July 2016

EMBEDDED SYSTEMS ENGINEERING

powered by **EECatalog**

Guiding Embedded Designers on Systems and Technologies

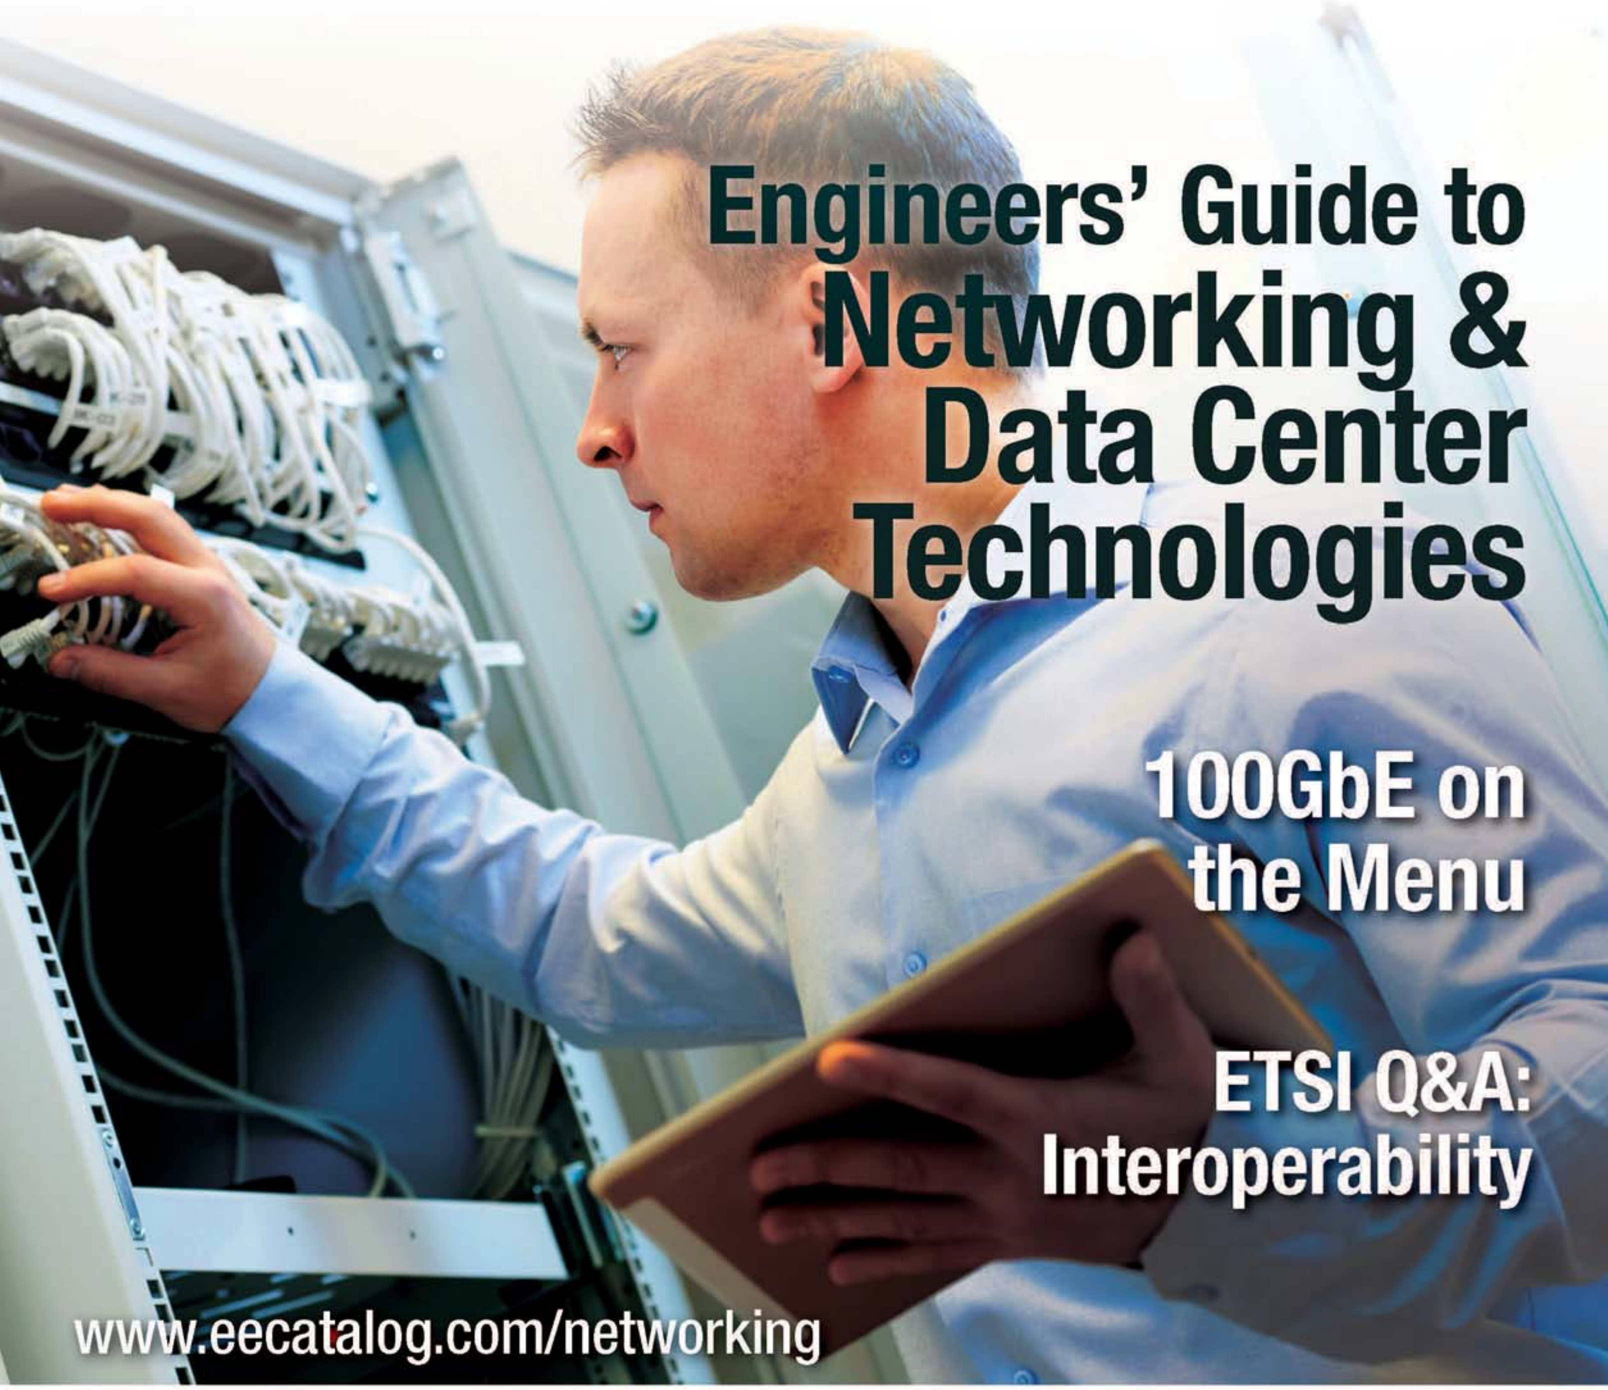

Engineers' Guide to Networking & Data Center Technologies

100GbE on the Menu

ETSI Q&A: Interoperability

www.eecatalog.com/networking

Sponsors

ADLINK TECHNOLOGY INC.

