REFERENCE ARCHITECTURE

Service Provider
Network Infrastructure and NFV

# Optimizing Virtual Customer Premises Equipment: Medium Enterprise

**Making the right choices about architecting vE-CPE empowers communications service providers to increase business agility, improve manageability, and cut costs while still meeting the business needs of medium enterprises.**

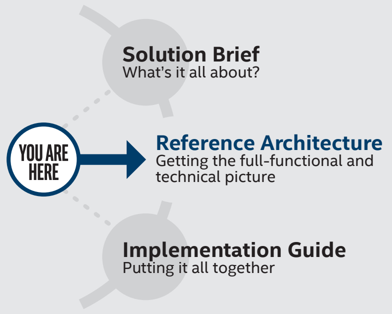

### What You'll Find in This Solution Reference Architecture

This solution provides a starting point for developing a vE-CPE infrastructure.

If you are responsible for:

- **Investment decisions and business strategy...** You'll learn how virtualization of customer premises equipment can help solve the pressing challenges facing CSPs today.
- **Figuring out how to implement vE-CPE...** You'll learn about the architecture components and how they work together to create a cohesive business solution.

## Executive Summary

Competition among communications service providers (CSPs) is fierce. Virtual-enterprise customer premises equipment (vE-CPE) is one way CSPs can react to market changes more quickly, improve the manageability of services and equipment, and minimize operational costs, thereby increasing their competitive edge.

vE-CPE can transform the connectivity landscape. But to unleash the power of vE-CPE, CSPs must understand the solution architecture and make the necessary decisions to build an appropriate architecture for their medium enterprise customers.

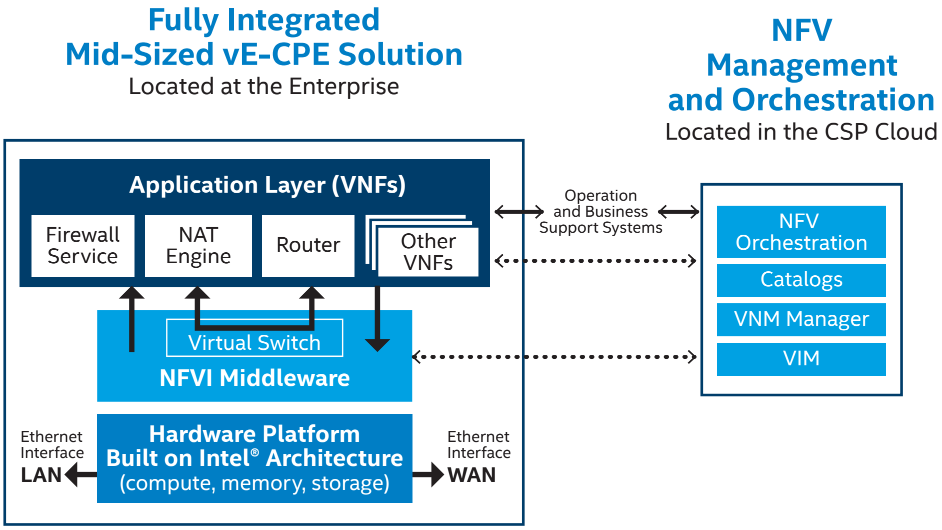

**Figure 1.** The vE-CPE solution architecture consists of hardware and interfaces, NFVI, the application layer (various VNFs), and management and orchestration.

## Table of Contents

**Executive Summary . . . . . . . . . . . . . . . . 1**

**Introduction . . . . . . . . . . . . . . . . . . . . . . . 2**

Overview of Network Function Virtualization . . . . . . . . . . . . . . . . . . . . . . . 2

Getting Started with vE-CPE . . . . . . . . . 3

**Solution Architecture . . . . . . . . . . . . . . . . 3**

Solution Overview and Benefits . . . . . 3

Solution Architecture Details. . . . . . . . . 4

Putting It All Together . . . . . . . . . . . . . . 8

**Architectural Design Considerations . . 8**

Security . . . . . . . . . . . . . . . . . . . . . . . . . . . 8

Initial Deployment of the Solution and Operational Support . . . . . . . . . . . 9

Out-of-Band Interfaces . . . . . . . . . . . . . 9

East-West Traffic Optimization . . . . . . 9

**Summary . . . . . . . . . . . . . . . . . . . . . . . . 10**

**Solutions Proven by Your Peers . . . . 10**

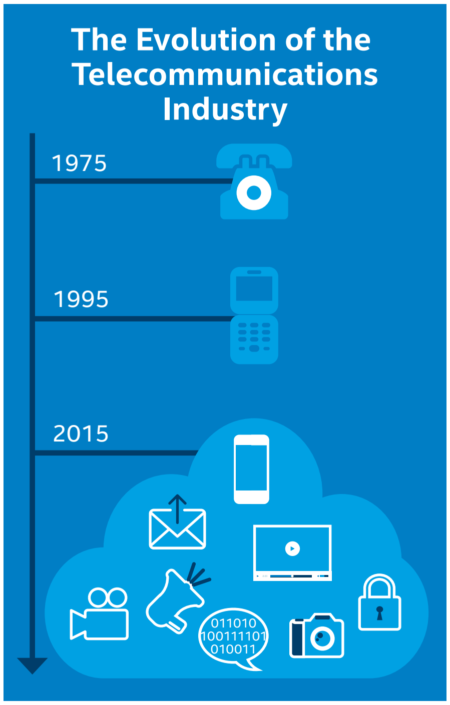

The Evolution of the Telecommunications Industry

## Introduction

Traditional approaches to enterprise customer premises equipment (E-CPE) using purpose-built solutions is expensive; deploying new services in a purpose-built environment can take more than a year—time that communication service providers (CSPs) cannot afford to take. To remain competitive and increase customer satisfaction, CSPs must digitally transform their businesses by establishing the following goals: quickly provision new services, improve the manageability of their services and equipment, and reduce costs.

### Overview of Network Function Virtualization

Recognizing the pitfalls of purpose-built hardware, forward-looking CSPs are achieving these goals through network function virtualization (NFV). The virtualized approach can reduce capital expenses through the use of standard high-volume servers (SHVS), also called commercial off-the-shelf hardware. Virtualization can also reduce operating expenses due to centralized management and orchestration (MANO). Intel® technology combined with technology from a wide variety of industry participants provides a foundation for virtual E-CPE (vE-CPE) that increases agility, simplifies management, and cuts costs for both CSPs and their customers. With industry-standard hardware and a wide selection of software applications that perform an array of virtual network functions (VNFs), CSPs can avoid vendor lock-in and significantly improve time to resolution when issues arise.

By decoupling the hardware and software components with a vE-CPE solution, CSPs can easily and affordably accommodate varying workloads and new services through network agility, flexibility, and scalability. Each of these workloads is a different VNF, such as deep packet inspection, firewall services, WAN acceleration, or a VPN.

All hardware and software components that build the environment where VNFs are deployed is referred to as the NFV infrastructure (NFVI). The NFVI can span several locations; the network providing connectivity between these locations is considered part of the NFVI. Figure 2 illustrates the typical VNF architecture.

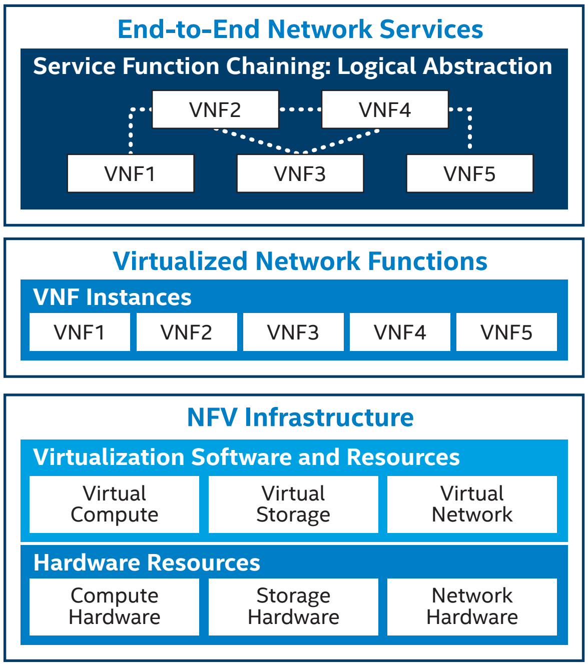

**Figure 2.** NFV decouples the hardware and software, enabling a flexible and scalable architecture that supports the quick launch of new services.

### Getting Started with vE-CPE

To unleash the business acceleration enabled by vE-CPE, CSPs must make informed decisions about platform design, software selection, integration of multiple VNFs, and maintenance and deployment strategies.[1] This reference architecture presents a proven vE-CPE architecture, identifying a range of architectural choices available to CSPs. These choices help CSPs build a vE-CPE solution optimized for a medium enterprise.

The following section describes each of the solutions' components and provides an example of a completed architected solution. Additional sections discuss architecture design considerations that CSPs should take into account as they begin to plan the vE-CPE architecture.

## Solution Architecture

vE-CPE solutions enabled by Intel technology and supported by components from industry participants give CSPs the ability to optimize their vE-CPE services to the specific needs of their customers. As shown in Table 1, vE-CPE solutions can be targeted for small, medium, and large customers using the appropriate network interface controllers and processors. This reference architecture focuses on the medium customer (50 to 500 users).

### Solution Overview and Benefits

The mid-sized vE-CPE solution is based on a foundation consisting of a mid-range hardware platform combined with middleware and an application layer (the VNFs) that are well suited for a branch office or other medium enterprise with 50 to 500 users.[2]

Sizing a vE-CPE solution is not based solely on the customer's number of users. For example, a small enterprise with five users may have a business need for five or more VNFs—a medium vE-CPE solution would be suitable in this scenario.

When sizing the architecture described in this document, Intel Solution Architects used traffic profiles and packet sizes that replicate the standard Internet traffic load and performed boundary tests for minimum and maximum packet sizes. Specifically, workloads found in the mid-sized enterprise consist of the following:

- VoIP traffic
- General Web traffic for browsing, email, and video streaming
- Corporate IT functions such as software maintenance and backup

**Table 1.** Small, Mid-Sized, and Large Options for vE-CPE Solutions

| | SMALL | MID-SIZED | LARGE |
|---|---|---|---|
| | **Intel® Atom™ Processor C2000 Product Family** | **Intel® Xeon® Processor D Family** | **Intel® Xeon® Processor E5 Family** |
| **WAN** | <100 Mbps<br>4 core | 100 Mbps to 1 Gbps<br>4–8 core | 1 Gbps+<br>8–16+ core |
| **LAN** | 1 Gbps | 1–10 Gbps | Multiple 10 Gbps |
| **User** | 10–50+ | 50–500 | 500+ |
| **Typical Number of VNFs** | 2–4 | 3–5 | 4–8 |

## Making the Right Decisions

The following questions summarize the key decision points CSPs should consider when designing a mid-sized vE-CPE solution. How to make these decisions is discussed in detail in the rest of the document.

### Network Interfaces

- Are two or four 1-Gbps interfaces needed?
- Is an OAMP (operations, administration, maintenance, and provisioning) interface required?

### NFVI Middleware

- Is the middleware open source, a commercial offering, or a reference base?
- What are the throughput performance requirements for the data plane?

### Application Layer

- Will VNFs come from a single vendor or from multiple vendors?
- Which VNFs are required?
- What VNF attributes are important (scalability, supportability, reliability, availability, and so on)?

### Hardware Platform

- What level of computing performance is needed?
- How much memory and storage is needed?

### Management and Orchestration Layer

- Is it interoperable with all the other layers?

### System Integration

- Who will put it all together?
- What is the level of subject matter expertise within the CSP's Operations and Field teams?

[1] This reference architecture assumes that the communication service provider (CSP) owns and operates the customer premises equipment (CPE). While it is possible for the enterprise to provide its own customer premises equipment and deploy the CSP's virtual network functions, we believe that the operational considerations of either the enterprise neglecting the service-level agreement (SLA) or the CSP having to accept an SLA for platforms that they have limited control over outweigh any potential benefits of enterprise-owned CPE for medium enterprises.

[2] While enterprises with fewer than 50 staff may also use the architecture in this solution, the initial hardware cost and infrastructure capability may be excessive and should be considered carefully. The small vE-CPE is based on a more economically priced Intel® Atom™ processor, which is generally more suitable for the business needs of smaller enterprises.

As illustrated in Figure 1, the vE-CPE solution architecture consists of the following primary building blocks (bottom-up view):

- Hardware platform (consisting of compute, memory, and storage)
- Network interfaces suitable for E-CPE workloads
- Appropriate NFVI middleware to support the core functionality and management of the platform, including a virtual switch (vSwitch)
- Application layer that comprises the necessary VNFs
- MANO layer

The first four of these components are located on the customer premises; the MANO layer (which is not officially part of vE-CPE solution but is necessary for resiliency and system configuration) is located in the CSP cloud.

Of course, once all the components are assembled, they have to be integrated. For a discussion of the various approaches to this task, see Putting It All Together.

## Solution Architecture Details

This section describes the various components—network interfaces, NFVI middleware, applications layer, hardware platform, and the management and orchestration layer—in more detail.

### Network Interfaces

As shown at the bottom of Figure 1 , the solution must have a minimum of two physical Ethernet interfaces: one facing the WAN and the other facing the LAN. By design, both the WAN and LAN interfaces are constrained to a maximum line rate of 1 Gbps.

Some vE-CPE environments may also require a separate OAMP (operations, administration, maintenance, and provisioning) interface. Intel offers solutions that can easily accommodate either two or four 1-Gbps interfaces in the mid-sized vE-CPE solution. If you need more than just WAN and LAN ports, or need a line rate higher than 1 Gbps, consider upgrading to a large vE-CPE design (see the sidebar Is It Time to Upgrade?).

### NFVI Middleware

Every vE-CPE solution includes VNFs that are coresident on a single host. Therefore, the NFVI middleware must be able to support the traffic flows between these VNFs (see Improving NFVI Performance) and the maximum (bidirectional) throughput of the data from the WAN to the LAN. Also affecting the NFVI middleware choice is whether the CSP is able to manage large-scale open source projects and possibly contribute to the code base, or whether it needs to rely on the ecosystem for system integration. Another consideration is the modularity of the VNFI. Technology is constantly changing, and infrastructures that can adapt to newer virtualization technology while still supporting legacy approaches can position a CSP to thrive in an ever-changing virtualized world.

#### Choosing an NFVI Middleware Provider

We recommend that a CSP with suitable skills or one that chooses to work through a system integrator use the current Liberty release (or a future release) of OpenStack*. The benefits of using OpenStack are two-fold:

- Its extensive and collaborative OpenStack community provides enhancements at a fast-paced cadence.
- Its open source nature can help reduce licensing fees, thereby decreasing the total cost of ownership of the entire vE-CPE solution.

Another NFVI middleware open source option is the Open Platform for NFV Project* (OPNFV*), introduced by the Linux* Foundation.

While this reference architecture makes specific design recommendations based on integration and test results, we also recognize that there are other options available, including the following:

- Commercial offerings, which include proprietary products such as VMware, Red Hat, or WindRiver, and “pure play” products such as Mirantis
- A reference base such as Intel® Open Network Platform (Intel® ONP)

The selection of NFVI middleware is unique, based on each CSP's resources, trained staff, and history with vendors and technology. There may also be business drivers such as commercial terms favorable to one choice over another. Each NFVI provider has strengths in areas that separate them from their competition. Table 2 provides a few examples of NFVI vendors; your Intel account team can provide further assistance, and the Intel® Network Builders program is also an excellent resource.[3]

### Improving NFVI Performance

In finalizing the choice of NFVI middleware, CSPs should consider the performance required for the data plane. Our lab tests, using simulated loads, have shown that maximum throughput is achieved by using an NFVI accelerated by Open vSwitch* (OVS) enabled with the Data Plane Development Kit (DPDK). For a detailed discussion, see East-West Traffic Optimization.[4]

While using OVS and the DPDK improves performance, it does require specific allocation of cores within the compute platform. Fundamentally there are two models of operation within the DPDK: run-to-completion and pipeline. We recommend using the pipeline model for vE-CPE solutions. More details on the DPDK can be found in the Intel Network Builders training materials or on the DPDK website.

Our lab tests also show that CPU pinning improves performance compared to implementations where the host OS is free to context-switch tasks between cores. With CPU pinning, each task (VNF) should be allocated to a fixed core. For example, in Figure 3, VNF1 and VNF2 are pinned to core 2, while VNF3 and administrative tasks are pinned to core 3.

CPUs that support Intel® Hyper-Threading Technology can be found in many 6th generation Intel® Core™ processors. These CPUs allow up to two workloads to be pinned to each core.

### Applications Layer

VNFs are at the heart of the vE-CPE solution. At a minimum, a vE-CPE solution will include the following VNFs: a firewall service, a Network Address Translation (NAT) engine, and a router. Additional VNFs will be required depending on the specific services being offered by the CSP to the customer.

[3] Inclusion in Table 2—or elsewhere in this document—does not indicate endorsement of any vendor by Intel. Many other choices are available; it is impractical to attempt to list them all.

[4] For detailed information about how Open vSwitch can enhance network performance, read "Open vSwitch* Enables SDN and NFV Transformation," at https://networkbuilders.intel.com/docs/open-vswitch-enables-sdn-and-nfv-transformation-paper.pdf.

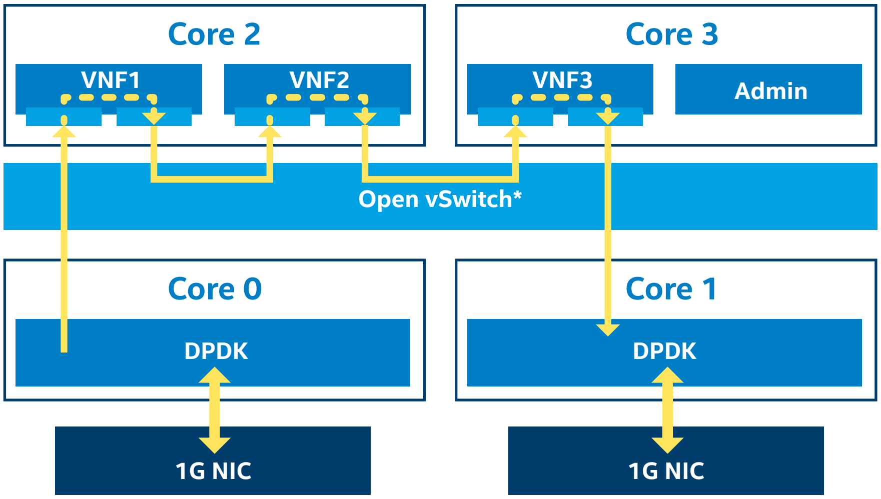

**Figure 3.** Pinning tasks to a specific CPU core improves performance.

**Table 2.** Examples of NFVI Suppliers

| Vendor | Notable Attribute |
|---|---|
| **Red Hat** | A common hypervisor provider |
| **VMware** | Experience with large commercial deployments |
| **Intel® Open Network Platform** | Reference architecture |
| **Open Platform for NFV*** | Integrated projects |
| **OpenStack*** | Extensive cloud support |
| **Mirantis** | Pure-play commercial offering |
| **WindRiver** | High availability features for the Teleco industry |

There are many choices when it comes to choosing the VNFs. When designing the system, CSPs must decide whether to choose integrated applications from a single vendor or assemble the end-to-end functionality from multiple vendors. Complete solutions from Cisco, Juniper, and Huawei, and other suppliers provide a vE-CPE solution out-of-the-box.

The service function chaining of VNFs provides the maximum performance but chaining VNFs from multiple vendors requires close collaboration with these vendors and a greater level of system integration. If the VNFs are sourced from one vendor, the system integration work may be reduced (see Putting It All Together), but this raises the risk of vendor lock-in at this critical solution layer. CSPs must balance the risk of vendor lock-in against simplicity and significantly accelerated time to market.

Which VNFs to include in the solution is primarily defined by the CSP's customer's current and future needs. Beyond that, CSPs should do careful research before choosing a VNF.

Security is a key aspect of VNF choice. The options are to place the vE-CPE workload (in particular, the VPN and firewall) on-premises or in the cloud. CSPs should soundly vet any cloud-located workloads, paying careful attention to any possible negative business impacts to the enterprise customer's governance or security requirements.

Reliability is another crucial consideration. When selecting a VNF, CSPs should assess its reliability and stability (lack of defects), along with the required level and ease of integration into the whole solution. Testing is often considered the process of verification and validation (is the function right and is this the right function)—but testing alone is not a measure of reliability. The quest is to establish a level of trust that is not quantifiable or measurable but results in assurance that the workload will perform as expected throughout a long life cycle.

Some additional aspects to consider include the following:

- **Scalability.** Can the VNF scale both up and down? Can you add capacity to a single VNF, or do you need to add capacity to all VNFs at once?
- **Supportability.** Does the VNF provider have an efficient software design? Does the provider have a reputation for responding quickly to issues and requests?
- **Availability.** Is a high-availability design required for current or future service offerings? Does the VNF feature a high-availability design when needed?

In addition, investigate the various technologies leveraged by the VNF. Network virtualization is an ever-evolving field, and many times it is best to choose VNFs that use advanced virtualization techniques such as containers and dynamic service function chaining.

### Hardware Platform

The workloads—that is, all the VNFs and the necessary throughput required of them—combined with the expected number of users and sessions help define the selection of the CPU and memory requirements for a vE-CPE solution. In addition to the VNF workloads, CSPs must also consider the workloads of the infrastructure (management interface, traffic processing in and out, and internal switching).

#### Choosing a CPU

Here's an example of the decision process over how many cores are required for a vE-CPE solution:

When selecting the CPU, the CSP should seek to maximize the performance while minimizing the overall system cost. As discussed earlier, CPU pinning reduces wasteful context switching during service chaining. The assumption therefore is that a specific core exists for each pinned workload. Since each core can sustain two threads, we can make the following recommendation:

- One core for each of the VNFs (assuming the three minimum VNFs previously discussed and also assuming CPU pinning)
- Two cores for the OVS and DPDK (one for each direction)
- One core for the OMAP interface
- One core for the host OS (recommended, but not required)

That's a total of seven tasks. Because hyper-threading is available on the reference architecture's recommended processors, the minimum number of cores on a CPU for this design requires a four-core processor such as the Intel® Xeon® processor D-1520. If high-end performance is required or additional VNFs may be added in the future, we recommend choosing an eight-core CPU such as the Intel® Xeon® processor D-1540. This processor is available in several configurations depending on power budget, core count, and interface options.

### Sizing the Memory

An optimized solution will appropriately size the memory to maximize the system capabilities and yet avoid over-purchasing. Our testing shows that for medium enterprise CPE workloads, 16 GB of memory is sufficient.

### Ensuring Adequate Storage

Storage typically is not a significant concern for vE-CPE solutions. 80 GB of solid-state drive (SSD) storage is usually adequate for the host OS, the NFVI middleware, and any VNFs, along with a reasonable amount of working storage for log files. We recommend using the Intel® Solid State Drive Data Center S3510 Series. SSDs are preferred over hard-disk drives for a number of reasons, such as ruggedness, access speed, and reliability. The removal of one spinning component reduces the number of physical parts that require mechanical movement to operate, thus removing one concern of possible physical failure.

### Sample Hardware Platform

Figure 4 shows a sample hardware platform for the medium vE-CPE solution. The platform assumes a 1U ("pizza box") form factor with 16–32 GB DRAM. As shown, the Intel® ANYWAN™ network processor series, or optionally the Intel® VINAX™ WAN interface components, can be added to support the desired WAN interface depending on the target deployment. Out-of-band (OOB) and legacy FAX (FXT) interfaces are optional. Security options include UEFI (Unified Extensible Firmware Interface) Secure Boot, Intel® Trusted Execution Technology, and IPsec software for the Intel® Communications Chipset 89xx Series.

## Management and Orchestration Layer

The MANO functions (shown on the right side of Figure 1) reside in the CSP cloud. Strictly speaking, the MANO functions are not actually part of the vE-CPE solution. Once the service is running in a steady state the MANO layer could theoretically be disabled—at least, until something breaks or the customer wants something reconfigured or changed.

But practically speaking, the MANO functions are necessary to the ongoing health of the vE-CPE solution, and are therefore discussed briefly here. Choosing products for MANO requires considering all the usual decision points: budget, functionality, interoperability, scalability, and so on. The following primary functions should be included:

- **NFV Orchestration** addresses capabilities for instantiating, managing, and chaining VNFs.
- **Catalogs** include the Service Catalog, Element Catalog, and VNF Catalog.
- The **VNF Manager (VNFM)** oversees life-cycle management of VNF instances and fills the coordination and adaptation role for configuration and event reporting.
- The **Virtualized Infrastructure Manager (VIM)** controls and manages the NFVI compute, storage, and network resources.

Many vendors offer MANO functions, so be sure to choose products that are interoperable with all the other components of the vE-CPE solution.

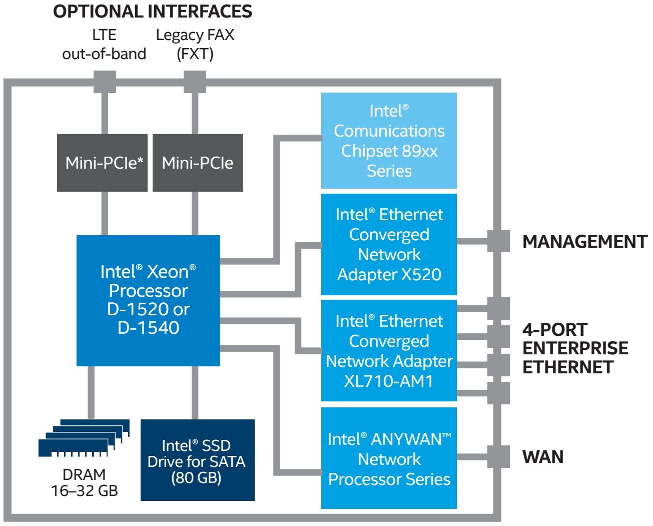
OPTIONAL INTERFACES
LTE out-of-band
Legacy FAX (FXT)
Mini-PCIe*
Mini-PCIe
Intel® Comunications Chipset 89xx Series
Intel® Ethernet Converged Network Adapter X520
Intel® Xeon® Processor D-1520 or D-1540
Intel® Ethernet Converged Network Adapter XL710-AM1
Intel® ANYWAN™ Network Processor Series
DRAM 16–32 GB
Intel® SSD Drive for SATA (80 GB)
MANAGEMENT
4-PORT ENTERPRISE ETHERNET
WAN

**Figure 4.** The complete medium vE-CPE solution's hardware platform accommodates LAN and WAN access, storage, and more.

### Putting It All Together

The previous sections have discussed the decisions associated with various components of a vE-CPE solution, such as network interfaces, compute and storage hardware, NFVI middleware, and applications (VNFs). CSPs may choose one of several approaches to make these decisions and then assemble all the components into a cohesive system. Deciding on which approach to choose can be affected by the CSP's regulatory environment, customer input, and business opportunities.

- **Engage with a well-known system integrator.** Many members of the Intel Network Builders program offer full-service, turnkey system integration services.
- **Partner with a hardware vendor.** The Intel Network Builders Program offers a wide choice of vendors.
- **Do-it-yourself.** CSPs with a strong background in system integration may choose to perform their own system integration. Note that the do-it-yourself approach requires a mature and dedicated internal support staff that will also provide the end-to-end support of the solution with minimal vendor assistance. In this case the CSP must choose an SHVS, including any necessary WAN interfaces to meet the needs of the physical connection at the WAN. This SHVS must, at a minimum, meet the specifications discussed in Choosing a CPU.

No matter which choice is made, the expectations of the end result are similar: a fully operational and fully supported vE-CPE solution that meets the needs of both the CSP and the customer.

## Architectural Design Considerations

Besides the overarching decisions about interfaces, hardware, NFVI middleware, VNFs, and MANO functions, there are other considerations when designing a vE-CPE solution. Not unlike purchasing a new car, there are standard configurations available, as well as various levels of options. With cars, some of these options are related to business drivers—a hybrid electric/gas option, for example. Other options are dictated by customer desires, such as a sunroof or fancy wheels. For vE-CPE, the same is true: some options relate to business concerns, while others relate to CSPs' operational differences.

The following sections discuss architectural design considerations relating to the following areas:

- Security
- How the solution is initially deployed at the customer's premises and then upgraded and maintained over time
- OOB interfaces
- Optimization of east-west traffic

### Security

This reference architecture provides a broad and general solution. However, we realize that security is an area that is particularly sensitive to specific CSP requirements. CSPs with high-performance security requirements may benefit from the addition of an optional accelerator to perform IPsec encryption and decryption. In these cases Intel® QuickAssist Adapters for servers may be an ideal addition to the system design. CSPs with less stringent security requirements probably will not need such an addition.

While using an accelerator is not part of the "pure" SHVS vision, practical considerations demand that CSPs address security in a way that meets their business and technical goals.

## Is It Time to Upgrade?

While the mid-sized vE-CPE reference architecture described in this document reflects the needs of a typical medium enterprise, there are situations where a higher-end vE-CPE solution may be advisable.

### Number of VNFs

The customer's number of users is not the only criterion that determines the sizing of a vE-CPE solution. The number of VNFs also comes into play. The typical medium enterprise uses two to five VNFs. If an enterprise has only 400 users but needs access to more than five VNFs, a large vE-CPE solution would be suitable.

### Ports and Throughput Requirements

Customers that desire additional WAN and LAN ports or need a line rate that exceeds 1 Gbps should consider stepping up to the large vE-CPE solution design. The Intel® Ethernet Controller I350 Family provides options for dual and quad interfaces in one package and supports line rates up to 1 GB/s.

## Initial Deployment of the Solution and Operational Support

One of the business drivers for virtualization of CPE is to reduce maintenance truck rolls. To reach this goal, a process must be in place to initialize the system, restart it when necessary, upgrade it occasionally, and manage the VNFs and the platform as a whole. There are several options:

- **Customer provides the bare metal hardware and handles initialization.** This approach may be taken in situations where the customs of the local jurisdiction are such that CSPs do not supply the CPE or when a multinational CSP may choose to locally source the SHVS to accelerate time-to-service initialization.

  Some medium enterprises may have technical staff suitably skilled at performing a PXE (preboot execution environment) boot over the LAN to initially install the software. However, CSPs cannot assume this is always the case, especially if they expect to scale the system to a large customer base. Also, the ability to download the entire operational image from the WAN may be constrained for a maiden platform. For example, the customer may not have a channel to accomplish this task initially. In general we recommend that CSPs avoid this approach and use one of the following options instead.
- **CSP or a third party provides a preinstalled bootable image.** In this case, the CSP provides the SHVS with all software, host OS, and NFVI middleware already installed. This approach makes initializing the system more manageable by the customer's non-technical staff, and it includes providing the native OS as part of the bootable image.
- **CSP provides a partial image.** We recommend a middle-ground approach, where the system arrives onsite with a minimum bootable image capable of reaching the WAN to pull any system updates. The vE-CPE solution can then access any additional images or updates on first connection to the WAN to complete the software image loading.

