A Guide to OCP for CLEVER CARRIERS

The Communications Service Providers' Guide to Open Compute Solutions for Network Equipment

Authors:
Doug Sandy, Chief Technology Officer, Embedded Computing
Todd Wynia, VP Communication Products

ARTESYN™
EMBEDDED TECHNOLOGIES

Executive Summary

The extensive use of virtualized environments by data centers and cloud providers has Communications Service Providers (CSPs) looking into ways to use virtualization and cloud technologies to reduce costs, increase efficiency, and improve service agility. For CSPs, such a shift requires a carefully orchestrated transition from the dedicated network appliances they use today to open systems based on more economical commercial off-the-shelf (COTS) servers and open software solutions.

Having different requirements than large-scale cloud providers and data centers complicates this transition. To implement new networking solutions such as Network Functions Virtualization (NFV) and Software-Defined Networking (SDN), CSPs must ensure that the core hardware platforms they deploy deliver the performance, reliability, serviceability, and regulatory and safety compliance their industry requires.

Artesyn is a long-time global leader in advanced networking designs, including but not limited to Advanced Telecommunications Computing Architecture (ATCA) products. More than 25,000 Artesyn systems are deployed in mobile and fixed networks globally. To help meet CSP needs for open solutions, Artesyn is taking a leadership position within the Open Compute Project (OCP) and European Telecommunications Standards Institute (ETSI) on developing standardized architectures for network equipment.

Artesyn has always embraced open ecosystems and over the last 15 years has been instrumental in promoting a variety of open standards including VME, cPCI, ATCA, and others. In this paper, we explore how Open Compute technologies and open “white box” hardware offer great promise for CSPs. We also consider the challenges the industry faces in making these technologies viable for the communications services industry.

The Move to Virtualization and Cloud Technologies

CSPs today are experiencing exponential growth in subscriber demand for secure, high-bandwidth applications. At the same time, they face increasing pressure to achieve greater network economies of scale while controlling the costs of network hardware upgrades.

Virtualization and cloud technologies offer promising solutions to these CSP challenges. Rapidly transforming the way companies and organizations process and store data, scale resources, and deliver services, virtualization and cloud technologies provide exciting opportunities to streamline operations and cut costs.

Virtualization saves money and improves efficiency by reducing reliance on expensive, dedicated (purpose-built) network appliances and their many different architectures and management interfaces

Artesyn's Centellis® OCP Platform is a rack solution that integrates servers, storage and top-of-rack switches, and offers improvements in density, availability, flexibility, scalability, serviceability, manageability, and ease of deployment versus traditional rack servers.

for operations, administration, maintenance and provisioning (see Figure 1). Using virtualization, CSPs can run many applications on more cost-effective, standardized, multi-vendor hardware that can be managed under a single management interface. The wide availability of such COTS hardware and open source software drives down costs and eliminates vendor lock-in.

Cloud technologies provide even greater operational flexibility, scalability and cost efficiency. Using these technologies, CSPs can launch new services in hours rather than months, and reduce latency by placing services closer to customers. CSPs can also leverage their well-established network-based businesses, local presences, customer relationships, and aggregator expertise to be strong competitors in the cloud provider space.

The Evolution to Open Network Architectures

Two recent cloud technologies, NFV and SDN, are driving industry-wide evolution toward open network architectures. These technologies offer exciting possibilities for making CSP networks more agile and responsive to the dynamic needs of today's traffic and services. By decoupling software and hardware, NFV and SDN enable CSPs to deploy services in software running on COTS hardware. This decoupling delivers many benefits, such as network-wide control and visibility, greater agility in services deployment, and CapEx/OpEx savings.

NFV enables consolidating network equipment types – traditionally hosted on proprietary hardware appliances – onto high volume COTS servers, switches and storage. By implementing network functions in software, NFV enables these functions to be moved to (or instantiated in) various locations in the network as needed without installing and maintaining new equipment. NFV is suitable for many data plane packet processing and control plane functions in fixed and mobile network infrastructures.

For example, NFV can be used for such applications as session boarder controllers (SBC), deep packet inspection (DPI), security appliances (firewalls, IDS/IPS, SSL VPNs, etc.), server load balancers, WAN acceleration, routers, gateways, and more.

