Solution Brief
Artesyn Embedded Technologies*
Telecom Solutions
Intel® Xeon® Processors

Different NFV/SDN Solutions for Telecoms and Enterprise Cloud

Networking solutions from Artesyn Embedded Technologies* and Intel address the rigorous requirements of the Telecom industry.

Introduction

Network functions virtualization (NFV) was envisaged to bring enterprise cloud concepts to the telecom world. However, one challenge is that enterprise class servers, so ubiquitous in the enterprise cloud, do not provide the compute density, I/O bandwidth, and carrier-grade 'hardness' required by many telecom applications. Addressing these requirements, Artesyn Embedded Technologies*, working with Intel, provides the carrier-grade-class server platforms needed to build an NFV node in a true telecom environment.

Addressing Service Provider Challenges

The telecom industry is at an inflection point, especially since the massive growth in subscriber data usage requires service providers to make substantial infrastructure investments. But since subscriber revenue is not growing at the same rate as data usage, simply adding more capacity and connectivity to their networks is not a viable option.

In addition, service providers face a challenge from over-the-top (OTT) content providers who supply services to end users without paying to use carriers' network resources. Again, there is no direct correlation between the volume of data transferred by OTT players and the revenue collected from subscribers.

On top of this, service providers must be very agile to compete with the OTT players serving the same subscribers, whose expectations are evolving, leading them to demand new content and services as quickly as possible.

For traditional telecom service providers, these challenges are converging and driving them to strongly consider networking solutions based on NFV and software-defined networking (SDN), which can help improve time to market for new services while lowering TCO,

NFV-based solutions *virtualize* network resources (e.g., compute, connectivity, storage, and supporting services) and promote a standard environment for network applications to run. In complementary fashion, network equipment based on SDN separates the data and control planes, making it easier to develop solutions that run on servers instead of expensive, proprietary hardware platforms. The principle is that networking solution vendors will develop software-based virtual network functions (VNFs) that can be deployed on general-purpose hardware powered by Intel® processors and Intel® Ethernet Controllers.

With NFV and SDN, a network can manage and orchestrate VNFs using a set of standard interfaces. It is also possible to configure networking resources to manage packet flows through the network and essentially send traffic to the right applications. This represents a sea change in how telecom applications and services are deployed. Historically, a new application required a ground-up design of hardware and software to produce a bespoke 'box' before it could be deployed in the network – a process that can take years to achieve. In comparison, the world

	Enterprise Cloud	Telecoms
Reliability	Internet quality	Carrier-grade (five 9's)
Performance	CPU and memory bound	I/O and memory bound
Memory view	Aggregated – all memory systems are treated similarly	Non-uniform memory access (NUMA)
Operating system services	Very important for applica-tions	Less important for data plane functions
Architecture perspective	Node-centric	Network-centric
Virtual machines (VMs)	Many small heterogeneous VMs	Few large VMs

Table 1. Computing Requirements and Attributes for Telecoms and Enterprise Cloud

of NFV/SDN enables applications to be designed to run on commodity hardware, which already exists in the network or is available as a commercial off-the-shelf (COTS) server. As a result, service creation is now an on-demand activity.

Different Requirements for Telecoms and Enterprise Cloud

The roots of NFV are in the enterprise, where it provided an attractive approach to creating, configuring, and destroying applications on demand to allow service creation with agility on a massive scale. However, transporting this model to the telecom industry poses some challenges because the computing requirements and attributes are different, as shown in Table 1.

In a typical enterprise cloud model, reliability is very important in terms of providing service, but continuous service is less important. For instance, if a virtual machine (VM) running a video server crashes, it just gets restarted, and the end user suffers a minor inconvenience of losing their video for a few seconds. In contrast, continuity of service is critical in telecoms. If a user dials a 911 emergency call to report their house is on fire, it is not acceptable to drop the call and require the user to call back.

