

The logo for WIND, featuring the word "WIND" in white, uppercase, sans-serif font on a red rectangular background.

AN INTEL COMPANY

# TITANIUM CLOUD VIRTUALIZATION PLATFORM

Glenn Seiler

Software Defined Infrastructure BU

30 Minutes

12 Content Slides

# Wind River Titanium Cloud

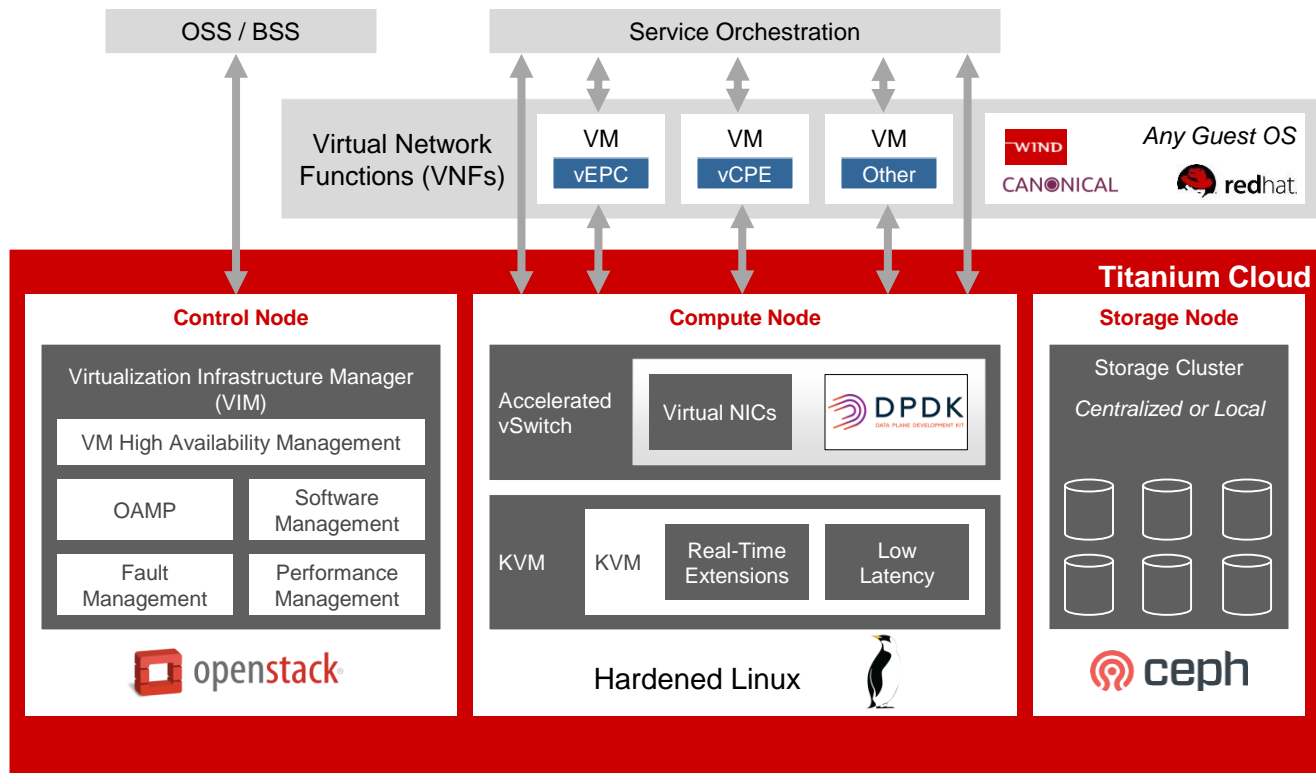
Titanium Cloud is a cloud and virtualization software platform designed for the rigorous demands of critical network services

- ✓ Complete cloud and virtualization software platform
- ✓ Support 3-layer decoupling
- ✓ Based on Open Source and Open Standards
- ✓ Fully Carrier Grade



