WHITE PAPER

Wind River and Intel Join Forces on Reference Design for Virtual Business Customer Premises Equipment

Communications services providers (CSPs) use business customer premises equipment (BCPE)—ranging from network interface devices (NIDs) to routers and firewalls—to connect enterprise customers to the telecommunications network and to deliver value-added services.

Analysts estimate that business customers spend more than $45 billion annually on wide area network data services and value added services requiring BCPE.[1] That market is expected to grow as CSPs offer new cloud, network security, and other services to enterprises looking to outsource their IT needs.

This market growth, though, accentuates the challenges CSPs have deploying BCPE and why many are looking to virtualize these devices. Multiple, separate BCPE appliances are required for each service, which can be costly, requires site visits for installation, and can have long lead times as the equipment must be ordered from the manufacturer for each customer.

Many CSPs would like to virtualize this BCPE paradigm using network functions virtualization (NFV),[2] which replaces discrete appliances with an Intel® architecture server and with software-based virtual network functions (VNFs) delivering the services. The results include lower network capital expenses, lower installation costs and faster time to market.

The reliability and performance of the virtualized BCPE (vBCPE) are critical to the acceptance of this new paradigm. One large service provider is exploring the potential of vBCPE to make a dramatic impact in its business – if it can develop a platform that delivers performance and carrier-class reliability. It is for this reason that Wind River®, in partnership with Intel, Intel® Network Builders ecosystem and key members of the Wind River Titanium Cloud™ ecosystem, have developed a reference design that it and other CSPs can use to build vBCPE platforms that are optimized for network-wide deployment.

[1] vCPE delivers opportunities for new business services http://searchsdn.techtarget.com/tip/vCPE-delivers-opportunities-for-new-business-services

[2] NFV Introductory Whitepaper: https://portal.etsi.org/nfv/nfv_white_paper.pdf (PDF download)

Table of Contents

Components of vBCPE

Traditionally, vBCPE marks the very edge of the network, often serving as the point of demarcation between networks owned by the CSP and those owned by the customer. In most cases CSPs own vBCPE and are responsible for repair and maintenance. The components of vBCPE include:

- Carrier-Grade Server: Intel architecture-based multi-core CPU-based server with memory and storage and with 1/10/40GbE network and telecom interfaces.
- NFV Software Platform: At a minimum, provides virtualization of the server, hosts virtual network functions (VNFs), and delivers networking services such as virtual switching.
- Management and Orchestration (MANO): MANO functions include orchestration of network services, management of VNF lifecycles, and infrastructure management, including controlling and managing NFV resources.
- Virtual Network Functions: Third-party software applications that provide value-added services, such as network routing, virtual private network, firewall, WAN acceleration, and others.

vBCPE Deployment Modes

Depending on the CSP and the service, there are three different ways to deploy vBCPE for service deployment:

Centralized or Telco Point of Presence (PoP) Deployment: In this deployment mode, servers and VNFs are installed in the telecom data center and services are delivered from this data center. This model is optimized for small and medium-sized businesses where CPE processing requirements are low and where there is an existing carrier Ethernet connection to the service edge. There is a significant up-front investment to build out this infrastructure, but it is the lowest cost deployment strategy because the shared server pools maximize the statistical multiplexing of compute loads for multiple customers.

Customer-Premises Deployment: When services are latency sensitive or require network security, it makes sense to deploy them entirely from the customer premises. This requires vBCPE server hardware with significant compute power, I/O, memory, and storage. With this infrastructure in place, CSPs can deploy significant, high-value services that make the most of this infrastructure. This deployment model requires up-front infrastructure investment, but scales with demand more cost effectively than the central deployment. Once the infrastructure is in place, CSPs can easily deploy new service-provisioning VNFs remotely without a need for a new appliance.

Hybrid Deployment: In many cases, a combination of deployment models will give the best new-service deployment flexibility, especially for medium-to-large-sized enterprises that need advanced security, WAN optimization, and other services. This model is designed for service deployment to a broad base of business customers ranging from SMBs to the largest enterprises.

A Firm Foundation: Choosing the Right NFV Software Platform

One of the critical concerns of service providers is that the virtual customer premises solution be reliable and manageable in a carrier application. From a software perspective, the foundation for carrier-class vBCPE is the NFV software platform. This is the software that virtualizes the underlying hardware and provides the network services to host and connect virtual network functions.

There is more than one path to choose from in selecting an NFV software platform: One could start with vanilla open source OpenStack* and build a solution from scratch, or use Wind River Titanium Server™—a complete and commercially available platform. Both share a common set of features—in fact, Titanium Server incorporates OpenStack—but both also have significant differences.

OpenStack*

OpenStack is an open source cloud operating system that is managed by the OpenStack Foundation,[3] is very popular in data center applications, and is being trialed as an NFV platform as well. OpenStack is made up of nine core projects that together provide a virtualization environment for compute, networking, and storage. In addition, there are optional services that add specific enhancements.

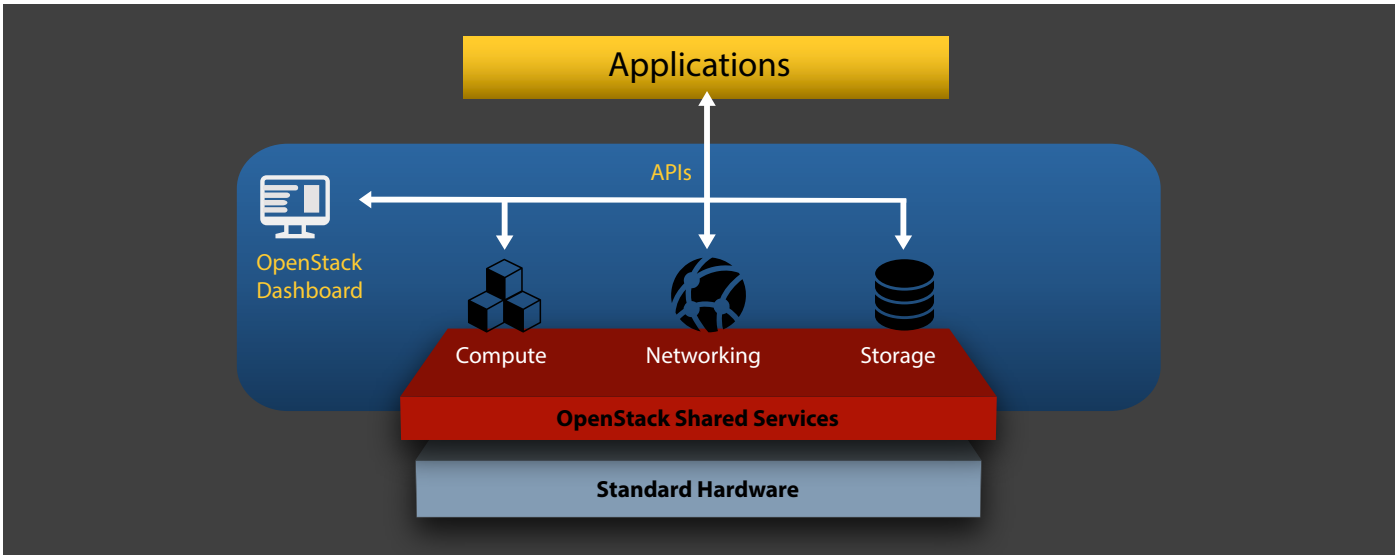

Figure 1: OpenStack Architecture

The core projects include:

Nova is OpenStack's compute virtualization project, providing the management of compute instance lifecycles. With Nova, large numbers of virtual machines can be deployed (spawned, scheduled, decommissioned) and managed network wide.

Swift is a storage system for unstructured data objects. Swift utilizes a RESTful, HTTP-based API to store and retrieve objects based on a unique identifier. The result is a very scalable and highly reliable storage system.

Cinder provides persistent block storage to VMs, a feature that improves data access speeds for applications where that's important.

Neutron provides the networking services for the VMs that are managed by Nova. It features an API that can define networks and attached devices from leading networking vendors.

[3] https://www.openstack.org

Keystone provides identity (authentication and authorization) services for other OpenStack services. Keystone maps endpoints with services permissions to manage service access.