ELMA Your Solution Partner

CONTENTS

EMBEDDED SYSTEMS ENGINEERING

Special Features

Product Showcases

Hardware

ENGINEERS' GUIDE TO NETWORKING & DATA CENTER TECHNOLOGIES 2016
www.embeddedsystemsengineering.com

Vice President & Publisher
Clair Bright

Editorial
Managing Editor
Anne Fisher
afisher@extensionmedia.com
Editor
Chris Ciufo
cciufo@extensionmedia.com
Senior Editors
Caroline Hayes
Gabe Moretti
Dave Bursky

Creative/Production
Production Traffic Coordinator
LS Jerrett
Graphic Designers
Nicky Jacobson
Simone Bradley
Media Coordinator
Jon Franke
Senior Web Developers
Slava Dotsenko
Mariam Moattari

Advertising / Reprint Sales
Vice President, Sales
Embedded Electronics Media Group
Clair Bright
cbright@extensionmedia.com
(415) 255-0390 ext. 15
Sales Manager
Elizabeth Thoma
415-255-0390 x17
ethoma@extensionmedia.com

Marketing/Circulation
Jenna Johnson

To Subscribe
www.eecatalog.com

Extension MEDIA

Extension Media, LLC Corporate Office
President and Publisher
Vince Ridley
vridley@extensionmedia.com
(415) 255-0390 ext. 18
Vice President & Publisher
Clair Bright
cbright@extensionmedia.com
Vice President, Business Development
Melissa Sterling
msterling@extensionmedia.com
Human Resources / Administration
Darla Rovetti

Special Thanks to Our Sponsors

ADLINK TECHNOLOGY INC.
ELMA Your Solution Partner

Embedded Systems Engineering is published by Extension Media LLC, 1786 18th Street, San Francisco, CA 94107.
Copyright © 2016 by Extension Media LLC. All rights reserved. Printed in the U.S.

ADLINK

Innovation of Network Appliance with Modular Industrial Cloud Architecture (MICA)

- Flexible modular design: computing, storage, network resources and IO modules can be deployed as needed
- High flexibility and scalability to optimize performance, cost and space

2U Network Appliance CSA-7200

- Intel® Xeon® Processor E5-2600 v3 Family platform
- High scalability with 8 Network Interface Modules
- Up to 64x 10GbE SFP+ ports with I/O intensive architecture
- Flexible storage interfaces: SATA, PCIe, M.2

4U Network Appliance CSA-7400

- Intel® Xeon® Processor E5-2600 v3 Family platform
- 8 CPU high-density platform
- 1.3Tbps total bandwidth, 128Gbps Fabric bandwidth
- Up to 720Gbps front IO density

ADLINK Technology, Inc.

▸ *Toll Free: +1-800-966-5200* ▸ *info@adlinktech.com* ▸ *www.adlinktech.com*

intel IoT Solutions Alliance Premier

100GbE on the Menu

A revision to PICMG 3.1 caters to the hunger for bandwidth that service providers, operators and consumers crave.

By Caroline Hayes, Senior Editor

The increase in cloud services and Internet traffic has placed considerable demands on bandwidth and capacity in recent years. As a result, data center developers have been forced to either invest in additional equipment or consider an architecture that can be scaled to add the capacity needed.

"PICMG 3.1 Revision 3 was drawn up to accommodate 100GbE. The open standard remains true to its values of preserving the ability to combine boards, switches and backplanes from multiple vendors to create 100GbE systems."

The dramatic increase in data traffic has been advanced by the proliferation of mobile devices. There will soon be more traffic originating from mobile devices than from PC-based ones. Last year, Cisco tracked this trend, in its Cisco Visual Networking Index (VNI) forecast. Data originating from non-PC devices, such as smartphones, tablets and TVs, grew from 40 percent in 2014 to a projected 67 percent in 2019.

There is expected to be an increase in all data traffic in that period, as Internet and cloud services are joined by IoT and Machine-to-Machine (M2M) data. Cisco believes that PC-originated data will grow at a Compound Annual Growth Rate (CAGR) of nine percent, significant, but dwarfed by tablet data traffic growing at a CAGR of 67 percent, smartphones at 62 percent and M2M at 71 percent CAGR. Another significant milestone by 2019, says Cisco, is that traffic from wireless and mobile devices will exceed that from wired devices. WiFi and mobile devices will account for 66 percent of IP traffic, compared to wired devices' 33 percent; in 2014 wired devices accounted for over half (54 percent) of traffic. In addition, an estimated 80 percent of telecommunications network traffic is video, which consumes a lot of bandwidth, and augments the case for 100G data transfer rates.

Figure 1: Doug Sandy, CTO PICMG, believes PICMG 3.1 R3.0 is a "robust alternative" for IT equipment.

100G CHALLENGES

To realize 100G operation, networks require Deep Packet Inspection (DPI) to process large data flows in real time. They also need support for Software Defined Networking (SDN) for design flexibility and network security.

Scalable systems are needed to support the network infrastructure as it grows, while upgradeable systems allow service providers to introduce services in response to customer demand.

Both SDN and Network Function Virtualization (NFV) allow features to be added to the networks, or to be reconfigured without upgrading a large portion of the hardware.

AdvancedTCA was introduced by the PICMG at the end of 2001, as a common hardware platform for computing and telecommunications equipment, with availability for central office applications. The intervening 15 years has seen industry progress from 10GbE to 40GbE and the IEEE 802.3bj-2014 standard, which adds 100 Gbit per second Physical Layer (PHY) specifications and parameters.

Incorporating 100Gbit backplane Ethernet (GbE) into the Advanced Telecom Computer Architecture (AdvancedTCA or ATCA) standard became the next step.

Figure 2: The Centellis 8000 40G/100G 14-slot ATCA system from Artesyn.

PICMG 3.1 Revision 3 was drawn up to accommodate 100GbE. The open standard remains true to its values of preserving the ability to combine boards, switches and backplanes from multiple vendors to create 100GbE systems.

“PICMG 3.2 R3.0 platforms provide a more robust alternative to operators who are not comfortable deploying standard IT equipment in their facilities,” explains Doug Sandy, Chief Architect/Lead Hyperscale Technologist for Artesyn’s Embedded Power Business, and Vice President of Technology / Chief Technology Officer of PICMG. “Some reasons might include backward compatibility with existing equipment, or more rugged requirements such as those found in traditional telecommunications central offices,” he continues, adding: “ATCA is also finding increasing adoption in military/aerospace applications.”

Sandy confirms that the revision, which incorporates 100-Gbit and 25-Gbit Ethernet into the AdvancedTCA platform was adopted last month and that the specification is ready for purchase from PICMG.

As well as backward compatibility, and multi-vendor interoperability, 100GbE operation throws up some technical challenges, for example, managing losses at 100G signalling rates, impedance control, which is achieved by limiting via stubs and crosstalk control, using trace geometries and grounding through connectors.

“Primary challenges were related to high speed signal integrity, interoperation between multiple vendors and backward compatibility,” says Sandy. “Artesyn Embedded Technologies’ Embedded Computing division spearheaded this work with a connector vendor and brought the initial concept to PICMG,” relates Sandy. “Through committee, collaboration (and lots of simulation work), the solution was expanded, improved and refined. The result is the specification that we have today.”

REALIZING 100GBE OPERATION

Artesyn offers 100G shelves, based on a QuadStar backplane architecture that has four switch blades or hubs, each fanning out to other cards in a rack. This increases the available bandwidth compared with dual-dual star architecture. Three switches can be active, and the fourth is in standby mode to offer 3+1 redundancy. The technology can be scaled up to 100G. This enables each blade to deliver up to 400G, or 4-Tbit per second aggregate system bandwidth in a non-redundant implement, and 300G bandwidth with redundancy.

An example is the ATCA-7480 QuadStar packet processing blade. It is based on two Intel® Xeon® E5-2600 v3 family processors providing up to 28 processing cores per blade.