Another aspect of carrier grade is deterministic performance, which can only be assured by having predictable resources available – CPU, memory, connectivity – to guarantee that an application instance will launch successfully. This leads telecom operators to carefully manage memory access, ensuring low latency applications are assigned the fastest memory in the system, like processor cache, and not slow hard drives. Conversely, enterprise cloud operators have an aggregated view of memory, meaning applications are normally not given preferential treatment with respect to memory.

The types of applications used in telecom are clearly distinct from enterprise cloud, not just in the resources they consume, but also in the types of services they run, as reflected in the different names: tele-communications and cloud computing. Telecom applications fall into a few general categories: voice, video, analytics, billing, user-oriented database, and deep packet inspection (DPI) at line rate, so the VMs are relatively large, but small in number. Cloud applications are numerous and diverse, requiring many small VMs.

Driven by the types of applications they support, cloud operators will choose commercial servers based on CPU performance and memory capacity, whereas telecom operators will be

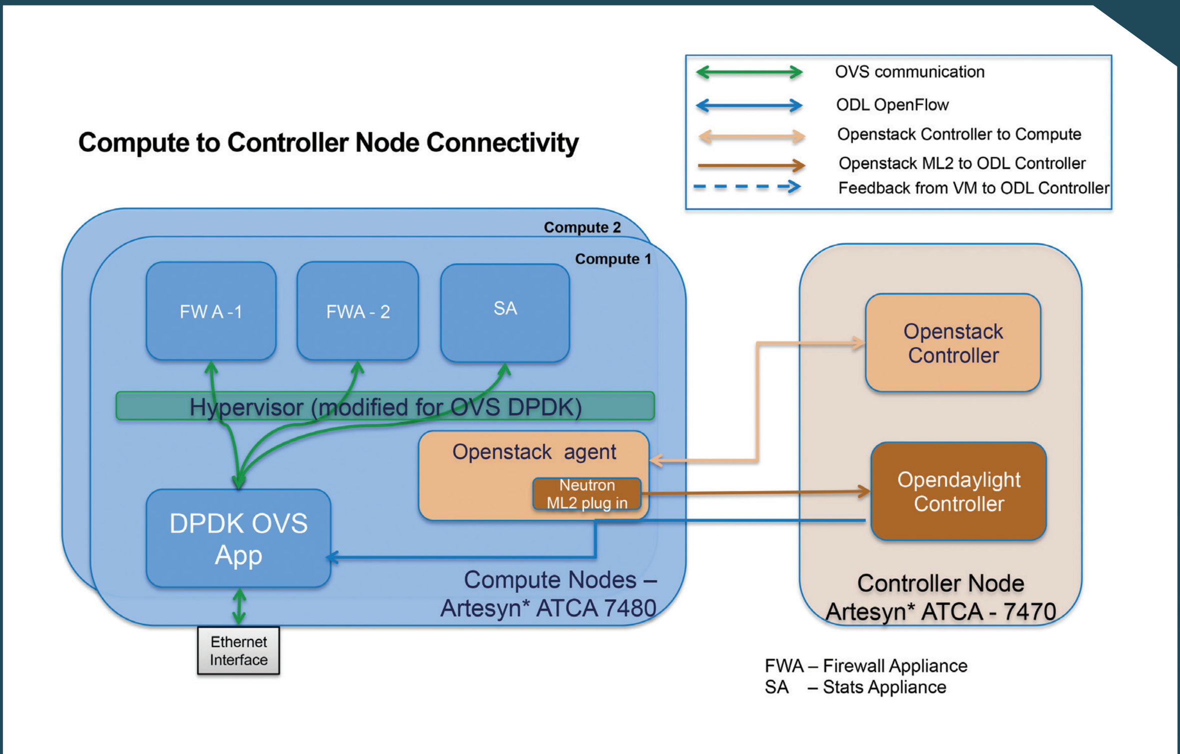

Figure 1. NFV and SDN Software Layers on an Infrastructure Node

focused on I/O bandwidth and memory size. A telecom server will typically have multiple 40 Gbps Ethernet ports, while a cloud server may have just a couple 1 Gbps Ethernet ports. In high-bandwidth telecom applications, it is equally important to ensure that I/O bandwidth and processing power are not squandered, in that packets should only be directed to the compute node that will actually consume them.