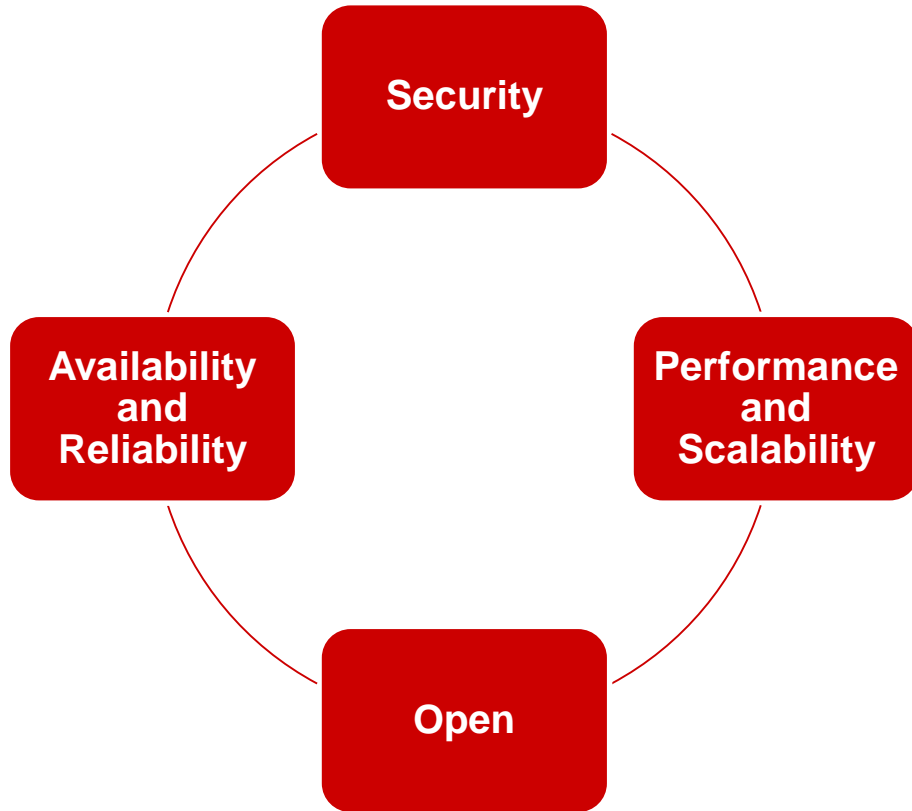
- Accelerate Time to Market with fully integrated and validated solution
- Lower OPEX through reduced operations complexity and remote and automated management
- Lower CAPEX through use of IT COTS hardware and software

# Titanium Cloud Architecture



- Containers or VMs on any guest OS
- ↑
- Add flexible storage cluster
- ↑
- Add Cloud Management VIM and middleware functions
- ↑
- Add high performance accelerated vSwitch
- ↑
- Add critical real-time virtualization enhancements
- ↑
- Based on standard open source components

# Complete Carrier Grade Cloud Platform



- **Performance and scalability**
  - Accelerated vSwitch: 20Gb/s guest throughput
  - 10µs interrupt-latency real-time virtualization
  - Scales from 1-node to 100s of nodes
- **Availability and Reliability**
  - Six nines (99.9999%) reliability at the platform level
  - Automatic detection and automatic recovery
  - Live migration of VMs with less than 150ms outage
- **Security**
  - Secure chain of trust from physical HW to VMs
  - Network-level AAA with secure identities
- **Open**
  - Open standards and open APIs
  - OPNFV testing and validation


# PERFORMANCE and SCALABILITY



# System Scalability


## Minimum-Footprint Edge Solution

Single server



## Highly-Available Edge Solution


Two servers



1:1 protected pair of servers

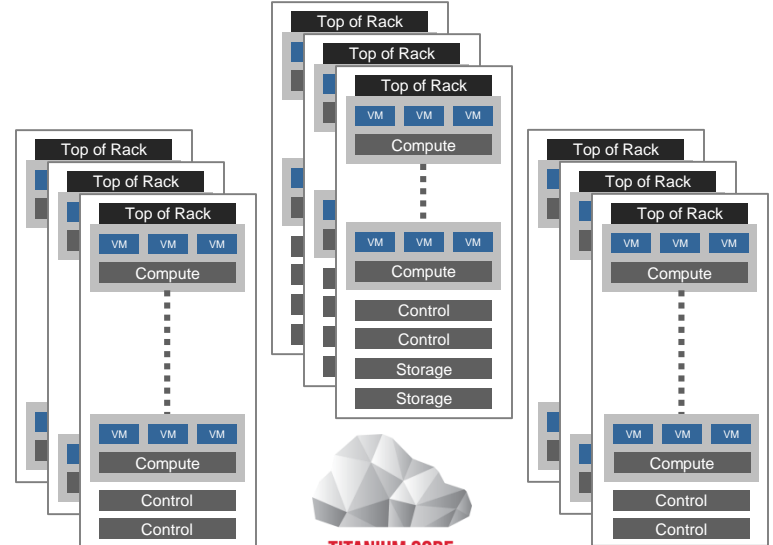
## Frame-Level Solution

4-100 servers



## Large-Scale Data Center Solution

Hundreds of servers



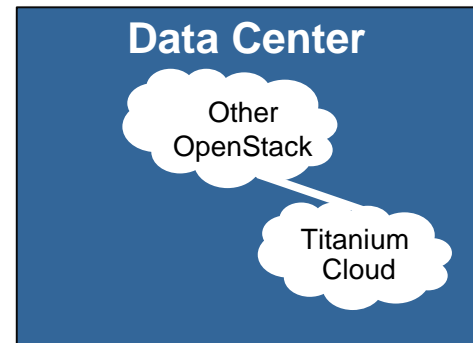
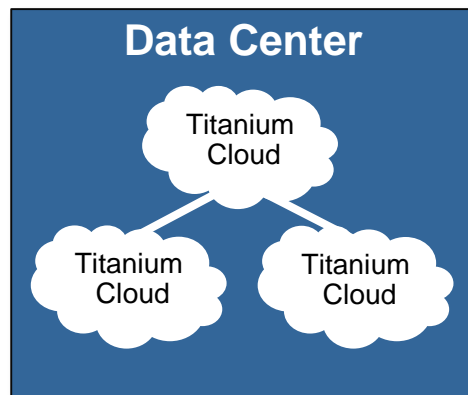
Multi-region cloud

# Multi-Region Scaling

**Region** = A discrete OpenStack environment with dedicated API endpoints that typically shares at least the Identity / Keystone service with other regions

Why multiple regions?

