vCPE, Performance Assured

Maintaining Performance-Critical Edge Functions for vCPE-Based Business Services & Data Center Connectivity

The virtual CPE is a real, tangible approach to simplifying business services deployment, cutting costly CapEx and OpEx, and eliminating unnecessary equipment at the customer premise. This white paper explores the vCPE - what it is, the components that make it work, and practical deployment options.

The vCPE • An Overview

With the promise of lower cost and faster deployment speeds, business services and data center (DC) connectivity providers are seeking to employ Network Function Virtualization (NFV) and, in-turn, reduce and simplify the physical hardware in the network. With the largest proportion of network elements located at the edge, the biggest savings in capital and operational expenditures (CapEx and OpEx) will be at the customer premise (CP). Unlike best-effort Internet connections, performance-assured connectivity requires assured performance to meet SLAs, and to differentiate products in a QoS-savvy market. The customer demands the same performance regardless of how connectivity is delivered – the transition to NFV-based services can't compromise performance or QoS visibility.

As managed services may traditionally employ a number of appliances at the CP – a router, network interface device (NID), firewall, IP VPN appliance, SIP PBX, etc. – Virtualized Network Functions (VNFs) can provide a competitive edge by replacing smany of these with their virtualized equivalents. Simplifying the CPE saves capital as well as OpEx related to installation, maintenance, troubleshooting and orchestration between appliances.

Service providers often refer to the virtualization of functions performed at the CP as a 'virtual CPE' strategy. Although approaches and definitions vary, it's clear that on-site, provider-owned equipment won't completely vanish into the cloud. Some form of CPE is required to provide service demarcation, OAM and QoS mapping functions, and to deliver what 'merchant silicon' and software can't deliver: precise packet time stamping for latency and delay variation measurements, scheduling for seamless test-traffic generation, traffic filtering and policing. Although establishing QoS, service bandwidth and filtered flows can be shifted to an aggregation or edge router, pushing 'traffic reducing' functions towards the core puts unnecessary load on the network. Extra bandwidth is consumed not only because traffic is transmitted then discarded – any dropped TCP packets require re-transmission, adding to the problem.

vCPE: Measured Benefits

Early adopters of vCPE strategies have documented substantial gains in operational and capital efficiency. Moving to virtualized routing (VRF) has been shown to introduce 30% cost savings, while virtualizing all L4-7 network functions delivers over 70% savings in CPE infrastructure cost.

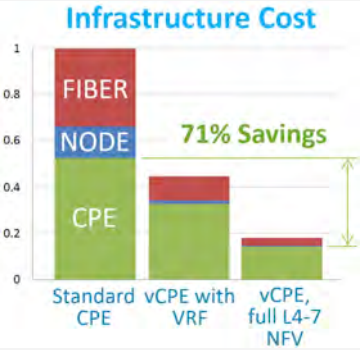

Source: Colt Technology Services, SDN & OpenFlow World Congress, Oct. 2014. Colt has used Accedian vCPE solutions since 2011, please see our press release at Accedian.com

The transition to NFV-based services can't compromise performance or QoS visibility

