SOLUTION BRIEF

Communications Service Providers
Universal Customer Premises Equipment

RAD vCPE-OS Overcomes Universal CPE Mass Deployment Challenges

Small footprint, integrated networking features, and support for a wide range of third-party VNFs mean communications service providers can deploy RAD vCPE-OS on servers powered by Intel® processors for simplified mass deployment.

Introduction: The Challenges of vCPE Mass Deployment

When network functions virtualization (NFV) and software defined networks (SDN) were first introduced, they held the promise of a true revolution—nimble networks, automated service rollouts, and general purpose, low-cost white box servers replacing dedicated appliances. Most importantly, customers would not be locked into an appliance vendor, allowing communications service providers (CommSPs) to mix and match elements they desire, be it a server, VNF, or an orchestrator. Virtualization at the edge was seen as particularly impactful for CommSPs as evidenced by the rapid adoption of software defined WAN (SD-WAN) solutions.

Mass deployment of these services remains a challenge. Popular NFV infrastructure (NFVI) software is resource heavy, which limits the range of servers it can be deployed on. RAD is an Intel® Network Builders ecosystem partner whose answer is the RAD vCPE-OS, an NFVI platform that has built-in networking features, needs minimal CPU and memory resources, and has features that can help CommSPs with mass deployment.

CommSPs Embrace Virtualization

In a strategy to improve network agility and reduce the cost of services, CommSPs have virtualized key parts of their network. Virtualization is important throughout the network, but the technology is especially powerful in remote locations at the edge of the network or on the customer premises. Not only does the service agility at the edge open up many new revenue and service opportunities, the ability to efficiently manage the service lifecycle remotely can reduce costs dramatically.

But mass deployment of virtualized servers has some lingering challenges that have made adopting this technology harder than expected. Some of these challenges include:

- Cost of hardware capable of hosting virtualized functions without compromising performance.
- The high license costs of brand-name VNFs.
- Integration with existing NFV managers and orchestrators.
- Secure device provisioning, including on-the-fly changes to service-chains and VM sizing.
- Diverse license management mechanisms that complicate service provisioning and make license renewing operations more complex.
- SLA assurance with enhanced performance monitoring, OAM, and diagnostics.
- Deployment over any access infrastructure, including legacy TDM, GPON, and xDSL.

Business Case for Mass uCPE Deployment

Despite these challenges, there is a very compelling business case driving edge compute virtualization. Some of the factors include:

Increasing revenue: CommSPs can increase revenue by using their edge virtualization infrastructure for value-added services, such as SD-WAN, IoT backhaul, and cloud/edge computing.

Improved agility and efficiency: Edge virtualization offers quick deployment of new functions/apps in a way that minimizes time, labor, and cost associated with buying, installing, and configuring services—especially when compared to maintaining, updating, and managing inventory for separate proprietary devices.

Security features and reliability: Robust security features and the flexibility to replace firewalls as security needs evolve both help protect against hacking and malicious attacks. For added reliability, the network can be made redundant through the use of wireless backup services.

Vendor ecosystem: Virtualized CPE utilizes a growing assortment of VNFs to tailor solutions for changing requirements. This wide range of VNFs, in addition to multiple options for connectivity and hardware, eliminate vendor lock in.

Reducing capital expense (CapEx): White box servers are widely available from competitive vendors in multiple form factors with varying compute, memory, and storage resources. This flexibility enables hardware cost optimization per service.

Reducing operating expense (OpEx): Virtualized edge servers can host multiple apps, thus simplifying installs and management and helping reduce total cost of ownership. These servers also feature automated routine network administration work, which reduces errors and manpower requirements.

Disaggregated Architecture

Virtualization delivers these benefits through network disaggregation, which separates the software that enables the service from the hardware appliance it runs on (see Figure 1). This architecture frees CommSPs to optimize products and solutions to meet their operational and business needs.