- Scalability of Titanium Cloud installations
  - Increase number of Nodes
  - Span geographies
- Inter-working / Federation with existing non-Titanium Cloud installations



# Low Latency Compute Profile

- New “Low Latency” Compute Node Profile
  - Selected at “Add Host” time
- Low Latency profile
  - Optimized for low-latency guest performance with the following tradeoffs
- Standard Profile, suitable for the **majority** of field use cases:
- Low Latency Profile best for **MEC & CRAN**
  - Host average latency  $2\mu s$  (3 sigma –  $4\mu s$ )
  - Guest average latency  $3\mu s$  (3 sigma –  $\sim 4\mu s$ )



Add Host

Host Info \* Board Management

Personality \* ⓘ  
Compute

From here you can add the configuration for a new host.

Personality Sub-Type \*  
Low Latency

Host Name

Management MAC Address \*

Cancel Add Host




# VM Performance and Scalability

- “Huge Pages” VM creation
  - Allocation of Huge Pages for VM delivers high performance for the VM
  - Support for 4K, 2M, and 1G memory pages
- “NUMA Optimized” VM creation option
  - NUMA affinity for enhanced VM performance
- “No Over Commit CPU” VM creation option
  - Guaranteed VM performance with “dedicated” CPU model
- Mixed “dedicated” and “shared” CPU models
- Ability to specify CPU models for VMs
  - In order to leverage advanced features of selected CPU architectures
- Ability to assign scheduling policy and priority to vCPUs of VMs

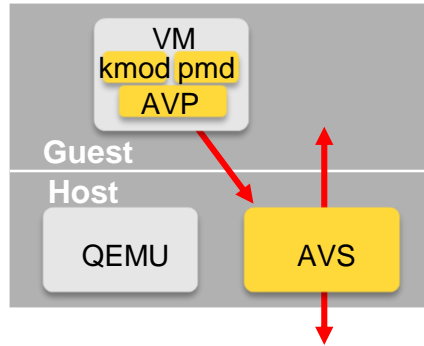
# Switching and Packet Performance

## Virtio Driver without vhost



- Standard VirtIO driver
  - NO changes required to VNF
  - NO recompile of VNF needed,
- SLOW performance (due to QEMU)

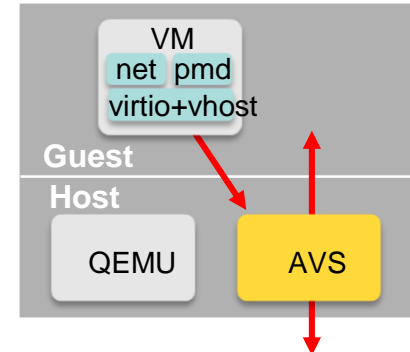
## WRS AVP Driver



- Uses open source AVP driver
- HIGH performance (KMOD / PMD: 8x / 40x 'VirtIO w/o vhost')
  - leverages high performance data transfer instr. set,
- SCALES for multiple Guests

