Software Defined Networking and Network Function Virtualization

Network and Telecommunications Community of Interest

FAA Industry Partnership / ACT-IAC FTI-2

Technology, Performance & Operations Working Group

Date Released: May 18, 2017

Synopsis

The Federal Aviation Administration (FAA) Telecommunications Infrastructure (FTI) network is an essential component for the day-to-day FAA mission. In this digital era where quick access to information is essential, the FTI network has grown significantly and is used to transport data, services, and information for the National Airspace System (NAS). Whether the follow-on network FTI-2 is completely on-premises, cloud-based, or a hybrid, it provides the foundation communication links that the FAA needs in order to run their applications and deliver services.

This paper examines Software Defined Networking (SDN) and Network Function Virtualization (NFV) technology as components of automating network and cloud service provisioning for enabling increased levels of accuracy in rapid services turn-up and restoration with a flexible, adaptable and scalable network infrastructure. While we are discussing both technologies in this one paper they are separate technologies that can complement each other but are not required to be deployed together. An example would be to deploy SDN to the current FAA fiber infrastructure to allow more efficient use of fiber for redundancy, without deploying NFV in the network. Conversely NFV can be deployed as network or customer premise devices without SDN control.

To illustrate the progress of the adoption of SDN/NFV technology, two Service Providers have provided input on the current status and potential future states for SDN/NFV technology in their networks. Both Service Providers examples show some of the benefits of deploying an integration of both technologies. The group expresses our gratitude for these two submissions.

American Council for Technology-Industry Advisory Council (ACT-IAC)

The American Council for Technology (ACT) – Industry Advisory Council (IAC) is a non-profit educational organization established to create a more effective and innovative government. ACT-IAC provides a unique, objective and trusted forum where government and industry executives are working together to improve public services and agency operations through the use of technology. ACT-IAC contributes to better communications between government and industry, collaborative and innovative problem solving and a more professional and qualified workforce.

The information, conclusions and recommendations contained in this publication were produced by volunteers from industry and government advisors supporting the objective of more effective and innovative use of technology by federal agencies. ACT-IAC volunteers represent a wide diversity of organizations (public and private) and functions. These volunteers use the ACT-IAC collaborative process, refined over thirty years of experience, to produce outcomes that are consensus-based. The findings and recommendations contained in this report are based on consensus and do not represent the views of any particular individual or organization.

To maintain the objectivity and integrity of its collaborative process, ACT-IAC does not accept government funding.

ACT-IAC welcomes the participation of all public and private organizations committed to improving the delivery of public services through the effective and efficient use of IT. For additional information, visit the ACT-IAC website at www.actiac.org.

Networks & Telecommunications (N&T) Community of Interest (COI)

The N&T COI mission is to provide clarity, impartial feedback, and points for consideration on networks and telecom issues identified in collaboration with the federal government and industry. The N&T COI provides a forum where government and industry executives are working together on key telecommunication issues such as interoperability, information sharing, communications architectures, wireless technologies, converged internet protocol based services, security, and continuity of service. The N&T COI established a working group to facilitate collaboration between government and industry on matters concerning the upcoming FTI-2 effort.

Disclaimer

This document has been prepared to contribute to a more effective, efficient and innovative government. The information contained in this report is the result of a collaborative process in which a number of individuals participated. This document does not – nor is it intended to – endorse or recommend any specific technology, product or vendor. Moreover, the views expressed in this document do not necessarily represent the official views of the individuals and organizations that participated in its development. Every effort has been made to present accurate and reliable information in this report. However, ACT-IAC assumes no responsibility for consequences resulting from the use of the information herein.

This paper was prepared by ACT-IAC after consultation with the Federal Aviation Administration. The information and opinions contained herein are those of the ACT-IAC and are not a reflection of any planned strategy or approach to FTI-2 by the FAA.

Copyright

Further Information

For further information, contact the American Council for Technology-Industry Advisory Council at (703) 208-4800 or www.actiac.org.

TABLE OF CONTENTS

LIST OF FIGURES

INTRODUCTION

The Federal Aviation Administration (FAA) Telecommunications Infrastructure (FTI) network is an essential component for the day-to-day FAA mission. In this digital era where quick access to information is essential, the FTI network has grown significantly and is used to transport different types of voice and data services, and other information for the National Airspace. It is anticipated that the FTI-2 network will be deployed on-premises in both the metro area as well as wide area network to provide communication links and other associated services that the FAA needs in order to fulfill their mission. Software Defined Networking (SDN) and Network Function Virtualization (NFV) are state of the art technologies to efficiently and effectively automate networking and cloud infrastructure providing adaptability, flexibility and scalability of service offerings.

SDN with network elements based on NFV, open network hardware and x86 platforms with Dynamically Reconfigurable - Field Programmable Gate Array (DR-FPGA) support integration of applications and networks for efficiency and performance which may be applied locally, in the metro area and Wide Area Network (WAN), at the data center, or in the cloud. Administering application services is difficult and sometimes impossible with today's static networks based on vendor specific network elements. For that reason, SDN and NFV are being deployed in various networks to innovate telecommunications by integrating networks with applications to reduce the administrative burden, while decreasing the time it takes to deliver new services. The current FTI network and National Airspace System (NAS) includes a large percentage of narrowband channelized Time Division Multiplexing (TDM) circuits and custom applications. (Note: This White Paper uses the term TDM in a limited context and not the more general use of the term in all serial frame and packet based connections.) To be most effective SDN controlled links are based on broadband capacity to support shared services and therefore are not usually associated with narrowband, channelized TDM circuits which usually support single fixed functions. SDN technology has not been widely applied to legacy applications because of lack of software transparency and difficultly in code modifications. However external application profiles can be associated with those legacy applications when they are initiated.

For many organizations, cloud-based application and data centers have been the early adaptors of premise based SDN with NFV due to the velocity of network changes, modest bandwidth requirements and availability of repurposed computing resources. FTI-2 is primarily a WAN acquisition so this paper will focus on the WAN aspects of SDN/NFV technology which is usually a more stringent case than on premise versions. This paper will first provide an overview of the SDN / NFV technology, then provide examples of how SDN/NFV is being deployed with two Service providers, and finally provide some analysis of potential deployment of the technology within FTI-2.

WHAT IS SOFTWARE DEFINED NETWORKING AND NETWORK FUNCTION VIRTUALIZATION?

SDN is a network architecture that separates control plane functions from the data plane migrating those functions which include network intelligence, management and policymaking

to a controller device or application. This architecture tracks what occurred in the Public Switched Voice Network (PSTN) narrowband services in the mid to late seventies with the advent and deployment of SS#7. The control messages were moved to a separate control plane that had north and south bound interfaces to control the circuit switches involved in the voice path as well as receive different command and policy data from associated databases.

For a number of years, packet switched data networks utilized in-band route updates to manage which paths packets traversed through the network. Individual packet routing is a computationally intensive process that does not scale well in large carrier and Internet backbone networks. It was observed that a vast majority of packets occurred in either one directional or bidirectional flows that could be switched through the network with layer 2 or 3 tunnels such as L2TP or MPLS. Label Switch Routing functionality was added to edge and core routers to tag flows and use that tag switch the packets without requiring route path determination at each hop in the path. To better manage an aggregate number of IP flows the OpenFlow protocol was developed and implemented as one of the first standards based SDN technologies. This protocol has been applied at all three layers of the communications stack with optical, electronic/digital and narrow or broadband frame switching, and IP flow switching.

The SDN architecture is dynamic, centrally manageable, cost-effective, and adaptable, making it ideal for the high-bandwidth, dynamic nature of today's applications. The control plane on a network device generally is anything that is needed in order to get routing, forwarding or switching working on that device; in other words, the control plane is the "brain" of the network device. Control plane management packets are destined to, or locally originated by, the network devices themselves. The data plane is called the forwarding plane for packet devices and networks; the forwarding plane forwards packets through a network device that is not destined for that device. In multi-layer networks the data plane also refers to the digital packet or frame transport layer as well as the underlying foundation optical transport layer.

