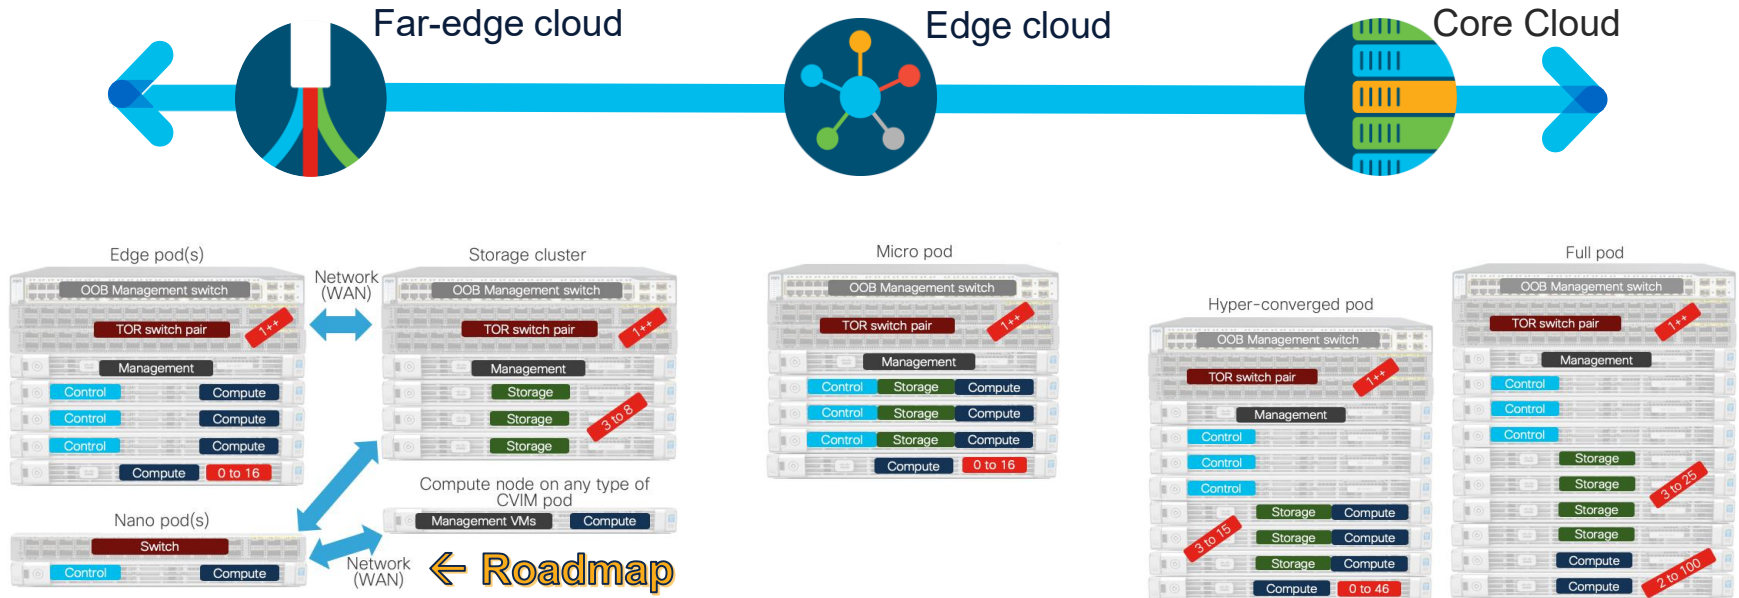
- Cloud Pulse
- Cloud Sanity
- Kloudbuster
- NFVbench
- VPP and networking-vpp
- MPLS EVPN
- Network Service Mesh
- Lots of patches

Result is Open and Usable Solution

- Use Open Source where possible
- Open API access
- Many different VNF's from multiple vendors have been deployed

CVIM POD Types

CVIM Pod Types



← Small optimized form factor with specialized hardware ... Large form factor with standard hardware →

CVIM For Core Cloud

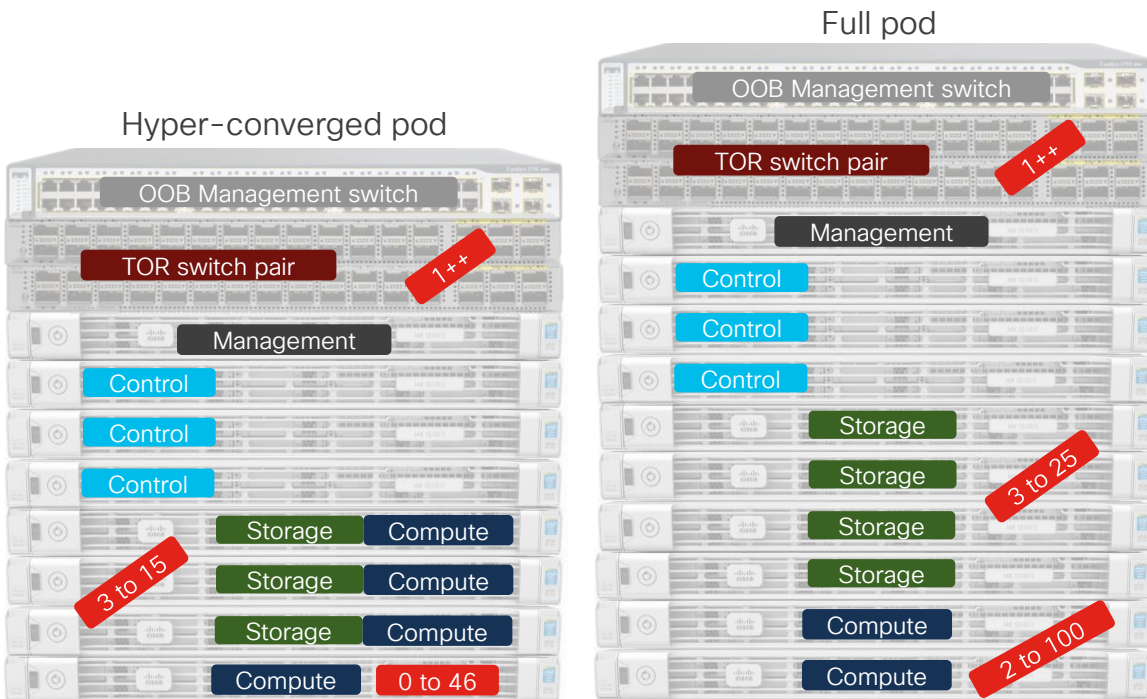


- **Full pod**


- Largest form factor
- Typically used for VNFs like vEPC, and other general IT applications

- **Hyper-converged pod**

- Second largest form factor
- Storage co-located with compute
- Suitable for use cases with moderate storage needs


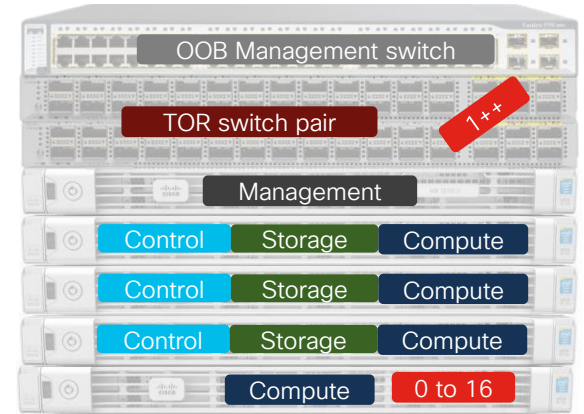


CVIM for edge cloud



- **Micro pod**

- Smallest form factor for independent cloud
- Control and Storage co-located with compute
- Typically used for small scale deployments, vRAN vCU, and MEC (Multi-access Edge Computing) use cases
- Also useful for small managed network services or SD-WAN type use cases
- Configurable minimal overhead for Linux, Control and Storage, all remaining resources available for workloads


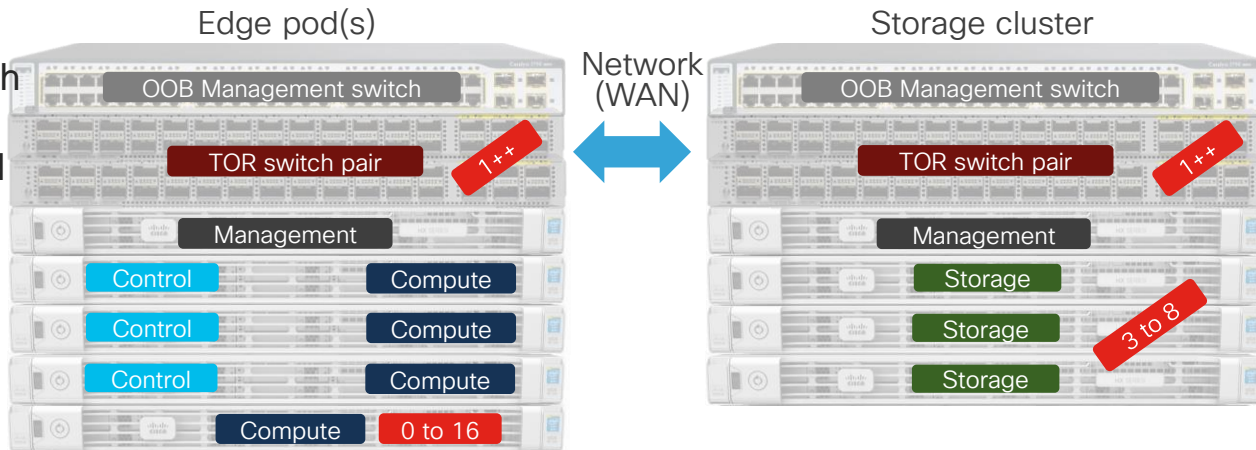


CVIM for far-edge cloud



• Edge pod

- Smallest form factor, with no volume storage, and image storage in shared storage cluster
- Control co-located with compute
- Configurable minimal overhead for Linux & Control
- Supports specialized FPGA NICs and Real-time kernel for time sensitive VNFs like vRAN vDU
- Typically used for remote deployments at small central offices that have limited power & space



CVIM for far-edge cloud


Roadmap



Nano pod(s)