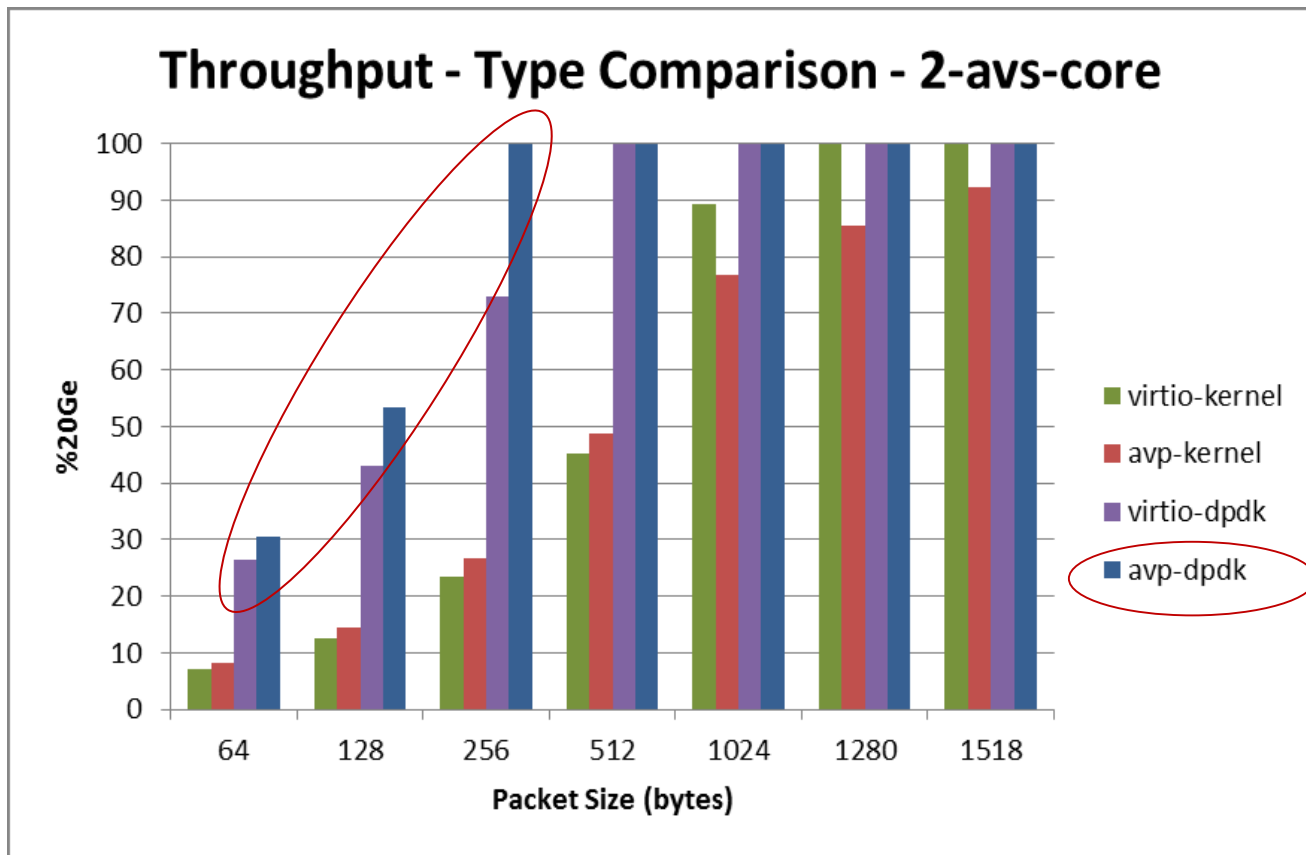
*avp-kernel*  
*avp-dpdk*  
*virtio-kernel*  
*virtio-dpdk*

## Virtio Driver with vhost



- Enhanced standard VirtIO driver
  - NO recompile of VNF
- Equivalent Performance
  - does NOT leverage high performance data transfer
- NOT scale for multiple Guests

# Industry Leading Packet Performance



## World class packet latency

- Predictable, deterministic performance with flexible profiles tuned to match use case

## High-performance switching for VM-to-VM traffic

- 40x performance of kernel-based vSwitching solutions
- No compromises from passthrough or SRIOV

# AVAILABILITY and RELIABILITY



# Meeting Telco Reliability Requirements

Titanium Server delivers true carrier grade reliability

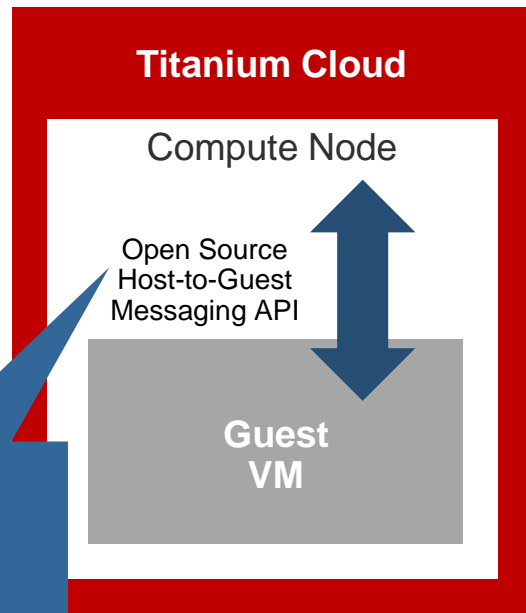
	Enterprise IT Platform Capability	Carrier Grade Cloud Requirements
Detection of failed VM	> 1 minute	< 1 second
Detection of failed compute node	> 1 minute	~ 1 second
Automated controller node failure detection & recovery	No support	Full Support
vSwitch performance	1-2 Gbps	Line rate with minimum core utilization
Network link failure detection	Depends on Linux distribution	50 ms
Live migration for DPDK-based VMs	No support	> 150 ms with DPDK

# Carrier Grade Availability APIs

- VM heartbeating and health checking
  - Simple validation of kernel scheduling
  - Application-specific health checks
- VM event acknowledgement / notification
  - Supported VM events: shutdown, migrate, reboot, pause, and suspend
- VM peer state notification and messaging
  - State change notifications of VMs within server group
  - Broadcast messaging mechanism between VMs
- VM guest scaling
  - Notification of add / remove of vCPU
  - Allows proper OS online / offline of vCPU