In telecoms, voice and video applications are not typically deployed on generic servers because both require specialized processing, and dedicated connectivity and I/O bandwidth. However, the software management interfaces to NFV management and orchestration layer may be generic even though the underlying hardware may not be just any COTS server.

Challenges

NFV and SDN network architecture concepts can help operators improve agility and reduce TCO by replacing expensive proprietary boxes with servers, but there is still work to be done. Both NFV and SDN were first adopted by enterprise-class data centers (aka 'the cloud'); however, enterprise-centric implementation of a non-standard framework is a barrier to short term adoption. Still, NFV and SDN do support de-facto, open-standard interfaces, like OpenFlow*, OpenStack*, and OpenDaylight* (ODL) shown in Figure 2, which can further standardization.

In fact, standardization of NFV has been undertaken by the ETSI Industry Specification Group, as well as aspects of reliability, availability, security, and performance required to transform NFV into a carrier-grade-ready platform. Making significant progress, the ETSI NFV group has been in session for over 18 months and recently voted to extend their work for another two years.

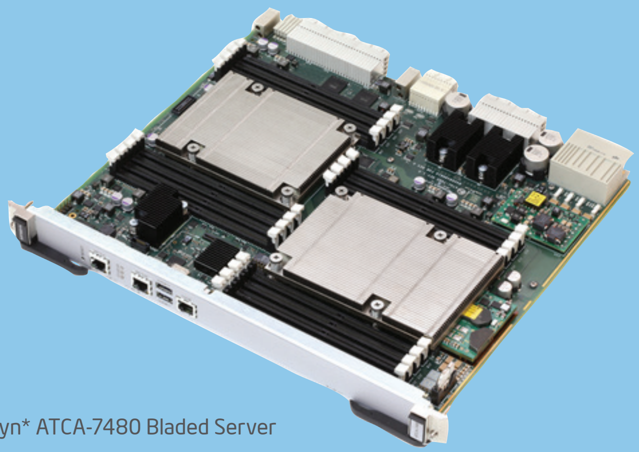

Figure 2. Artesyn* ATCA-7480 Bladed Server

Solution

A telecom solution in an NFV/SDN environment needs to employ a high-compute-density, high-bandwidth, highly-available hardware platform. The environment should be 'standard', but the hardware must be hardened.

This is where the high-performance ATCA platforms from Artesyn Embedded Technologies come into play. The Artesyn* ATCA-7480 bladed server, pictured in Figure 2, utilizes dual, from the 14-core Intel® Xeon® processor E5-2600 v3 product family, running in conjunction with 40 Gbps Intel Ethernet Controllers to provide 4 x 40 Gbps backplane connectivity per blade. Residing in an Artesyn Centellis* ATCA chassis, up to 10 blades can be deployed as a single NFV node with up to 280 Intel® processor cores and a total of 1 Tb/s of I/O connectivity.

Each Intel processor in the system is enabled with the software components required for it to be a true NFV-managed compute node:

- Highly-optimized Linux* operating system
- KVM hypervisor, providing a VM environment for application deployment with predictable resource availability

- Intel® Data Plane Development Kit (Intel® DPDK) Accelerated Open vSwitch* to maximize efficiency of packet delivery to individual VMs
- The local OpenFlow and OpenStack agents required for management via standard protocols and interfaces

For this design, Artesyn took advantage of the Intel® Open Network Platform Server Reference Design (Intel® ONP Server Reference Design), including the Intel DPDK Accelerated Open VSwitch, Intel Ethernet Controllers with powerful flow control features, and Intel® Solid-State Drives (Intel® SSD) to boost performance.