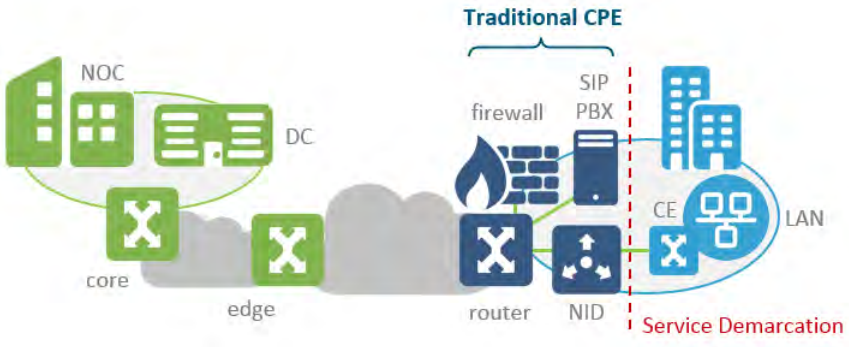

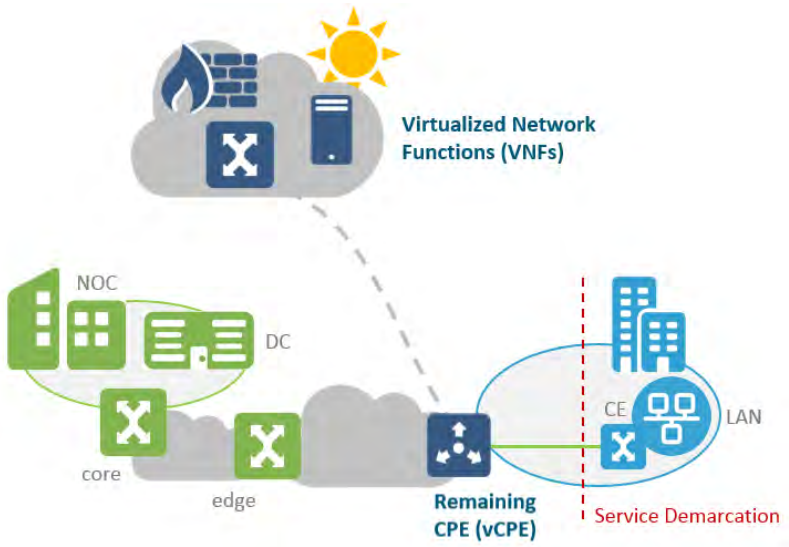

In general, the vCPE implies virtualizing as many CPE functions as possible, replacing racks of equipment with the minimum hardware required to establish, monitor and maintain the service. Areas targeted for virtualization include most layer 3-7 functions (e.g. NAT, routing / VRF, DPI, etc.). With these centralized in provider-owned NFV Infrastructure[1], the service provider gains significant economy of scale, simplified management, extensible shared compute resources, increased reliability, and many other advantages.

vCPE Architectures

The most common vCPE approaches seek to virtualize layer 3 networking and higher layer functions, separating the IP demarcation point from the layer 2 edge - which remains at the customer premise. Offloaded functions can be assigned to other network elements such as aggregation or edge routers, or may be deployed as Virtual Network Functions (VNFs) hosted in aggregation sites, x86 blades in core routers or centralized in data centers. With routing functions displaced from the customer premise, layer 2 broadcast domains extend directly to the customer equipment (CE); it's therefore critical that Ethernet OAM and layer 3 performance monitoring (e.g. TWAMP) extend to the vCPE to maintain last-mile visibility.

Ideally, the vCPE should also be able to map VLANs, set layer 2 & 3 flow priorities, filter and police outbound traffic to ensure performance-sensitive application requirements are respected between the vCPE and VNF-hosting locations (the NFV Infrastructure, NFVI). With VNFs handling many network functions, providers realize the benefits of accelerated, simplified service chaining, application deployment and policy enforcement. If the vCPE also includes automated provisioning and Service Activation Test (SAT) functionality, the provider can also benefit from truck roll-free customer self-install and automated SLA validation.

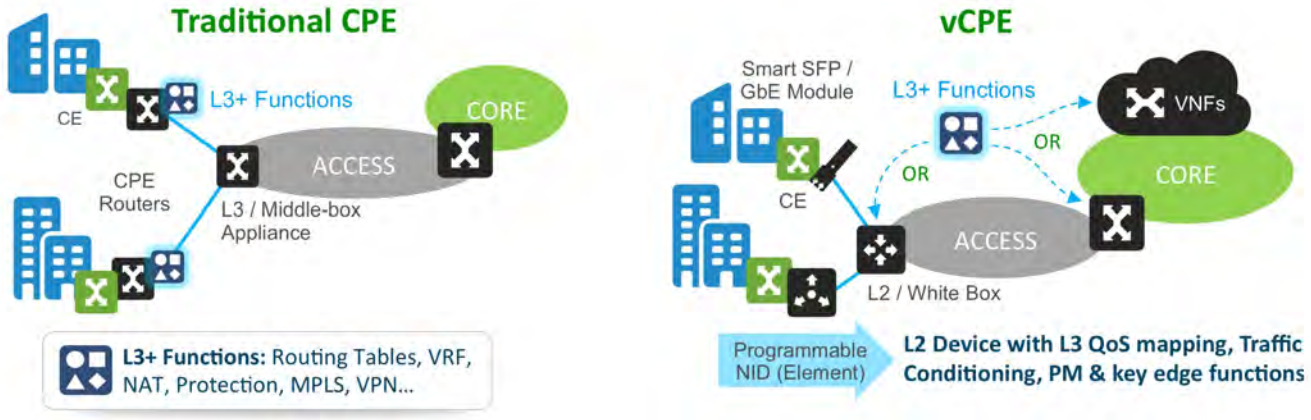

[1] NFVI: a datacenter or device capable of hosting virtual machines (VMs) and therefore VNFs)

Where to Virtualize?

NFVI Location & Performance Considerations

VNFs can be hosted pretty much anywhere a provider has equipment: from the vCPE itself to network elements or servers in the access, edge or core, to centralized, large-scale data centers. In general, the further the NFVI is from the customer, the higher the latency (and the further VNF-related traffic is carried through the network). As a result, the ideal of location for a given VNF will depend upon performance and telemetry requirements. As an example, a VoIP call control VNF is relatively latency insensitive so may be located in a data center without affecting quality of experience (QoE). Functions directly interacting with traffic (e.g. shaping, DPI-based routing) need to be conducted as close as possible to the CP.

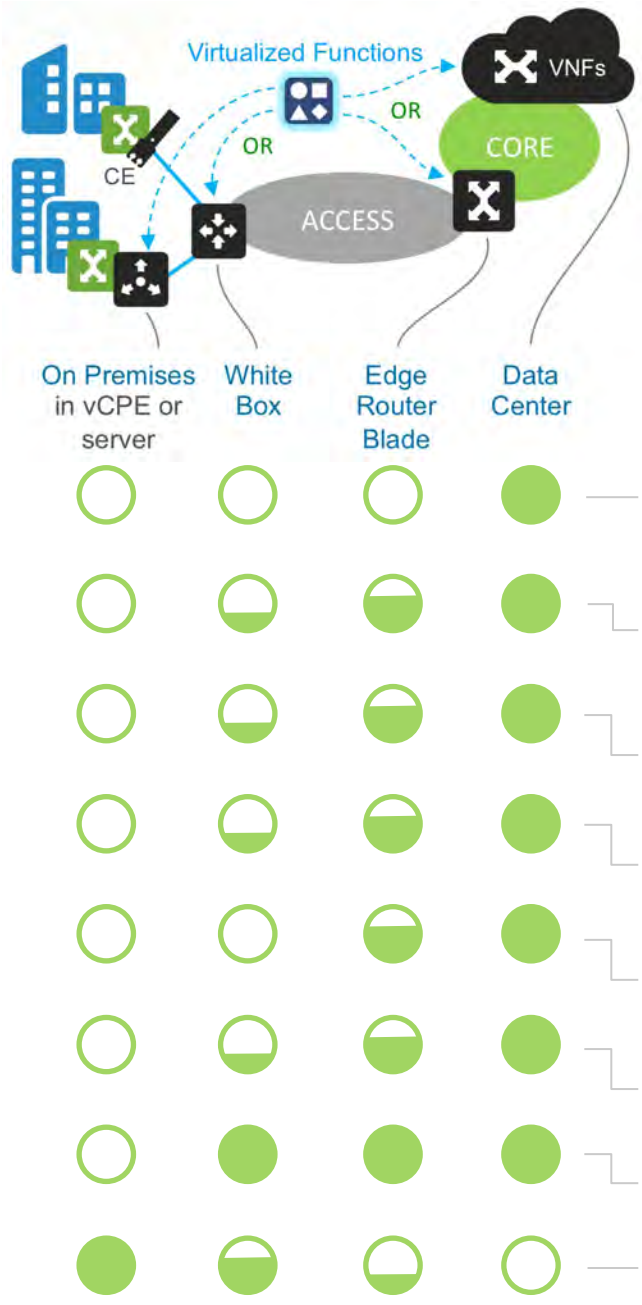

In general, most virtualized functions have little effect on service performance. Network Address Translation (NAT), firewall, and other session-based applications are well suited to centralized NFVI[2].

In General NFVI on premise is a step backwards vs. centralizing

Overall, it's a better choice to centralize NFVI in data centers and locations closer to the core, when application requirements permit.

When choosing an NFVI location, key considerations include:

Scalability: fixed compute power available in vCPEs, standard servers, or x86 NFVIs on blades in routers don't scale with Moore's law without hardware upgrades, favoring the elastic compute capabilities of data centers.

Capacity: If compute capacity is fixed, is there enough available for a given VNF? Will adding more VNFs to the same host affect the performance of others?

Security: physical access at the customer site is a liability, and deploying many NFVI nearer to customers are more difficult to secure and control than centralized NFVI.

Simplicity: fewer NFVI are easier to maintain, manage, upgrade and optimize. Chained services are also more responsive and easier to provision when their VNFs are co-located.

Reliability: large data centers offer full resiliency, load balancing, location diversity and many other aspects smaller NFVI can't offer.

Operational Expenditure: fewer truck rolls to fewer sites drives down cost, favoring the use of centralized resources when possible.

Footprint: hosting NFVI at a customer site contradicts many benefits of NFV, while unnecessarily consuming power and space.

Latency: VNF responsiveness – also dependent on compute speed and service chain path length

[2] See final page for a primer on security functions, including firewall localization

vCPE Options

From Programmable Elements to Virtualized Instrumentation

Accedian offers traditional network interface devices (NIDs, a.k.a. Performance Elements) and NFV-powered Modules to address common vCPE installations. Based on a programmable FPGA architecture, both physical and VNF-assisted form factors are ideal for low-latency services and precision SLA assurance, covering layer 2 and 3 demarcation functions, full support for layer 2 service OAM (Y.1731), layer 3 TWAMP (RFC-5357) performance monitoring and per-flow utilization metering. They also support integrated RFC-2544 & Y.1564 Service Activation Testing with multi-flow traffic generation, line-rate loopback, service test sequence automation and reporting.

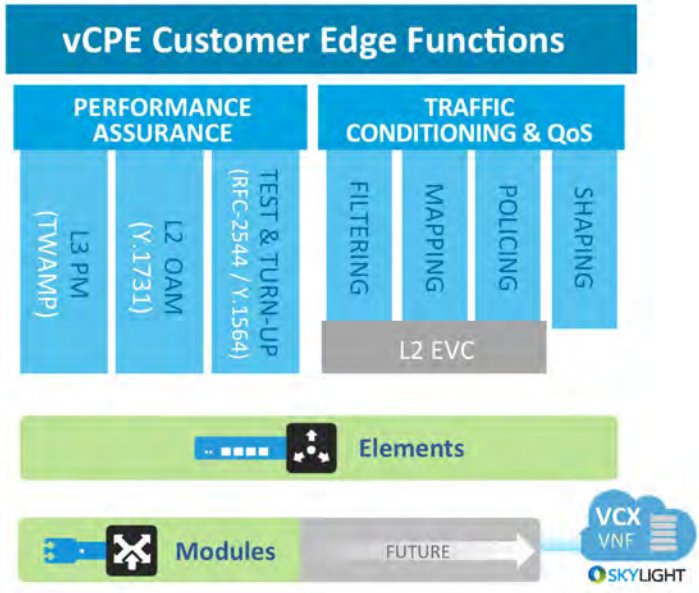

Network Performance Elements (NIDs), ideal for multi-tenant installations, also include layer 2-4 traffic filtering, QoS mapping, and MEF 10.3 certified hierarchical bandwidth policing (H-QoS). With multiport, switch-free aggregation, the industry's fastest micro-shaping and integrated remote troubleshooting tools, Accedian Performance Elements have been a preferred CPE solution for nearly a decade.

Nano smart SFP and GbE ant Modules employ NFV to offer the same suite of performance assurance features. When orchestrated by the SkyLIGHT™ VCX – a virtualized performance assurance VNF – Modules also offer layer 2-4 traffic conditioning functions. With the smallest possible footprint, Accedian Performance Modules are the most cost-efficient, performance-optimized vCPE solution for direct fiber and broadband business service connections.

A mix of Modules and Elements can be seamlessly managed, right-sized to each site's unique service assurance requirements. Both feature Plug & Go™ auto provisioning, realizing the highest-possible operational efficiency with customized workflow automation, tailored to a providers' existing deployment practices.

Step out of the virtual world – contact our solution experts to tailor a solution to your particular application.

Primer: Security Matters

Security is of increasing concern. The change to new network architectures invite the chance to revisit implementation of encryption and firewall functions. Although not necessariliy the role of vCPEs, security applications can seamlessly coexist. Here are some considerations for implementing these key functions as part of a vCPE roll-out.

Encryption Implementation & the vCPE

IPsec, widely available in the Linux Kernel, and included in the IPv6 standard, is commonly implemented at the customer premise using an IP VPN appliance, or on servers and desktops where traffic originates. IPsec is not generally supported by vCPEs. Being at the service demarcation point, the vCPE must be capable of processing full line-rate traffic. Wire speed IPsec encryption requires dedicated hardware – it cannot be implemented in a simple COTS / x86 module in the vCPE without introducing unwanted latency and significant packet loss under high-load conditions.

An alternative approach is to implement MACsec, layer 2 encryption over the 'first hop' into the access network, then apply IPsec thereafter if required. MACsec is simpler to implement, while offering the same Advanced Encryption Standard (AES) protection as IPsec. Being a lighter-weight protocol, a vCPE with an ASIC or FPGA capable of fill line-rate packet processing can perform the encryption without compromising performance.

Encryption can be performed on either side of the vCPE (service demarcation) - by the provider or the customer - without impairing vCPE functions. Traffic conditioning (H-QoS, flow regulation, filtering, priority and VLAN mapping), aggregation and performance monitoring functions only require packet header access. Traffic encrypted in 'transport mode', where only the payload is encrypted, provides full access to the header. When 'tunnel mode' encryption is performed (e.g. for IP VPNs) the entire packet is encrypted and a new header is added. If key traffic classification identifiers (e.g. CoS, DSCP, VLANs, etc.) are mapped to the new header, this more complete form of encryption is also transparent to the vCPE.

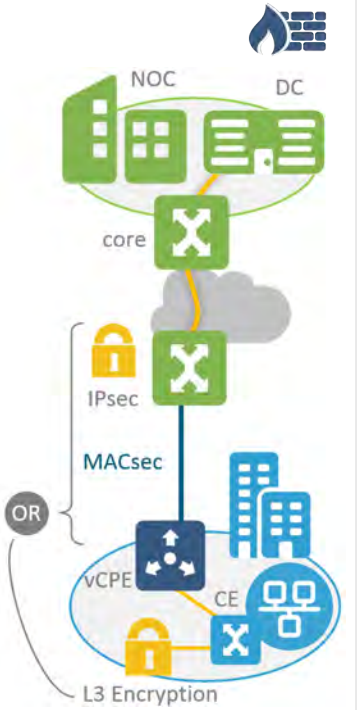

No Need for Firewall On-Site

Firewalls

Base firewall functions of identifying connection status, and allowing or denying access to a private network, are session based. Since sessions are end-to-end connections, blocking a session at any point along the packet path prevents the connection from being established.

As an example, if YouTube traffic is 'blocked', a user will be unable to play a video. The control packets required to start streaming (pressing the play button) cannot communicate with the video server, preventing the session from ever starting. Said another way, the firewall is not discarding full flows of YouTube packets as it encounters them, which would consume bandwidth over an access link – it simply prevents them from ever being transmitted to begin with.

The result that the location of the firewall is unimportant, counter to the misconception that it should be provided on-site as a 'last line of defense'.

A firewall is equally effective when implemented in a service providers data center, with an edge or core router, or any other convenient location. Firewalls take on a number of forms, sometimes grouped with more advanced higher-layer functions including load balancing, deep packet inspection and intrusion detection.

Like the firewall, these complementary applications can also be deployed as virtual network functions (VNFs). When these releated VNFs are hosted in the same datacenter, management, scalability and service chaining are simplified.