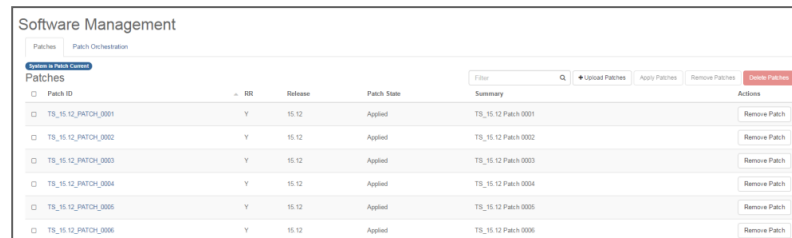
Available on  
GitHub

- Specification
- Reference implementation



# Carrier Grade Software Patching

- Ability to deploy software patches
  - Manually
  - Automatically through patch orchestration tool (Titanium Edge SX included)
- Support for patching of host OS, NFVI components
  - Apply, remove, query patch status per node
- Patches may or may not require reboot of impacted compute nodes
  - If required, guest VMs are live-migrated off of node
- In-service patching
  - Through support of rolling patches across all affected nodes
  - Patch orchestration tool provides “one-button” application of patch across system
- Comprehensive patch status
  - For each node
  - At the system level
- Patch dependencies ensure consistency of patches applied to system



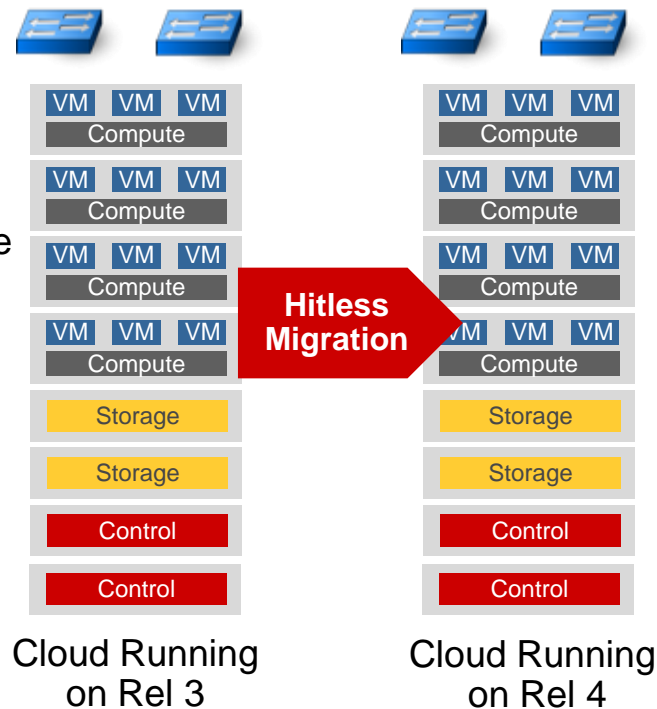
The screenshot shows a 'Software Management' interface with a 'Patches' section. It features a table with columns for Patch ID, RR, Release, Patch State, Summary, and Actions. All listed patches are in the 'Applied' state.

Patch ID	RR	Release	Patch State	Summary	Actions
TS_16_12_PATCH_0001	Y	16.12	Applied	TS_16_12 Patch 0001	Remove Patch
TS_16_12_PATCH_0002	Y	16.12	Applied	TS_16_12 Patch 0002	Remove Patch
TS_16_12_PATCH_0003	Y	16.12	Applied	TS_16_12 Patch 0003	Remove Patch
TS_16_12_PATCH_0004	Y	16.12	Applied	TS_16_12 Patch 0004	Remove Patch
TS_16_12_PATCH_0005	Y	16.12	Applied	TS_16_12 Patch 0005	Remove Patch
TS_16_12_PATCH_0006	Y	16.12	Applied	TS_16_12 Patch 0006	Remove Patch

## Hitless Software Upgrades

# Carrier Grade Software Upgrades

- Integrated end-to-end solution
  - Automated; low number of steps
  - No additional hardware required for upgrade
- Rolling upgrade across nodes
  - Managing API compatibility between nodes at N and N+1 release
  - Live migrating of hosted applications that support it; otherwise cold migrating
- Manages upgrades of all platform software
  - Host OS changes
  - New / upgraded OS packages
  - New OpenStack release
  - New / upgraded Titanium Cloud features