Glance is a tool for providing access to virtual machine disk images during VM provisioning. This core service stores and retrieves these images.

OpenStack optional services include:

- Horizon – graphical dashboard to manage OpenStack
- Ceilometer – telemetry services
- Heat – orchestration
- Trove – database
- Sahara – elastic map reduce
- Ironic – bare-metal provisioning
- Zaqar – messaging service
- Manila – shared filesystems
- Designate – DNS service
- Barbican – key management
- Magnum – containers
- Murano – application catalog
- Congress – governance

OpenStack provides a comprehensive set of features and functionality, but is optimized for the data center and is far from being considered "carrier grade," with the reliability and manageability features needed for CSP networks.

Wind River Titanium Server CPE

Wind River Titanium Server CPE is part of the Wind River Titanium Server product portfolio. It gives CSPs an NFV software platform that easily deploys on a pair of commercial off-the-shelf servers, providing a fully integrated, cost-effective carrier-grade NFV foundation.

As part of the Titanium Server software family, Titanium Server CPE shares the carrier-grade design of the software and also offers compatibility with VNFs from companies that are part of the Titanium Cloud ecosystem. Some of the essential elements of Titanium Server CPE software platform include:

Carrier-Hardened Linux*: Titanium Server CPE is based on a trusted, cloud-targeted Linux distribution, leveraging open source contributions and adding availability, security, clustering, and performance enhancements that make it appropriate for use in production carrier networks.

Real-time Kernel-Based Virtual Machine (KVM): This hypervisor software is part of Linux. KVM facilitates the creation of virtual machines when used on Intel architecture-based processors or others that support hardware virtualization extension. Wind River has made significant improvements to KVM that reduce interrupt and timing latency for more predictable performance.

Accelerated vSwitch: A high performance, feature rich vSwitch for protected, inter-VM networking that leverages the Data Plane Development Kit (DPDK) architecture. This accelerated vSwitch vastly improves packet processing between VMs and from the network controller to VNFs over competing vSwitching technologies. These performance improvements allow more VMs to be run per core, increasing VM density and lowering operating expenses.

Carrier-Grade OpenStack Plug-ins: Titanium Server CPE has integrated OpenStack for cloud-computing functionality, but has added key reliability, usability, manageability, and availability extensions required for CSP networks.

Virtualization Infrastructure Manager (VIM): The VIM's features include rapid detection of, and automated recovery from, VM and host node failures; VM resource allocation and management; and the rapid, automated migration of VM workloads to new servers upon server failures.

Firewall and Secure Access: Controller-functions in the software are protected from hackers through a built-in firewall. Management of the system can also be secured through the authenticated role-based access, enforcement, integrity, and confidentiality features. The built-in service chaining capabilities also can be used in conjunction with a third-party firewall VNF for added security.

VNF Network Acceleration: VNF performance enhancements include support for Virtio drivers for enhanced network and disk performance. Open source accelerated virtual port (AVP) drivers and kernel loadable modules (KLMs) are supported if additional performance is needed.

Selecting the Right NFV Software Platform

So does it make more sense to build vBCPE on vanilla OpenStack or Titanium Server CPE? OpenStack is a highly acclaimed open source cloud computing software with many of the components necessary for the virtualized environment needed for the vBCPE.

Indeed, OpenStack offers a compelling solution and Titanium Server CPE leverages OpenStack. But several CSPs have asked some critical questions about going to market with an "off-the-shelf" OpenStack-based vBCPE.