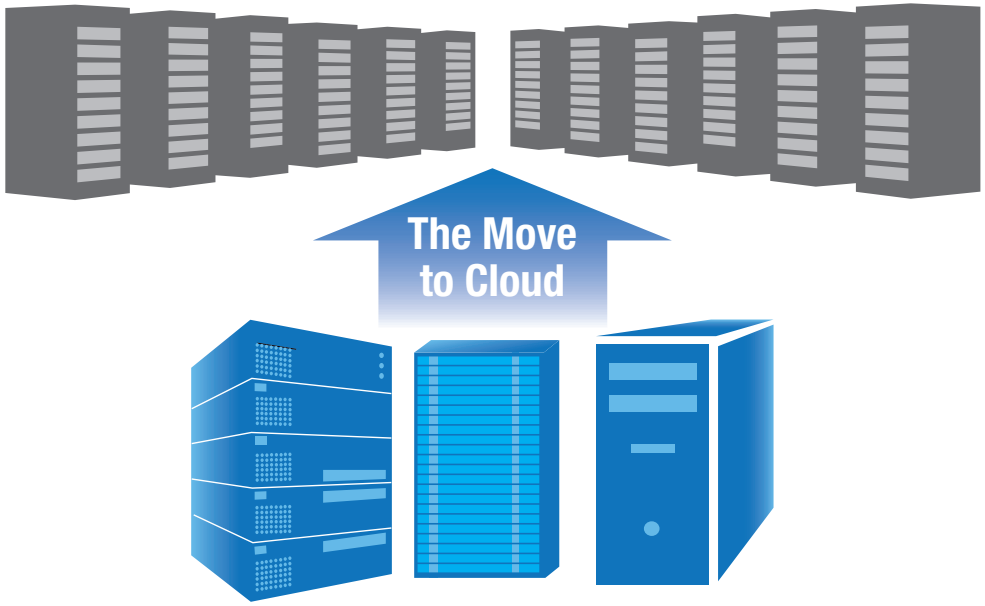

- Single purpose hardware
- Unique management interfaces
- Vendor lock-in
- Months to deploy new services
- Multiple actions per server
- Single management interface
- COTS and vendor independence
- Hours to deploy new services

Figure 1.
Virtualization enables CSPs to run many applications on more cost-effective standardized, multi-vendor hardware. Deploying cloud technologies on this COTS hardware enables greater agility in service delivery.

SDN separates control and data planes to provide a centralized controller and global view of the network. Enabling external applications to program the network, SDN optimizes the on-demand nature of cloud services and enables more efficient use of data center network, compute and storage resources.

As complementary technologies, NFV and SDN support the overall goal of programmability for the network in a carrier cloud environment using COTS hardware. CSPs can use NFV and SDN to gradually transform their networks into layered and distributed clouds enabling elasticity and optimization of data transport, delivery of new services, and continuously user-selectable quality of experience (QoE).

Development and Expansion of Open Compute Technologies

Many CSPs are looking to adopt COTS hardware and deploy NFV and SDN for an increasing amount of their operations. According to industry sources, by 2017, approximately 10 to 12 percent of AT&T's and Verizon's total next-generation network equipment CapEx will be enterprise COTS-based. The potential savings of COTS computing and networking hardware is also creating great interest in the latest data center innovation: Open Compute technologies.

Pioneered and promoted by the Open Compute Project (OCP), these technologies focus on the most efficient and economical ways of scaling COTS computing infrastructure. Founded by Facebook, OCP's original objective was to guide the design from the ground up of the most cost-efficient data center infrastructure possible.

Facebook sought a new rack-scale architecture that would use generic servers to make its new data centers as low cost and efficient as possible. The company viewed traditional 19-inch racks as falling short in several areas:

- 19-inch rack solutions each have a unique form factor, I/O, and management system, locking a company into a single supplier
- 19-inch solutions are typically individually powered, increasing cabling and power costs and limiting the ability for centralized management
- 19-inch solutions have minimal real estate for increasing functionality

> CSPs showed early interest. As early as November 2013, Frank Frankovsky, OCP chairman and president at the time, announced in a blog active collaboration "on the development of more than 30 potential contributions covering most of the network hardware stack and even some of the network software stack." Recognizing the need to address the industry-specific needs of CSPs, a group dedicated to the CSP industry formed within the OCP in October 2014. Their goal is to provide a set of recommendations or specifications. In a blog on taking an open approach to SDN, Omar Baldonado, manager of software infrastructure engineering at Facebook said, "We started it [OCP] for servers, storage, data center designs, and it's only recently that we've started it for networking. Those initiatives have all shown a lot of interest. Over 150 companies are involved with Open Compute, and we've had contributions from a lot of companies in the networking phase already."

Facebook wanted a new rack-scale solution in which all servers would be identical no matter what company manufactures them. Servers needed to be powered, plugged into the rack, and cabled the same. Determined to remove anything that didn't contribute to efficiency, Facebook even had manufacturers remove server faceplates, choosing to handle the regulatory EMC shielding at the facility level instead of server or rack level.