Compute node on any type of CVIM pod

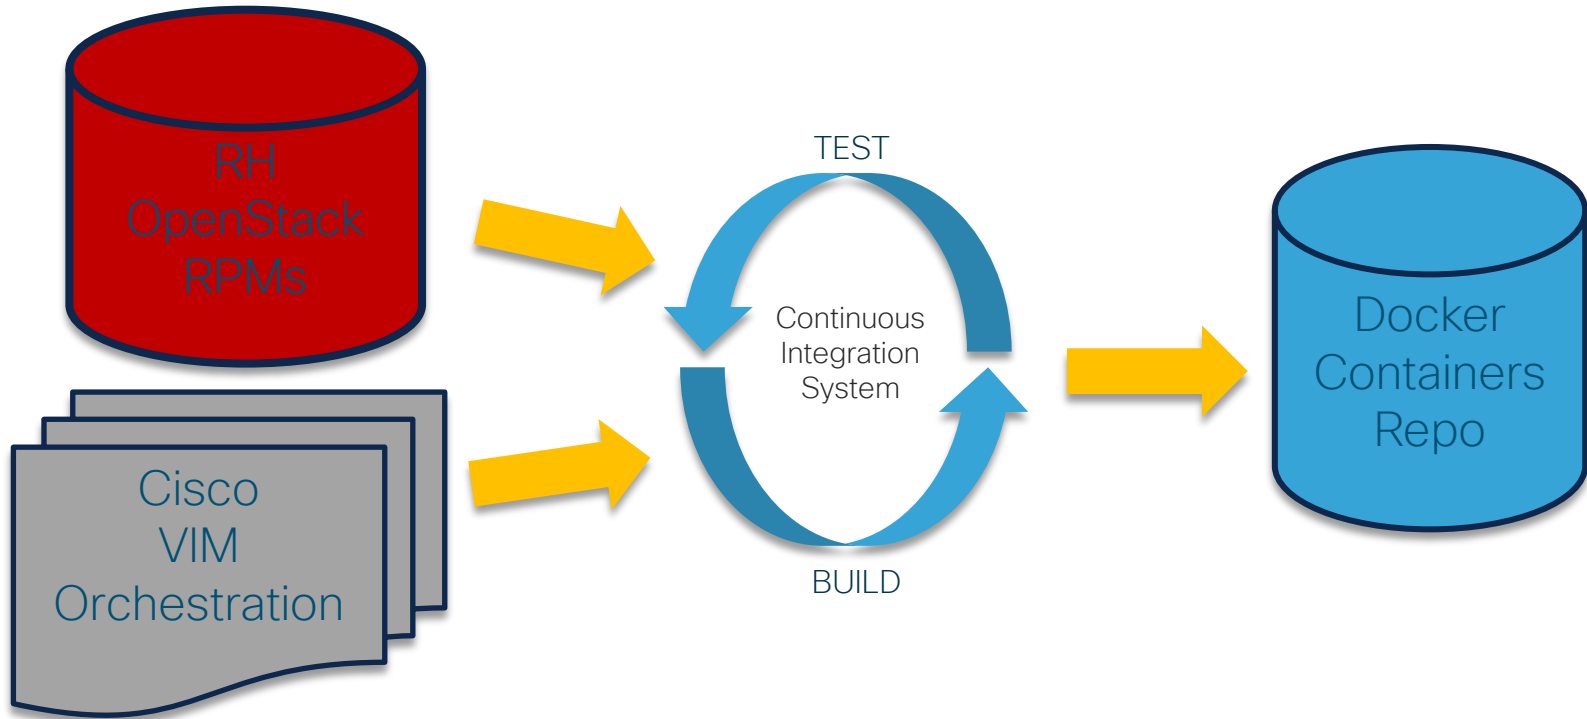


- **Nano pod**

- Single server form factor, with no volume storage, and image storage in shared storage cluster
- Management node functionality in in a VM running on some other cloud or as a nano node running mgmt. role.
- Similar to Edge pod, control co-located with compute
- Configurable minimal overhead for Linux & Control
- Supports specialized FPGA NICs and Real-time kernel for time sensitive VNFs like vRAN vDU
- Typically used for sparsely populated remote deployments and/or cloud-on-wheels for disaster relief networks

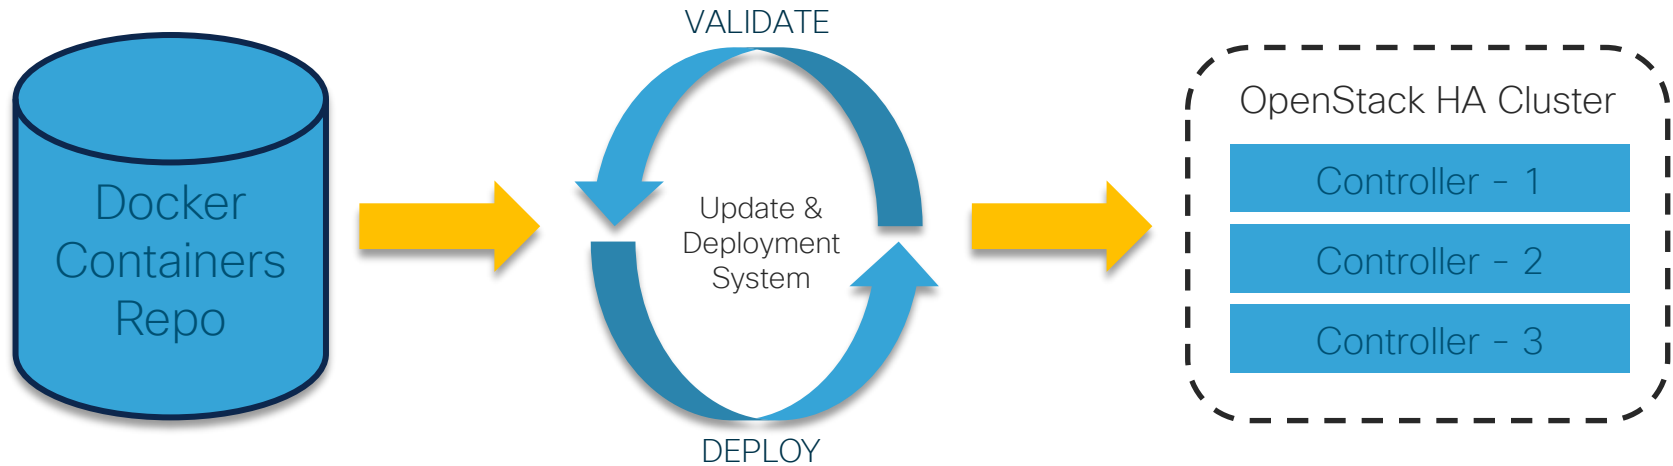
CVIM Deployment

Building CVIM



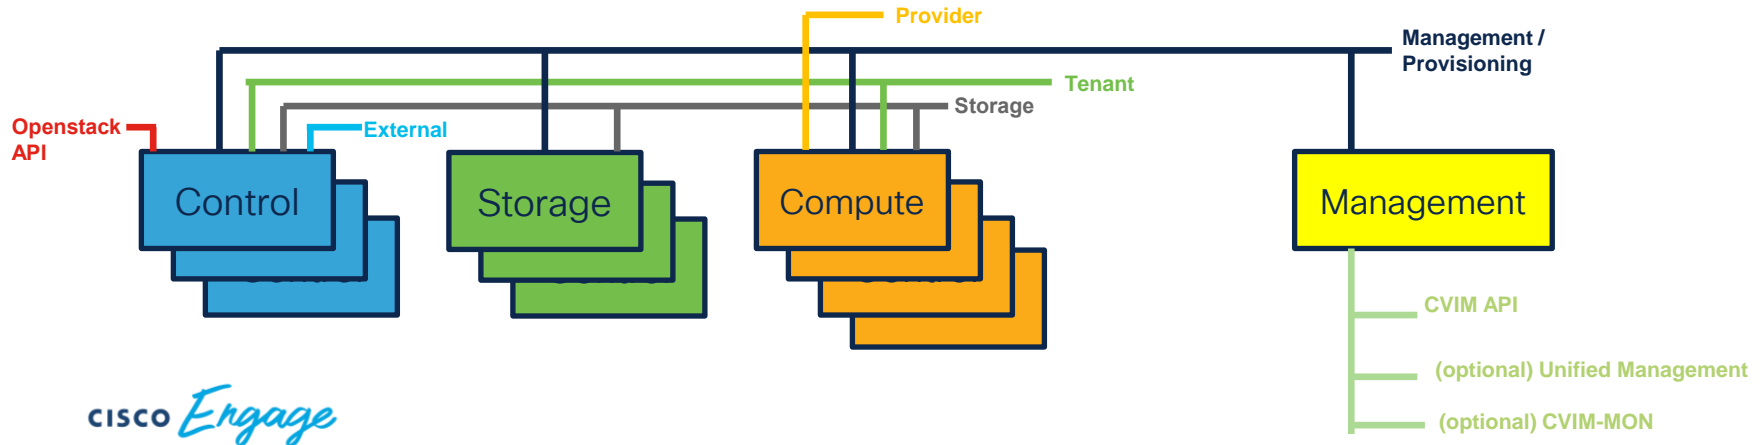
Deploying CVIM

Container-based Atomic Deployment and Update




CVIM Network Segmentation

- API – OpenStack API end point for managing/using the NFVI
- External – Link to world beyond the cloud via OpenStack virtual routers (L3 agent)
- Management/Provisioning network – PXE boot, management and OpenStack inter-service communication
- Provider – Link to existing infrastructure networks
- Tenant – Inter VM traffic via OpenStack tenant networks
- Storage – Ceph data replication traffic, access to Netapp, Cinder volume access
- CVIM API – Provides CVIM API, External access into CVIM mgmt. node, requires access to Openstack API



Design Options

- Hardware Choices



- VIM configuration (at install time or reconfigure options)

- | | | | |
|--|--|---|--|
| <input type="checkbox"/> Auto Backup | <input type="checkbox"/> Neutron base mac-address | <input type="checkbox"/> Cobbler | <input type="checkbox"/> CVIMMON |
| <input type="checkbox"/> Enable Esc Priv | <input type="checkbox"/> Enable Read-only OpenStack Role | <input type="checkbox"/> Enable TTY LOGGING | <input type="checkbox"/> ES_REMOTE_BACKUP |
| <input type="checkbox"/> External LB VIP FQDN | <input type="checkbox"/> External LB VIP TLS | <input type="checkbox"/> Heat | <input type="checkbox"/> Horizon Aliases |
| <input type="checkbox"/> Install Mode | <input type="checkbox"/> Inventory Discovery | <input type="checkbox"/> Ironic | <input type="checkbox"/> LBAAS |
| <input type="checkbox"/> LDAP | <input type="checkbox"/> MGMTNODE EXTAPI REACH | <input type="checkbox"/> NETAPP_SUPPORT | <input type="checkbox"/> Networking |
| <input type="checkbox"/> NFVbench | <input type="checkbox"/> NFVI Monitoring | <input type="checkbox"/> Pod Name | <input type="checkbox"/> PROVIDER VLAN RANGES |
| <input type="checkbox"/> Registry setup settings | <input type="checkbox"/> SNMP | <input type="checkbox"/> SRIOV CARD TYPE | <input type="checkbox"/> SSH Banner |
| <input type="checkbox"/> Swiftstack | <input type="checkbox"/> Syslog Export Settings | <input type="checkbox"/> TENANT VLAN RANGES | <input type="checkbox"/> TORSwitch Information |
| <input type="checkbox"/> Vim Admins | <input type="checkbox"/> Vim LDAP Admins | <input type="checkbox"/> VMTP | |

- Openstack Flavors -> CPU pinning, anti-affinity, huge pages, Thread isolation, NUMA nodes...