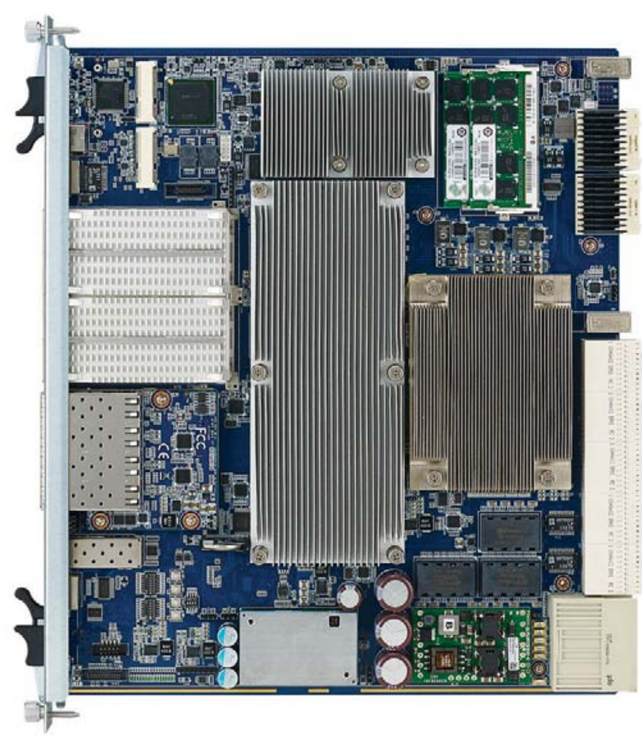

Figure 3: Advantech’s ATCA-9223 uses the Intel Atom C2000 processor for programming and Virtex-7 FPGAs for packet processing on 100GbE ports.

Together with connector manufacturer, ERNI, the company has developed connector and backplane technology for 100GbE connectivity in an AdvancedTCA shelf. Following the launch of the Centellis 8000 14-slot systems (Figure 2) for high availability applications, the company added the Centellis 8840 AdvancedTCA open standard server, with 100G AdvancedTCA technology integrated into a Network Equipment Building System (NEBs)-ready platform that will accept 40G and 100G blades, as they become available.

Also available for 100G operation, Advantech targets network security in carrier and large enterprise data center networks, with the ATCA-9223, 100GbE AdvancedTCA hub blade (Figure 3).

Like Centellis, it is designed for use in both SDN and NFV. According to the company it optimizes traffic flows and load balancing across clusters of AdvancedTCA node blades. It uses an onboard Intel Atom™ Processor C2000 for programmability, with a Broadcom BCM56150 base fabric switch with 70-Gbit per second switching for PICMG 3.1 GbE backplane connectivity. Two Xilinx Virtex-7 690T FPGAs connect two 100GbE and eight 10GbE for inline processing of packets between external ports and blades on the fabric interface.

Caroline Hayes has been a journalist, covering the electronics sector for over 20 years. She has worked on many titles, most recently the pan-European magazine, EPN.

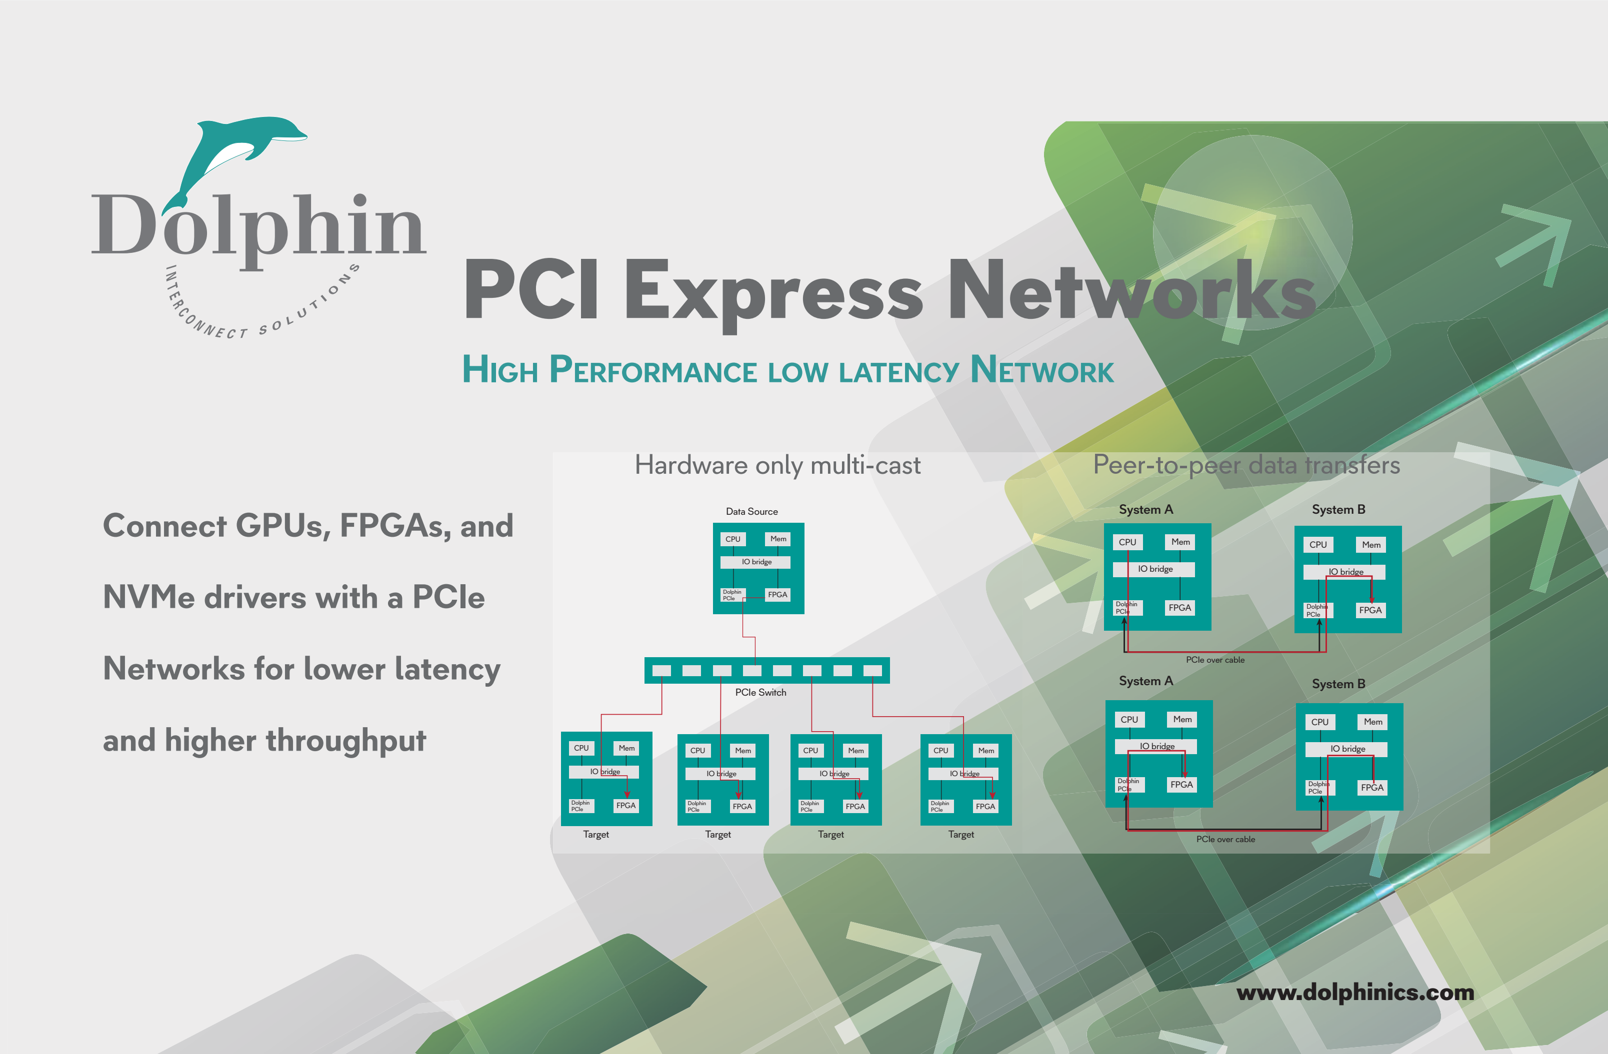

Increased Interoperability: Q&A with Don Clarke, ETSI

NFV's objectives, why Telecommunications infrastructures require rigorous specifications, and more.

By Anne Fisher, Managing Editor

Don Clarke, Chairman of the NOC (Network Operator Council), ETSI NFV Industry Specification Group.