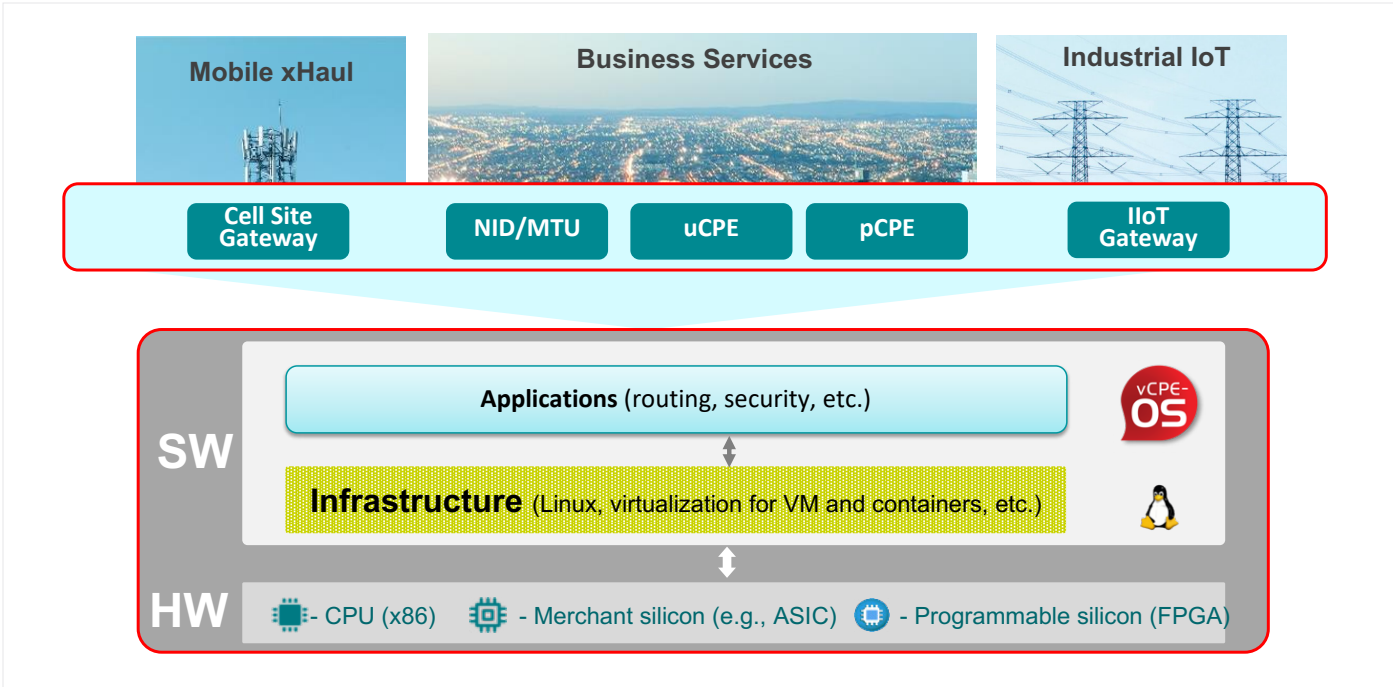

Figure 1. Benefits of disaggregation.[1]

Network disaggregation, and NFV in particular, opens the door for CommSPs to make their technology validation and procurement cycles more agile, and allows them to introduce new services faster, becoming less dependent on their traditional network equipment providers (NEPs). In addition, NFV is changing not only pricing models, but the entire supply chain, with CommSPs developing their own software.

Centralized vs. Decentralized Approaches

CommSPs can opt to locate their virtualized functions anywhere in the network, including at the customer edge, network edge, and/or cloud depending on cost, performance, or application requirements (see Figure 2). In cases where the virtualized functions are centralized, it is possible to reduce costs at the customer premises by utilizing lower cost hardware without NFVI since VNFs are not needed. This is typically referred to as pCPE, where "p" stands for programmable. pCPE differs from traditional appliances in that it typically shares the same management API and networking capabilities as uCPE to simplify its integration into the network.