When the network control and forwarding plane functions are separated in the network, network elements have a North bound application interface and a South bound network device control interface. The network device control plane becomes directly programmable from a central controller and enables the underlying infrastructure to be abstracted for both management and service delivery purposes. This standard abstraction can increase the level of automation of a network like FTI-2 network and reduce the level of manual provisioning for network services. For example, when a new network device needs to be added to the network, a centralized controller can automatically identify the device and provide a configuration for the new device. This model provides exceptional levels of visibility into the network as well as planning and business intelligence. Some of the more relevant benefits of SDN are:

- **Programmable:** Control and underlying data planes are directly configurable because the control packets are removed from the data forwarding and switching planes. It allows rapid configuration and management changes, in-depth and distributed security, and optimization of resources quickly via real time

management Application Programming Interfaces (APIs). DevOps can rapidly develop and modify both network management and end user applications to increase the efficiency and effectiveness of the underlying dynamic network.

- **Agile:** Separating the control plane from the packet services, digital and optical transport layers allow administrative and end user applications flexibility in optimizing packet and data flows through the network infrastructure. For example, once a packet flows source and destination ports are identified, single hop digital or optical pipes can be dynamically instantiated to allow packets to avoid a number of intermediate hops.
- **Centralized management:** Network intelligence is logically centralized in Administrative Domain based SDN controllers that maintain a complete view of their network. This provides the network administrators a single pane of glass interface to manage and visualize their network infrastructure. UNI and NNI interfaces support end to end signaling from user through one or more network providers to the destination network.
- **Open standards-based and vendor-neutral:** Open standards simplify network management and operation because instructions are provided over standard protocols by SDN controllers instead of vendor-specific applications and protocols.
- **Virtualization:** NFV extends the benefits of virtualization, currently residing in data centers, to networks, increasing resource flexibility, utilization and reducing infrastructure costs and overhead.
- **Application Centric: This benefit r**-educes Total Cost of Ownership (TCO) for end users and service providers by dynamically matching user requirements with tailored service provider capabilities without stranding bandwidth and other resources. This matching allows total end to end system optimization matching computing, data transport requirements and provider capabilities on all segments of the path. Also matching is accomplished by using business-relevant policy models across multiple service providers, networks, servers, storage, security, and services.

TRADITIONAL NETWORK MODEL

Traditional Networks have the control plane tightly coupled to the network device that results in minimal application programmability for network devices (**Figure 1**). For example, In IP networks IS-IS and BGP protocols pass link availability data in band for intra and external ASN links.

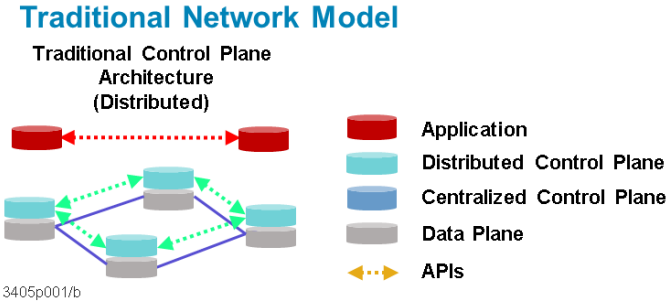

Figure 1. Traditional Network Model (Courtesy of SP1)

CENTRALIZED CONTROL PLANE MODEL

The forwarding hardware uses a control protocol like OpenFlow. OpenFlow, in this example (**Figure 2**), is the agent used by the controller to communicate with the forwarding and switching network components.

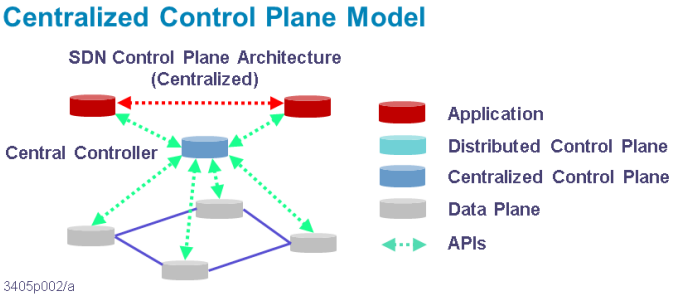

Figure 2. Centralized Control Plane Model (Courtesy of SP1)

HYBRID CONTROL PLANE MODEL

A hybrid approach to SDN allows customers to take advantage of existing investments in platforms while elevating more sophisticated tasks for handling by the controller. Recently, network device vendors have been adding Application Program Interface (API) support to their OS exposing underlying functions and hardware. The API specifies how software components interact and Graphical User Interface (GUI) components control the APIs. This allows customers to use the distributed control plane to perform routine convergence and local repairs while allowing this centralized controller to run global and tactical optimization with resiliency against failures (**Figure 3**.)

The hybrid version of SDN is a lower risk implementation as it also provides a fall back if the SDN central controllers or the control plane reachability are compromised. If the controllers were rendered useless for any reason, the advantages of the central controllers would be lost, but the underlying network would still continue to operate with the traditional brain of the networking device operating as usual.

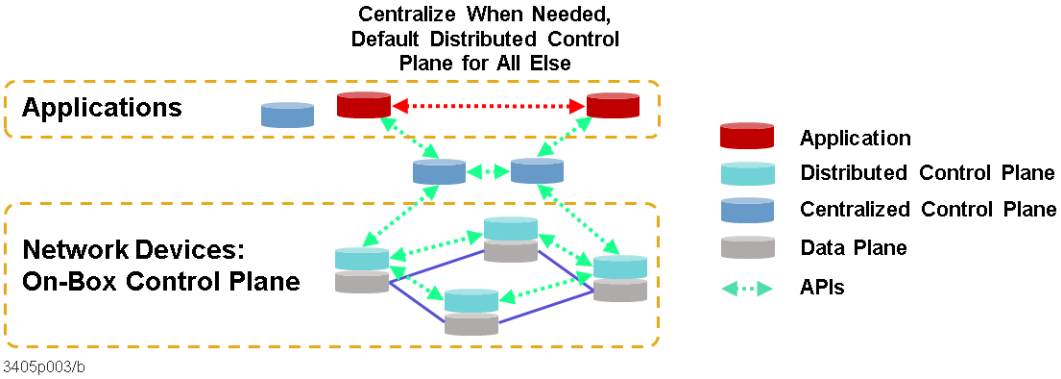

Figure 3. Hybrid Control Plane Model (Courtesy SP2)

SDN removes a layer of complexity by automation and central management encouraging and simplifying required network changes. Standardized SDN and NFV are still in the early stages, but by 2022, SDN and NFV will be more mature technologies and much more prevalent. There are many industry organizations that are working towards a common SDN platform by creating

network standards to make networks interoperable and programmable. We would expect the industry to be much further along with standards by the time FTI-2 is implemented.

(Note: In both of these illustrations, the data plane is the packet services plane and the digital services and optical data planes are not illustrated for simplicity and ease of understanding.)

NETWORK FUNCTION VIRTUALIZATION

The current versions of NFV applications and devices were developed to reduce rack space and cabling in data centers that were consumed by a variety of deployed, single purpose network appliances. Historically NFV were implemented in a common set of hardware platforms with different communication interfaces in which the box's personality was determined by the software loaded in the system. There were different levels of abstractions for the hardware including non-standard and industry standard virtualization technology and straight multi-tasking, interrupt drive embedded operating systems. These systems would have narrow focused ASICs to accelerate functions that were not suitable for software only applications. ASICs are commonly used for frame and packet processing on high speed (1 Gbps and greater speed) links.