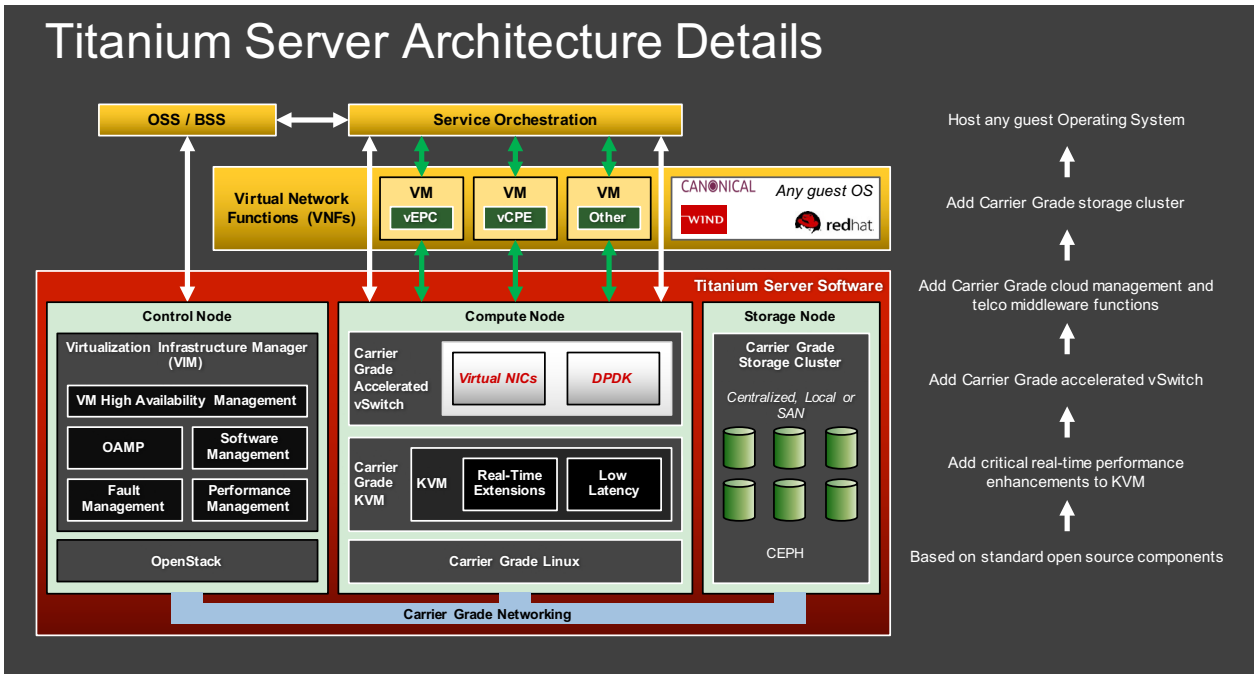

Figure 2: Wind River Titanium Server Architecture

In a presentation at the 2015 SDN and OpenFlow World Congress, a BT representative laid out several key challenges for CSPs' use of OpenStack[4]:

1. Controller Scalability: In carrier tests, one OpenStack controller managed about 500 compute nodes, which is limiting when considering that a CSP could have hundreds of thousands of CPEs to manage in their network. This scalability is a challenge when a network comes back after an outage and all of those customer premises devices simultaneously signal the network to connect.

2. Service Chain Agility: OpenStack lets CSPs set up service chains to direct data flows to VNFs in a specific order, but these service chains are hard to modify. Adding another VNF requires disconnecting the interface, setting up the new service chain and then reconnecting. In BT's vBCPE test networks, some VNFs locked up completely when added to a service chain, requiring them to be reinstalled to work.

3. Internet Security: In carrier tests, using OpenStack in a DSL environment meant making too many data exceptions in the firewall for it to have a high degree of integrity. Additionally, in OpenStack, each VNF has a serial number that is used to determine connection order. This makes it difficult to verify that LAN and WAN connections are properly made, which can jeopardize the effectiveness of services such as firewalls.

4. Version Interoperability: vBCPE applications are spread over such a large customer base that it's a given that there will be multiple versions installed on the network. OpenStack backward compatibility is a challenge.

5. Binding Virtual NICs to Virtual Network Functions: To ensure deterministic behavior of the vBCPE, it's important to ensure that the correct VNF interface is always connected to the correct virtual NIC, especially after an interface has been disconnected and then reconnected. Testing with an off-the-shelf OpenStack distribution, however, reveals that in some cases the connections are restored incorrectly and that in others the VNF locks up.

6. Start-up Storms: When a link is cut then subsequently restored, hundreds or thousands of compute nodes will then simultaneously attempt to attach to a centralized controller. Testing shows a standard OpenStack controller has insufficient resiliency to cope with this scenario. It can become overloaded and not recover.

[4] BT Threatens to Ditch OpenStack; Light Reading, Oct. 14, 2015

In the work done by the team developing the vBCPE reference design it became clear that there were more challenges with the controller architecture of OpenStack in this application. It's important to remember that OpenStack was designed primarily for data center/IT environments, which are significantly different from the carrier network with tens of thousands of remote sites.

Separating the controller from the compute node not only impacts scalability, as noted above, but also offers challenges when a controller fails or network conditions change and bandwidth is more limited. Either of these scenarios can mean that management traffic does not reach the compute node, or that data must be throttled at either end of the network (call-gapping) to deal with the data overload on the link between the controller and the compute node.

The architecture of Titanium Server CPE overcomes these controller challenges through an integrated controller-compute node architecture that distributes the controller function out to the remote sites.

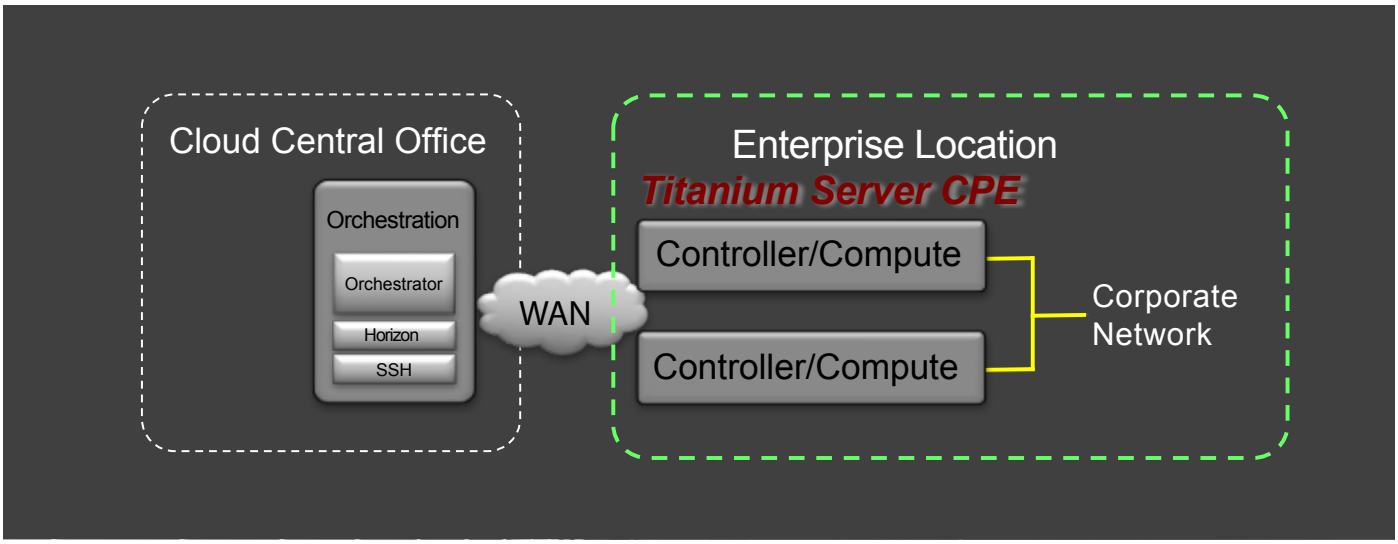

Figure 3: Integrated controller-compute node

Titanium Server CPE's no single point of failure architecture ensures a backup controller is always available in hot standby mode, able to take over in the event of a failure of the primary controller.

Developing a vBCPE Reference Design

Given these challenges with OpenStack, Wind River engaged members of its Titanium Cloud ecosystem to develop a reference design that would offer the functionality and feature set needed by CSPs.

The vBCPE reference design was based on Wind River Titanium Server CPE running on Intel architecture-based hardware and included VNFs from companies who are a part of the Wind River Titanium Cloud ecosystem. The elements of the reference design include:

Intel

The reference design is optimized for servers based on a wide range of dual-socket Intel® Xeon® processors, which provide power and virtualization support for medium to large enterprise CPE. The Intel Xeon processor D-1500 SoC product family for midrange routers, network appliances, and security appliances can also be used for entry level and medium enterprise applications. The reference design also makes use of high-performance Gigabit Ethernet and 10 Gigabit Ethernet controllers based on the Intel® Ethernet Controller 82599 and Intel® Ethernet Controller XL710 families.

ADLINK* Compute Systems for Mobile Edge Computing (MEC) and Edge vCPE

ADLINK Server for Extreme Outdoors (SETO) is a dual Intel Xeon processor-based, carrier-grade MEC server that can be

placed in harsh, outdoor environments (e.g., radio towers).[5] Latency and optimization of vCPE functions are critical to businesses implementing specific services requiring low latency and speed. ADLINK also provides a distributed common platform, Modular Industrial Cloud Architecture (MICA), which is a reusable vCPE appliance system that can be placed in data centers and central offices and implemented on an OCP-Telecom based platform. These systems are pre-integrated with Titanium Server software and have been fine tuned for the performance needed for vCPE applications.

ADVA* Ensemble Orchestrator*

The ADVA Ensemble Orchestrator is an ETSI MANO NFV Orchestration platform that supports end-to-end network service lifecycle management across industry-leading cloud platforms. It also provides full VNF lifecycle management including a built-in VNF Manager for on-boarding and control as well as integrated VNF/Service operations and analytics.

Check Point* vSEC for NFV

Check Point vSEC for NFV is an advanced threat prevention VNF allowing CSPs to offer cyber-threat protection services. vSEC for NFV protects enterprise customers from internal and external threats using cloud security protection and management. The Check Point vSEC for NFV offering includes advanced multi-layer threat prevention with a comprehensive management platform as well as logging and reporting capability.

Riverbed* SteelHead*

Riverbed SteelHead is an application acceleration and WAN optimization solution that increases the performance of all applications including on-premises, cloud, and SaaS, across Hybrid WANs (MPLS, private VPN, and public Internet). SteelHead provides increased visibility into application performance and end-user experience and the ability to ensure business performance service level agreements (SLAs) through an application-aware approach based on centralized business intent-based policies.

Brocade* 5600 vRouter*

The Brocade 5600 vRouter is a layer 3-7 router that is purpose-built for high-performance NFV implementations. It is designed with carrier-class performance, reliability, and features. Leveraging Brocade's vPlane Technology, the vRouter is able to achieve 10+ Gbps performance per physical core.

The resulting reference design can be seen in Figure 4 below:

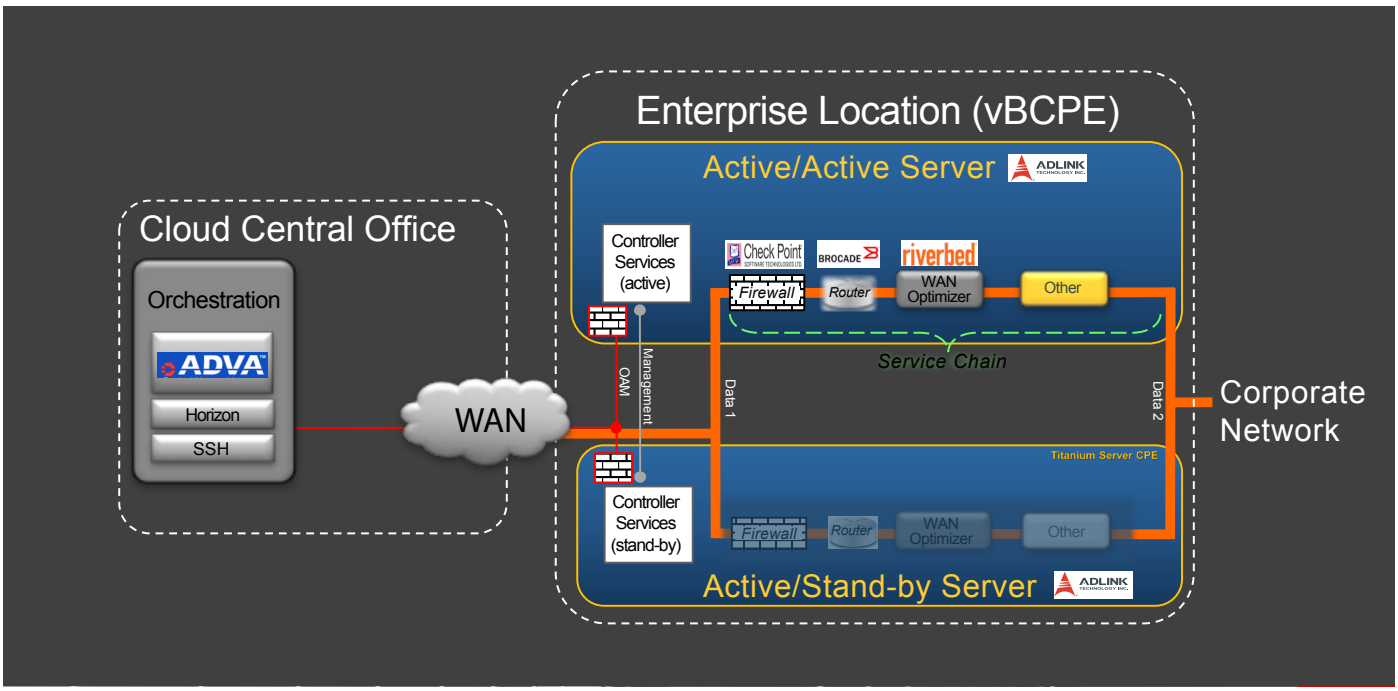

Figure 4: Wind River High Reliability vBCPE Reference Design

[5] See http://www.adlinktech.com/mobile_edge_computing/ and http://www.adlinktech.com/PD/web/PD_detail.php?cKind=&pid=1577

As depicted, Wind River Titanium Server CPE provides controller services, accelerated virtual switching, as well as compute, storage, and networking virtualization on the same platform, which can be powered by multi-core Intel Xeon processors. The design supports 1GbE, 10GbE, and 40GbE network connections.

VNFs from Brocade, Check Point, and Riverbed are in a service chain that directs data first through the Check Point firewall in order to protect the network from malicious traffic. The firewall protection also covers operations, administration, and management (OAM) data traffic, and the reference design provides authenticated management access.

All of the NFVi functionality and the VNFs are mirrored on the stand-by server providing 99.9999% availability with sub-second VNF failure detection and recovery, 50 msec network failure detection, and controller failover.

Conclusion

vBCPE's offer tremendous cost and service provisioning advantages for CSPs, but must match the performance and reliability of the discreet CPE devices they are replacing. A key component to this carrier-class reliability is selecting the correct NFV software platform. Wind River has teamed up with Intel and its VNF partners to develop a reference design that combines performance, reliability, scalability, and manageability required for this application.

About Wind River

A global leader in delivering software for intelligent connected systems, Wind River offers a comprehensive, end-to-end portfolio of solutions ideally suited to address the emerging needs of IoT, from the secure and managed intelligent devices at the edge, to the gateway, into the critical network infrastructure, and up into the cloud.

Wind River technology is found in nearly 2 billion devices and is backed by world-class professional services and award-winning customer support.

About Intel

Intel (NASDAQ: INTC) is a world leader in computing innovation. The company designs and builds the essential technologies that serve as the foundation for the world's computing devices. As a leader in corporate responsibility and sustainability, Intel also manufactures the world's first commercially available "conflict-free" microprocessors.[6] Additional information about Intel is available at newsroom.intel.com and blogs.intel.com and about Intel's conflict-free efforts at conflictfree.intel.com.

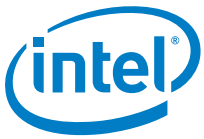

[6] Conflict free" and "conflict-free" means "DRC conflict free", which is defined by the U.S. Securities and Exchange Commission rules to mean products that do not contain conflict minerals (tin, tantalum, tungsten and/or gold) that directly or indirectly finance or benefit armed groups in the Democratic Republic of the Congo (DRC) or adjoining countries. We also use the term "conflict-free" in a broader sense to refer to suppliers, supply chains, smelters and refiners whose sources of conflict minerals do not finance conflict in the DRC or adjoining countries.