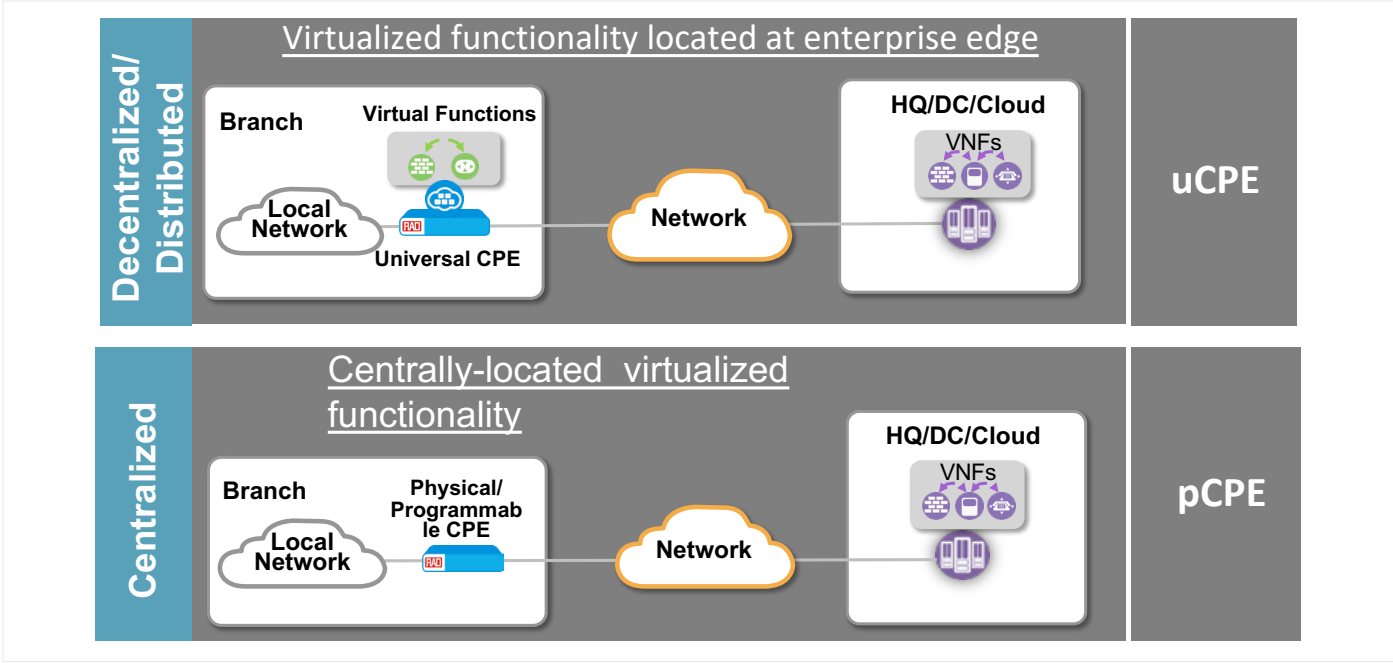

Figure 2. Comparison of centralized vs. decentralized approaches to virtualization.

RAD vCPE-OS Provides Flexible NFVI Platform

RAD has designed its vCPE-OS to facilitate mass vCPE deployment on generic white box hardware. This is achieved primarily by improving performance and built-in capabilities while minimizing compute/memory costs and operational complexity. Specifically vCPE-OS:

- Utilizes vector packet processing (VPP) to ensure high performance forwarding irrespective of packet size, thus reducing expensive cores while addressing future throughput requirements.
- Requires a very small memory footprint of only 2 GB to minimize compute/memory resources and costs.
- Incorporates essential networking capabilities, such as routing, NFVI, and service level assurance (SLA) that are needed in nearly all vCPE use cases.

Table 1 below provides more details on how RAD vCPE-OS addresses key mass deployment challenges.