There are two popular types of virtualization, Virtual Machines (VM) and containers which are light versions of VMs with less available functionality that require less processing power and fewer system resources. Heavy compute and database applications are not virtualized nor are heavy I/O applications as the virtualization process adds unacceptable amounts of overhead to each operation. Another aspect of clusters of computers in data centers is modular network functions can be distributed around the data center, lessening the load on any one instance and allowing siphoning off capacity from a number of systems without effecting performance or throughput.

If the network device applications require the real-time servicing of system events (hardware interrupts) such as a full packet had just been received, real time embedded operating systems are used without virtualization because of the rapid response required.

In the past, hardware acceleration could only be achieved with narrowly focused ASICs which would require daughter boards or mother board swap outs to add, modify or increase the functionality. Dynamically Reconfigurable Field Programmable Gate Arrays (DR-FPGA) are firmware/software configurable hardware which allows NFV to be deployed with both network device application virtualization as well as appropriate hardware virtualization/acceleration. The smallest form factor NFV would be a small card/enclosure with a DR-FPGA, which can include a full microprocessor on chip, with static memory and I/O interface hardware in separate chips.

In the selection of technology for a particular NFV the hard limits of Space, Weight and Power (SWaP) and environmental issues are usually the deciding factor. A number of non-NFV communications devices will be smaller in form factor than their NFV replacements and may consume less power especially cards in chassis. Those devices located in ATC towers or other

exposed locations may have to be protected against nearby lightning strikes, wide temperature variations and low quality power and ground.

In situ programming of an NFV device requires a directly connected communications link that can be used to upload code and restart the device before it can be used for its intended function. The other option is the device has to be removed from its normal location to a programming station to update the application/firmware/hardware. A number of large wide area networks utilize Frame Relay as an Out of Band Administrative (OOBA) network for updating, monitoring and controlling their network devices. OOBA networks are especially useful when software and other problems develop on the network device isolating it from the data or normal control plane of the network.

The FAA networks have a wide variety of network devices and communication interfaces. NFV within the constraints already discussed, can allow a smaller common set of hardware and software platforms to be utilized to provide a variety of the required services throughout the network. It is envisioned that a mixture of NFVs, with different configurations can be utilized at various points and interfaced in the network infrastructure. Newer purpose built network devices can be upgraded with Open Network technology to allow control and integration into SDN controlled environments. Inline Network Devices may require sufficient performance to use Open Network based purpose built devices supporting routing, switching and security device requirements.

There is a spectrum of Open Network/NFV devices in the marketplace with additional items introduced frequently. Some classes of devices include:

- Medium to higher performance compute servers with multi-application virtualization or single medium performance virtualization with 1 or 2 Gbps optical interfaces
- Medium performance compute devices with single application virtualization with 100Mbps or 1Gbps optical interface
- Small form factor device based on DR-FPGA technology for individual service virtualization
- High performance, medium form factor compute servers with one or more DR-FPGAs for hardware virtualization supporting 1 or 2 10Gbps optical interfaces
- Medium form factor, Open Network controlled CWDM or DWDM NID with 100 Gbps customer side and up to 1.2 Tbps network side
- Medium form factor, Open Network controlled G.709 termination NID
- Edge Network aggregation router, Open Network controlled with greater than 1,000 virtual route instances

Because the copper plant in large segments of enterprise and service provider's networks has reached end of life, it is anticipated that copper delivery systems will be replaced by fiber optic based systems with significantly increased bandwidth. In general, these aggregation points including multiple 10, 40, 100 and 1,000 Gbps optical based bandwidth delivery systems

utilizing CWDM and DWDM platforms with Open Network interfaces which already allow dynamic upgrades to device functionality and capability.

Once older network components are virtualized, it provides the ability to take network functions in and out of service quickly and scale within performance of the system limits. NFV-like technology is attractive because it reduces the number of single function appliances in enterprise networks to a minimum of different scalable platforms.

While SDN/NFV increases the complexity of FAA and Service Provider networks, it allows much more flexible and cost effective matching of user requirements with Service Provider capability. Service Providers can better allocate and manage network resources and move toward intelligent networks that automatically optimize network performance against current resources and turn up additional lambdas and other capabilities as needed. These capabilities also allow the FAA to respond to necessary network changes quickly since network functions will be automated and streamlined using a dynamic service delivery approach.

These SDN and NFV capabilities can be delivered to the FAA from a combination of MSPs, contractors and Service Providers. The next sections outline examples of activities and technology adoption of two Service Providers in the commercial marketplace.

(Editors Comment: 1. Current NFV mantra suggests that commodity x86 boxes have sufficient computing, I/O and storage capability to reduce the technology refresh rate in network edge devices by use of these devices as Virtual Customer Premise Equipment (vCPE). It has been industry experience that the refresh rate on x86 boxes is two to three years while the refresh rate on network devices is three to five years. There is a competition between the White Box (common x86 hardware), Black Box (common or specific network devices with hardware assist) with open APIs and the traditional network device manufacturers that have been making software controlled and upgradable devices for over ten years now and have also added API interfaces to their equipment. Another option for all of this equipment is to integrate in mission applications, hardware and requirements to minimize the total number of devices deployed to small and medium size FAA locations.
2. There is a real need to model and understand network reliability with SDN and NFV in the mix. Complexity and reliability theorems state the larger the number of devices and interactions the more likely there will be an error, break, component or system failure. Most existing fixed or multi-function devices are fully designed, engineered and tested for network roles and are simple enough to have a high reliability factor after any break-in periods. SDN adds additional protocol stacks to the mix, while NFV may add a complete underlying operating system and a myriad of other components and devices.

These are some factors to consider when operationalizing SDN and NFV in production environments.)

Service Provider Prospective — Service Provider One (SP1)

(Editor's Note: This section of the document has been provided by SP1 who is solely responsible for its content. This is an illustration of one vendor's view and implementation of SDN/NFV technology and no vetting or endorsement of this section is implied by its inclusion here.)

Introduction

Changing Requirements for Networking

Current networks deploy a large number and increasing variety of proprietary integrated network appliances. Launching new network services requires another set of proprietary appliances and a long-planning and deployment process to certify the new devices, create new skills and competencies, and find appropriate space and power. As appliances rapidly reach the end of life, the plan-procure-design-certify-integrate-deploy-retire cycle starts again, consuming more resources and up-front costs. SP1 is the operator of one of the largest and most reliable global networks, and foresees rapid and fundamental changes in customer needs, technologies, and best practices for operating networks in the near future. These developments mean that the traditional model for purchasing and managing network services will change as networks evolve.

The network infrastructure now has to cope with massive increases in data traffic. Fifteen petabytes of new data are being produced on a daily basis[1] and all these bits and bytes need to be stored, moved, and accessed at a moment's notice. Sixty-two percent of all Internet traffic will be video traffic in 2015 and this number is only going to grow, jumping to 76 percent by 2018[2]. This, clearly, has serious implications for how network bandwidth is allocated and managed. The infrastructure also has to accommodate the growing numbers of endpoints. By 2018, 1.7 billion mobile users[3] will be connecting to the network and the vast majority of these users will be carrying more than one device and using multiple mobile productivity and collaboration apps.

Cloud-based workloads are also increasing as companies continue to migrate their infrastructures to the cloud. Forty-six percent of enterprise customers anticipate total IT delivery through the cloud by 2016[4] and this means more mission-critical workloads traversing the network than before. All these trends are creating a set of unique challenges for the network infrastructure, which now needs to be intelligent, dynamic, and on-demand so that the network can respond and adapt as quickly as business needs change. The network must enable

[1] Source: CIO Magazine, The New Face of Storage: Cloudy, 2012
[2] Source: Cisco Visual Networking Index, 2014
[3] Source: Strategy Analytics Mobility Report, 2014
[4] Source: IDC Cloudtrack Survey, 2012

organizations to introduce new services rapidly, provide better service and support to their customers, and respond quickly to business and market opportunities.