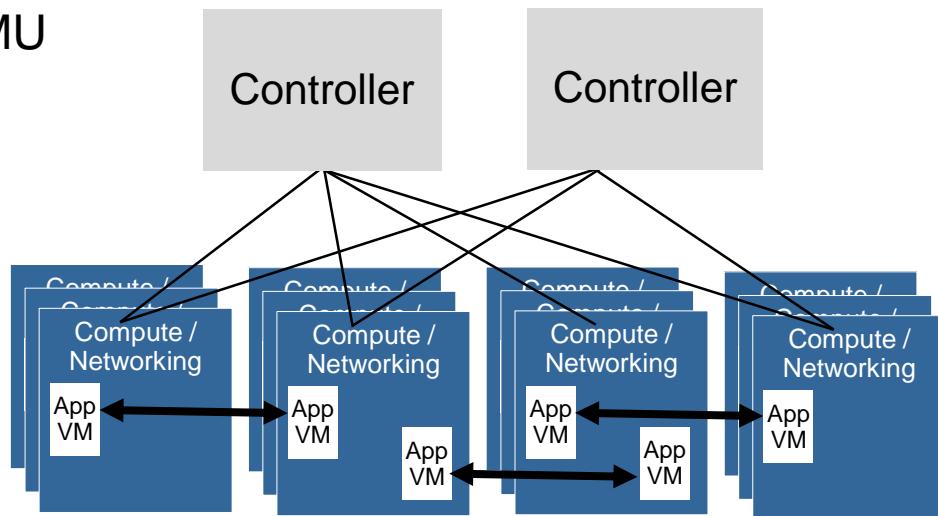


**Hitless Software Upgrades**



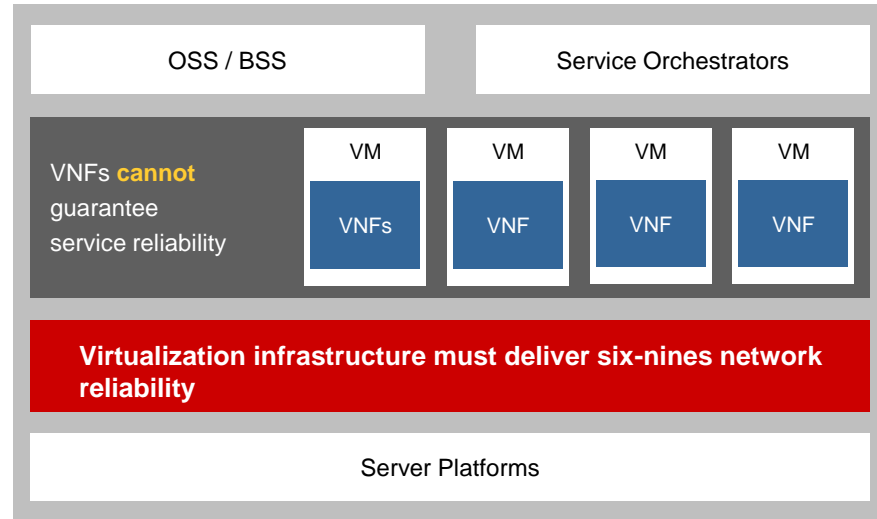
# Carrier Grade VM Management

- Automatic VM recovery on KVM / QEMU failure
- Automatic VM recovery on guest OS / application failure
- Live migration of VMs
- Graceful VM shutdowns, migrations, pause, etc.
- Support for persistent replicated virtual disks for VMs
- Support for VM server groups



# Application-Level High Availability Is Not Sufficient

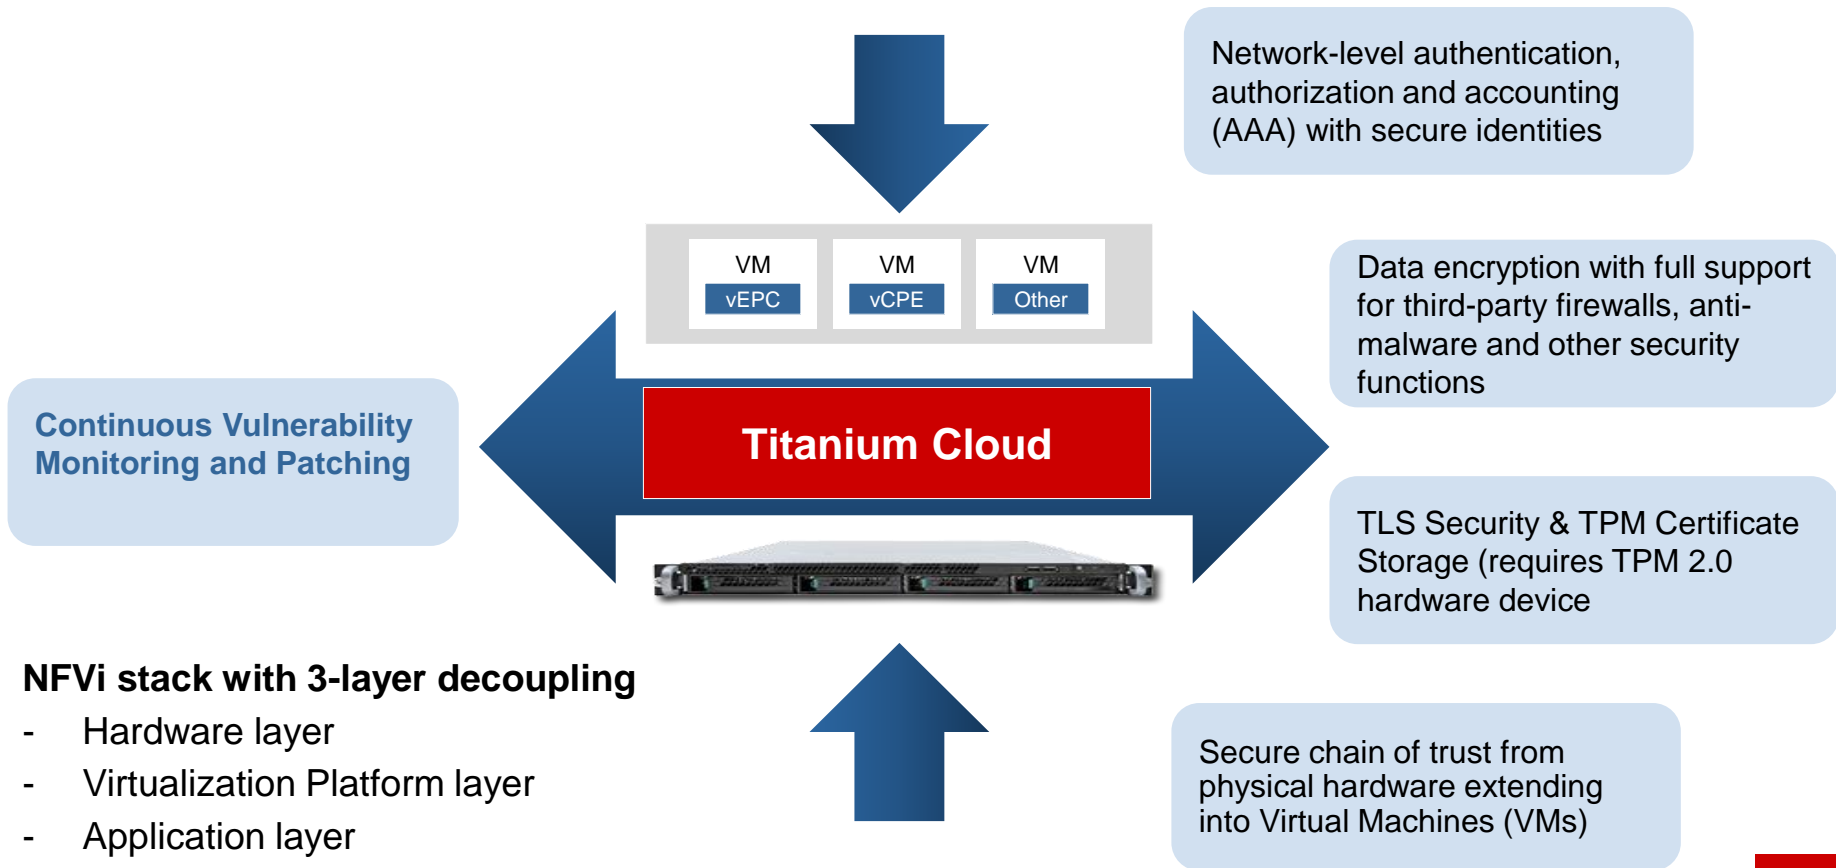
- Several approaches exist for application-level HA
  - Active / Active
  - Active / Standby
  - N-way Active with load balancing
- None of these meet system-level requirements for reliability or resiliency
  - No awareness of underlying system resources: could deploy both instances on same server
  - No guarantee of deterministic, consistent VNF performance (e.g., NUMA awareness)
  - No awareness of service chains
  - No capability of automated recovery from system-level failures
  - No support for platform-level security



# SECURITY



# Holistic Security for a NFVi Platform



# Carrier Grade Security

## ■ Integrity ✓

- Critical process monitoring and recovery on Titanium Cloud nodes
- Resource monitoring on Titanium Cloud nodes
- Monitoring and recovery of Titanium Cloud node connectivity
- Titanium Cloud controller node program store authentication
- Tenant data stored on private closed network

## ■ Access ✓

- Network firewall on external OAM interface
- Role-based access control
- Secure password enforcement
- Password aging
- Restricted access to root account and root cmds
- Auto-logout of in-active user sessions
- External LDAP integration—keystone

## ■ Confidentiality ✓

- Secure keyring database for storage of encrypted passwords
- ACL filters for authenticity of connectivity to hosted VMs
- QOS for protection of connectivity to hosted VMs
- TLS Security & TPM Certificate Storage (requires TPM 2.0 hardware device)

## ■ Host environment ✓

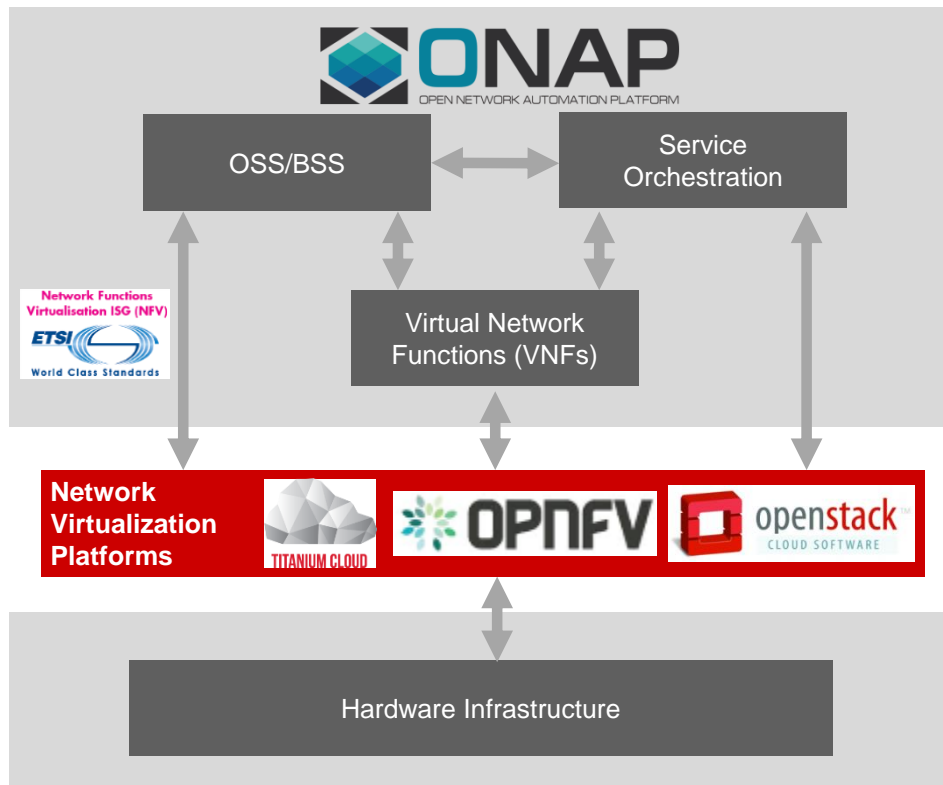
- UEFI Secure Boot & Cryptographically Signed Images
- User and group permissions
- Chroot jail
- Virtual TPM for Secure Guests
- Process group isolation

# OPEN



# Open = Interoperable and Compatible = No Lock-In

In-depth technical validations through Titanium Cloud ecosystem



- In-depth compatibility and interoperability validation
- Comprehensive test plans and validation reports with OPNFV tests
- Strict compatibility with essential components; OpenStack, Ceph, etc
- RESTful APIs for High Availability and other functions

# Demonstrated Open and Fully-Interoperable Platform at OPNFV Plugfest

- **Titanium Core: the only NFVI platform that demonstrated full interoperability**
  - Complete virtualized nanocell and EPC
  - Collaboration with China Mobile, Huawei, Red Hat and others
- **Demonstrated on complex, real-world configurations**
  - Raisecom virtual EPC and virtual security gateway VNFs hosted on Titanium Core
- **FuncTest** (base system functionality testing)  
**Yardstick** (infrastructure verification) testing





# Wind River Active in Key NFV Open-Source Projects



## Active in maintenance of Yocto Linux project

- Linux-yocto kernel and tooling
- Meta-cgl, meta-cloud-services, meta-openembedded, meta-security, meta-selinux, meta-virtualization, meta-zynq

**Contributed real-time KVM to many related open-source projects**



## Leading contributor to Nova, SR-IOV and other core projects

- Wind River is ranked in top 10% to the Nova Compute project
- Wind River is ranked 86th out of 318 contributors to OpenStack

**Wind River is focused on critical telecom-focused enhancements to Nova**



## Active in critical workgroups

- Contributing our HA APIs to HA Workgroup
- Compliance & certification workgroup
- OPNFV validation and testing requirements (C&C, Dovetail)

**Proven Vendor Interoperability at OPNFV PlugFest**

# SUMMARY



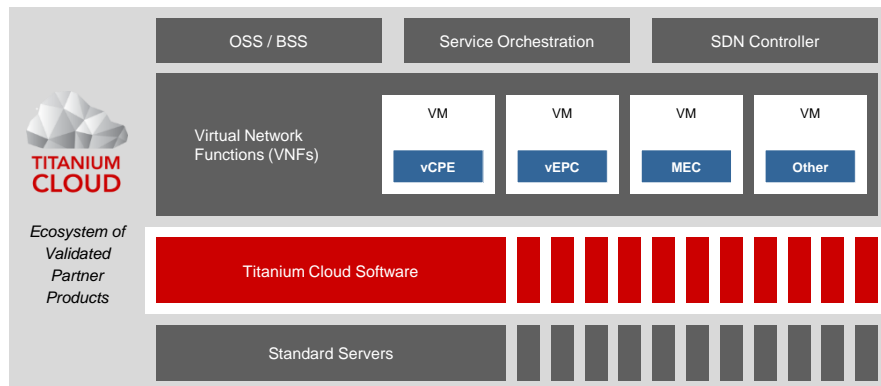
# Titanium Cloud Solves Key Telecom Problems

Adding unique value by solving real-world telecom and carrier grade issues

- ✓ No Vendor Lock-in with fully Open platform
- ✓ Industry-best performance; latency, packet through put, failover
- ✓ Fully scalable architecture from 1 to 100s of nodes
- ✓ Best Carrier Grade reliability in the industry
- ✓ Secure Cloud platform for every Telecom use-case



# TECHNICAL BENEFITS ARE CLEAR....



- **Secure, integrated telco cloud platform**
- **Runs virtualized services with carrier grade reliability**
- **Seamless scalability from 1 node to 100s**
- **Simplified installation, commissioning and maintenance**
- **Compatible with Open Standards**
- **Accelerated VNF performance**
- **Extensive ecosystem**

.... So how do we quantify the business benefits?