CHALLENGES	RESOLUTION	RAD SOLUTION
Hardware Costs	Disaggregated Architecture	pCPE/uCPE hardware optimized to application requirements
Expensive RAM	Minimize OS resource consumption	RAD vCPE-OS consumes only ~2 GB RAM
Performance while minimizing CPU/Core costs	High performance forwarding plane with acceleration	VPP based forwarding with SR-IOV, Intel® QAT and FPGA acceleration options
Minimize VNFs required	Embedded capabilities at no extra charge	vCPE-OS incorporates enterprise router, performance monitoring, and an optional vFirewall VNF
VNF licensing costs	Pay-as-you-go, but also eliminate vendor lock-in	Simple VNF onboarding and replacement
Hardware selection and integration	Variety of pre-certified white boxes	vCPE-OS can run on any white box with easy hardware selection menu
Orchestrator selection and integration	Versatile API	vCPE-OS supports OpenStack REST API and NETCONF/YANG

Quick time-to-management with zero-touch provisioning (ZTP)	Domain orchestrator	RADview option for fast time to market with full lifecycle orchestration
Installation costs	Eliminate/minimize truck rolls	ZTP simplifying install/RMA and VNF onboarding processes
Troubleshooting	Performance monitoring and diagnostics	Active network, compute, and VNF PM and resizing
Brownfield support	Universal access options	RAD plugware includes TDM, GPON, CE2.0, Timing

Table 1. uCPE mass deployment challenges and solutions.

Part of RAD's Service Assured Access (SAA) solutions, vCPE-OS runs on any white box server and is pre-loaded in RAD's vCPE servers. vCPE-OS is designed with performance and functionality that can scale with the server's processing power.

vCPE-OS provides consistent "look and feel" with a variety of servers to ensure cost-effective deployment per use case:

1. Centralized VNFs: In a pCPE use case, vCPE-OS can provide high performance networking when deployed on a cost effective server with only two CPU cores powered by Intel Atom® processors.
2. Decentralized VNFs: Beyond networking, vCPE-OS can also feature NFVI designed to host and service-chain multiple VNFs on multi-core servers powered by Intel Atom or Intel® Xeon® D processors.
3. For very high performance applications, RAD vCPE-OS can make use of Intel® field programmable gate array (FPGA) hardware acceleration functionality built into servers that utilize Intel Xeon D processors.

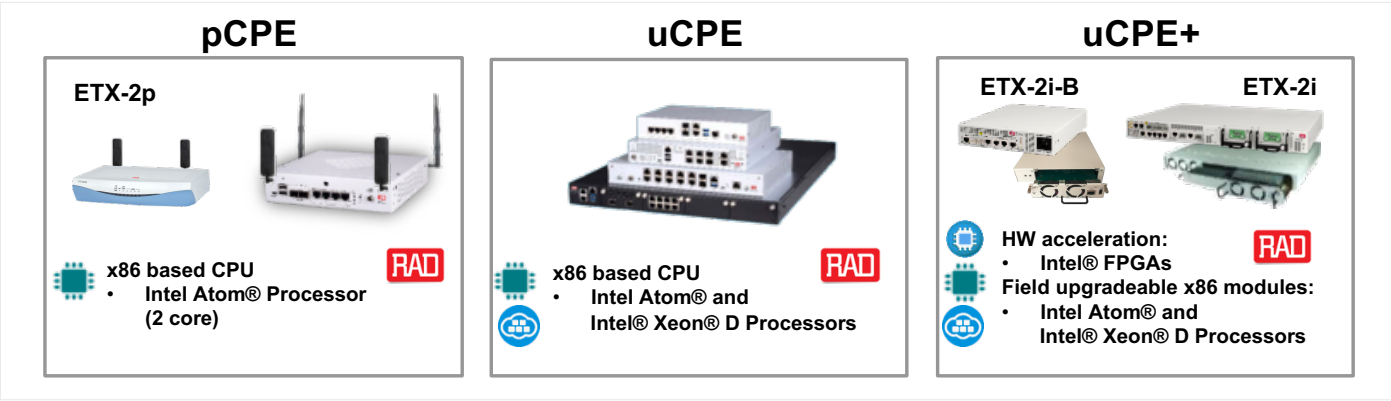

Figure 3. RAD vCPE-OS works with servers at a variety of cost-effectiveness levels.