Compute & Network Technology Overview

CVIM and VNF Compute Configurations

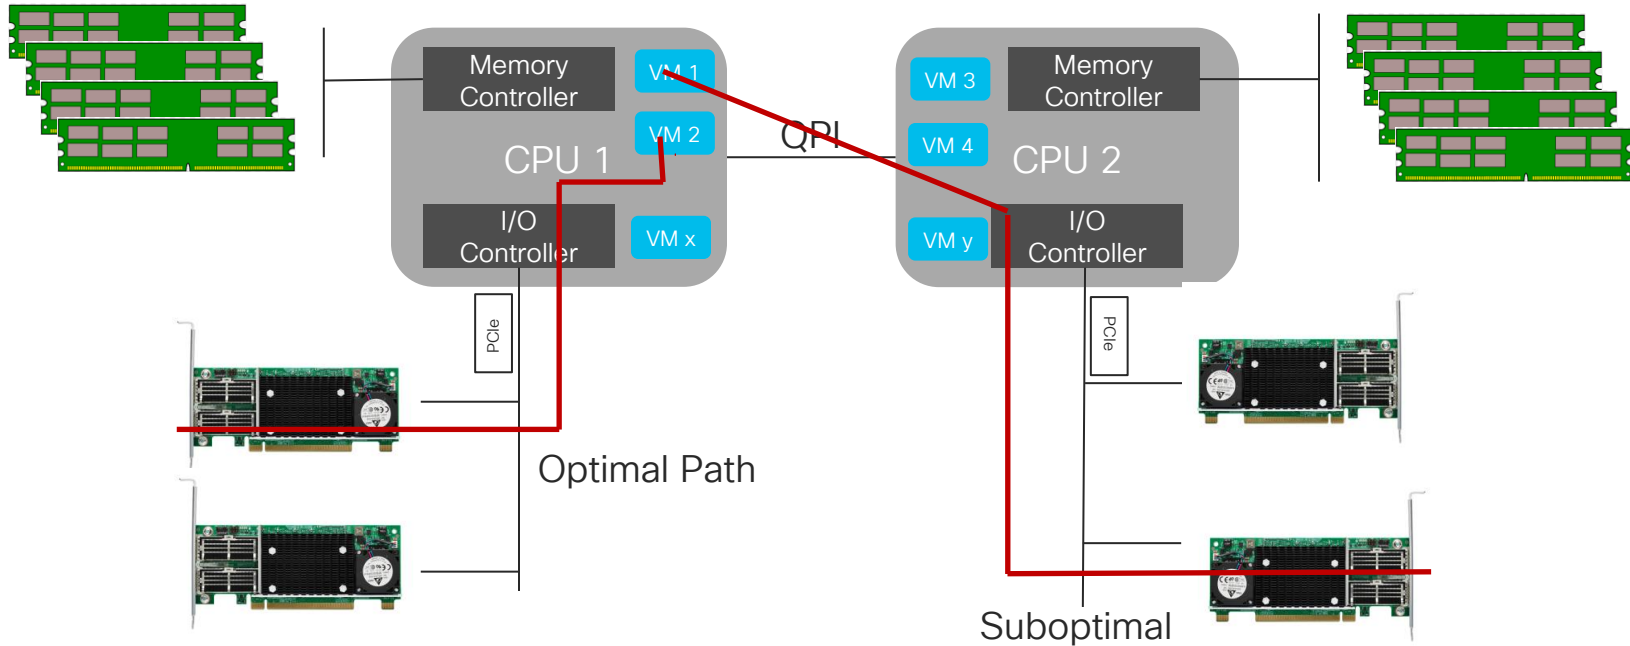
CPU Pinning and Thread Isolation are defined in Flavors.

Update Flavor Metadata ✕

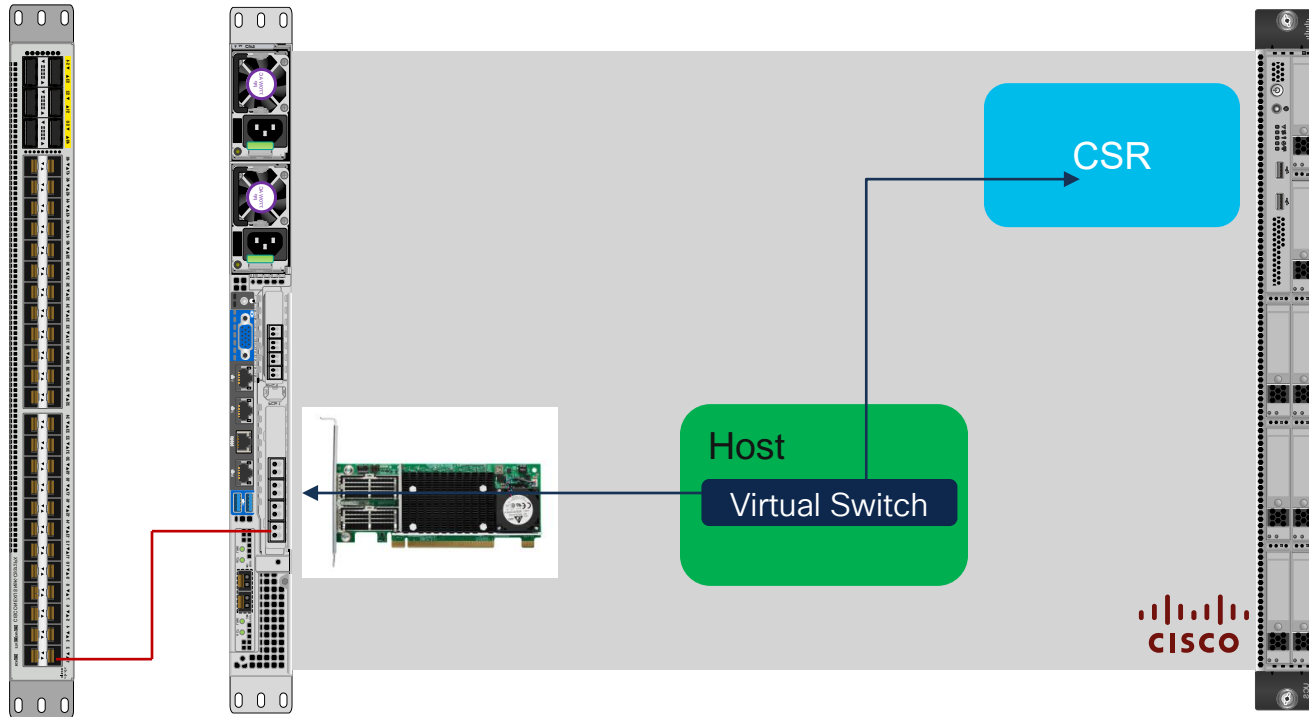
You can specify resource metadata by moving items from the left column to the right column. In the left column there are metadata definitions from the Glance Metadata Catalog. Use the "Custom" option to add metadata with the key of your choice.

Available Metadata	Filter	Q	Existing Metadata	Filter	Q
Custom <input type="text"/>			hw:cpu_policy	dedicated	-
No available metadata			hw:cpu_thread_policy	isolate	-
			hw:numa_nodes	2	-