Today, the goal of the OCP is to spark a collaborative dialogue and effort among peers on OCP technology, collectively developing the most efficient computing infrastructure possible. Project focus includes addressing servers, storage, networking, hardware management, Open Rack (a rack standard), data center design, and certifications for solution providers.

An ETSI Open Compute group, the ETSI NFV Industry Specification Group (ISG), is paying close attention to OCP. Focused on NFV standards, this group addresses virtualization technology for consolidating network equipment types on COTS high volume servers, switches and storage. Seven of the world's leading telecoms network operators initiated the ETSI NFV ISG. They have been joined by over 200 companies, including network operators, telecoms equipment vendors, IT vendors, and technology providers. Their goal is to provide guidance through documentation.

Working on defining requirements and architecture, as well as addressing technical challenges, the ETSI NFV ISG delivered approximately 15 documents in Release 1 at the end of 2014. Release 1 covers management and orchestration, compute domain infrastructure, hypervisor domain, network domain infrastructure, and virtualized network function architecture. An important decision in 2014 was to work on recommendations for NFV infrastructure (NFVI) Node Architecture, including management functions implemented in software. Other publications outside the release plan will address security issues as well as performance and portability best practices.

Where Existing Open Compute Solutions Fit Today in Carrier Networks

Currently, OCP is best suited for traditional data center applications. While CSPs are showing interest in OCP architecture and its advances, their needs are much different and more complex than companies like Facebook that run narrowly focused hyperscale data centers. Consequently, at this time CSPs are primarily looking at OCP for handling applications such as billing services, ISP applications, and hosting capabilities for enterprise services.

As for COTS in general, results from a December 2014 published survey of CSPs conducted for Artesyn by Capital Research, a division of Markinetics, show CSPs are in the early deployment stage with COTS in the network. CSPs are using it for mobile network functions such as load balancing, IP edge routing, intelligent gateways, and compact evolved packet core (EPC). As for data plane NFV solutions, the research found that CSPs have concerns about their performance. In addition, because every CSP feels their network is strategic to their very existence, most are reluctant to take chances on new concepts.

Where Open Compute Solutions Fall Short in Carrier Networks

The differences between hyperscale data centers like Facebook's and traditional CSP data centers make current OCP implementations unsuitable for many CSP applications. To satisfy the needs of CSPs, OCP must meet the industry's higher standards in a number of areas. While OCP solutions for CSPs may not need to completely meet the Network Equipment-Building Standard (NEBS) – the most common set of safety, spatial and environmental design guidelines applied to U.S. telecommunications equipment – they do need to factor in the unique needs of CSPs to maintain reliable metropolitan to nationwide networks.

An Open Compute implementation for CSPs needs to address many or all of the following requirements:

1. **Co-location.** CSPs need Open Compute solutions that can co-locate with existing data center equipment. To avoid problems with radio interference, radio emissions must be controlled at the server/rack, not the facility level as currently implemented by OCP.

2. **Seismic Protection.** CSPs need greater seismic protection against earthquakes. To ensure the continuing operation of equipment and vital communications during a seismic event, all server racks and enclosures must comply with one or more of the industry standards for seismic protection.

3. **Front Cabling.** CSPs see front cabling as a maintenance obstacle. Having each server individually cabled makes node replacement a lot of manual work and increases the chance for human error. CSPs also need faster interconnect speeds and greater scalability than the front cabling used in OCP designs. Instead of 1 GbE and 10 GbE connections CSPs want the ability to scale from 10 GbE to 25 GbE, 40 GbE, and even 100 GbE connections to meet their higher capacity-per-node requirements. These higher speed connections will enable the cost-efficient handling of data-demanding applications and line-rate functions such as firewalls, switching, routing, and security appliances.

4. **Availability.** CSPs require greater service availability than ordinary cloud providers to maintain regulatory compliance (e.g., FCC regulations in the U.S.), service level agreements (SLAs), and QoE. High availability is particularly important for physical layer applications such as transport. CSP levels of availability and reliability are currently not supported by today's Open Compute solutions.

5. **Software.** Current cloud and virtualization deployments use an assortment of open source components to create non-standard software frameworks for managing application instances. CSPs will require true standard software interfaces with the reliability to match the needs of network applications, such as those under development by the ETSI NFV group. In addition to application resources, CSP will require a rack-level and a data center-level interface for management of power and mechanical systems such as fans, and other components.

