PREPARING FOR THE CONVERGED FUTURE

Proof of concept from Intel, Casa Systems, and Kyrio drives virtualized fixed and mobile access

Today's service providers are focused on providing optimal quality of service (QoS) to their customers. In our future 5G world, those customers will not only need to be connected 24/7, they will also need fast, agile services, access to massively increased bandwidth capacity, and different levels of service for different types of customers and use cases. At the same time, service providers are also looking to create new revenue-generating opportunities to keep up in an increasingly competitive marketplace.

Given these opportunities and today's escalating demands, companies are asking themselves an important question: Is my network ready? Specifically, is it equipped to support multiple best-in-class wired and wireless access technologies? A growing number of companies are finding that they are not adequately prepared and are exploring ways to:

- Virtualize and simplify their networks
- Move access closer to the end user
- Leverage resources across multiple access technologies

ACCESS NETWORK EVOLUTION

As the industry changes, the pressure on service providers continues to grow. In addition to the calls for greater capacity and the ability to differentiate and provide advanced performance classes to their traffic, these companies must also analyze and extract value from all of the data and do it over various physical access mediums. To complicate matters further, all of that must be done cost-effectively.

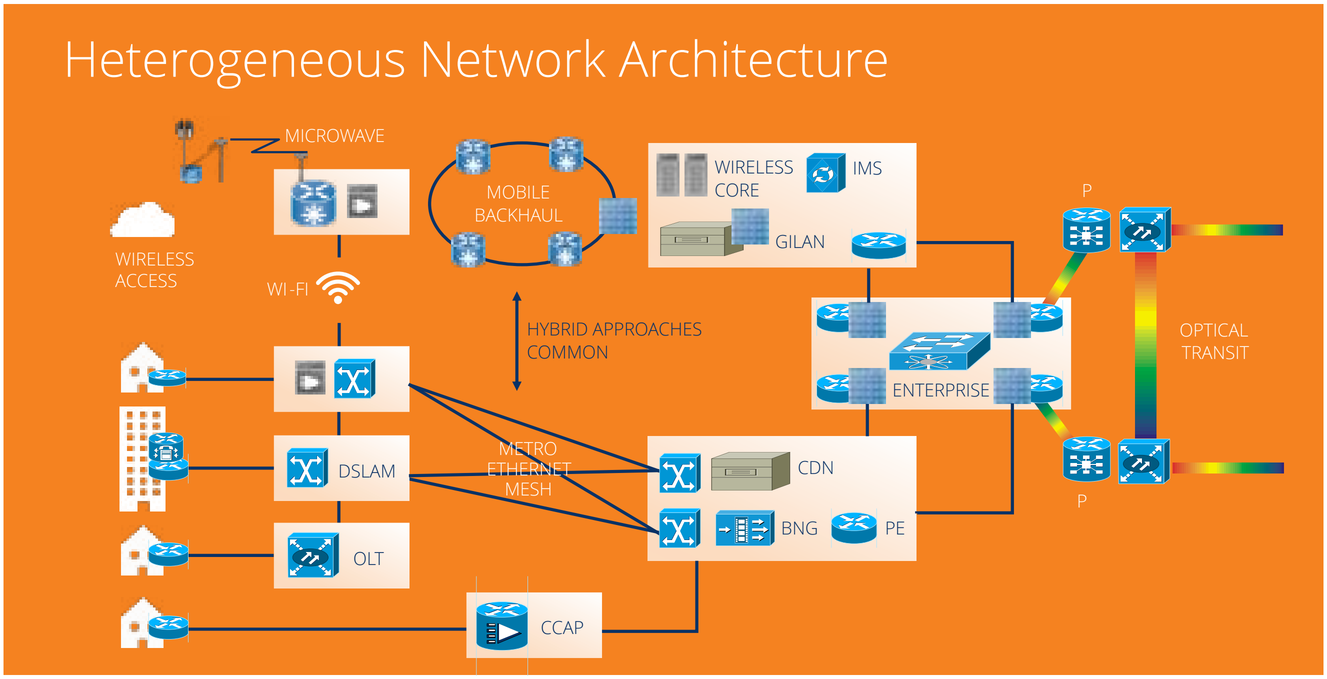

In the view of many, these goals can largely be met by the virtualization of network functions and the development of software-defined infrastructure in general. Such an approach moves companies from fixed-function appliances with fixed capacity to virtual network functions written in software and running on commercial off-the-shelf (COTS) servers and switches.