SP1's Software-Defined Network Architecture

To respond to these changing customer needs, SP1 is transforming its network architecture from the current state to a future state where network services and common infrastructure are used, provisioned, and orchestrated in a cloud-like model similar to that already well-established in data centers. This will offer advantages of available and scalable on-demand, software-driven infrastructure and configuration, and increased use of automation to create and manage services.

SP1's software-defined network initiative relies on two enabling technologies — SDN and NFV. While SDN is an approach to intelligent networking in which network control is decoupled from the underlying data plane and is directly programmable, NFV decouples network functions from dedicated hardware devices.

Taken together, these technologies make it possible to implement changes in the network more rapidly and without having to remove and re-deploy physical infrastructure and the associated capital investment. These services will increasingly become cloud-centric workloads. Because the essence of networking is now provided through software, it becomes possible to provide innovative new capabilities, roll out new services expeditiously, and flexibly scale up or down the network as demand changes.

SP1's Software-Defined Network Architecture Security Approach

A comprehensive, unified approach to security is essential across both SDN and NFVs. Today, the SP1 SDN-enabled network takes full advantage of the existing Providers' security measures and in the near future it will introduce additional capabilities, such as dynamic security controls and software-defined security configurability. The new features will provide a mechanism to respond to security threats in near real time, mitigating the risks to Provider and its customers' networks.

- **Existing security measures:** The SDN controllers, the virtual machines and servers hosting SDN controllers as well as the shared infrastructure are protected with the existing advanced security policies, proven security measures, best practices and architecture that protect the SP1 network from various types of intrusions and vulnerabilities. These existing mechanisms, offering best-in-class, enterprise-grade protection include:
  - *Identity and Access Management* -- prevents malicious hacking by providing secure access management via multi-factor authentication, access control encryption, and historical log view.
  - *Threat Management* -- provides malware mitigation through monitoring and detection of security threats (DDoS attacks, malware *data*, network/application

exploits and malicious network events), analysis of security events to identify real and unknown threats, and alerts on detected security threats.

- *SP1's Trusted Software Program* -- provides a highly secure software development process, software security testing, software configuration controls, and code validation process to help ensure vulnerability-free software implementation. The SP1 Trusted Software Development and Life Cycle Management Program will also manage the transformation to a highly secure SDN architecture. It will address software vulnerabilities including vendor-supplied, internally developed, and open source code. Key security controls to be introduced by this program will include highly secure but agile software development process, software security testing, software configuration controls, and tool-based code validation.
- *Portal Security Tools* -- provide comprehensive security assessment during portal design, development, test, and post-deployment phases. The focus is on identifying common vulnerabilities that could be introduced with code changes, developing code in accordance with accepted security policies, conducting various code scans, and remediating security violations.

The new SDN architecture will augment the existing security practices with new security features, including dynamic security controls and software-defined security configurability.

- **Dynamic security controls:** The elastic virtual infrastructure will enable this capability for the SDN controller. One of the benefits of deploying SDN controllers in a cloud platform is that, if a hacker infects one of them with a malware, the compromised SDN controller virtual machine (VM) can be quickly detected, isolated, shut down, quarantined, and replaced by another dynamically instantiated SDN controller VM. The threat then can be quickly resolved by applying security patches to fix the code vulnerability. Similarly, if a hacker gains access to the SDN controller, the role-based access control function will activate and limit the hacker to performing least-privileged activities that should not significantly disrupt services while SP1 quickly works on shutting down the threat. The virtual infrastructure will also enable dynamic security controls for VNFs. All the features of the elastic cloud platform that protect against a failure also protect VNFs against a potential security breach or attack. Just like with a compromised SDN controller, a compromised VNF VM can be quickly isolated and replaced by another instance of VNF VM. If an entire site is compromised, all the VNFs at that site can be replicated to other sites. Because of the redundant distributed design, this may be necessary only in extreme cases.
- **Software-defined security configurability:** In the near future, the SDN controller will also be able to dynamically modify security measures and policies. For example, in response to an incident, the SDN controller will dynamically adjust the security gateway rules to block malicious traffic or a DDoS attack. This will provide a near

real time mechanism to address emerging security risks and prevent service disruptions.

Transforming the Network

To achieve this software-defined transformation, SP1 is simplifying the network by consolidating previously separate layers within the network architecture and bringing the network edge closer to the customer. In parallel, SP1 is implementing an SDN controller into the network architecture to automate network provisioning and orchestrate changes across devices, locations, and services. This is the key to simplifying their ordering and provisioning systems infrastructure, thereby cutting out manual steps and making it possible to turn up new services or changes to existing services in near real time. With SDN-enabled automation implemented in the access and transport network, orders can flow electronically from the customer's request directly to the pertinent network elements, with service activation times measured in minutes.

In addition, SP1 is moving network functions from hardware to software elements that can be dynamically instantiated on a common infrastructure, when and where needed. Initially, we are transforming applications that support existing monolithic control plane elements such as route reflectors, DNS servers, and DHCP servers. Overtime, a broad range of network edge and middle box functions are expected to migrate to NFV infrastructure, including broadband network gateways, IP edge routers for services like IP-VPN and Ethernet, and load balancers and distributors. Because these elements do not typically need to forward large aggregates of traffic, their workloads can be distributed across a number of servers, adding an elastic capability that is unattainable in a monolithic model. This transformation will enable SP1 to create a completely different customer experience that is dynamic, responsive, self-managed, and on-demand (**Figure 4**).

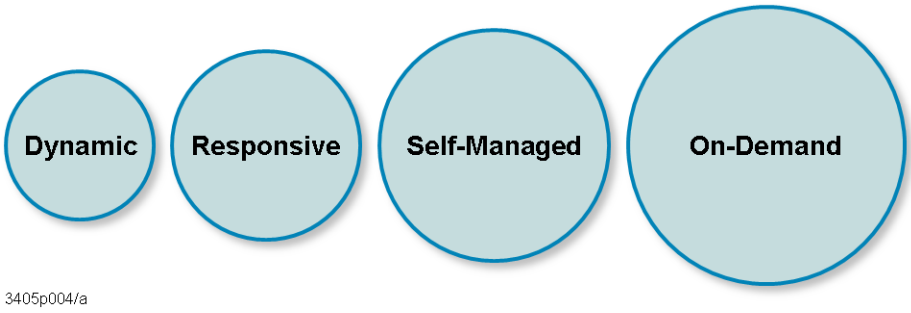

Figure 4. Transforming the Network (Courtesy SP1)

SP1's SDN Approach

From Service Provider's perspective, SDN promises to add value and greatly improve the way networks are currently managed, not only from the perspective of the carrier, but also the consumer. High-level SDN controller languages made accessible via the emerging SP1 architecture will simplify configuration, ease the introduction of nimble policy control, reduce

errors, and enable more real-time changes in the network. Using controller-to-network interface standards like NETCONF, for example, will make it easier to not only manage and maintain networks, but also to incrementally introduce improvements in a non-disruptive manner.

SDN in this network is used for two basic classes of functions: traffic management and network configuration. TheSP1 SDN traffic management implementation is based on a fully distributed design with two levels of control — basic and global. The basic level of control (routing layer) leverages the very mature and highly scalable MPLS architecture that offers end-to-end Class of Service (CoS), fast restoration, and resiliency options. The global level of control (SDN control layer) augments the existing MPLS architecture, using network and application analytics to enable traffic optimization and new services and capabilities such as bandwidth on demand, dynamic traffic redirection, diverse routing, and just-in-time provisioning.