The European Telecommunications Standards Institute (ETSI) develops Information and Communications Technologies standards deployed worldwide for fixed, mobile, radio, broadcast and Internet. This role naturally makes this standards organization the holder of a key role in the development of Network Functions Virtualization (NFV) technologies. Don Clarke, chairman of the Network Operator Council group in the ETSI NFV Industry Specification Group recently responded to e-mailed questions from EECatalog about the first NFV Plugtests Event organized by the ETSI Center for Testing and Interoperability, which will be held from January 23 to February 3, 2017, and other data center and virtualization topics. Edited excerpts follow.

"Such virtualized infrastructures need to be managed end-to-end, which requires new standards and new tools."

EECatalog: What should our readers be aware of regarding the NFV Plugstests Event being held at the beginning of next year [January 23 to February 3, 2017]?

Don Clarke, ETSI: ETSI Plugtests are an essential source of feedback to our standardization activities, which allow us to validate and improve the quality of the specifications as they are being developed.

The first NFV Plugtest focuses on testing relevant ETSI Network Functions Virtualization (NFV) capabilities over a number of combinations of NFV infrastructure, Management and Orchestration (MANO) solutions and Virtual Network Functions (VNFs) provided by the industry and Open Source projects. This focus allows ETSI to evaluate and increase the interoperability among vendor and open source implementations.

Besides being a source of essential feedback for ETSI NFV, the NFV Plugtest is also a great opportunity for the industry and open source projects to learn how the rest of the NFV ecosystem uses their implementations.

EECatalog: Many open source communities have emerged to drive NFV implementation. Are standards still needed?

Clarke, ETSI: Open source is an excellent way for the common elements of an implementation to be created collaboratively, and for vendors to focus their individual commercial efforts on capabilities built on top of open source. But Telecommunications infrastructures require rigorous specifications to ensure interoperability and to support legacy services that are deployed at massive scale. Telecommunications networks must also meet numerous regulatory requirements including support for critical national infrastructures. Current open source governance models do not provide these guarantees. Ideally there is a model where Standards Development Organizations (SDOs) developing specifications work more quickly and hand-in-hand with open source communities.

ETSI NFV has led the way in converging and specifying operator requirements (38 operators are involved) and the ETSI NFV work is widely referenced by the industry including open source communities. ETSI consequently established the Open Source MANO (OSM) group in February 2016 to deliver an open source NFV MANO stack using best-in-class open source workflows and tools to ensure rapid development and delivery. The activity is closely aligned with the evolution of ETSI NFV and provides a regularly updated reference implementation of NFV MANO. OSM enables an ecosystem of NFV solution vendors to rapidly and cost-effectively deliver solutions to their users.

EECatalog: How would you say embedded virtualization differs from that used for data centers and enterprise IT networks?

Clarke, ETSI: I prefer to use the term Network Functions Virtualization (NFV). The objective of NFV is to use IT and Cloud techniques, including virtualization and management and orchestration, but to identify and specify additional requirements that will enable these technologies to be used to create "carrier grade" network solutions inside cloud environments. In this context, "carrier grade" means the ability to assure deterministic bandwidth, jitter and latency, and to enable configurations that can deliver the appropriate level of reliability and availability for the services being delivered via the virtualized infrastructure.

In addition, network operators require cloud infrastructures to be "distributed," that is, extending beyond the data center. For example, instances of cloud infrastructure could be physically located in the access network, and even in the end user premises. Such virtualized infrastructures need to be managed end-to-end, which requires new standards and new tools.

EECatalog: What are some examples you have seen of embedded developers putting virtualization to innovative use?

Clarke, ETSI: We are seeing the early application of NFV to enable high-performance software implementations of network functionality previously only possible using hardware devices for such tasks as routers, firewalls and security monitoring. Implementing these functions purely in software enables automation and faster deployment, including customer self-provisioning.

EECatalog: How do you expect virtualization where the need for real-time response is also involved to look five years from now?

Clarke, ETSI: Achieving automation is key. There is still a lot of work to do to enable network operators to fully automate network design, provisioning and operations. Currently virtualized networks need a lot of manual intervention to design and deploy. This is why early NFV deployments are often in conventional data center environments where existing tools can be used. A key area of focus is to converge information modeling approaches across the industry to minimize complexity and simplify tooling and skill requirements. A collaborative multi-SDO effort is underway to do that.

EECatalog: What technology developments are you keeping especially close watch on?

Clarke, ETSI: The emergence of container technology as an alternative to virtual machines is of high interest. Containers are more resource efficient and faster to deploy than virtual machines, but there is more dependency on the host operating system version, which needs to be taken into account to ensure interoperability.

Today, commercial VNFs are often based on hardware appliances that have been re-purposed to run in a cloud environment. Such re-purposing can be inefficient in use of resources, so we are interested to see VNFs designed from the ground up to be more resource efficient and more optimized for automated deployment and operations.

Elma Electronic

High Speed Networking Solutions

If proven performance in demanding applications is what you look for in an Ethernet switch, then Elma has the answer. Our Ethernet and PCIe switches and routing engines perform in a wide range of applications. Address your networking requirements with the broadest and most technologically advanced family of networking products on the embedded market today. Elma offers solutions for high speed networking in defense, industrial, medical and telecommunication systems. Based on the latest silicon available from Marvell, our Interface Concept line of products have been serving critical roles in multiple defense programs as well as key commercial installations for over twelve years. Un-managed and managed models support basic switching up to the latest layer 3 routing protocols. Our Switchware™ graphical user interface (GUI) makes network management simple with it's easy to use interface which allows for network optimization, analytics and troubleshooting. All Ethernet switching products are supported through our Horsham, PA offices by one of the most experienced staffs in the industry. Customization is available at low volumes. Look to Elma's Interface Concept line of networking products for the speed, scalability, reliability and security your systems require.

FEATURES & BENEFITS

- 12 year track record of solid performance, reliability and longevity in demanding environments
- US based sales and support organizations enable fast response time
- Low volume customs including port location, count, type and GUI enhancements
- Custom packaging and backplane designs along with thermal analysis for system platforms
- Chassis systems or board solutions with complete integration and system design support

TECHNICAL SPECS

- 3U, 6U, mezzanine and custom chassis solutions including SFF and 19" rackmount
- VPX, VME, cPCI, cPCI serial, mezzanine and PCI/104 versions
- Air and conduction cooled models, -45°C to +85°C with optional conformal coating
- 1Gig, 10 Gig and 40Gig links with copper and fiber port options and front or rear I/O for up to 28 ports on a 6U board
- Layer 3 routing with full protocol support such as VLAN services, Quality of Service, multi-cast options, priority queues, bandwidth management, link aggregation, user authentication and more.

AVAILABILITY

10-12 weeks ARO, varies by model

APPLICATION AREAS

As the backbone of the IoT revolution, Ethernet based IP network communications facilitates high speed data and voice communications for signal intelligence gathering, engine control systems, enterprise data transfer, security systems, transportation systems, imaging systems, positioning systems, data storage and robotics equipment in the defense, manufacturing, medical and telecommunication industries.

CONTACT INFORMATION

ELMA
Your Solution Partner

Elma Electronic
235 Gibraltar Rd
Horsham , PA 19044
USA
Tel: 1(215) 956-1230
www.elma.com

Innovation of Network Appliance with Modular Industrial Cloud Architecture (MICA)

- Flexible modular design: computing, storage, network resources and IO modules can be deployed as needed
- High flexibility and scalability to optimize performance, cost and space

2U Network Appliance CSA-7200

- Intel® Xeon® Processor E5-2600 v3 Family platform
- High scalability with 8 Network Interface Modules
- Up to 64x 10GbE SFP+ ports with I/O intensive architecture
- Flexible storage interfaces: SATA, PCIe, M.2

4U Network Appliance CSA-7400

- Intel® Xeon® Processor E5-2600 v3 Family platform
- 8 CPU high-density platform
- 1.3Tbps total bandwidth, 128Gbps Fabric bandwidth
- Up to 720Gbps front IO density

ADLINK Technology, Inc.

▸ *Toll Free: +1-800-966-5200* ▸ *info@adlinktech.com* ▸ *www.adlinktech.com*

intel IoT Solutions Alliance Premier