6. **Cooling.** The use of small cooling fans in OCP servers is a maintenance and reliability obstacle for CSPs. Small fans tend to be noisier and less reliable. The use of thousands and tens of thousands of them in a data center creates both maintenance and noise issues. For CSPs, rack-based rear cooling with larger, quieter, more efficient fans makes more sense and helps reduce OPEX.

7. **Interoperability.** CSPs have traditionally purchased single-purpose system for each function from a specific vendor. To scale, CSPs added more of the exact same equipment from the vendor. CSPs are not accustomed to relying on COTS servers from a mixed supplier environment for their networked functions. To make the transition to OCP, CSPs need greater assurance through certification that these systems will work together.

8. **Certification.** CSPs will want carrier grade requirements demonstrated and certified on Open Compute technologies targeting the carrier cloud, particularly high-volume COTS servers for NFV solutions. In particular, servers and technologies must be proven capable of delivering the required low-latency performance to support CSP SLAs.

How the Industry Will Address CSP Requirements

Working with OCP and ETSI, Artesyn sees real opportunity to help CSPs economically scale out their networks using COTS equipment from the data center through the metro service edge (core). As these organizations work to help data centers move to Open Compute technologies, manufacturers and suppliers in the CSP space such as Artesyn can work in parallel, defining and developing a more networking-centric version of Open Compute technologies.

Focusing first and especially on the high-touch Layer 4 through 7 services targeted for NFV, there's great opportunity for a networking-centric Open Compute platform to support a significant portion of central office and data center application needs (see Figure 2). Layer 4-7 services, sometimes referred to as the upper layers, support end-to-end communication between a source and destination application and are used whenever a message passes from or to a user.

The list of potential Layer 4-7 services is long in cloud computing and SDN infrastructure because these services can be built largely independent of the underlying network. For example, load balancing, WAN acceleration, and virtual security appliances can all be virtual applications. From an immediate benefits point of view, some of the earliest benefits could be achieved by NFV implementations hosted in the central office where services are delivered over the carrier Ethernet infrastructure.

A next step is to use Open Compute technologies designed for CSPs at the network edge to enable virtual switches and virtual load balancers. Using virtual switches and load balancers would allow these resources as well to be elastic. Through SDN, software on Layers 4-7 could then have control over the network on the lower layers (L2/L3) or IP edge. This control would generate added value by giving applications running on Layers 4-7 the necessary visibility over network layer resources to make resource provisioning more efficient.

While layer 1-4 functions such as transport or edge routing may not be suited for NFV deployments, a carrier-focused version of OCP could allow these functions to be deployed within the same infrastructure as the upper layers. This would enable a common equipment practice and management interface across the full spectrum of CSP deployments.

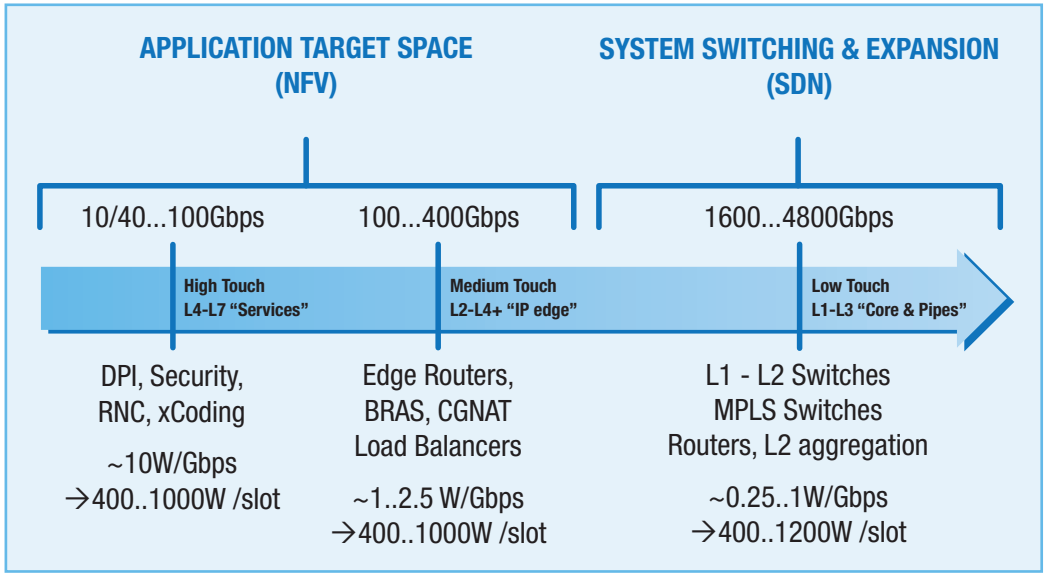

Figure 2.
Artesyn and other CSP providers are working together to drive Open Compute solutions that address the high-and medium-touch layers of carrier networks.