## Out-of-Band Interfaces

Some CSPs will require OOB management using a wireless interface, such as a 4G LTE (Long Term Evolution) interface. To accommodate this interface, the bare metal must have the ability to support a wireless modem and SIM card. The Intel® XMM™ 7160 slim modem provides LTE/4G/3G/2G access and may be an ideal addition. Living up to the name "long-term evolution," in the near future carriers will be offering 5G access, so be sure to check with your Intel account team for the latest information on 5G.

If the CSP is using 4G, the system will need to be provisioned through a source at the carrier. This option comes with some additional complications. For example, the system may be located in a RF-free or significantly diminished RF signal zone. CSPs can easily test for this by having the customer verify the availability of a sufficiently strong signal through the use of a mobile device on the same carrier network in the intended space. The customer can simply make a voice call and test Web browsing from the same physical location using a 4G LTE handset on the carrier's network on which the CPE is to be installed. If a more advanced verification is necessary, the customer can perform a speed test using one of the many available resources such as Ookla.

## East-West Traffic Optimization

When multiple VNFs are coresident on the platform there is a significant amount of traffic that flows from VNF to VNF. This "east-west" traffic by design is chained together to provide the service, which is called service function chaining. State-of-the-art VNFs use service function chaining at Layer 2 within the IP switching.

One option would be for all VNF traffic to exit the platform and be reflected back by a physical Top of the Rack (TOR) switch to the same physical port with an IP destination of the next VNF. But this approach is not optimal because since the traffic is east-west within the CPE, it makes sense that this traffic should remain within the CPE as it moves from VNF to VNF.

A better approach is to use a vSwitch, such as OVS. Many NFVIs today support some level of OVS. Using a vSwitch, traffic that is chained from one VNF to another VNF on the same physical machine remains within that machine and does not egress to a TOR only to be reflected back to the originating port. This approach reduces complexity and improves performance. In fact, our testing has shown that OVS accelerated by DPDK provides a near 12x performance increase compared to native OVS. Therefore, choosing NFVI middleware that supports OVS and DPDK will yield better performance.

## Summary

Technology, especially the field of virtualization, epitomizes the adage “The only thing that is constant is change.” Therefore, this reference architecture provides a snapshot in time of the relevant technology and industry players in the field of vE-CPE. Like a game of leapfrog some players may be “in the air,” and this snapshot will certainly change over time.

That said, the recommendations given here represent the best-known practices for CSPs seeking to optimize their vE-CPE solutions for medium customers. These best-known practices range from choosing the right network interfaces, compute resources, NFVI middleware, VNFs, and MANO functions, to deciding how to handle system integration.

## Solutions Proven by Your Peers

vE-CPE empowers CSPs to increase business agility, improve manageability, and cut costs while still meeting the business needs of medium enterprises. This and other solutions are based on real-world experience gathered from customers who have successfully tested, piloted, and/or deployed these solutions in specific business use cases. Solution architects and technology experts for this solution reference architecture include:

- **Larry Horner,** Solution Architect, Industry Sales Group, Intel Corporation
- **Ron Whitfield,** Marketing Manager, Sales and Marketing Group, Intel Corporation

Intel Solution Architects are technology experts who work with the world’s largest and most successful companies to design business solutions that solve pressing business challenges.

To find the best solution for your organization, contact your Intel representative or visit **intel.com/communications**.

## Learn More

This solution reference architecture complements product documentation and is part of an entire solution kit of content that is full of key insights and learnings:

- **Solution Brief:** “Realizing the Value of Virtualizing Customer Premises Equipment”

You may also find the following resources useful:

- **Intel® Network Builders:** networkbuilders.intel.com
- **Intel® Open Network Platform and open standards:** intel.com/content/www/us/en/communications/network-infrastructure-open-source-open-standards.html

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

Intel technologies’ features and benefits depend on system configuration and may require enabled hardware, software, or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer, or learn more at intel.com.

Software and workloads used in performance tests may have been optimized for performance only on Intel® microprocessors. Performance tests, such as SYSmark* and MobileMark*, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit intel.com/performance.

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

0716/JWAI/KC/PDF ♲ Please Recycle 334200-001US