The SDN control layer creates a global view of the entire network topology by gathering information in near real-time from the routers and using this information to identify additional traffic optimization opportunities based on the knowledge of the global state of the network. As opportunities are identified, the SDN control layer will communicate with the network nodes to update their routes. In the unlikely event that the SDN control layer fails for any reason, the existing distributed traffic engineering and fast restoration mechanisms will continue to work for existing services, which will be appropriately rerouted.

The SDN controllers are virtualized network functions distributed in the shared infrastructure to provide the necessary level of resilience. Distribution ensures that, even if one instance of SDN controller is compromised, the controller functionality is still available from the other SDN controller instances deployed in the shared infrastructure throughout the network.

Ensuring SDN's Reliability

An SDN controller can potentially be compromised in two ways — by failure of the controller itself and/or by security compromise. The Service Provider's software-enabled architecture helps address those potential failures.

Reliability is built into the SP1 SDN architecture in a number of ways.

- **SDN Control Layer Redundancy:** The global level of control (SDN control layer) has redundancy. The SDN controllers are virtualized network functions distributed in the shared infrastructure to provide the necessary level of resilience. Distribution ensures that, even if one instance of SDN controller is compromised, the controller functionality is still available from the other SDN controller instances deployed in the shared infrastructure throughout the network.
- **Quick Recovery Time:** Just as it is important to have very low failure rates, it is also important to have minimum downtime when a failure occurs. The SDN controller resides inside a virtual machine. If a virtual machine fails, a new replica can be

created in near real-time. The combination of distributed design and near-real-time replication means downtime is greatly reduced.

- **Availability of Spare Capacity:** The infrastructure leverages a cloud platform, which means there is ample pre-provisioned spare capacity. The cloud platform is engineered with sufficient capacity to add new SDN controllers when needed. Engineering accounts for maintaining high throughput even under failure scenarios.

The SP1's NFV Approach

The second enabling technology that SP1 is introducing into their infrastructure is NFV. NFV separates the network functional software, called VNFs from the underlying hardware platform. This represents either a one-for-one mapping of an existing network function or, in the future, a combination of new and existing network functions.

NFV makes use of virtualization technology to place various network functions such as routers, switches, gateways, WAN accelerators, VPN concentrators, DDoS protection, firewalls, and tunneling elements onto industry-standard high-volume servers, switches, and storage that can be located in data centers, network Points of Presence (PoPs), or customer premises. The virtual devices can be moved to, or instantiated in, various locations in the network as required, without the need for installation of new equipment. VNFs will also distribute data services in a way that will eliminate multiple layers/tiers of appliances and provide more flexibility in hardware redundancy.

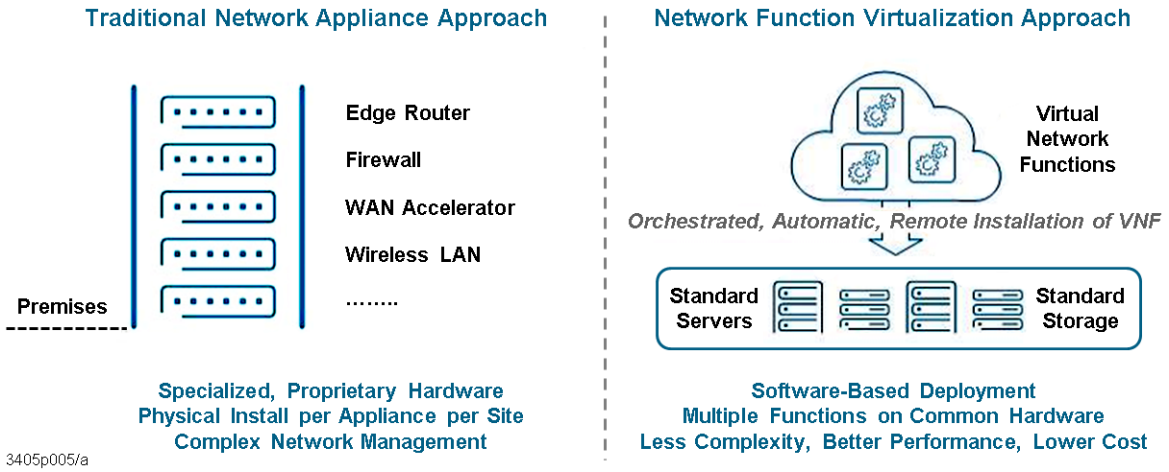

Figure 5. SP1 NFV Approach (Courtesy of SP1)

NFV is important because it means that the FTI-2 network may no longer need to purchase dedicated hardware devices, which simplifies the deployment of network services and reduces both CapEx and OpEx. Having this flexibility allows Agency IT departments to respond in a more agile manner to changing network service demands and implement on-demand usage-based services to meet their business needs.

Ensuring NFV's Reliability

VNFs running in a shared infrastructure can potentially be compromised in two ways — by some kind of failure or capability degradation of the infrastructure and/or by a security attack or breach. SP1's software-enabled architecture helps address those potential failures. While we will always be constrained by physical assets in the network, virtualization offers an opportunity to create high-availability solutions through distribution and replication.

Reliability in the NFV environment is ensured in a number of ways:

- **Reduced Service Restoration Time:** The software-enabled, virtualized infrastructure provides on-demand elasticity so that a failed VNF can be quickly restored. If a VNF fails, a new VNF can be created at the same site in a very short period of time, while redundant VNFs at other sites are continuing to provide service. If needed, VNFs at a failed site can be easily moved or replicated at another available site, ensuring reliability.
- **Distributed Service Deployment:** The software-enabled virtualized infrastructure is designed to be geographically distributed. As part of the automated provisioning and orchestration process, VNFs are replicated at multiple sites and therefore can withstand site failures with minimal service interruption. The number of replicated sites can be specified and adjusted to meet customer service quality requirements.
- **On-Demand Capacity:** The cloud-based nature of virtualized infrastructure helps address network service overloads that result from traffic surges, failures, and malicious attacks. If an overload occurs, cloud sites enable quick creation of additional capacity to meet VNF traffic demands, allowing VNFs to be quickly moved, restored, or replicated as needed.

CARRIER PERSPECTIVE — SERVICE PROVIDER TWO (SP2)

(Editor's Note: This section of the document has been provided by SP2 who is solely responsible for its content. This is an illustration of one vendor's view and implementation of SDN/NFV technology and no vetting or endorsement of this section is implied by its inclusion here.)

We have seen that much has recently been written and discussed about how agencies can benefit from SDN and NFV capabilities. Though it sounds complex, but in reality, this technology is simply a concept of IT services convergence.

SDN and NFV technologies make it easier to bring network, cloud, data center, and applications together, allowing customers to have more control over their services. The following is an overview of what this technology trend means for the FAA.

- **Automation:** NFV helps us to be more efficient and provide the best possible FAA experience. As we focus on selling communications solutions and consulting services on top of our network, hosting and cloud services, more automation and fewer systems will enable us to simplify our processes and provide more innovative services.
- **Self Service:** Greater automation allows for self-service capabilities with more flexible purchasing and bundling options. The FAA will have more web- and phone-based options to purchase what services they need, when they need it, and be able to view the status of their service requests whenever they want to. Bringing our network and the cloud together and making FAA products easy to access, understand and buy, is what self-service is all about.
- **Programmable:** Programmable means paying for what the FAA uses and how much you use it. Agencies can easily add network capacity by managing IT in their own FCS cloud. When demand returns to normal levels, the FAA should be able to scale back. This is more efficient for the FAA, avoids network slow-downs or lack of server capacity, and saves the agency money because in a programmable environment, you only pay for what you use. More applications and more flexibility are wins for the FAA and the carrier.

SP2 is aggressively working to transform its existing core network into an SDN- and NFV-based environment with plans to have full global virtualization coverage in its IP core network and data centers by 2018. When fully implemented, SP2will be able to integrate the delivery of our network, hosting, cloud, and IT services through a self-service application marketplace and on-demand environment, creating a customer experience others will find difficult to replicate.

This company has already built its NFV platform in 36 network and data center locations in 7 countries, with plans to expand into a total of 44 locations by the end of the year. SP2 also was one of the first providers to use NFV technology to virtualize security features such as firewalls and a Content Delivery Network (CDN) to distribute video through its company's TV service,

providing quicker delivery of these services. In the coming months, the Service Provider will begin deploying a range of virtualized data and voice infrastructure services, including virtualized Multi-Protocol Label Switching (MPLS) routers and Customer Premises Equipment (CPE) to enable dynamic delivery of scalable services for customers.

SP2 is actively replacing hardware with software over its recently launched software services backbone to quickly deliver new services to enterprises and small and medium businesses through virtualized functions. This platform will bring SDN and NFV functions to core and metro networks, and new next-generation CPE, bringing the company's networking, cloud, hosting, and IT services into an integrated offering that will be accessed through a portal and set of APIs.

SP2 began implementing network virtualization and automation in 2013 when it partnered with five vendors to develop an open environment through common API structures, creating interoperability and enhancing the creation of applications. SP2 is working with multiple hardware partners in the industry to deliver a next-generation Operations Support Systems (OSS) and Business Support Systems (BSS) solution through a single, reliable transport infrastructure which can rapidly fulfill customer demands and create the support needed for implementing NFV and SDN services. SP2 plans to have full global virtualization coverage in its IP core network and data centers by 2018.

Considerations and Potential Benefits in applying the fiber optic based SDN/NFV Technology to FAA Networks

(Note: This section is primarily written from a Carrier perspective with certain editorial corrections and not incorporating other points of view. It is being presented for discussion purposes only and no verification or endorsement of these comments is implied by being included in this paper.)

The current FTI network is based on a private optical backbone infrastructure managed by the current contract holder. The Layer 1 infrastructure consists of leased optical wavelengths terminating in optical switches in the 21 cities hosting FAA ARTCCs. This backbone supports the Operational IP (OP-IP) network that provides Layer 3 services to NAS applications. Locations in the 21 metropolitan areas access the backbone network over optical fiber rings. Remote locations employ leased circuits to access the backbone network. Core and aggregation routers for the backbone network are owned and operated by the contractor at co-location facilities in the 21 cities.

The OP-IP backbone employs a dual-core architecture; each core consists of a distinct set of core routers utilizing a completely isolated routing protocol.

The current architecture evolved from requirements that could not be met by current shared service offerings. These included end-to-end visibility and control of diverse circuit paths in the backbone, protection from system-wide outages resulting from routing table corruption, and rapid restoration and recovery from failures in the network. The FAA is interested in how the

application of SDN/NFV technologies can be applied to shared service offerings in the future to mitigate these current shortcomings. Specifically, the following topics are explored in this section.

- The potential for SDN/NFV functionality to provide the required availability and survivability in the FTI-2 OP-IP backbone through use of SDN control of the optical and packet services plane and backbone with central visibility and real time control.
- The potential for reduced access costs in a shared SDN/NFV backbone network through the deployment of tiered layers of packet services, and dynamically reconfigurable digital and optical transports allowing switched recovery of services and not stranded backup or optical ring bandwidth capacity
- The potential for virtualized CPE or other NFV technology at remote sites to simplify the operations and maintenance functions at these sites and reducing restoration times as well as the need for physical site visits.
- The potential for NFV technology to reduce the complexity and frequency of technology refresh cycles.

BACKBONE AVAILABILITY/SURVIVABILITY

This section explores the topic of backbone network availability and survivability from different perspectives.

Background

A private network (dedicated capacity on a shared fiber or fiber in a shared bundle)) provides the network operator control over key parameters of the network backbone. These parameters include the number of diverse paths out of each core router, the physical diversity of network links, the number of diverse paths between any two core routers, and the alternate routing table between any two core routers. In current shared networks, customers have virtual routing tables at the edge controlling connectivity within their domain, but the core connectivity and routing in the service provider's network is common to all customers and not visible to any customer.

Most if not all carrier class routers and switches are not affected by any customer traffic. While NFV can allow virtual routers created on shared platforms to be dedicated to single customers this limits performance and reliability. SDN, with its central control and visibility, can provide topologies and routing strategies unique to a customer. If for example an Optical Virtual Private Network is implemented by a carrier, then provisioning of the resources can be performed by the customer without carrier personnel intervention.

Carrier Perspective

Although SP1 plans on leveraging NFV technologies to virtualize certain core networking elements such as MPLS Provider Edge (PE) routers, these components will continue to operate as shared, multi-tenant virtual network devices. At this time, there are no plans on that

provider's technology roadmap to support deployment of dedicated, virtualized core network elements such as PE and P routers for individual customers. Other service providers' edge routing devices allow hundreds to thousands of Virtual Routing Instances to be supported simultaneously for better customer isolation.

System Integrator Perspective

One of the reasons that the FAA chose a system integrator as the prime on FTI was because of the tens of thousands of devices and hundreds of applications that is their primary responsibility to operate and maintain. Enterprises, systems integrators, ISPs, Cloud providers and information companies work with a host of different carrier and service providers for geographically dispersed optical networks for both CONUS and OCONUS. As stated above, FTI is a virtual private network on shared underlying optical and IP networks that is managed and operated by the prime contractor. This backbone is actually a set of leased optical wavelengths from multiple carriers (approximately ten carriers). Nationwide access to the backbone requires the prime contractor to lease access services from hundreds of carriers around the nation (~ 200) which is normal for a CONUS wide network implementation.

There are the traditional access providers like Local Exchange Carriers and Competitive Local Exchange Carrier (CLECs) as well as a new breed of fiber access provider with good fiber coverage but in limited areas around the country. The prime contractor can be the primary point of contact to these new fiber carriers to develop fiber to the site either as the primary fiber route or as the redundant route. Site bandwidth can be shared across both links to verify functionality in the loop.

One option for the prime contractor and the FAA would be to transition T-carrier services to fiber delivery and put automated mini Digital Access and Cross Connect systems at various access PoPs, co-location facilities and carrier hotels to groom and switch the T-carrier based services for the FAA until those services can be retired or transitioned to other transport types.

The widespread deployment of Dense Wavelength Division Multiplexing (DWDM) systems based on Photonic Integrated Circuits (PICs) has dropped the cost and size of DWDM terminals while the bandwidth have increase several orders of magnitude. Today for example there is a 1 RU (rack Unit) whose customer interface is 12 either 100 Gbps Ethernet or Carrier Ethernet and the network side is a DWDM, coherent super channel with 10 tightly spaced wavelengths providing 1.2 Tbps (Tera bits per second) of bandwidth. By stacking these systems together one fiber can carry 26.7 Tbps of data.

It is envisioned that the FTI-2 backbone could include switched lambdas (colors of lasers) on a switched DWDM foundation under SDN control with fiber to the premise for a number of sites including the rural ones. As the PICs are integrated into fiber in the loop systems and cheap fiber provides sufficient bandwidth and can be strung on telephone poles, trees or direct buried

the cost for remote site to be connected is dropping expanding the footprint of these rural carriers.

Access Cost Reduction

Background

Because of FAA reliability and redundancy requirements in each airspace, the current FTI has two point-of-presence serving the facilities in the metropolitan areas of the Air Route Traffic Control Centers (ARTCCs). In many cases, however, remote sites must lease inter-LATA (Local Access Transport Area) circuits to access the backbone network. With a multiple carrier implementation, there can be a Point-of-Presence (PoP) in each LATA, potentially reducing the leased circuit costs. As the cost of fiber optic access has allowed wide spread residential deployment, aggregating circuits on to redundant fiber systems can further reduce both deployment and operational costs.

System Integrator Perspective

The FTI backbone actually has two PoPs in every airspace, and they are usually in the same LATA. The two PoPs are located near the hub of each airspace, i.e., the ARTCC facilities. Nearby facilities such as TRACONs (Terminal Radar Approach Control) and ATCTs (Airport Traffic Control Tower) also use these two PoPs in many cases. The issue is in the thousands of small sites that potentially have to cross LATA boundaries to reach the backbone PoPs that need a more cost-effective solution from potentially alternative carriers or technology.

The question of what effect SDN/NFV has on the architecture and number of core and satellite nodes and how that affects costs depends on the status of network equipment in the different vendors' service offerings. Additionally, the quantity of remote FAA facilities will always make ubiquitous coverage challenging.

Carrier Perspective

By utilizing a carrier-centric, network-based approach for its connectivity needs, the FAA will have the ability to leverage all available network access PoPs) meet-me-rooms, carrier hotels and cable vaults, supported by multiple service provider and hence potentially reduce associated network access costs. Both Service Providers (SP1 and SP2) intend to continue to rollout SDN and NFV capabilities across their network platforms over the course of the next several years. To date, SP1 has enabled SDN capabilities in more than 1,300 wire centers across 100 LATAs within its multi-state footprint. Expansion of these capabilities beyond this and other provider's footprint will be subject to general availability and timing of SDN with UNI and NNI capabilities by other Interexchange Carriers (IXCs), Competitive Local Exchange Carriers (CLECs), Local Exchange Carriers (LECs) and alternative carrier both domestically as well as internationally.

O&M at Remote Sites

Background

Radio, navigation and other vital systems at unmanned remote sites must perform a variety of functions and can contain multiple pieces of hardware, which may vary from site to site. Maintenance and restoration activities often require a site visit, which is both costly and time consuming. Remote sites often do not have clean power or grounds and have limited data communications capability which may have to traverse radio or satellite links. Typical virtualized devices are typically not hardened and may require rotating mass storage which tend to be too fragile for this environment. With low speed circuits connecting the sites to the network, this does not allow any substantial remote access for re-programmability. Hardware abstracted NFVs may be tested in this environment but with the current states of technology it may not be sufficient for these tasks.

Hardware Provider Perspective

Over the past several years there have been two trends under taken by network equipment vendors. One is the combining of several network and security functions along with orchestration software on a common platform and, the other, in the data center is disaggregating these platforms into distributed services on multiple computing resources in those same data centers. These traditional and new vendors are also developing virtualized CPE with hardware acceleration to support a lower cost medium capacity devices for enterprises and network edge applications.

For device management, network monitoring and control, these systems are deployed with APIs that interconnect with higher level orchestration software to allow rapid configuration changes, validation and turn-up of services. The FAA has developed a number of standard and one-off configurations for their different types and classes of CPE which means a very large number of unique configurations required for the operational network. These configurations are developed and tested at different FAA or contractor facilities around the country and potentially modified during installation and service turn-up. These current processes assume a standard hardware/firmware/software configuration that has been thoroughly tested and validated in the field.

A substantial amount of the equipment installed today allows for remote configuration, but is typically done on a box-by-box basis. NFV promises to allow for a centralized configuration at the controller level which will automate CPE box configuration for carriers that deploy the same box with minor configuration changes to a number of customers. This new method would likely result in lower Operations & Maintenance (O&M) costs for them.

Since standards based SDN is a new development in carrier networks, actual use cases with a balanced risk approach and associated savings versus consequences are not readily available.

As we get closer to the implementation of FTI-2, there may be a clearer picture of the tradeoffs of risk between the promised benefits of NFV in O&M, weighed against the potential risks.

Carrier Perspective

(***Editor's note: This is a generic view of service provider furnished CPE devices. These devices can be equally provided and managed by contractors, Managed Service Providers or GFE. No endorsement of any of these methods is implied by this section.***)

From SP1's perspective, the use of NFV will eliminate the need for proprietary/purpose built hardware appliances and enable networking component/functions such as routers, switches, firewalls, load balancers, WAN accelerators, content delivery systems, and many other network functions to run as software on virtual machines.

The ability to run these network functions on appropriately sized servers or other types of NFVs, could potentially result in cost savings, which the network and managed service providers can pass on to enterprises. NFV-based services will facilitate the shift from purpose built hardware appliances to virtual CPE, which can be provisioned, maintained and restored remotely in minutes and without the need to dispatch technicians to customer site, and in general will cost significantly less than dedicated hardware.

SP1 is currently in the Control Introduction (CI) phase of its universal CPE (uCPE) platform launch which consists of several form-factors of generic x86 server hardware that will be capable of supporting two or more VNFs that will be downloadable from the vendors Cloud environment. The supported VNF Software Categories will include the following – Router, Firewall, WAN Acceleration, Remote Access, Wireless LAN, Load Balancer, Visibility, Session Border Controller, and Domain Name System/Dynamic Host Configuration Protocol (DNS/DHCP).

The initial launch of these capabilities will be based on the carrier managed model and will include SP1 support with configuration, fault management, trouble ticketing, reporting, maintenance and MACD (Move, Add, Change and Delete) of the uCPE and NFV platform. Customer managed option will follow later.

System Integrator Perspective

There are at least two major challenges to overcome with this concept:

1. The FAA culture — FAA operational personnel are planning the movement away from narrowband channelized TDM network services to optic and IP-based network services. Over 90 percent of the services are still TDM based. It will be difficult to validate virtualized solutions of any sort, including CPE. Normal review and due diligence processes on the part of the FAA will and should involve FAA operational personnel so they will be completely aware that SDN/NFV capabilities are being used. When incidents occur on the network today, a great deal of emphasis is placed on performing a post-mortem to understand what caused the incident and processes are put into place to prevent it from happening again. In a software, controlled network and many

virtual environments, it is likely that operational incidents will occur as a result of software or protocol issues that do not directly involve FAA and/or contractor personnel. Testing, software and protocol validation as well as integration testing is one of the keys and keeping the complexity of the system manageable is another. These types of incidents need to be identified prior to service turn-up and will cause FAA operational personnel a great deal of concern and even more scrutiny will be placed on testing and post-mortem activities to ensure the these and other types of incident are not repeated.

2. When the mission is, critical and involves public safety, it is important to balance operational concerns with technology benefits. In the past, one of the most important functions for operations is "break-fix," i.e., restoral of service when things break.

   A lesson learned from FTI is that the introduction of new technology can provide benefits, but unfamiliarity with the new technology can often take more time to restore service when it breaks.

   The new SDN paradigm is to monitor and track services especially when they begin to degrade and switch to alternative facilities before an outage occurs. Network links, nodes, access connections can be monitored in real time to alert the network control function to switch around a node or link and signal operators to ascertain the problem so that it can be repaired before the next incident.

Technology Refresh

Background

Typically, the hardware in a private network is owned by the end user or the service provider (carrier or integrator), both in the backbone and on the customer premise edge. Technology refresh requirements can be generated by the need to upgrade the capacity or performance of components or by the End-of-Life designation of the component provider. In either case, the transition to a new set of equipment is both expensive and disruptive to network operations. Software only updates in a virtual environment will not increase the performance of the hardware, it will allow dynamic reallocation of network device and application functionality.

Hardware Provider Perspective

As has been discussed, NFV has the potential to reduce the need for some technology refresh by utilizing a set of computer platforms with hardware assist to perform multiple functions such as routing, WAN acceleration, and security. The ability to combine these functions on a small number of platforms can reduce the number of stocking units and the associated logistics costs.

The actualized benefits on the hardware refresh cycles are still a work in progress as SDN deployments have just begun. These benefits may be somewhat limited by a few physical

factors. Expectations of having average cycles of hardware refresh exceeding 7 or 8 years would also have a potential side effect of having more component failures as we get later into the tech refresh cycle. There is also the possibility that changes in the hardware/software (Moore's Law) may provide opportunities for FTI-2 to deploy new hardware that could provide significant technology improvement in other areas such as security.

For example, a platform that is deployed in 2022 may have a processing capacity of x, and new security software developed in 2027 may require processing capacity of x plus y. In order to provide the enhanced security to the network, it may require a newer platform.

System Integrator Perspective

The technology refresh benefit is significant. It does help to reduce the potential for operational disruption and associated cost of technology upgrades. Most upgrades can be accomplished from the Network Operations Center/Security Operations Center (NOC/SOC) without the need to dispatch field technicians, enabling a minimum effort approach to technology refresh and providing a significant benefit to the FAA should the refresh be chosen.

Carrier Perspective

Service provider networks with SDN enabled functionality with customer based orchestration portals and UNI interfaces to the carrier network allows customer service turn-up, turn-down and moves, adds and changes without requiring carrier intervention. This decreases the cycle time of order, install and turn-up while reducing customer required and carrier resources. Experience has shown that over time the major requirement is for increase performance which requires hardware assists or improved hardware performance. With the appropriate DR-FPGA based NFV capabilities combined with the SDN are not only expected to reduce network installation/provisioning times and simplify network management activities for network DevOps and management teams, but also increase operational agility and reduce Total Cost of Ownership (TCO) for the carrier as well as improve capability at reduced cost to the enterprises.

SDN/NFV Operational Issues

SDN control like capabilities are not new but with OpenFlow, OpenNetwork and OpenStack there are standardized architectures and interfaces that have been put into production over the last five years. Pre-standard SDN like technology has been applied to narrowband channelized T-carrier's systems in the past for switching and grooming of T1, T3 circuits on SONET carriers.

It is envisioned that Managed Service Providers, System Integrators and customer network teams can utilize carrier and other Service Provider customer ordering portals for service turn-up and down, SDN based configurations Moves, Adds and Changes portals to control WAN optical, MPLS and Carrier Ethernet capabilities, features and functions. This of course is prerequisite on more carriers implementing these portals or only utilizing carriers that current implement these portals. Alternative, SD-Wan and other providers with UNI interfaces to

customers and NNI interfaces to other carriers will be able to deploy SDN services more rapidly than others.

SDN and NFV are essentially not applicable to TDM. The FAA estimates show that a very high percentage of the network relies on TDM. Until these applications and services are replaced with IP, the benefits of SDN will not be available. Some of these TDM locations will implement transition routers; routers which can support both TDM and IP. These devices can work in an SDN hybrid mode as described in the introduction section of this paper.

The other potential limitation of deploying SDN will be the FAA's risk adversity. The NAS has historically taken a very cautious approach to new technology and SDN is still a nascent technology at this point. As we look forward toward 2022, the technology is likely to be more mature but may not yet be in the mainstream phase of the technology lifecycle.

OBSERVATIONS AND SUGGESTIONS

Fiber based networks under SDN control, augmented by different NFV technologies, will enable the building and operation of flexible, scalable and user controlled and configurable high bandwidth networks. The new standards based SDN allows Government customers to setup regional or CONUS wide networks using standard protocols like OpenFlow, Generalized Multi-Protocol Label/Lambda switching employed in multi-vendor networks.

While the SDN will continue to leverage various LSP (Label Switched Path) architectures, the addition of a global control layer will enable traffic and networking optimization which are efficiently coupled to computing and storage resources. New services and capabilities, such as bandwidth on demand and near real time provisioning SDN controllers, which are geographically distributed and potentially virtualized, will enable rapid network reconfigurations based on intelligent network decision-making. The SDN architecture helps ensure reliability by providing redundancy in the SDN control, real time path computations through various levels of network transport and capacity (optical, electrical, frame or packet) allowing rapid recovery in times of bandwidth surges or network failures. Recovery scenarios include switching around failed, routing or switching devices, link failures and wider high level IP plane failure.

Anticipated SDN control benefits to the FAA include:

- Rapid Service creations, deployment, failure recovery, path and capacity switching between facilities
- Utilization of GMPLS-UNI and NNI interfaces to control network links end to end optically, at transport or MPLS/IP layer
- If T-carrier circuits are transported over fiber than switching and rerouting of DSx services over the fiber backbone

Impediments to current deployments of SDN include:

- Lack of maturity of OpenFlow and a variety of related protocols that do not have a number of required elements baked into the protocols as of early 2017
- The Open Network Foundation was started in 2011and there is insufficient operational experience to have well defined products, procedures and failure modes
- In these early stages, operational risks are probably higher than traditional security risks and these include:
    - SDN multi-layer control plane with two to three layers for the service provider and two to three customer layers for MSP and end user control.
    - SDN control plane failure modes including maintaining existing services until the control functions are back online

- Inadvertent customer initiated Denial of Service Attack either programmatic or not returning allocated capacity
- SDN, NFV are added layers of software control stacks which are difficult to prove correctness of code, validation, certification and failure modes. More limited functionality in hardware platforms usually bound the parameters of the device to be tested and allows certification with different software loads.

Within data centers, NFV makes use of virtualization technology to place various network functions into industry-standard high-volume devices that can be instantiated in various locations without physical deployment, installation, and integration of new equipment. Different NFV technologies hold promise to achieve more flexible, scalable and adaptable customer controlled network backbone devices and lower speed network interface devices.

With any new technology, there will be some learning curve issues with SDN/NFV. SDN will allow parts of the network to be configured from a single point, which presents a different type of risk profile than is present in today's FTI network. An unchecked mistake at the controller level would have more significant consequences than a mistake at the device level. As discussed, the promised benefits such as decreased costs in operations and maintenance, traffic optimization, and more rapid network reconfiguration should outweigh some of these risks.

In the opinion of the authors this technology should be deployed in FAA development, test and evaluation networks to gain familiarity and operational knowledge of the technologies. A group of stakeholders will be required to build out a realistic evaluation of the network.

Authors & Affiliations

John L. Lee, Lead – Internet Associates, LLC
David Garbin, Noblis Corporation
Mark Graham, Harris Corporation
Don George, CenturyLink
Steve Dempsey, Cisco Systems, Inc.
Luis Garcia, Cisco Systems, Inc.
Marek Jantac, AT&T Corporation

ACRONYMS

Acronym	Definition
API	Application Program Interface
ARTCC	Air Route Traffic Control Center
ATC	Air Traffic Control
ATCT	Airport Traffic Control Tower
BSS	Business Support Systems
CAPEX	Capital Expenditure
CDN	Content Delivery Network
CO	Central Office
CoS	Class of Service
COTS	Commercial Off the Shelf
CPE	Customer Premises Equipment
DDoS	Distributed Denial of Service
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name System
FAA	Federal Aviation Administration
FCS	Fiber Channel Standard
FTI	FAA Telecommunications Infrastructure
GMPLS	Generalized Multi-Protocol Label/Lambda Switching
GUI	Graphical User Interface
IP	Internet Protocol
LAN	Local Area Network
LATA	Local Access and Transport Area
LEC	Local Exchange Carrier
MPLS	Multi-Protocol Label Switching
NAS	Network Attached System
NETCONF	Network Configuration
NNI	Network to Network Interface
NFV	Network Function Virtualization
O&M	Operation and Maintenance
OPEX	Operational Expenditure
OS	Operating System
OSS	Operations Support System
PoP	Points of Presence
PSB	Programmable Services Backbone
SDN	Software Defined Networking
SDP	Services Delivery Point
SLA	Service Level Agreement
TCO	Total Cost of Ownership
TDM	Time Division Multiplexing
UNI	User to Network Interface

VNF	Virtual Network Functions
VPN	Virtual Private Network
WAN	Wide Area Network