Artesyn plans to take a major role in OCP and ETSI in helping to design an Open Compute rack architecture that addresses CSP needs and concerns such as reliability, OpEx, equipment co-location needs, cooling, and NEBS compliance. Work is ongoing with specification work and pilots planned for 2015 and products and standards available sometime in 2016.

Some of the specific innovations for CSPs include:

- Integrated optical interconnects supporting 10/25/40/50/100 GbE for massive bandwidth scalability
- Use of hot-swap sleds in trays for carrier-grade serviceability
- Industry-standard software infrastructure designed to support deployment of COTS applications within an SDN/NFV managed environment including OpenStack, OpenFlow, and NETCONF
- Fewer, more highly efficient, quiet, reliable hot swap fans
- Use of highly efficient processors and other components
- Network-centric hardware design and certification addressing everything from low latency requirements to regulatory and safety compliance
- Rack management infrastructure to manage and monitor everything from the electromechanical components to the payload and switches

Conclusion: Working Towards a CSP-Standardized Architecture

CSPs are embarking on an exciting period of business transformation. The ability to use high volume COTS servers to implement cloud technologies, such as NFV, will help them reduce CapEx and OpEx, unleash new flexibility and elasticity in their operations, and radically improve their time to market for new services.

While OCP provides an excellent solution for the enterprise data center, CSPs require a higher grade of hardware platform designed to meet their more challenging needs for low-latency performance, bandwidth scalability, reliability and serviceability, and regulatory and safety compliance. Considering how critical and specialized their equipment is, the shift requires a carefully orchestrated transition.

Working with the OCP and ETSI, companies like Artesyn are rapidly developing solutions that will meet CSP requirements with standardized architectures. Through a performance-optimized solution that maximizes data flows to virtualized applications while maintaining five nines reliability, CSPs in the future should easily and confidently be able to implement Open Compute solutions. These solutions will enable Open Compute not only for their enterprise data center needs, but also for NFV solutions and a range of new innovative services that will help them better compete in the cloud provider and communications industries.

About Artesyn Embedded Technologies

Artesyn Embedded Technologies is a global leader in the design and manufacture of highly reliable embedded computing solutions for a wide range of industries including communications, military, aerospace and industrial automation.

Building on the acquired heritage of industry leaders such as Motorola Computer Group and Force Computers, Artesyn is a recognized leading provider of advanced network computing solutions ranging from application-ready platforms, single board computers, enclosures, blades and modules to enabling software and professional services.

For more than 40 years, customers have trusted Artesyn to help them accelerate time-to-market, reduce risk and shift development efforts to the deployment of new, value-add features and services that build market share.

Artesyn has over 20,000 employees worldwide across nine engineering centers of excellence, four world-class manufacturing facilities, and global sales and support offices.

© Copyright 2015 Artesyn Embedded Technologies, Inc., All rights reserved.

Trademarks
Artesyn Embedded Technologies, Artesyn and the Artesyn Embedded Technologies logo are trademarks and service marks of Artesyn Embedded Technologies, Inc.
Intel®, the Intel logo, and Intel® Xeon® are trademarks or registered trademark of Intel® Corporation or its subsidiaries in the United States and other countries.
All other product or service names are the property of their respective owners.
Reproduction of this material in any manner whatsoever without the express written permission of Artesyn is strictly forbidden.

Notice
While reasonable efforts have been made to assure the accuracy of this document, the authors assume no liability resulting from any omissions in this document, or from the use of the information obtained therein. The authors reserve the right to revise this document and to make changes from time to time in the content hereof without obligation of the authors to notify any person of such revision or changes. Electronic versions of this material may be read online, downloaded for personal use, or referenced in another document as a URL. The text itself may not be published commercially in print or electronic form, edited, translated, or otherwise altered without the permission of the authors. It is possible that this publication may contain reference to or information about products (machines and programs), programming, or services that are not available in your country. Such references or information must not be construed to mean that the companies intend to announce such products, programming, or services in your country.

Limited and Restricted Rights Legend
If the documentation contained herein is supplied, directly or indirectly, to the U.S. Government, the following notice shall apply unless otherwise agreed to in writing by Artesyn. Use, duplication, or disclosure by the Government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data clause at DFARS 252.227-7013 (Nov. 1995) and of the Rights in Noncommercial Computer Software and Documentation clause at DFARS 252.227-7014 (Jun. 1995).

A Guide to OCP for Clever Carriers - version 1