NUMA: Resource Locality



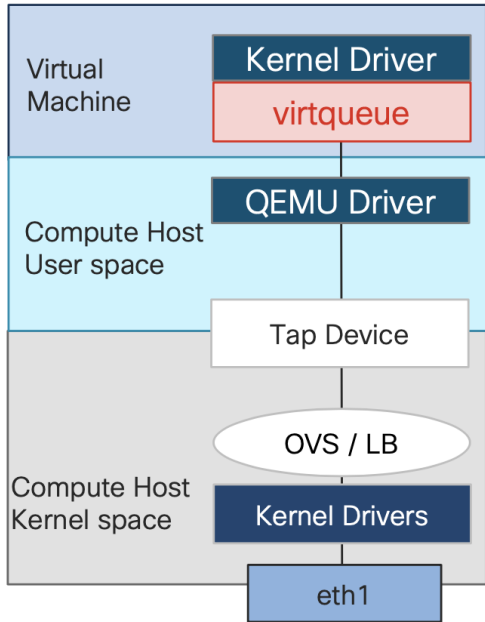
Virtual Switch



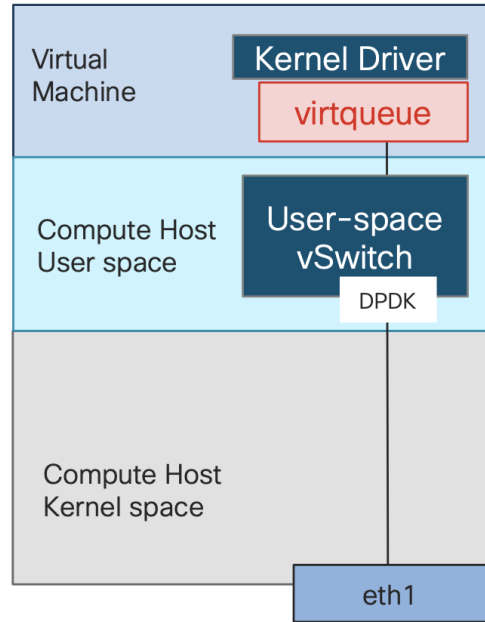
SR-IOV



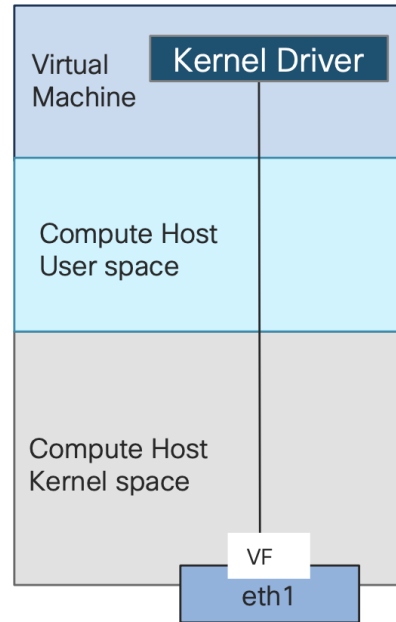
Networking Technologies



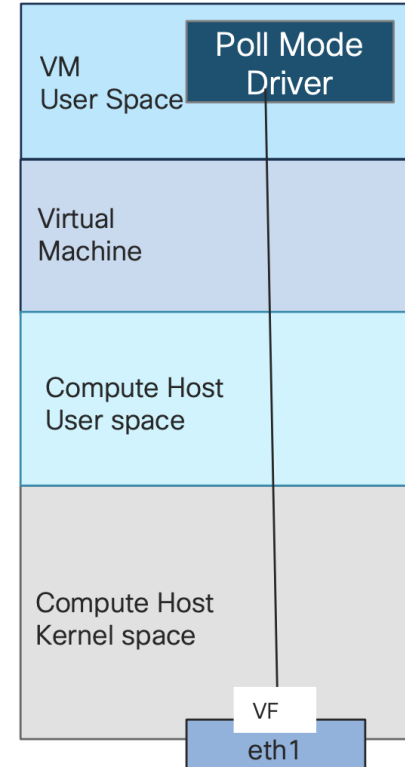
Kernel Space



User Space

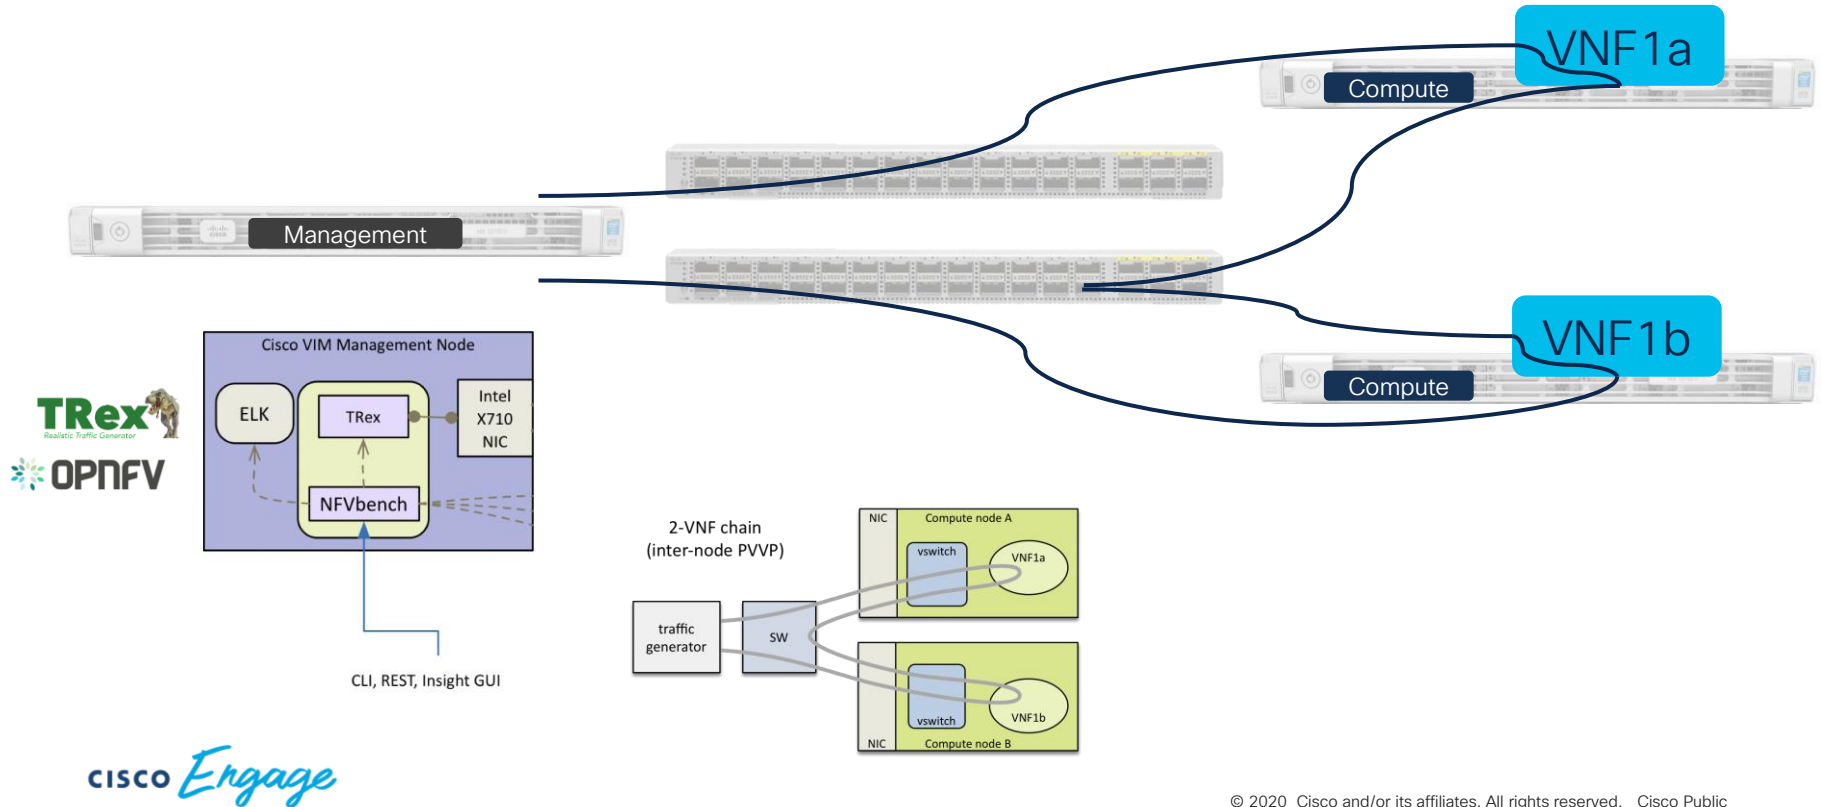


SR-IOV



SR-IOV

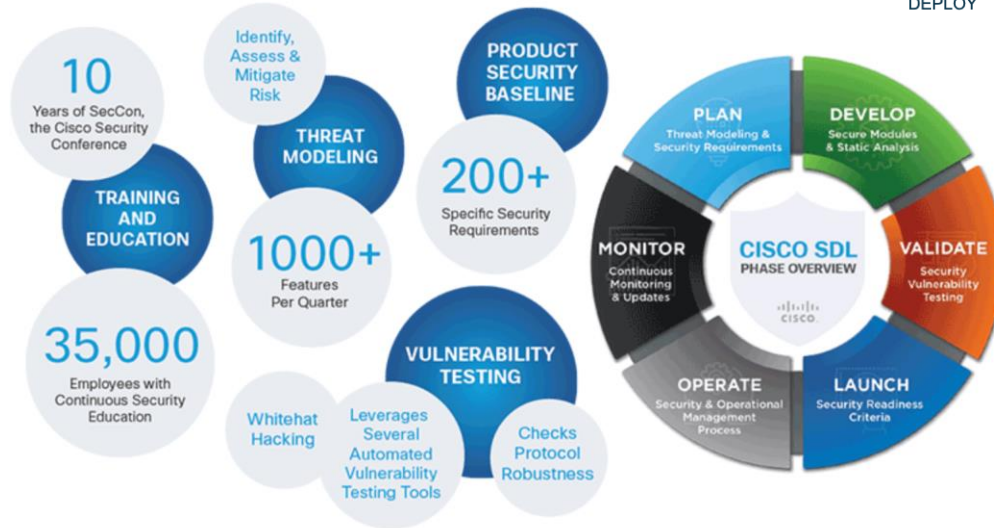
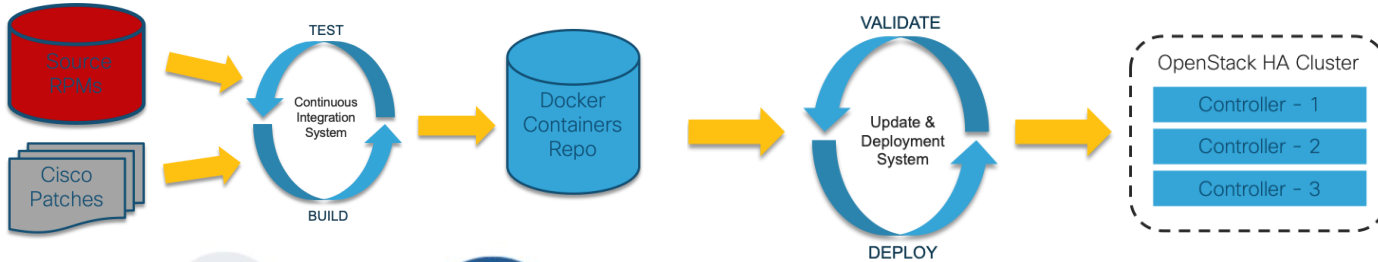
NFVBench: Performance Testing Made Easy



Performance Numbers

- Suggested Cisco live session-> BRKSDN-2411



Security



- SELinux everywhere
- No unused services running
- Host based firewall
- Network segmentation
- ...

HA architecture and more ...

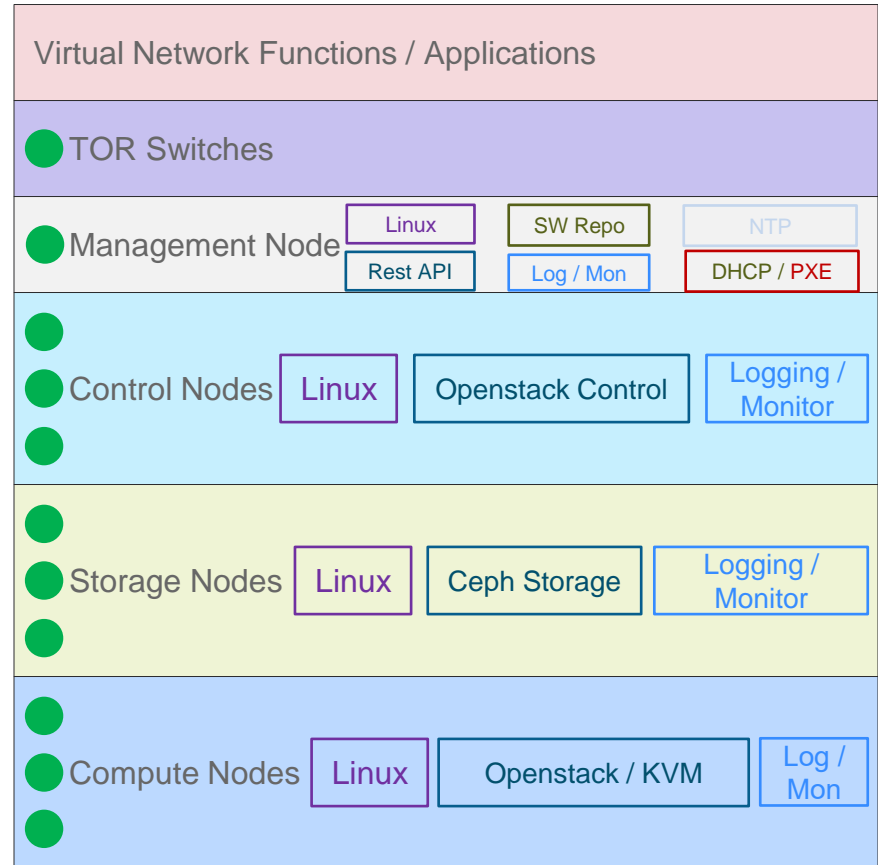
- Control node
- Storage cluster
- NIC port HA
- ToR HA
- Mgmt backup and restore
- VNF HA



CVIM Software Update

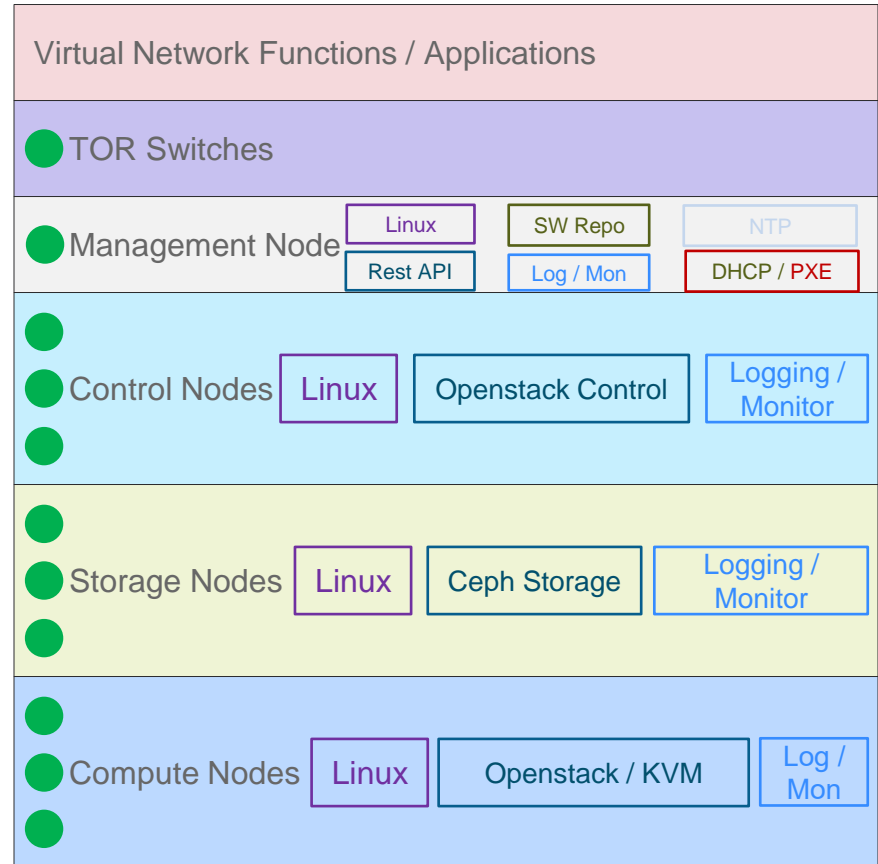
Software Update

- OSS system (or human) triggers CVIM lifecycle manager (one command)
 - Every version corresponds to a single file that is downloaded to management node.
 - Update will download all updated packages, check integrity, and install where needed.
 - If an issue is detected, system is automatically rolled back to state prior to update start.
 - If update includes kernel update, controllers and storage nodes will be rebooted one by one to ensure control and storage remains up.



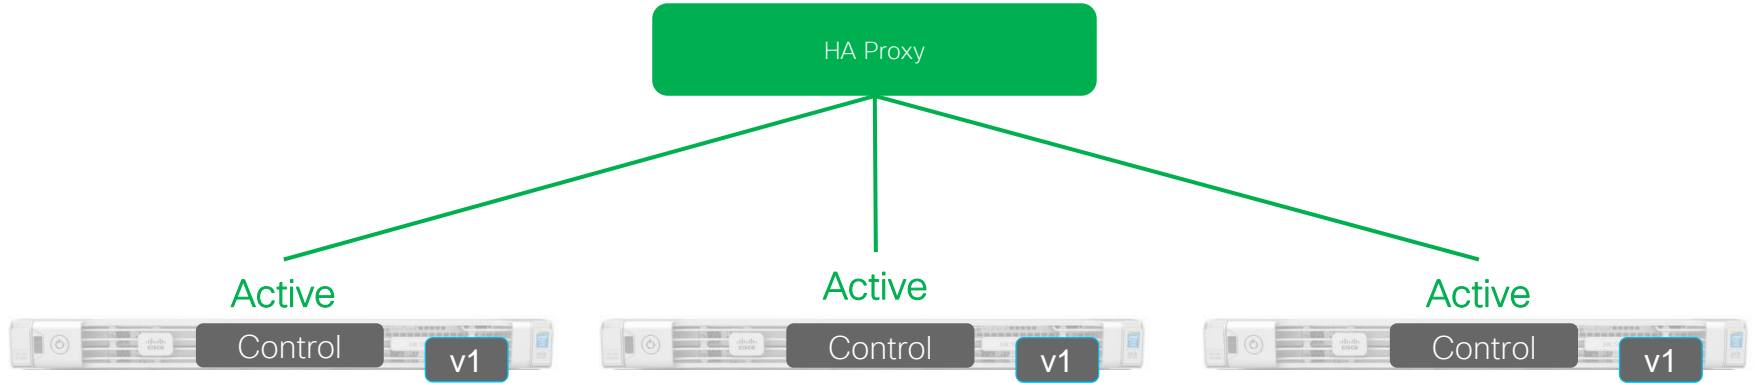
Software Update

- OSS system (or human) triggers CVIM lifecycle manager (one command)
 - If update includes kernel update, then only compute nodes with no running VNF's are automatically rebooted. Compute nodes with active VNF's will be put in reboot pending so operator can determine appropriate reboot time. This usually depends on the HA mechanisms of the VNF's.
 - If update succeeds, operator needs to commit the update before additional changes to the CVIM pod can be made. This to ensure consistency.

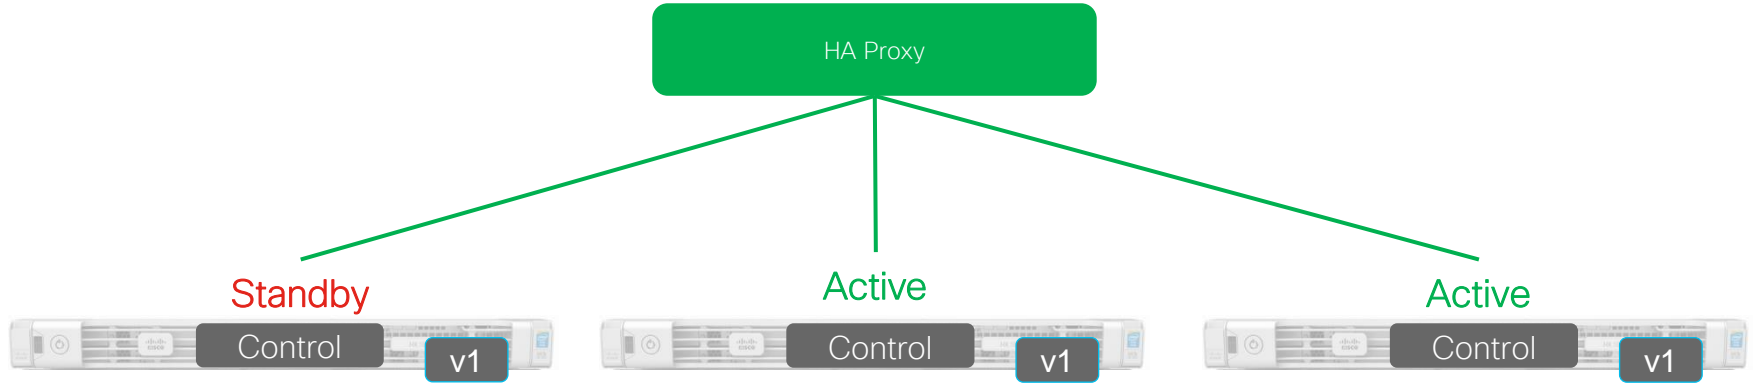


Software Update: Rolling updates

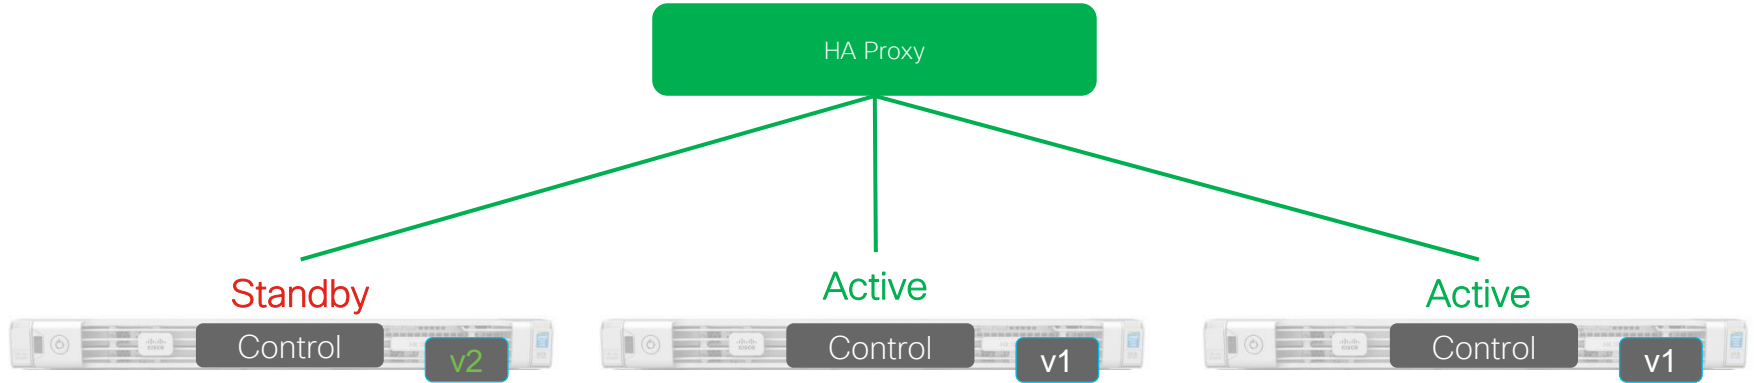
- Rolling updates ensure control plane remains up.



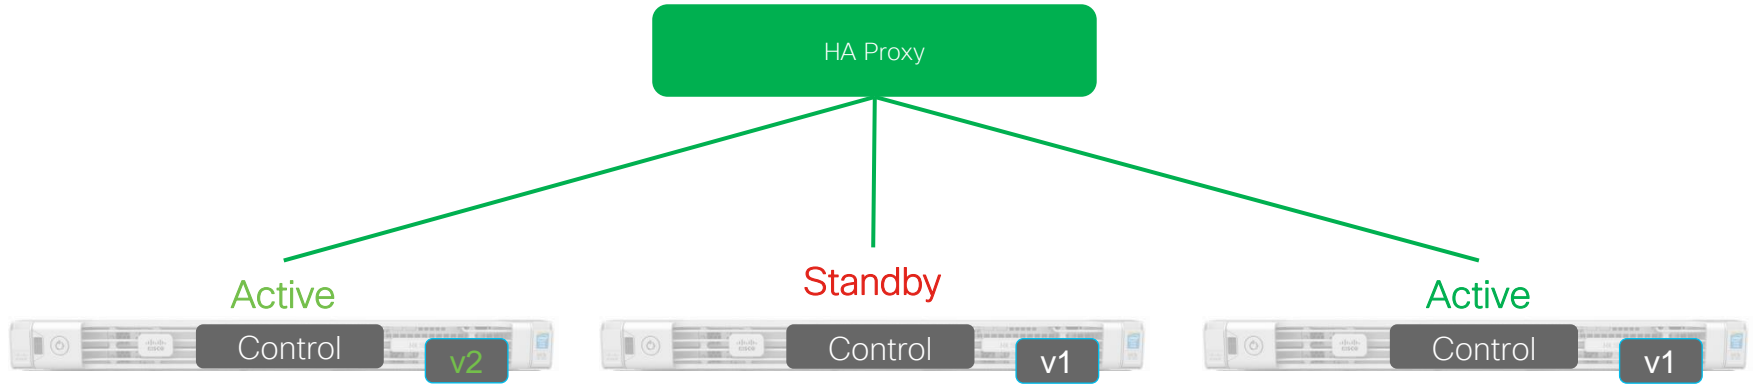
Software Update: Rolling updates




Software Update: Rolling updates



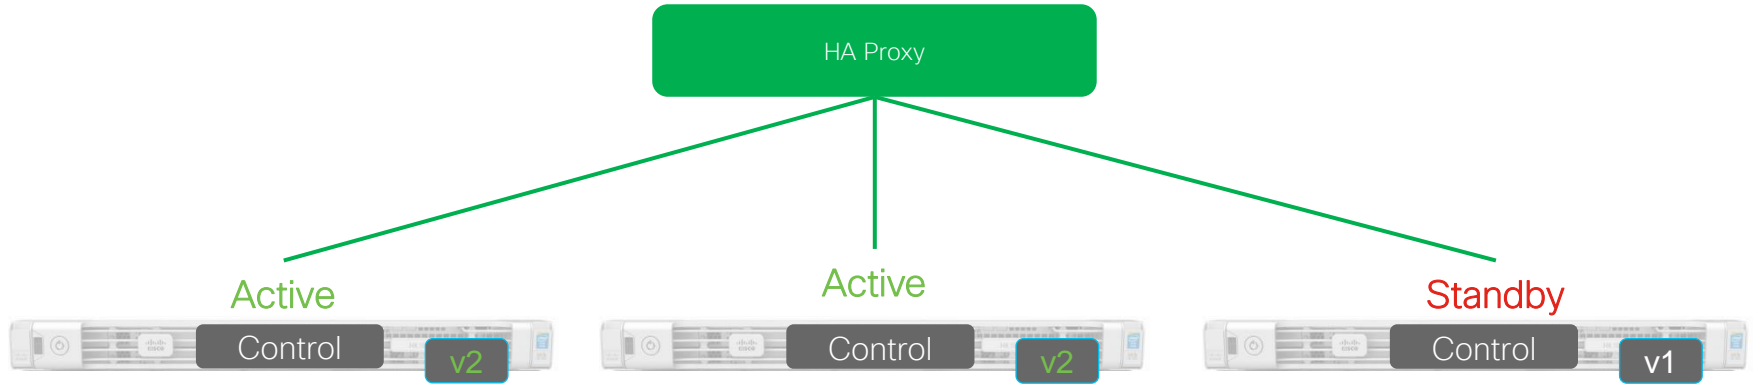
Software Update: Rolling updates



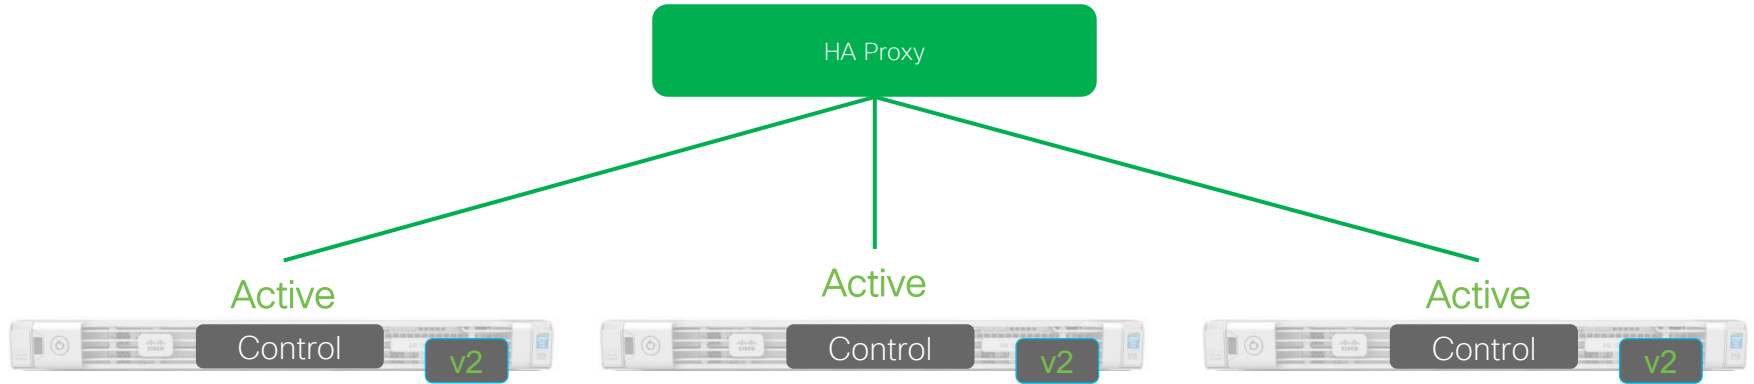
Software Update: Rolling updates



Software Update: Rolling updates

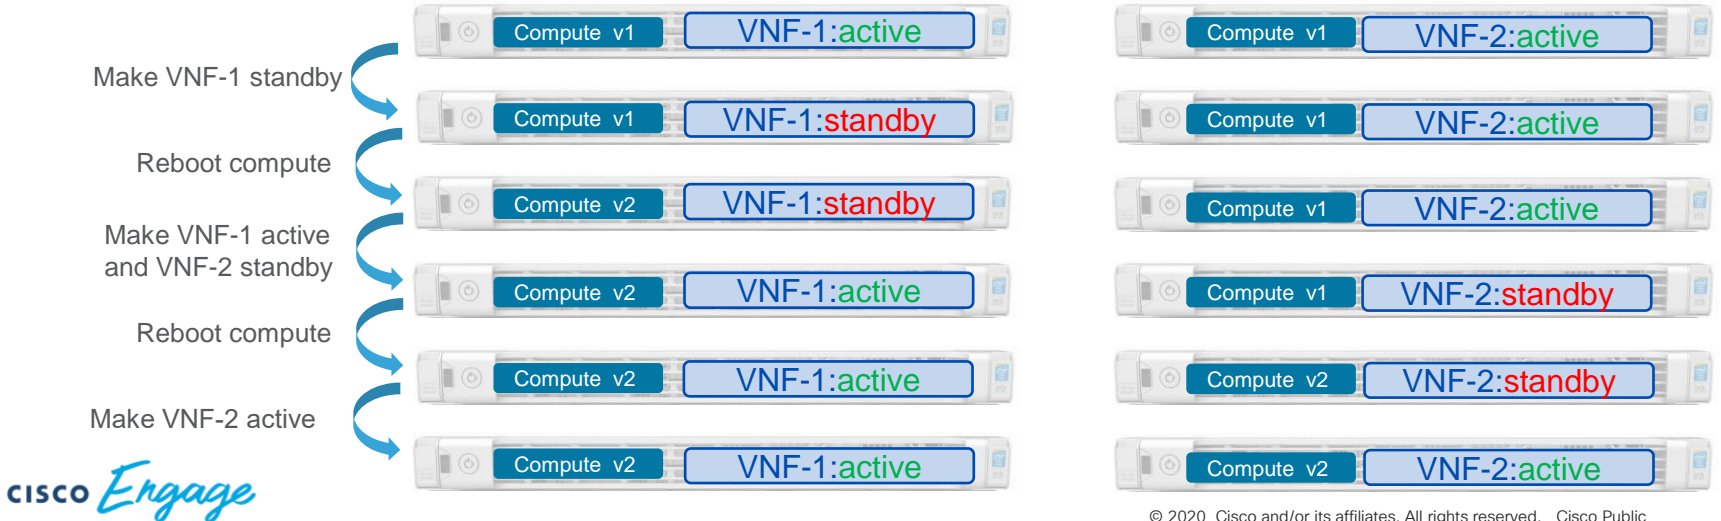


Software Update: Rolling updates

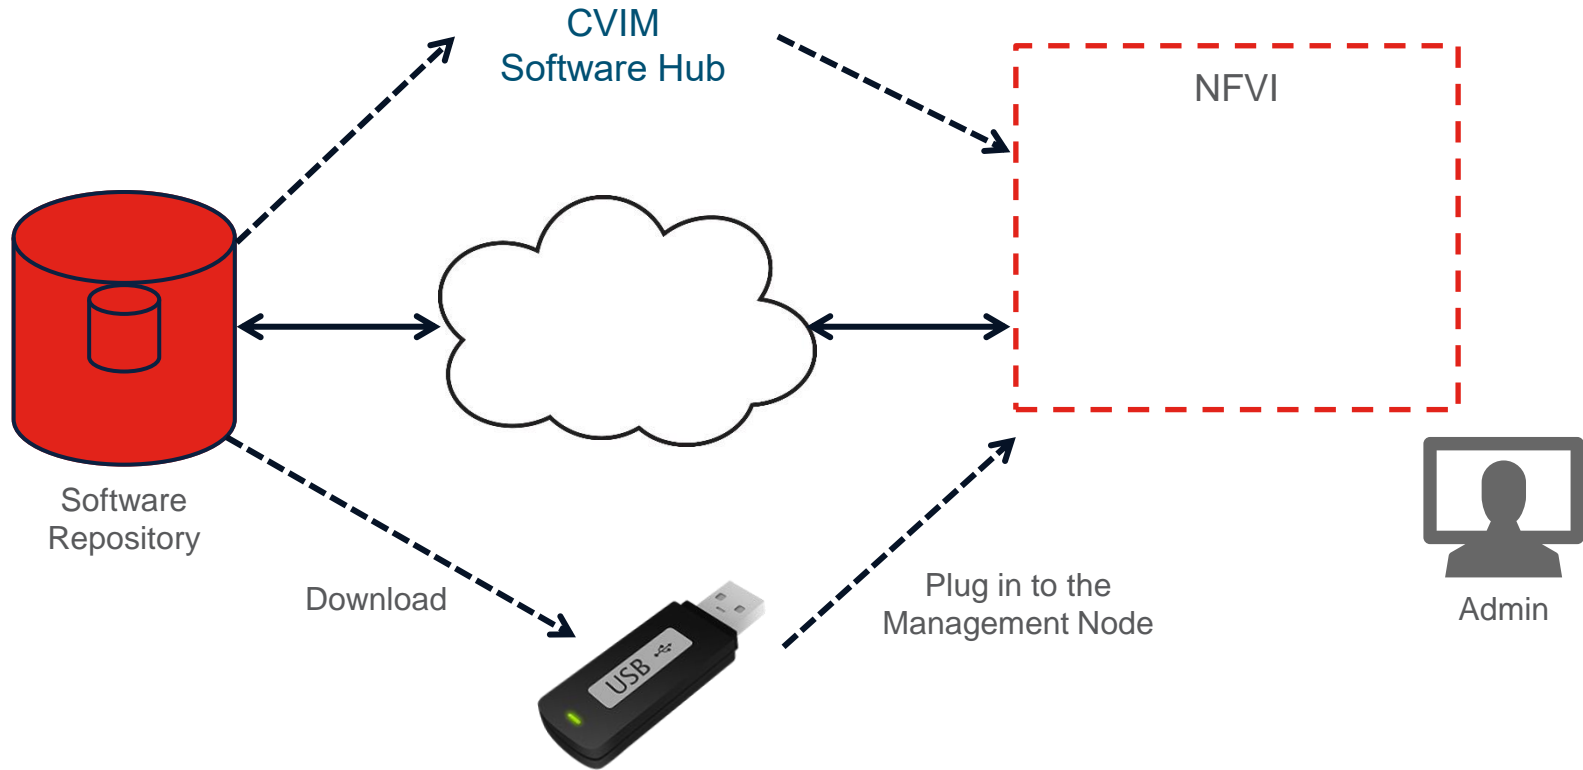


Software Update: Compute nodes

- If kernel is updated, only compute nodes with no running workloads are automatically rebooted.
- Other nodes are put in “pending reboot”, operator decides when to reboot.


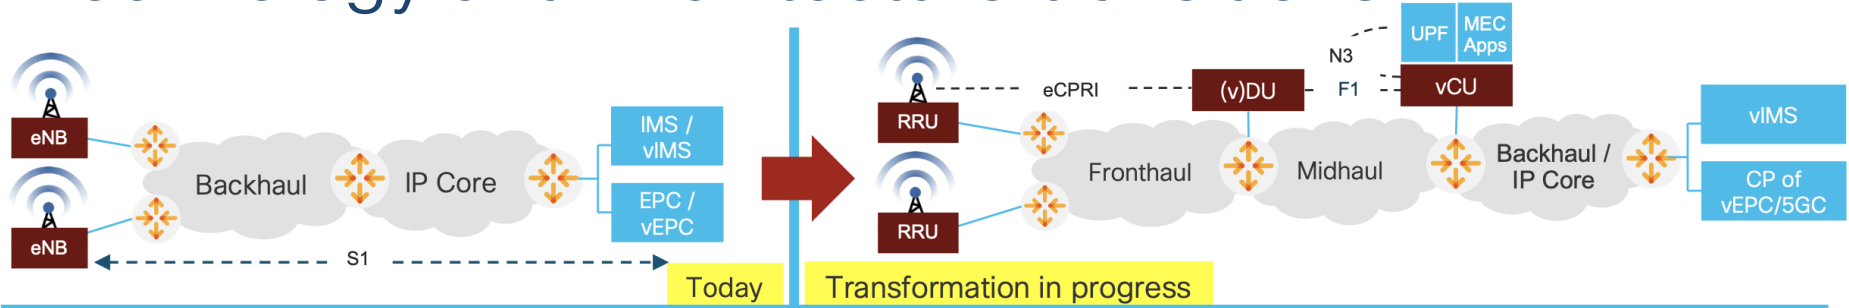


CVIM Offline & Online Install / Update / Upgrade



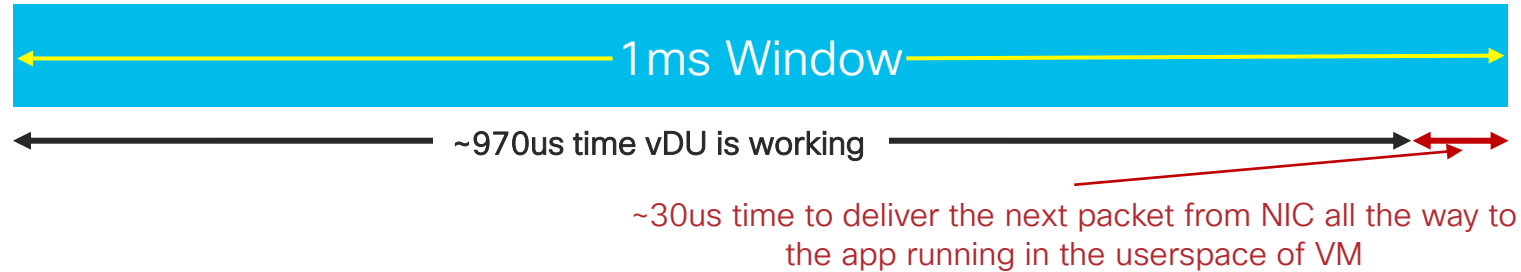
5G Transition

Technology and Architecture transitions



(Hard) Real Time Virtualization in CVIM

- Enabler for vRAN (vDU workload)



Scaling the Edge: Central Control?

3 cores gained (2% gain)
Resiliency reduced

Edge pod



1 core gained (2% gain)
Resiliency reduced

Nano pod



Core DC



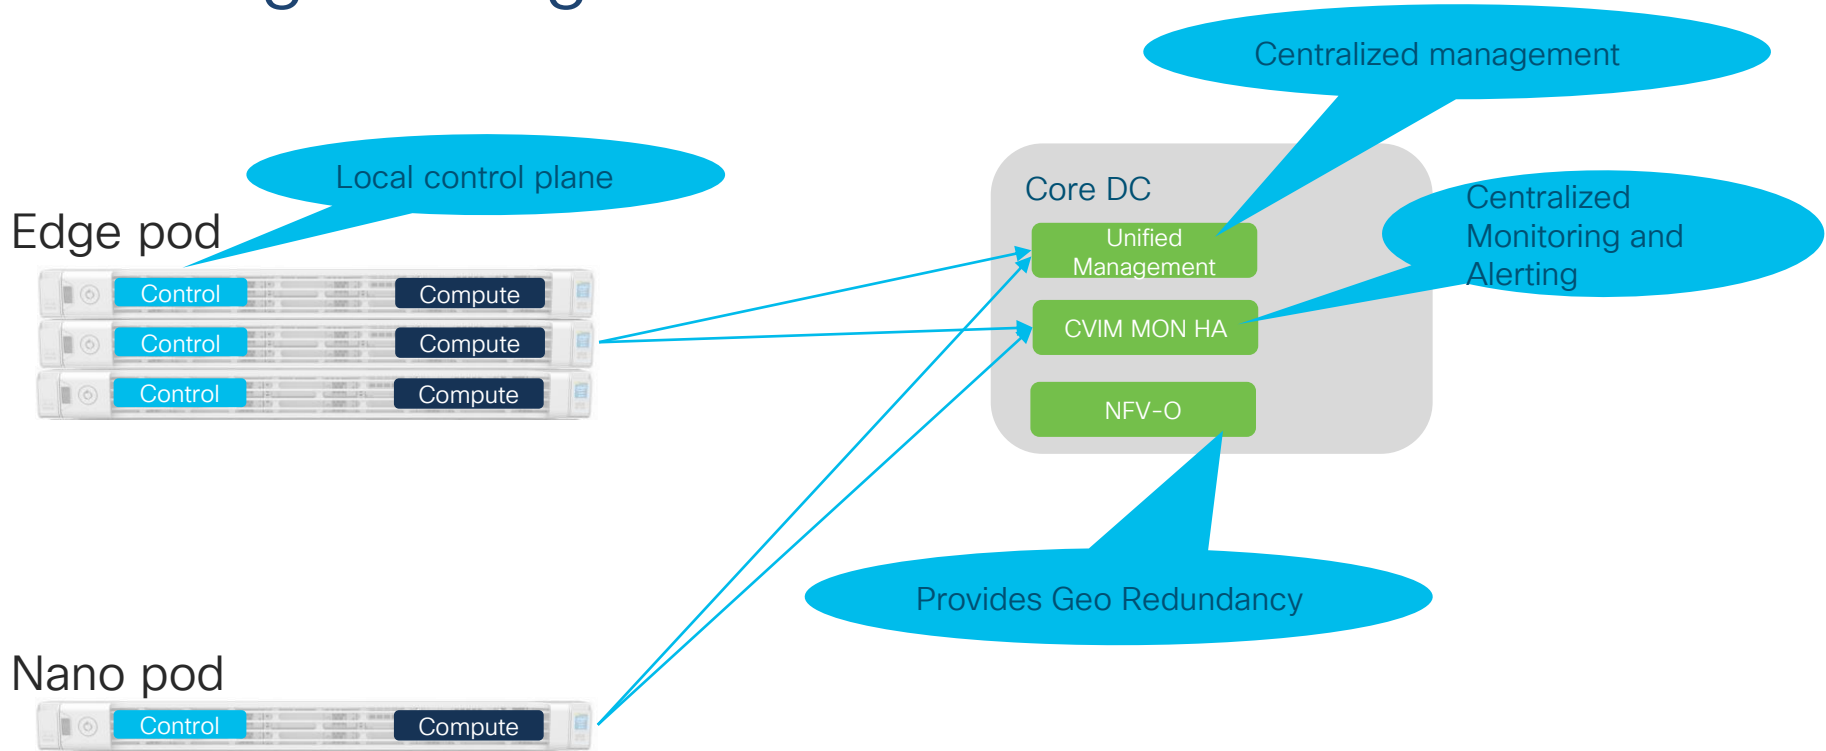
- Compute scaling limited by messaging bus bottleneck
- Additional complexity to maintain resiliency

Control Plane suffers from:

- Random delay
- Packet loss
- Out of order messages
- Unpredictable reliability
- ...

Huge
Blast
Radius

Scaling the Edge



Moving workloads?

Move VNF from compute-1 to compute-3




OSS/BSS




NFV-O




VNF-M



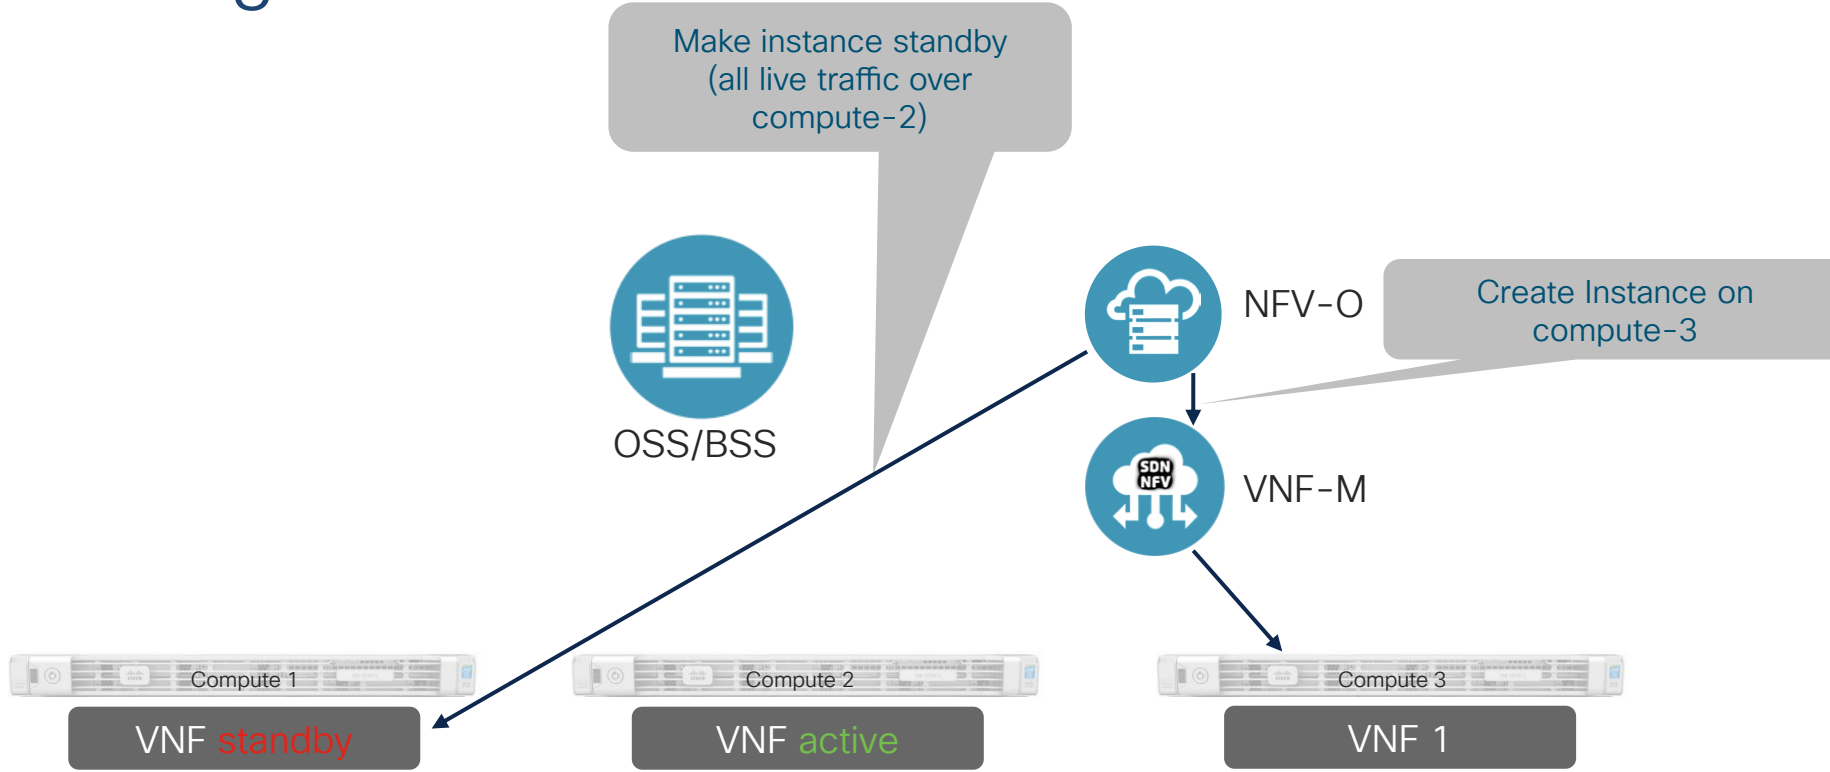
VNF active



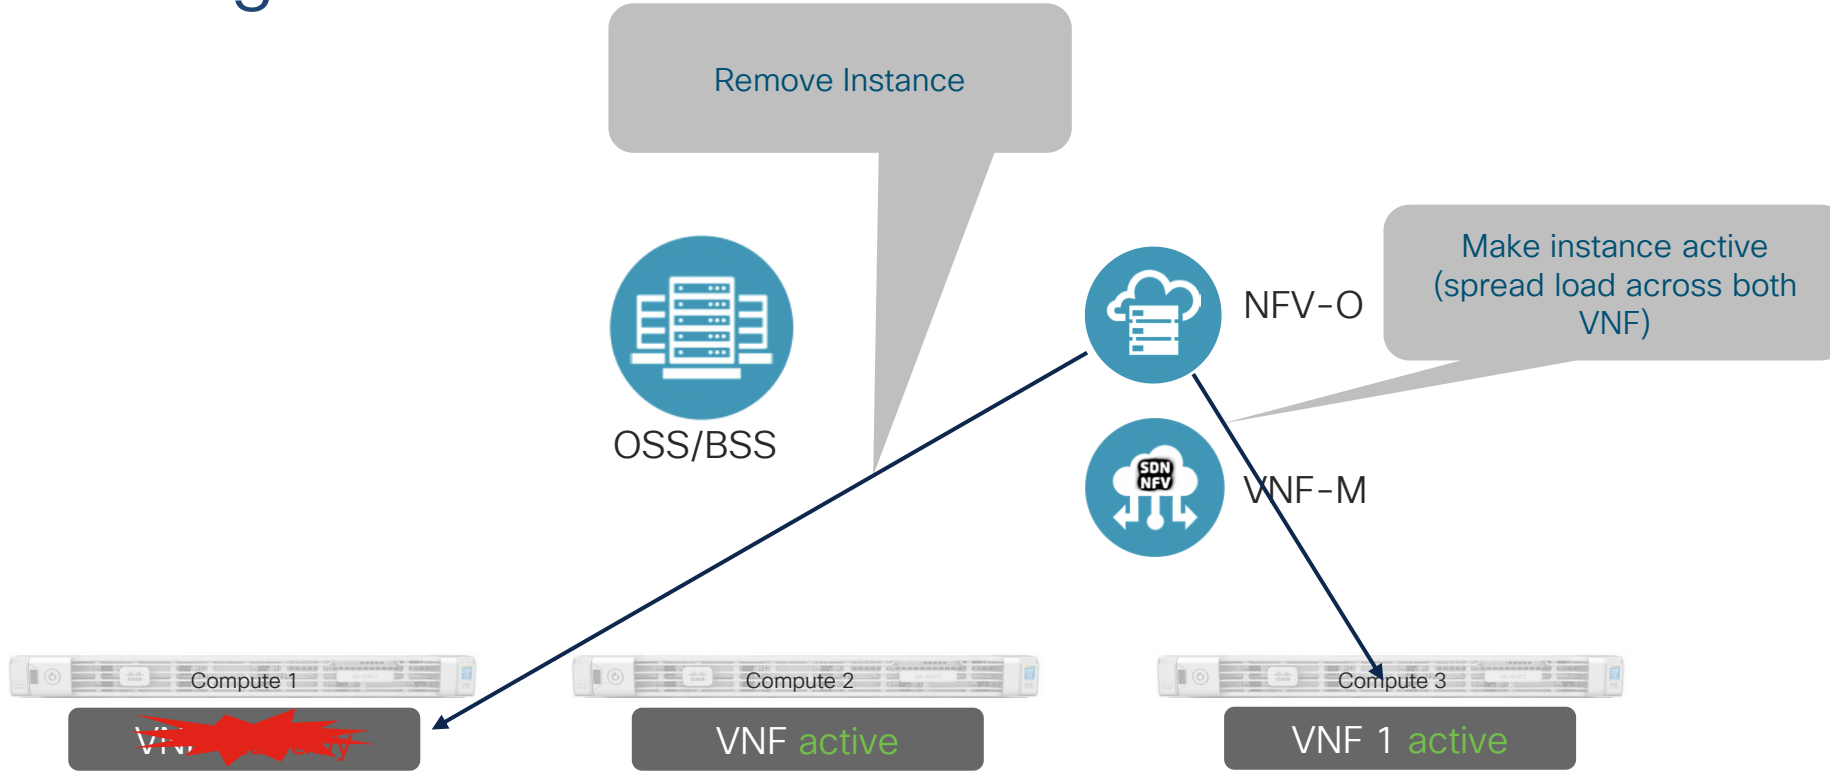
VNF active



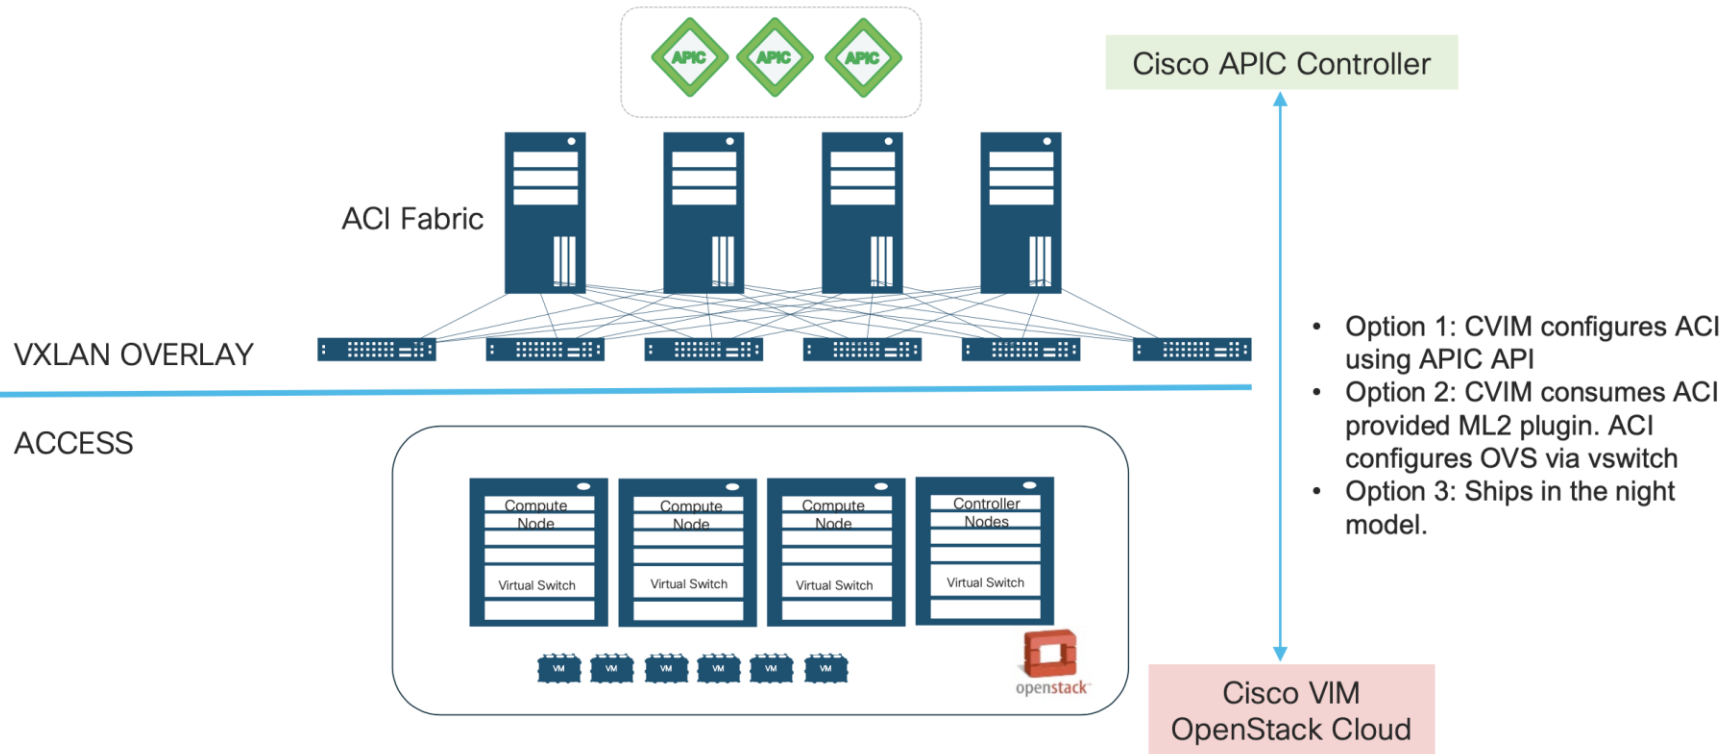
Moving workloads?



Moving workloads?



Cisco VIM Integration with Cisco APIC



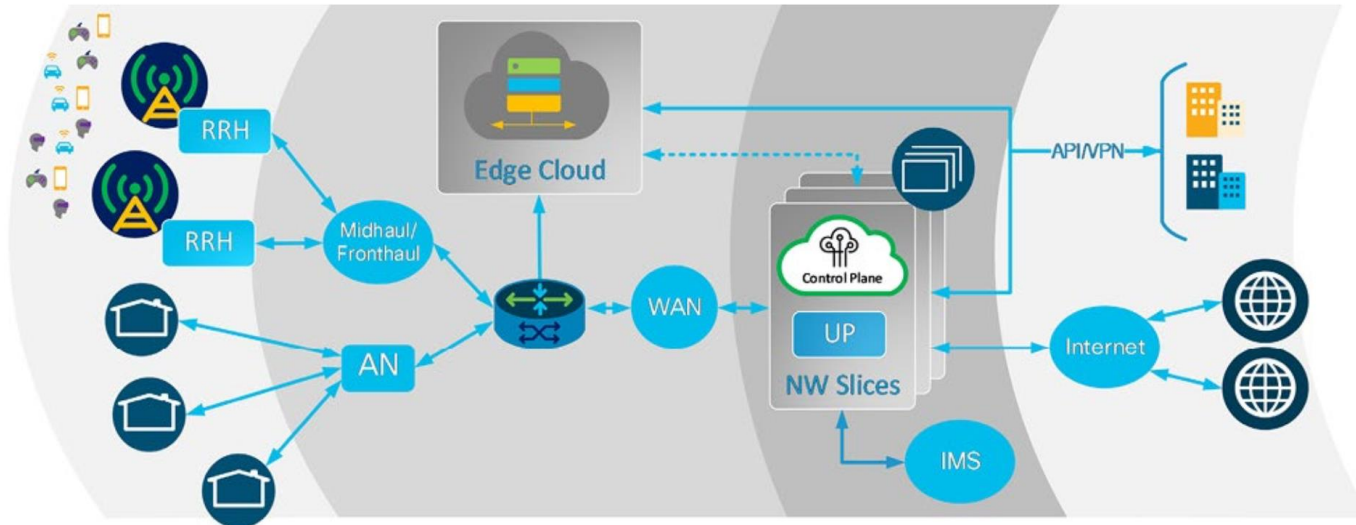


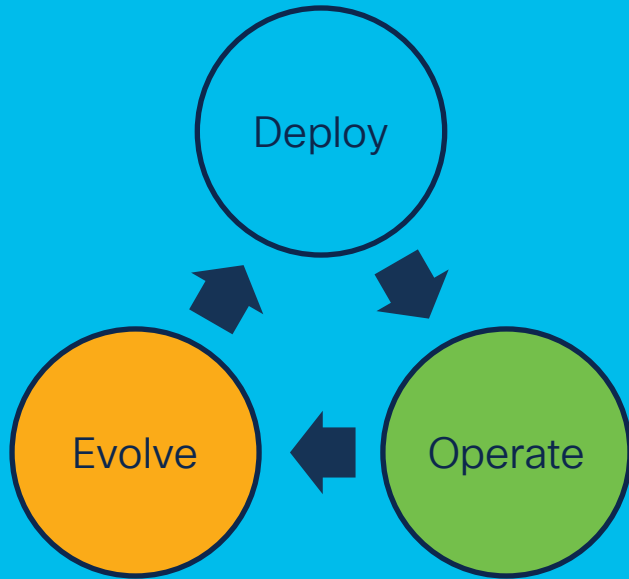

Future

CISCO *Engage*

Future Use Cases

- New requirements from emerging 5G and Multi-access Edge (MEC) use cases:
 - Lower latency – URLLC (Ultra Reliable Low Latency Communications), AR/VR
 - Higher scale – mMTC (massive Machine Type Communications), IoT
 - Higher bandwidth – eMBB (Enhanced Mobile Broadband)
 - More programmability – Network Slicing





- Customers need special capabilities in telco cloud platform
- What matters is not just the deployment, but the whole lifecycle
- CVIM is all about automation and completely lifecycle management
- CVIM is evolving to enable “Cloud Native” NFV

Thank you

CISCO *Engage*





Possibilities

#CiscoEngage