The principal aim is to increase flexibility and service agility, with the secondary goal of improving total cost of ownership (TCO). Part of the formula is bringing down the cost per bit over the network, while equipping operators with enhanced flexibility when it comes to responding to market demands.
By exchanging a hardware-based cadence to one driven by software, users can greatly accelerate the speed of innovation.

Of course, there will always be a certain amount of purpose-built infrastructure. But it is becoming increasingly clear that important benefits can be achieved when progressively adding more network function virtualization (NFV) and software-defined networking (SDN)–based deployments over time. This evolution becomes an important tool for driving network scale and agility for service providers.

THE INTEL VISION FOR NFV

Intel sees the enormous potential in NFV solutions. That is why the technology leader is focusing time, attention, and resources on helping accelerate its adoption everywhere. Intel is working to enable an access network that is coax, twisted pair, fiber, and wireless, all deployed by the same service provider. In Intel's view, today's existing networks, which are currently deployed alongside each other, can be converged to reduce real estate requirements, lower cost, and boost flexibility.

Getting there starts with partnership across the ecosystem. It also goes beyond technology, Intel® chips, platforms, and hardware. A successful transformation means rethinking business processes, developing skillsets in the workforce, and helping develop an end-to-end value chain for service providers and network operators. Intel is further supporting this vision by contributing to various open source projects, including Open V-Switch, DataPlane Dev Kit, and various efforts under the OPNV banner.

WORKING TOGETHER TO TRANSFORM NETWORKING

In 2017, Intel brought its vision to two innovative partners, Casa Systems and Kyrio, a wholly owned subsidiary of CableLabs. The three companies collaborated on a unique proof of concept (PoC) that employs a virtualized, converged, and distributed network. The PoC is built around an open source-based platform running software from Casa Systems and Kyrio that brings together several access-related virtualized network functions (VNFs) into the same management domain.

By combining virtual CCAP (vCCAP), vWAG, virtual evolved packet cores (vEPCs), and other VNFs running within an open-source network functions virtualization infrastructure (NFVi), providers can achieve a highly scalable, programmable, and automated infrastructure. They can also be expanded for other operator needs through flexible and generalized SDN/NFV principles. Excited by the results, Intel, Casa Systems, and Kyrio unveiled the PoC to great interest at the 2017 CableLabs Summer Conference.

A CLOSER LOOK: THE CONVERGED BROAD ACCESS PLATFORM

Wi-Fi first and other unlicensed spectrum-based offerings, such as the citizens broadband radio service (CBRS), enable cable service providers to use the existing hybrid fiber-coaxial (HFC) infrastructure for backhaul offerings to MNOs as well as for branded mobility services. That means those same service providers can deploy converged access and core VNFs, which are needed for both fixed and mobile or wireless convergence on a common virtualized environment.

There are significant benefits that come from using a common, standards-based infrastructure, including the hardware, the software, and the management system in the head end and the data center for both fixed and wireless services. By taking this route, service providers can reduce the cost to reuse transport for wireless as well as fixed broadband.

Other opportunities and benefits are made possible as we move toward 5G. Service providers are interested because such functionality can be deployed closer to where the users actually are, meaning on-premise or at the edge of the network. In the end, an access-agnostic multiservices core will enable differentiation in terms of converged services, unified policy enforcement, and lower total cost of ownership.

MULTISERVICES VIRTUALIZED CORE: USE CASES

Advancing state-of-the-art, any-access networks promises to serve an expanding range of industries and applications, from the enterprise to hospitality, healthcare to the Internet of Things (IoT). With IoT, specifically, we are seeing a large-scale deployment for a mobile scenario, though it could also be attached to a fixed environment.

There are also advantages to be enjoyed when it comes to critical emerging applications such as vehicle-to-vehicle and vehicle-to-infrastructure communications. These and other new uses increasingly require ultra-low latency, high reliability and high speed, commonly only made possible by successfully positioning the solution closer to where the communications occur. With enhanced mobile broadband, users will be able to tap into higher data rates and higher-density mobility, as well as seamless convergence wherever they are.

A POC THAT LIGHTS THE WAY FORWARD

The PoC introduced at the CableLabs summer conference is garnering considerable industry interest because it addresses today's expanding functionality demands by illustrating a converged network service. The PoC demonstrates a voice-over Wi-Fi (VoWiFi) call with LTE handoff, leveraging an R-PHY node for backhaul of the small cell and a vCCAP, vePDG, and vEPC running in one software build. In this case, the R-PHY, mounted on a strand in the field, is connected to the Intel® Puma™ cable modem in a home or office. Casa Systems provided its Apex Small Cell to emulate an environment with CBRS for indoors and a licensed spectrum outdoors. Users can be latched into Wi-Fi first, move into a small cell environment that is a licensed or CBRS spectrum, and then roam back down to the Wi-Fi.

Wi-Fi Call via Casa Systems' R-PHY Node and Cable, Mobile, Wi-Fi VNFs in Kyrio OPNFV-based NFVI

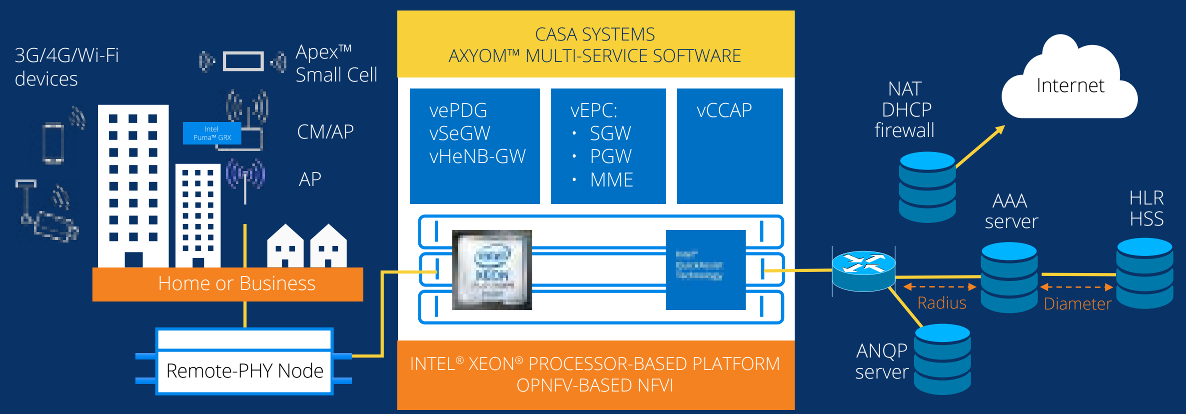

Casa Systems' Axyom multi-service platform, based on the Intel® Xeon® processor, uses fixed infrastructure to deliver Wi-Fi-first services

With regard to software, the PoC connects the R-PHY to a vCCAP core. Casa Systems' vCCAP and vEPC are installed in an OpenStack cloud running on a standard x86 server platform provided by Intel.

The net of this unique design is that the PoC is able to showcase Wi-Fi-first scenarios, illustrating how users can roam into a cellular network and then out into a Wi-Fi network using the cable backhaul that exists in the infrastructure.

CASA SYSTEMS TECHNOLOGIES

Axyom™ Ultra-Broadband Cloud

Casa Systems provided not only the Apex Small Cell for the LTE portion and the DA node (R-PHY node) for the backhaul, but also key fixed and mobile VNFs in the Axyom Ultra-Broadband Cloud. All the VNFs from Casa Systems (vEPDG, vEPC, and vCCAP in the case of this PoC) make use of Axyom's common software framework, including the management capability and the security and routing functionality.

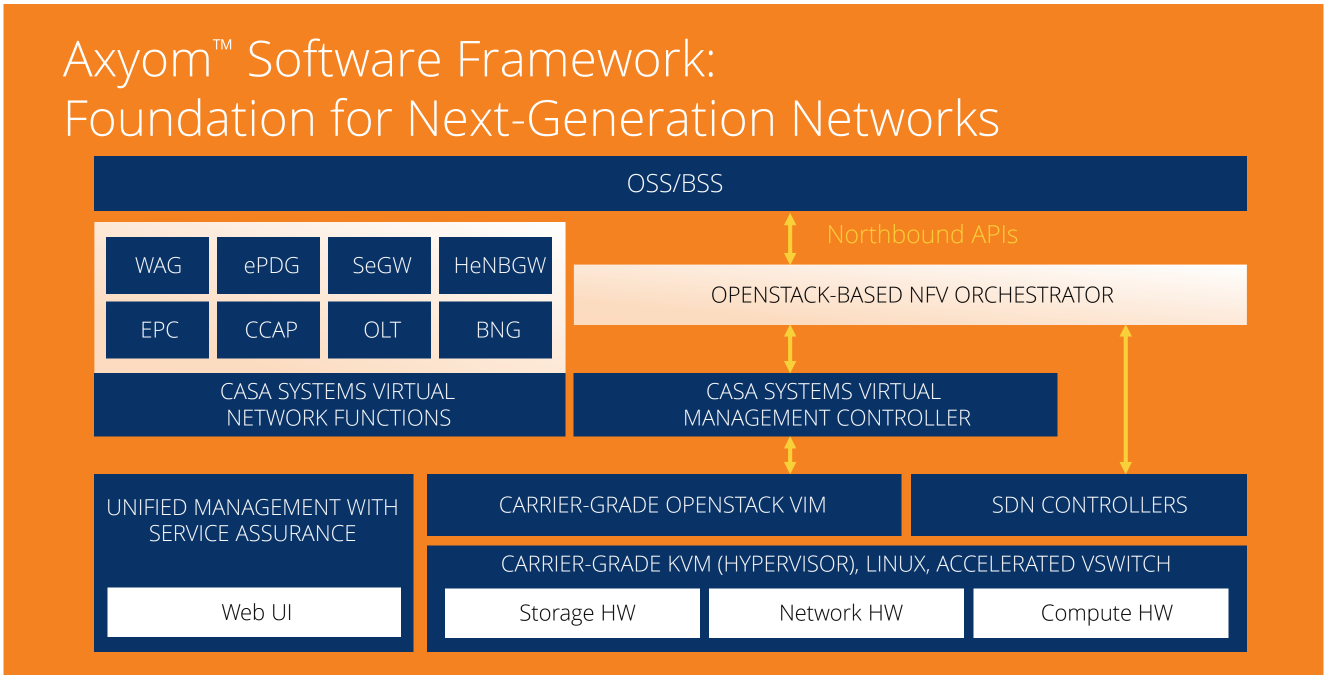

The VNFs in the PoC exist within an OPNFV OpenStack cloud. Installation was accomplished with an early release of CableLabs' SDN/NFV Application Platform and Stack™. The SNAPS™ package adds an installer and test facilities to OpenStack, providing an open-source cloud platform for implementer development and operator deployments.

The SNAPS virtual infrastructure manager controlled the operation of the VNFs in the PoC system. There is also the Casa Systems Virtual Management Controller. This serves as the virtual network functions manager (VNFM) and can work with an orchestrator to allow users to manage the VNF, whether at the central core network, at the edge, or on premise. That means users can start making the move from a proprietary network entity to a virtualized entity to a 5G slicing mechanism.

POC KEY INSIGHTS

The PoC developed by Intel, Casa Systems, and Kyrio represents a real system stood up with real hardware at the Kyrio NFV Interop lab. That collaboration and ultimate integration of the demo led to significant insights sure to help accelerate understanding and advancement of similar efforts.

NFVI/installer choice

COTS vs. OS matched to internal expertise and VNF requirements

The first hurdle to overcome with such a project is determining what you are going to do for a virtualization infrastructure. Are you going to go with COTS hardware and an open source DIM? Or perhaps you are considering a vendor-provided infrastructure that is either purpose built or OS-matched or

OS-tuned hardware? This is an important decision because choosing open source, for example, involves embracing software and application development—and implies sufficient in-house expertise regarding software development projects, tool chains, and version control, among other issues. If you do not possess the appropriate background or skills, you will need to partner with the right third-party provider for support and guidance.

NFV infrastructure

Stack sizing

This project used six COTS servers for the stack, one for a build server to install OpenStack, a controller node, and five compute nodes. Given the aim of the project, the team did not require high availability, which would have occupied three controller nodes, leaving just two compute nodes. The takeaway is that it is important to think and plan for the size of the stack needed. This could vary anywhere from a desktop if the project is a small-scale experiment to several racks of servers to support a software application.

Network interface cards (NICs) and switch speeds

The PoC relied on a 10 GB switch and 10 GB network interface cards (NICs). This proved to be about as low as the team could realistically go as OpenStack demands sizable networking resources to image the nodes, install OpenStack, and handle internode communication for the cloud and operation. The lesson learned here is that 10 GB is the minimum, with 40 GB and up being preferred if you are a carrier and considering deployment.

CPU pinning

During the integration, the team learned that simply assigning a certain number of cores to a flavor in OpenStack and then building a virtual machine (VM) on that flavor does not guarantee those cores are dedicated and isolated to that VM. It is necessary to go down in to the Linux OS and group the CPUs using NUMACTL and then go back up into OpenStack to make sure that the groupings are assigned to the VMs and are dedicated and isolated cores.

Four layers of networking

Networking got complicated in this project. First, there is the environment in which the system is going to live, which would include IT-controlled switching, routing, gateways, and the assignment that goes with that. Second, upon setting up the stack there are server nodes that are running a Linux OS, each one with NICs with subnets that must be assigned correctly. Third, there are Open V-Switch and OpenStack, which lay bridging interfaces on top of the NIC interfaces with subnets that must be configured correctly. Finally, in the fourth layer, there is the OpenStack project cloud with subnets configured for the VMs, which need network configuration and management.

In combination, these layers add up to a sizable amount of networking that needs to be done to make sure that the packets get where they need to go, both inside and outside the stack. Unless you bring the requisite expertise with all of those layers, it may require putting together a cross-functional team composed of IT, software developers, system administrators, and others to promote the necessary synergies and desired results.

Automation and orchestration

For the project, the cloud infrastructure was OPNFV-aligned OpenStack deployed by the SNAPS installer, which has been contributed to the open-source community. Casa Systems' software applications were deployed via Heat templates.

The lesson here is that if you wish to actually deploy in trial or attempt a commercial deployment, it is absolutely necessary to have orchestration and complete automation. There is no way to do a real deployment by hand. You must have an orchestration layer integrated. This is why the PoC team has undertaken additional work to add an orchestration layer to the PoC.

PROJECT CHALLENGES

Networking issues

Virtual large area network (VLAN) vs. Virtual Extensible LAN (VXLAN) and provider networks

The PoC team used an extreme summit 670 10 GB switch, which provided the necessary connectivity and speed. However, it did not offer native VXLAN support. It turned out that a network requirement of one of the VMs in the core was layer-2 connectivity to the radio-frequency interference (RFI) node, the remote PHY device (RPD). Something outside the cloud project needed layer-2 connectivity into one of the VMs in the cloud, a requirement that would have been made easier with VXLANs.

DHCP, port, and protocol security

By default, OpenStack disables dynamic host configuration protocol (DHCP) traffic from within the project cloud to anything outside. It assumes that OpenStack will handle IP addressing. But when trying to run a DHCP server as a network application in a project cloud, it is necessary to configure the OpenStack security policy to enable DHCP communication. It is important to take a close look at the VNF communication requirements and then explicitly and correctly tailor the OpenStack security policy to meet those requirements.

Layer 2 vs. layer 3

This refers again to the situation where a network element outside of the cloud needs layer-2 connectivity to a VM in the cloud. As noted earlier, there are different ways to address this challenge, depending on the requirements and what the hardware and NFVi will support.

Maximum transmission unit (MTU) size

The MTU for Ethernet is 1,500 bytes. DOCSIS 3.1, which was the access network used in the PoC, specifies an MTU size of about 2,000 bytes. That required the team to go into the stack and increase the MTU size. The process proved to be more complicated than anticipated. The first step was to set the MTU size at the tore switch. It then had to set it on the Linux nodes of the OpenStack hosts and in Open V-Switch and OpenStack. Finally, the team had to go into the VMs within the project cloud and set the MTU size there as well.

Installer capabilities and provider networks

For this project, the team started out using the OPNFV Apex Installer, which is the default installer for the new release of OpenStack, only to learn that it would not actually permit the assignment of provider network IPs to VMs within the cloud. These are the external IPs that facilitate communication from inside to outside the cloud.

That constraint led the team to move from the OPNFV Apex installer to CableLabs' SNAPS package and installer, which does permit assignment of provider network IP addresses directly to VMs. While likely specific to this installation, the issue does underscore the sort of hurdle you can encounter when standing up a system in a particular environment.

Monitoring and metering

The tools and resources available in OpenStack to monitor and meter both virtual and physical resources are limited. The demo team used the command line in OBS and Docker and the different OpenStack modules, such as Nutron, and looked at VMs to determine processes and process help. The team is setting up Wire Shark on VMs within the cloud to check connectivity within the cloud. Mirror ports on the top of rack (TOR) switch are being established to see what is coming in and out of the project cloud. While pleased to make it work, the team realized that if deployment is the goal, monitoring and metering must be a part of the system architecture.

POC POSITIVE OUTCOMES

1. CableLabs' SNAPS is a working NFV infrastructure manager

Integrating the PoC system within the SNAPS environment demonstrated that SNAPS is a viable platform for network function virtualization. Notably, it is open source code that is aligned with ETSI/OPNFV. It provided a stable virtualization platform for the PoC, and with suitable configuration and hardening could serve as a commercially deployable virtualization platform.

2. Kyrio NFV Interop Lab is an efficient development environment

The Kyrio NFV Interop Lab proved itself to be a worthwhile option for facilitating this kind of multivendor collaboration. The lab proved it can provide public IP access into the systems, sufficient tools and permissions to remote engineers, and facilitation and project management functions for implementation.

3. Partner-identified advantages

Partners Casa Systems and Intel also believe the Kyrio NFV Interop Lab greatly benefited the project, providing the sort of safe haven necessary for truly effective collaboration. CableLabs is rigorously vendor neutral with respect to technology development. That commitment helped the team focus on the goal of the PoC, the engineering, and the creation of a system that actually worked, and do it within a very compressed time frame of six to eight weeks.

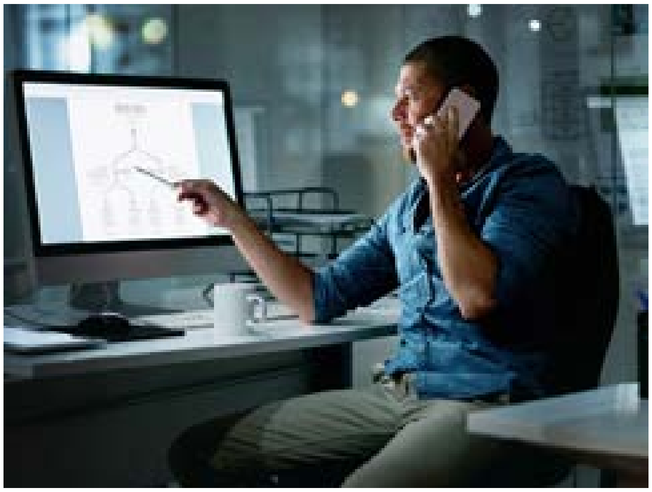

THE PROMISING RESULTS OF THE POC

The goal of the joint PoC was to show a working end-to-end solution that contained both physical and virtual elements to demonstrate the promise of virtualized converged networks. But what makes this type of system exciting and part of the bigger selling point of an SDN and NFV approach is that once this type of software infrastructure has been established to orchestrate and deploy the VNFs, including the vCCAP, vEPC, and vePDG, the same infrastructure and process can be used to deploy any type of VNF.

If a network needs a change, updates to the ability to assist can be made in real time. If users need more capacity on a particular access technology (e.g., DOCSIS), more vCCAP cores can be spun up. If those users need a CBRS connection or Wi-Fi capabilities and capacity in the system, the VNF can be added into the system to make that possible. And when the needs diminish, downward adjustments can be made and new software workloads brought on. In short, systems of this kind free users to unlock the potential of an existing infrastructure.

The success of and excitement around the demo have led to conversations with operators from both the technical and operations disciplines as well as at the executive level. That interest has continued beyond the summer conference, prompting a move toward adoption and commercialization. This includes a promising field trial, two additional collaboration projects to implement orchestration for the vCCAP and mobile core, and a pair of potential operator engagements.

The PoC represents new thinking and inroads on the way to developing VNFs and all the supporting software. But it is also about taking virtualization, distributed access, and mobile core and creating real systems and advancing our understanding of how those new systems can be deployed on real networks.

ACCELERATING THE PATH FORWARD

The call for network transformation is here as a massive influx of devices and use cases puts tremendous pressure on your network infrastructure. How do you cope with the volume and complexity of that data? How do you enable the next generation of services? Intel is working with partners like Kyrio and Casa Systems on advancing the idea of network functions virtualization and software-defined networking toward cloud-ready networks.

These new networks are going to be scalable enough to cope with the heightened demands, the escalating traffic, and the varying levels of performance and QoS that will be needed for different types of traffic. Readying your network and doing so in a cost-effective manner will require new levels of automation, easier and faster development and deployment of new services, and trusted security. Demos like the one stood up by Intel, Kyrio, and Casa Systems are helping show the way.

To learn more about the partners and their technologies, visit their websites:

intel.com

kyrio.com

casa-systems.com

Intel, the Intel logo, Intel Puma, and Xeon are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.
*Other names and brands may be claimed as the property of others.
© Intel Corporation
© 2018 Casa Systems. All rights reserved.
Casa Systems and the Casa Systems logo are registered trademarks of Casa Systems, Inc.