RAD vCPE-OS vs OpenStack/OVS-DPDK

In order to maximize uCPE benefits, vCPE-OS offers a streamlined OS that replaces bulky OpenStack compute agent and controller without compromising support for OpenStack REST API. This approach maintains interoperability with existing orchestrators while providing the following performance and resource optimization benefits:

Bi-Directional Throughput: Eliminating performance bottlenecks is critical for white box viability, therefore:

- vCPE-OS provides high-performance routing and network address translation (NAT) in front of service function chains (SFC), so any VNF in the chain can be independently replaced—thus avoiding vendor lock-in.
- vCPE-OS provides high-performance VPP-based forwarding plane so that SFC throughput depends only on VNF performance.

Reduced RAM Requirements: RAM is a significant white box cost component, therefore:

- vCPE-OS streamlines the OS by removing irrelevant OpenStack components to significantly reduce its footprint.
- By freeing up RAM requirements, vCPE-OS maximizes resources available to host VNFs.

RAD vCPE-OS combines powerful networking capabilities with virtualization for hosting SD-WAN and any other value-added virtual network function (VNF) applications from any vendor. RAD's vCPE-OS is interoperable with open source management platforms, and easily integrates with standards-based SDN controllers, orchestrators, and operations/business support systems (OSS/BSS) from major providers.

Featuring a comprehensive management and security suite, vCPE-OS is unique in its convergence of key OS components of NFV infrastructure, including a KVM hypervisor, Open vSwitch, integrated router, and other embedded networking capabilities as well as integrated drivers for LTE and Wi-Fi. In addition, it enables seamless integration of RAD's plugware functions, such as operations, administration and management (OAM), packet timing, and support of TDM, DSL, and PON connectivity.

RAD vCPE-OS Specifications

- Comprehensive management suite: NETCONF/YANG, CLI, Syslog, alarms, and more
- Wide range of security tools: SNMPv3, SSH, SFTP, Access Control, TACACS+, RADIUS
- Integrated networking capabilities, more secure tunneling/VPN
- LTE, Wi-Fi support
- Performance monitoring: TWAMP, ICMP echo, UDP echo
- Single IP for networking and virtualization management
- NFVI monitoring: VM resources, reports, self and on-demand troubleshooting tools

Conclusion

Resource heavy NFVI can cause challenges for the network-wide deployment of uCPE-based services by restricting CommSPs to using expensive, high-performance uCPE servers in their network. RAD's alternative is RAD vCPE-OS which is a low footprint NFVI solution with built-in core networking functions that CommSPs can standardize on for a wide range of services. This platform simplifies mass deployment of uCPE, making it easier for CommSPs to offer a full range of edge-deployed virtualized services.

About RAD

As a global telecom access solutions and products vendor, RAD is committed to enabling service providers and critical infrastructure operators to evolve any service over any network using disaggregated architecture (DA). By keeping at the forefront of pioneering technologies and engaging in co-innovation with our customers, we strive to help service providers move up the value chain at a pace that is right for them, while offering their end-customers and network operators added value—be it in network edge virtualization and vCPE, Industrial IoT, or 5G xHaul. RAD is a member of the $1.5 billion RAD Group of companies, a world leader in communications solutions. More information is available at www.rad.com.

About Intel® Network Builders

Intel Network Builders is an ecosystem of infrastructure, software, and technology vendors coming together with communications service providers and end users to accelerate the adoption of solutions based on network functions virtualization (NFV) and software defined networking (SDN) in telecommunications and data center networks. The program offers technical support, matchmaking, and co-marketing opportunities to help facilitate joint collaboration through to the trial and deployment of NFV and SDN solutions. Learn more at http://networkbuilders.intel.com.

[1] Figures provided courtesy of RAD.