In addition to the Intel processor-based server blades, Artesyn's ATCA-F140 hub blade is enabled with FlowPilot* – an application to build in intelligent packet flow balancing between the payload blades. FlowPilot ensures the blades and cores only receive the packets they will consume, which is a significant advantage because processing nodes no longer have the overhead associated with examining and forwarding packets not destined for them. Figure 3 illustrates how FlowPilot uses flow pinning based on a 5-tuple to steer traffic flows to the right CPUs and cores.

The Centellis platform is enhanced further by the availability of specialized processing blades for DPI and media transcoding. These can sit alongside ATCA-7480 blades in a wide variety of configurations, providing an extremely flexible environment.

The Artesyn Centellis platform is a purpose-built, telecom service provider solution for an NFV infrastructure node:

- It provides a true carrier-grade platform, built on the industry-standard of ATCA
- It is an incredibly compute-dense platform with extremely high interconnect bandwidth, a key requirement of tele-communications applications
- It has the processing scalability, flexibility, and heterogeneity that allows a carrier to deploy the right platform for their applications, safe in the knowledge they can expand and customize their system configuration in the future

As a member of the Intel® Network Builders alliance, Artesyn Embedded Technologies benefits from Intel's dedicated technical support, early chip samples, and Intel product roadmaps.

Conclusion

NFV and SDN are exciting network architecture concepts that are expected to bring enormous benefit to the telecom industry by decreasing TCO and improving the time to market for new services. Agility in the development and deployment of services will allow telecom service providers to maximize the potential of their networks, as well as ease future scale-out.

The ETSI ISG for NFV is working to produce a standard for NFV, based largely on existing concepts from the enterprise cloud using open standards (e.g., OpenStack, OpenFlow, etc), while paying close attention to the availability needs in the telecom world.

This work has to be underpinned with standard, carrier-grade hardware as a basis for the NFV infrastructure layer. Satisfying the special requirements of telecom operators are the Artesyn Embedded Technologies' Centellis ATCA platform and ATCA-7480 Server Blade with NFV management and orchestration layer software.

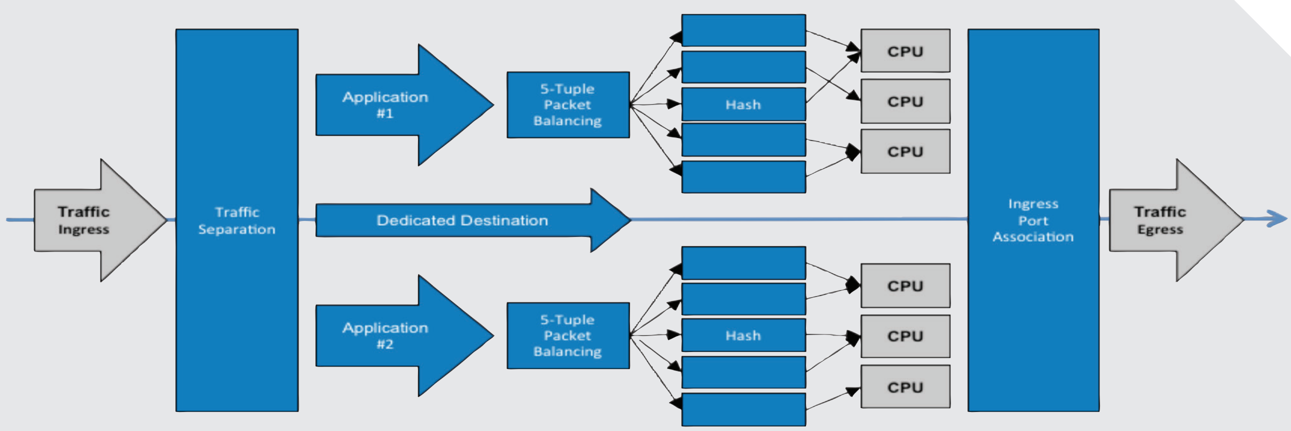

Figure 3. Artesyn* FlowPilot* Balances Traffic among CPUs on Multiple Server Blades

For more information about telecom solutions from Artesyn Embedded Technologies, visit **www.artesyn.com/embedded**

Solution Provided By: