FAA Telecommunications Infrastructure - 2

Security Considerations

Networks and Telecommunications Community of Interest (COI) FAA Telecommunications Infrastructure - 2 (FTI-2) Technology and Performance Subcommittee

Date Released: May 18, 2017

Synopsis:

The FAA obtains approximately 25,000 telecommunications services under the existing FAA Telecommunications Infrastructure (FTI) contract that expires September 2017. The FAA, in collaboration with industry, is planning for the subsequent program, known as "FTI-2", to provide all telecommunication services nationwide.

This White Paper was developed in consultation with FAA for the purposes of offering the Federal Aviation Administration (FAA) considerations pertaining to emerging security best practices and technologies that can be leveraged for use on FTI-2.

The paper focuses its analysis and recommendations on the following areas: security governance, FTI-2 technology opportunities, enterprise opportunities, security considerations and implications.

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 and Telecommunications (N&T) COI

The ACT-IAC N&T COIs mission is to provide clarity, impartial feedback, and points for consideration on networks and telecom issues identified in collaboration with the Government Advisory 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 – the replacement for the current FTI contract. The FTI-2 Working Group is comprised of three committees, each with an industry chair and government advisor, who engage with FAA, agency users, and industry to create a body of knowledge to support the government in the FTI-2 effort:

- Acquisition Strategy
- Technology, Performance and Operations
- Transition and Implementation

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 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

LIST OF TABLES

1.0 PURPOSE

This White Paper was developed at the request of the ACT-IAC FTI-2 Technology and Performance Working Group for the purposes of offering the FAA considerations pertaining to emerging security best practices and technologies that can be leveraged for use on FTI-2.

Assessments are also provided with respect to governance and policies that exist within the FTI program framework today, a discussion of prevailing security approaches used on the program and how well they align to the direction of emerging technology, and identify unique enterprise opportunities designed to improve the effectiveness of the security solution for FTI-2.

2.0 GOVERNANCE

Network security is inherently a discipline that must be defined, monitored and enforced by comprehensive policies, directives, processed and guidelines that govern the means, ways and methods followed to engineer, establish and operate a network security architecture. Clear, unambiguous and enforced governance is essential to maintaining the proper, appropriate and approved level of risk within the architecture, and to assure operational objectives and requirements are continuously met.

Governance of the FTI security architecture is derived by the E-Government Act of 2002 and Title III of the Act, entitled the Federal Information Security Modernization Act (FISMA), amended in 2014. Provisions of FISMA are implemented through policies and guidelines established by the National Institute of Standards and Technology (NIST). One of its more important guidelines is Publication 800-53, entitled Security and Privacy Controls for Federal Information Systems and Organizations, Revision 4, released in April 2013.

It is essential that the FTI-2 contract include full compliance with the latest revision of NIST 800-53 in effect at time of award and allow for the ability to expand requirements as new standards and guidelines are released.

The FTI program network is assigned a Federal Information Processing Standard (FIPS) Publication 199 classification Moderate-impact system; meaning that the loss of confidentiality, integrity, or availability of FTI assets or services could be expected to have a serious adverse effect on FAA organizational operations, organizational assets, or individuals. Controls used to define the security architecture on FTI therefore are selected based on this Moderate categorization as set forth in NIST 800-53. This level of categorization is also considered appropriate for the FTI-2 contract.

When considering prevailing FAA information or network security policies, there are several cornerstone policies that must continue to be kept current and strictly enforced during the FTI-2 contract period. These include:

- Information Systems Security Program (FAA Order 1370.82A)
- Implementation of IS Requirements (FAA Order 1370.114)
- Boundary Protection Policy (FAA Order 1370.116)
- IP Addressing Requirements (FAA Order 1370.117)
- Security Control Assessment/Authorization (ISS Authorization Handbook)

For the important policies, the following recommendations are provided for consideration:

- Information Systems Security Program (FAA Order 1370.82A)
  - Last revised in 2006.
  - Recommend revise terminologies ((e.g. Computer Security Incident Response Center (CSIRC) now Cybersecurity Management Center (CSMC))
  - Recognize the National Airspace System (NAS) Cyber Operations (NCO) and provide policy relative to security incident reporting responsibilities for NAS internal network and non-NAS FAA network domains.
  - Recommend define policy relative to the reporting of security events by contractor-managed FTI-2 network provider Security Operations Centers (SOC) to CSMC or NCO as appropriate.

- Implementation of IS Requirements (FAA Order 1370.114)
  - Last revised in 2012
  - Recommend where possible that legacy systems currently waivered from using authorized NAS Infrastructure Services be migrated or replaced with systems capable of using NAS Infrastructure Services to mitigate risk to the NAS.
  - Recommend FAA reassert governance over programs and services that allow the use of encryption of voice and/or data traffic as they bypass existing security controls.
  - Recommend policy be defined that prohibits remote access to NAS internal network elements. Should the FAA wish to establish a Remote Management Access Gateway (RMA) the policy should address it and define policy relative to its purpose, use and access control authority.
  - Alternatively, policy should be defined to allow remote access to non-NAS network elements for disaster recovery and restoration activities only using AO-approved methods and means.

- Boundary Protection Policy (FAA Order 1370.116)
  - Last revised in 2012
  - FAA enforcement must be vigilant and consistent, especially in regards to placement of hosts within the DMZ (internal and external) components of approved gateways. Ideally, all programs should follow the preferred DMZ hand-off approach for flows through the NESGs (see diagram below). Deviations to policy DMZ hand-offs increase risk to the NAS.
  - Suggest adding anti-Malware Protection as capabilities resident within FAA-approved gateways. The use of Anti-Virus is still relevant but not a completely protective feature.

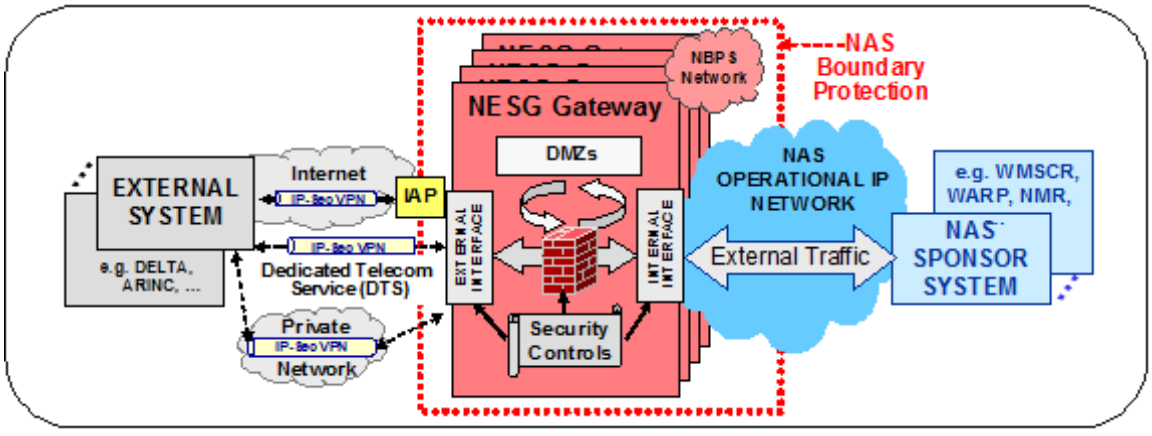

Source: Harris Corporation

Figure 2.0-1. NAS Boundary Protection.

- IP Addressing Requirements (FAA Order 1370.117)
  - Last revised in 2014
  - Current version is relevant

- Security Control Assessment/Authorization (Information System Security (ISS) Authorization Handbook)
  - Revised annually
  - Facilitates implementation of the Risk Management Framework (RMF) and security authorization processes within the FAA
  - Recommend the FTI-2 contract requirements include the latest revision of the "Handbook" and establish a process to facilitate revising requirements as appropriate as Handbook revisions are released. The Handbook is based principally on provisions of NIST 800-53, Revision 4, however the existing FTI contract requirements include compliance with NIST 800-53, revision 2.

Governance from a FTI-2 Contracts perspective could benefit from a consolidation of FAA Security requirements within the FAA Program Management Organization (PMO) that captures and maintains the latest set of standard requirements, policies and processes that all programs must comply with. These standard and consistent set of security contract requirements should reflect the latest federal standards (FISMA, NIST, and FIPS) as well as latest orders, directives and policies issued by the FAA Chief Information Security Officer (CISO) and Air Traffic Organization (ATO) Cybersecurity organizations. Deviations from these standard set of contract requirements would be approved by appropriate security authorities, CISO, Information System Security Officer (ISSO), and Authorizing Official Designated Representative (AODR) to ensure any risks derived are evaluated and accepted by the agency.

Internal to FTI-2 program execution, governance also applies to the security approach used to oversee engineering, operations and maintenance of the FTI network. Security is obviously an essential component to the FTI service construct and by its nature should be embedded within the program team to guide and oversee all facets of service delivery from concept development, engineering design, testing and evaluation, operational integration and cutover, as well a lifecycle operational support through decommissioning. While it is necessary that the security functions be established with separation of duties and functions apart from engineering and operations reporting directly to program management to remove conflicts of interest (real or perceived) , it cannot be separated by contract boundary away from these functions. Security operators must have the direct access, influence and oversight to all elements of the FTI program and its processes to ensure that policies are consistently and verifiably applied, and that the posture of the architecture remains appropriate to the level of approved risk.

3.0 FTI-2 TECHNOLOGY OPPORTUNITIES

The following sections provide consideration regarding prevalent emerging technology opportunities that have been considered for possible use on the FTI-2 contract.

3.1 SHARED AND WIRELESS INFRASTRUCTURES

The current FTI network is based on a private optical backbone infrastructure managed by Harris. The Layer 1 infrastructure consists of leased optical wavelengths terminating in optical switches in 21 metropolitan areas hosting FAA Air Route Traffic Control Centers (ARTCCs). This backbone supports the Operational IP (OP-IP) network that provides Layer 3 services to NAS applications. Locations in the 21 metropolitan areas across the U.S. 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 Harris at colocation facilities in the same 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 requirements 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.

FTI-2 will face many technology transitions over its projected lifecycle that will offer options to use an increasing number of shared and wireless media to transport FTI services. Two of these transitions will likely be the use of cloud-based networks, carrier ethernet and last-mile wireless solutions. As these technologies are evaluated, Security managers must play an integral role in defining an appropriate approach that meets FISMA/NIST regulations and FTI-2 contract requirements. Some specific considerations are provided below.

3.1.1 Encryption

One of the most important recommendations that can be made for FTI-2 is the need to establish and maintain a comprehensive enterprise-wide encryption management strategy. This is especially important as FAA data expands to cloud-based domains that host NAS and non-NAS systems and services. To this end, FAA Order 1370.103, FAA Encryption Policy last released in 2008 should be updated to reflect the latest developments in encryption techniques and technologies, including those that apply to cloud-based providers and other virtualized systems.

3.1.2 MACsec

Media Access Control Security (MACsec) is an industry-standard security technology that provides secure communication for all traffic on Ethernet links and should be considered for use on FTI-2. MACsec provides point-to-point security on Ethernet links between directly connected nodes and is capable of identifying and preventing most security threats, including denial of service, intrusion, man-in-the-middle, masquerading, passive wiretapping, and playback attacks.

3.1.3 MACsec Overview

In today's diverse workplaces, consultants, contractors, and even guests require access to network resources over the same LAN connections as regular employees, who may themselves bring unmanaged devices into the workplace. As data networks become increasingly

indispensable in day-to-day business operations, the possibility that unauthorized people or devices will gain access to controlled or confidential information also increases.

The best and most secure solution to vulnerability at the access edge is to use the intelligence of the network. IEEE 802.1X provides port-based access control using authentication, but authentication alone does not guarantee the confidentiality and integrity of data on the LAN. While physical security and end-user awareness can mitigate threats to data on an IEEE 802.1X-authenticated LAN, there may be situations or locations (such as remote offices or publicly accessible areas) in which the LAN needs additional protection. When additional protection is needed, network devices need to enable data confidentiality and integrity on the LAN by using MAC Security (MACsec). Defined by the IEEE 802.1AE standard, MACsec secures communication for authorized endpoints on the LAN.

MACsec uses several protocols, including:

- **Extensible Authentication Protocol (EAP)** — The message format and framework defined by RFC 4187 that provides a way for the supplicant and the authenticator to negotiate the EAP authentication method and MACsec association
- **EAP Method** — Protocol that defines the authentication method-that is, the credential type and how it will be submitted from the supplicant to the authentication server using the EAP framework; for MACsec, the EAP method must be capable of generating keying material to export a master session key (MSK) to the supplicant and authentication server
- **MACsec Key Agreement (MKA)** — Protocol that discovers MACsec peers and negotiates the keys used by MACsec; MKA is defined in IEEE 802.1X-2010
- **Security Association Protocol (SAP)** — A pre-standard key agreement protocol similar to MKA
- **EAP over LAN (EAPoL)** — An encapsulation defined by IEEE 802.1X for the transport of EAP from the supplicant to the switch over IEEE 802 wired networks; EAPoL is a Layer 2 protocol
- **RADIUS** — Essentially the standard for communication between the switch and the authentication server-the switch extracts the EAP payload from the Layer 2 EAPoL frame and encapsulates the payload inside a Layer 4 RADIUS packet; RADIUS is also used to deliver keying material to the authenticator

3.1.3.1 MACSEC Benefits and Limitations

MACsec offers the following benefits on wired networks:

- **Confidentiality** — MACsec helps ensure data confidentiality by providing strong encryption at Layer 2.
- **Integrity** — MACsec provides integrity checking to help ensure that data cannot be modified in transit.
- **Flexibility** — Network admin can selectively enable MACsec using a centralized policy, thereby helping ensure that MACsec is enforced where required while allowing non-MACsec-capable components to access the network.
- **Network Intelligence** — Unlike end-to-end, Layer 3 encryption techniques that hide the contents of packets from the network devices they cross, MACsec encrypts packets on a hop-by-hop basis at Layer 2, allowing the network to inspect, monitor, mark, and forward traffic according to your existing policies.

Although MACsec offers outstanding data security, it has limitations that should be addressed by the FTI-2 design if required:

- **Endpoint Support** — Not all endpoints support MACsec.
- **Hardware Support** — Line-rate encryption typically requires updated hardware on the access switch.
- **Technology Integration** — Enabling MACsec may affect the functions of other technologies that also connect at the access edge, such as IP telephony. Understanding and accommodating these technologies is essential to a successful deployment.

MACsec provides secure communication on wired LANs. When MACsec is used to secure the communication between endpoints on a LAN, each packet on the wire is encrypted using symmetric key cryptography so that communication cannot be monitored or altered on the wire.

MACsec was primarily designed to be used in conjunction with IEEE 802.1X-2010. IEEE 802.1X provides port-based access control using authentication. An IEEE 802.1X-enabled port can be dynamically enabled or disabled based on the identity of the user or device that connects to it.

While many use cases of MACsec apply to LAN environment between hosts and LAN switches, MACsec is also widely deployed in WAN environments. MACsec can be deployed between Customer Edge (CE) routers to CE routers or CE routers to Provider Edge (PE) routers over either layer 2 DOT1Q trunk or access link. It can also be deployed over layer 3 point-to-point routed ethernet connections. In addition, MACsec can be implemented via either ethernet physical interface mode or DOT1Q sub-interface mode in either a layer 2 or layer 3 design. Layer 3 information is encrypted when MACsec is enable over a point-to-point layer 3 link.

When MACsec is applied on both the uplink and the downlink, the MACsec sessions are completely independent. Moreover, while all traffic is encrypted on the wire, the traffic is in the clear inside each switch. This feature allows the switch to apply all the network policies (quality of service [QoS], deep packet inspection, NetFlow, etc.) to each packet without compromising the security of the packet on the wire. With hop-by-hop encryption, MACsec secures communication while maintaining network intelligence.

3.1.4 Packet Inspection – IPS

Like many other networks around the world, FTI-2 will be facing intruders and attackers who could come from both outside and inside the network. They can launch denial-of-service (DoS) attacks or distributed denial-of-service (DDoS) attacks; attack network connections; and exploit network and host vulnerabilities. At the same time, Internet worms and viruses can spread across the world in a matter of minutes. There is often no time to wait for human intervention so the network itself must possess the intelligence to instantaneously recognize and mitigate these attacks, threats, exploits, worms, and viruses.

IPS is an inline, deep-packet-inspection-based feature that enables sensor to effectively mitigate a wide range of network attacks. While it is common practice to defend against attacks by inspecting traffic at the data centers and corporate headquarters, it is also critical to distribute the network-level defense to stop malicious traffic close to its entry point at the branch or telecommuter offices.

It is suggested that FAA consider the use of integrated technologies that some equipment vendors are beginning to offer regarding packet inspection by more multi-functional devices (e.g. ASA firewalls offering IDS capability) to expand the IDS footprint to ensure all appropriate FTI dataflows are inspected.

3.1.5 Packet Inspection – Firewall

Stateful Packet Inspection (SPI) works at the network layer and examines some basic information contained within the packet, such as the packet header (address, port information) and also determines if the packet belongs to a valid session. That information is used by a basic firewall to determine if the packet should be allowed to enter the network or be blocked. Firewalls using SPI also check to see what connections have been established from the inside of the network to the Internet, using that information to determine if there is an open connection related to the packet before allowing the packet to traverse the firewall and into the internal network. If the packet fails to meet any of the basic requirements set forth by the firewall, it will be rejected.

Deep Packet Inspection (DPI) looks at not only the header and footer of a packet, but also examines the data part (content) of the packet searching for illegal statements and predefined criteria, allowing a firewall to make a more informed decision on whether or not to allow the packet through based upon its content. DPI delves into the data content of the packet, which allows additional determinations to be made before that packet can travel into the network.

DPI accomplishes that by dissembling incoming packets, examining the payload (data), comparing that data with defined criteria, and then re-assembles the packet for transmission (or rejection). When examining the payload, DPI engines can also employ signature matching, stealth payload detection and numerous other security capabilities.

3.1.6 Physical Layer Segmentation

Refers to separation of two networks at the physical layer, meaning that there is a change or disruption in the physical transmission medium that prevents data from traversing from one network to another. An example of physical layer segmentation could be users in marketing and engineering department use two physical switches to connect to corporate network even they're sitting in the same area. Therefore, the terminology of "air gap" is a physical layer segmentation method.

3.1.7 VLAN (Layer 2) Segmentation

The role of providing access into a LAN is normally reserved for an access layer switch. A virtual local area network (VLAN) can be created on a Layer 2 switch to reduce the size of broadcast domains, similar to a Layer 3 device. VLANs are commonly incorporated into network design making it easier for a network to support the goals of an organization. While VLANs are primarily used within switched local area networks, modern implementations of VLANs allow them to span WLANs, MANs, and WANs over technologies such as OTV, VPLS…etc.

As the backbone of FAA, FTI-2 provides essential transport services to all FAA programs. VLANs will allow FTI-2 to separate and inter-connect each program's specific and common infrastructure to give necessary segmentation between program flows to meet security requirements.

3.1.7.1 VLAN Security Benefit

Security is one of many benefits VLANs provide. Groups that have sensitive data are separated from the rest of the network, which decreases the chances of confidential information breaches. Within a switched internetwork, VLANs provide segmentation and organizational flexibility. VLANs provide a way to group devices within a LAN. A group of devices within a VLAN communicate as if they were attached to the same wire. VLANs are based on logical connections, instead of physical connections.

VLANs allow an administrator to segment networks based on factors such as function, project team, or application, without regard for the physical location of the user or device. Devices within a VLAN act as if they are in their own independent network, even if they share a common infrastructure with other VLANs. Any switch port can belong to a VLAN, and unicast, broadcast, and multicast packets are forwarded and flooded only to end stations within the VLAN where the packets are sourced. Each VLAN is considered a separate logical network, and packets destined for stations that do not belong to the VLAN must be forwarded through a device that supports routing.

3.1.7.2 VLAN Security Vulnerabilities

There are a number of different types of VLAN attacks in modern switched networks. The VLAN architecture simplifies network maintenance and improves performance, but it also opens the door to abuse. It is important to understand the general methodology behind these attacks and the primary approaches to mitigate them.

VLAN hopping enables traffic from one VLAN to be seen by another VLAN. Switch spoofing is a type of VLAN hopping attack that works by taking advantage of an incorrectly configured trunk port. By default, trunk ports have access to all VLANs and pass traffic for multiple VLANs across the same physical link, generally between switches.

3.1.7.3 VLAN Best Practices

A good security practice is to separate management and user data traffic. The management VLAN, which is VLAN 1 by default, should be changed to a separate, distinct VLAN. To communicate remotely with a FTI-2 switch for management purposes, the switch should have an IP address configured on the management VLAN. Users in other VLANs would not be able to establish remote access sessions to the switch unless they were routed into the management VLAN, providing an additional layer of security. Also, the switch should be configured to accept only encrypted SSH sessions for remote management.

All control traffic is sent on VLAN 1. Therefore, when the native VLAN is changed to something other than VLAN 1, all control traffic is tagged on IEEE 802.1Q VLAN trunks (tagged with VLAN ID 1). A recommended security practice is to change the native VLAN to a different VLAN than VLAN 1. The native VLAN should also be distinct from all user VLANs. Ensure that the native VLAN for an 802.1Q trunk is the same on both ends of the trunk link.

DTP offers four switch port modes: access, trunk, dynamic auto, and dynamic desirable. A general guideline is to disable auto-negotiation. As a port security best practice, do not use the dynamic auto or dynamic desirable switch port modes.

Finally, FTI-2 will continue to carry voice traffic with stringent QoS requirements. If voice devices are on the same VLAN as general users, each tries to use the available bandwidth without considering the other device. To avoid this conflict, it is good practice to use separate VLANs for IP telephony and data traffic.

3.1.7.4 Native VLAN Function and Best Practices

The native VLAN on each Layer 2 trunk port is VLAN 1 typically and it cannot be disabled or removed from the VLAN database. The native VLAN remains active on all access switch Layer 2 ports. The default native VLAN should be properly configured to avoid several security risks: worms, viruses, or data theft. Any malicious traffic originated in VLAN 1 will span across the access-layer network. With a VLAN-hopping attack, it is possible to attack a system that does not reside in VLAN 1.

As a best practice, the best way to mitigate this security risk is to implement an unused and unique VLAN ID as a native VLAN on the Layer 2 trunk between the access and distribution switches. For example, configure VLAN 801 in access switch 1 and in the distribution switch. Then change the default native VLAN setting in both the switches. Thereafter, VLAN 801 should not be used anywhere for any purpose in the same access-distribution block.

3.2 Software Defined Networking (SDN)/Network Function Virtualization (NFV)

Software Defined Networking (SDN) or Network Function Virtualization (NFV) are emerging concepts that are getting well-deserved attention in the industry for the potential they have to innovate telecommunications by integrating networks with applications to reduce the administrative burden, while decreasing the time it takes to deliver new services. Its attraction is driven by these new ways to look at how networking and cloud solutions can be automated, made more efficient, and more scalable.

Security best practices will continually drive to striking the balance between risk and operability. Industry, including producers and consumers, is always eager to create more flexibility, scalability and efficiency within a network infrastructure, fostering the development and proliferation of such advancements like SDN and NFV. To ensure the balance is not lost in the risk versus operability scale, these solutions should carefully apply appropriate security measures within the virtualized infrastructure to compensate for the inherent risks they create.

For example, security architects should develop designs to monitor and inspect the baselines of traffic behavior within the virtualized infrastructure at so-called east/westbound interfaces between the different layers of controllers thus enforcing both intra and inter-domain security policies. Another concern is the affect security attributes impose on the operability of controllers supporting the SDN/NFV functions, such as TLS, or more advanced mechanisms to avoid eavesdropping, man-in-the-middle and DoS attacks on the control plane.

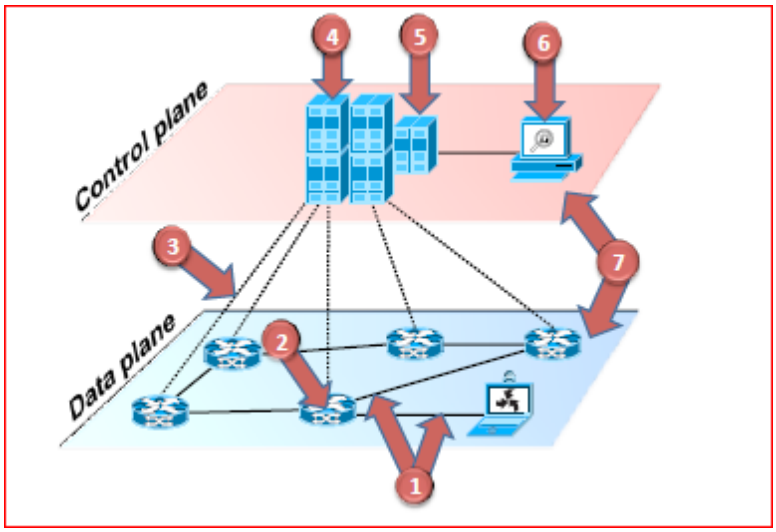

Source: Software-Defined Networking: A Comprehensive Survey. Diego Kreutz, Member, IEEE, Fernando M. V. Ramos, Member, IDDD, Paulo Verissimo, Fellow, IEEE, Christian Esteve Rothenberg, Member, IEEE, Siamak Azodolmolky, Senior Member, IEEE, and Steve Uhlig, Member, IEEE

Figure 3.2-1. SDN/NFV Control Plane Architecture.

Different threat vectors have already been identified in SDN architectures. While some threat vectors are common to existing networks, others are more specific to SDN, such as attacks on control plane communication and logically-centralized controllers. It is worth mentioning that most threats vectors are independent of the technology or the protocol (e.g., OpenFlow, POF, ForCES), because they represent threats on conceptual and architectural layers of SDN itself. The figure below shows the threat vectors involved.

While these threat vectors can be appropriately countered by a well-designed, well placed and integrated security architecture, it highlights the need to ensure that the security engineering effort is seamlessly woven in to the FTI-2 program engineering effort to ensure that the operabilities achieved do not create undue risk exposure to the agency.

SDN SPECIFIC VS. NON-SPECIFIC THREATS

Threat vectors	Specific to SDN?	Consequences in software-defined networks
Vector 1	no	Open door for DDoS attacks.
Vector 2	no	Potential attack inflation.
Vector 3	yes	Exploiting logically centralized controllers.
Vector 4	yes	Compromised controller may compromise the entire network.
Vector 5	yes	Development and deployment of malicious applications on controllers.
Vector 6	no	Potential attack inflation.
Vector 7	no	Negative impact on fast recovery and fault diagnosis.

Source: Software-Defined Networking: A Comprehensive Survey. Diego Kreutz, Member, IEEE, Fernando M. V. Ramos, Member, IEEE, Paulo Verissimo, Fellow, IEEE, Christian Esteve Rothenberg, Member, IEEE, Siamak Azodolmolky, Senior Member, IEEE, and Steve Uhlig, Member, IEEE

Figure 3.2-2. Main Threat Vectors Associated with SDN.

3.3 CLOUD COMPUTING SERVICES

There are three distinct models for cloud computing service – infrastructure-as-a service (IaaS), platform-as-a-service (PaaS), and software-as-a-service (SaaS). Each offers a different balance of control and responsibility between the host and the client business, making it more or less appropriate to a particular set of requirements. Ownership of cloud components vary depending on the model desired as depicted in ***Figure 3.3-1*** below.

The model chosen to operate within therefore dictates the domain boundaries for authorization and the demarcation points for event collection and correlation by Security Information and Event Management (SIEM) tools and the placement of security appliances needed to monitor and inspect traffic, perform access control, anti-malware, logging, auditing and more. The determination of the model chosen will depend on the degree of freedom the FAA is willing relinquish to a Cloud Service Provider (CSP). In all cases, the CSPs should provide the

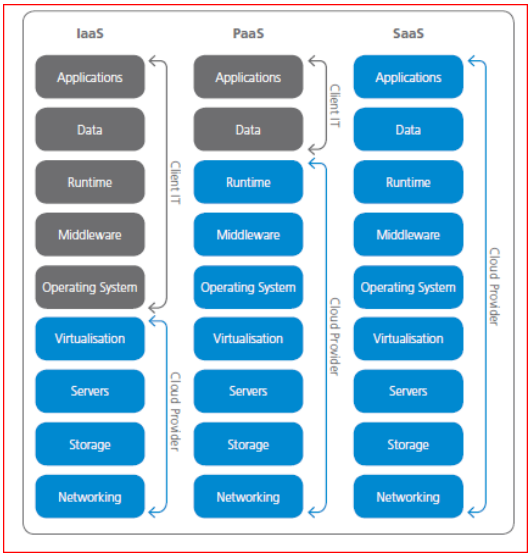

Source: Section 3.3 (Page 12): Harris Corporation

Figure 3.3-1. Cloud Computing Model Variability of Ownership.

same levels of security controls for their applications and infrastructure as if it was the FAA managing their environments.

There needs to be compatibility and collaboration between the FTI-2 program and CSPs, if used, to ensure that the same oversight measures followed for access control, authentication, boundary control, identification and authorization, monitoring & reporting, cryptography and information assurance.

Connections from FTI-2 infrastructure to a FEDRAMP-authorized CSP should be used as a boundary point for placement of firewalls and Intrusion Detection Systems (IDS) that properly separate and define the delivery points of services as well as the scope of security responsibility.

3.4 VOIP Traffic vs Latency and QOS (Establish a Voice Gateway)

FTI-2 should establish security requirements for the National Airspace Enterprise Gateway (NESGs) to accommodate the emerging needs for sharing VoIP traffic with external partners that traverse through the NESGs. Today, the NESG design does not support the latency and Quality of Service (QoS) policies that will need to be met supporting NAS Voice System (NVS) services. Requirements will have to be defined and imposed to address the prioritization and quality of critical IP voice services.

One possible solution is to establish an Enterprise Voice Gateway (EVG) with the use of Session Border Controllers within the NESG infrastructure that handle critical voice services independently of critical data services. Creating a shared environment for both voice and data services with the NESG architectures will have to be carefully defined and designed to ensure the appropriate RMA values are achieved. One such option to consider is the use of VLAN separation as previously mentioned.

4.0 ENTERPRISE OPPORTUNITIES

4.1 INTRA-AGENCY SIEM/EVENT DETECTION, CORRELATION AND REPORTING

FTI currently uses a Security Information and Event Management (SIEM) solution that provides a holistic view of the security status of all relevant IT systems, and integrates security into existing management processes and workflows. The SIEM solution collects, normalizes, aggregates, and filters millions of events from thousands of assets across the network into a manageable stream that is prioritized according to risk, exposed vulnerabilities, and the criticality of the assets involved. These prioritized events can then be correlated, investigated, analyzed, and remediated, providing situational awareness and real-time incident response time.

The SIEM solution receives data from multiple FTI assets on the NAS and Mission Support networks. Combining information from NAS and Mission Support provides a more comprehensive view of security events, allowing FTI to correlate information across both networks.

The SIEM also has a bi-directional feed between the FAA Cyber Security Management Center (CSMC) and FTI Mission Support, and a one-way feed from FTI NAS to CSMC.

There is also a one-way feed from FTI to NAS Cyber Operations (NCO) for select events. Authorized NCO users can also run queries from the SIEM user interface.

When planning intra-agency SIEM communications, agencies should consider which assets and/or networks to include, whether there are specific fields that are needed for reporting, the expected volume of information, and storage capabilities. This information helps determine the best architecture for forwarding and storing events, ensuring that the SIEM capabilities are fully utilized.

4.2 NETFLOW DATA ANALYTICS

Using NetFlow telemetry and contextual information from the FTI-2 network infrastructure, a network security analyst will be able to exam any suspicious activity, user identity and application information from a single pane of glass. With this information, the analyst can decipher the correct next steps to take concerning the threat in a timely, efficient, and cost-effective manner for advanced cyber threats such as:

- **Network Reconnaissance** — The act of probing the network looking for attack vectors that can be exploited by custom-crafted cyber threats
- **Network Interior Malware Proliferation** — Spreading malware across hosts for the purpose of gathering security reconnaissance data, exfiltrating data, or creating back doors to the network
- **Command and Control Traffic** — Communications between the attacker and the compromised internal hosts
- **Data Exfiltration** — Export of sensitive information back to the attacker, generally via command and control communications.

A Netflow data analytic solution delivers detailed visibility into user activity, enabling network operators, security administrators and datacenter personnel to determine within seconds who is responsible for and affected by events anywhere across the network. Administrators can simply search the user name or IP address associated with the event from management console and

the system returns the appropriate real-time information. User-centric monitoring capabilities also allow network and security teams to run flow queries and reports based on actual user names versus just IP addresses. In addition to pinpointing responsible users, the solution can simplify the identification of other users affected by an event, so that quarantine and corrective actions can be taken sooner. This information is invaluable for combating advanced attacks and insider threats, as well as for improving incident response and forensic investigations.

Netflow data analytic solution also uses a combination of deep packet inspection (DPI) and behavioral analysis to identify applications and protocols in use across the network no matter if they are plain text or use advanced encryption and obfuscation techniques. It also gathers packet-level performance statistics at a fraction of the cost of traditional probe-based devices, playing a key role in troubleshooting both security incidents and application performance problems.

For FTI-2, many communications will occur in virtual environment. Because virtual-machine-to-virtual-machine (VM2VM) communications inside a physical server cannot be monitored by traditional network and security devices, this lack of visibility complicates problem identification and resolution. For virtual environments with limited system resources, virtual Netflow solution also enables operators to see the same detailed traffic statistics for their virtual networks as they can for their physical environments, effectively eliminating the blind spots associated with virtualized infrastructure. The solution will be able to capture vital traffic statistics to address multiple virtualization challenges, including gaining virtual network topological and location awareness, securing virtual networks, demonstrating compliance, controlling VM sprawl and tracking virtual machines when they are moved such as via VMware VMotion.

4.3 Deep Packet Capture/Forensics

Deep packet capture is used to store and view raw data packets passing through a network. It can provide granular historical information about security events. One of the primary considerations when planning for deep packet capture is determining how long to store the data, which is dependent upon network volume and the amount of available storage. For busier networks, it can be costly to keep more than a few days of data, and running queries may take an excessive amount of time. Lighter network traffic might allow keeping months of data with the same amount of storage.

One strategy is to use NetFlow data to establish a baseline for what is normal in the network, and then make use of full packet capture when anomalies are observed. Since most useful information is found in packet metadata, the metadata could be stored for longer periods and the full capture for shorter periods.

5.0 OTHER SECURITY CONSIDERATIONS

5.1 COMPLIANCE AND VULNERABILITY ASSESSMENT AND RESPONSE

NIST Special Publication (SP) 800-37, as amended, Guide for Applying the Risk Management Framework to Federal Information Systems: A Security Life Cycle Approach provides a disciplined and structured process that integrates information security and risk management activities into the system development life cycle. The RMF steps include:

- **Categorize** the information system and the information processed, stored, and transmitted by that system based on an impact analysis.
- **Select** an initial set of baseline security controls for the information system based on the security categorization; tailoring and supplementing the security control baseline as needed based on an organizational assessment of risk and local conditions.
- **Implement** the security controls and describe how the controls are employed within the information system and its environment of operation.
- **Assess** the security controls using appropriate assessment procedures to determine the extent to which the controls are implemented correctly, operating as intended, and producing the desired outcome with respect to meeting the security requirements for the system.
- **Authorize** information system operation based on a determination of the risk to organizational operations and assets, individuals, other organizations, and the Nation resulting from the operation of the information system and the decision that this risk is acceptable.
- **Monitor** the security controls in the information system on an ongoing basis including assessing control effectiveness, documenting changes to the system or its environment of operation, conducting security impact analyses of the associated changes, and reporting the security state of the system to designated organizational officials.

Baseline compliance and vulnerability assessments and response activities are key to a successful monitoring program for any network security enterprise, especially FTI. To be effective, full accounting of all network elements, applications, Operating Systems (OS) and appliances should be continuously updated to ensure accuracy of the baseline architecture. It is this baseline architecture that should be subjected to regular and continuous (nominally monthly) scanning using prominent scanning tools (e.g. Tenable Nessus) and proactive compliance management tools (e.g. Tenable Security Center) to assess the degree of baseline compliance against known baselines (e.g. CIS) and the latest vulnerability scan plug-ins scripts.

5.1.1 Patching

Recently, the FAA Information System Security (ISS) Program Office introduced new Patching cycle timelines to be followed by NAS systems according to their FIPS-199 impact level (FTI is a FIPS Moderate system), the position of the device or application within the network, and the criticality of the vulnerability provided (based on CVE scoring). Its release was predicated on the FAA's desire to create a standardization for patching that all NAS programs should follow. ***Table 5.1.1-1*** is provided below.

Table 5.1.1-1. FAA ATO Patching Timeline Guide.

FIPS-199 Impact	Criticality of Vulnerability	External (Public Facing) Including Networking Assets	Mission Support/Admin	NAS Gateway	NAS Internal
High	Critical	30 days	30 days	60 days	120 days
	High	30 days	30 days	60 days	150 days
	Medium	60 days	60 days	90 days	180 days
Moderate	Critical	30 days	60 days	60 days	120 days
	High	30 days	60 days	60 days	150 days
	Medium	60 days	90 days	120 days	180 days
Low	Critical	30 days	60 days	60 days	120 days
	High	30 days	60 days	60 days	150 days
	Medium	60 days	120 days	120 days	180 days

Source: FAA

Note: Per FAA ATO, the figure above shows what is a suggested patch interval, based on FIPS, Criticality of patch, and Exposure, which is contained in the FAA SAR, and was tailored using the Agency's ISS Handbooks as a starting point and is subject to change.

While the need to establish a clear patching implementation policy is valid, the challenge with this patching timeline is that complexities exist within the framework of the FTI process that includes time-delaying factors beyond the control of the FTI vendor as well as the need to manage the vulnerability assessments more narrowly for network elements that are more externally facing than others that are not. For example, devices that support FTI operational services require the release of FAA controlling authorities for all services that rely on that device to function in order to implement a patch update.

To ensure the effectiveness and success of a patch management program on FTI-2, FAA should consider simplifying and defining patch implementation requirements for security-related (vulnerability) patches made to network elements. As vulnerabilities are assessed for their applicability, degree of impact, and operational risk to implement, they are assigned a priority to test and apply. Depending on the number of devices involved, the degree of coordination needed for FAA FTI service release and the complexities with the patch itself (e.g. bug-free releases), the variability of time to fully implement can be considerable. All these factors should be taken in to account to develop a reliable and achievable patch management process for FTI-2.

5.2 Penetration Testing/Red-Blue Team Exercises

Penetration testing (also called pen testing) is the practice of testing a computer system, network or Web application to find vulnerabilities that an attacker could exploit. Pen tests can be automated with software applications or they can be performed manually. Either way, the process includes gathering information about the target before the test, identifying possible entry points, attempting to break in, and reporting back the findings.

Because of the criticality of the NAS infrastructure supported by the FTI program, pen testing on the operational network itself is not recommended. As a practical alternative, pen testing of lab environments similarly configured to that of the operational production network is a valid choice and should be done minimally once annually. A Plan of Actions and Milestones (POAM) should be assigned for any action items and lessons learned to be addressed prior to the next pen testing event.

5.3 PROTECTION OF PII DATA

Policies governing the Protection of Personally Identifying Information (PII) can be found in NIST Special Publication 800-122, Confidentiality of Personally Identifiable Information (PII) and FAA Order 1280.1B, Protecting Personally Identifiable Information.

PII is any information about an individual maintained by an agency, including any information that can be used to distinguish or trace an individual's identity, such as name, social security number, date and place of birth, mother's maiden name, or biometric records; and any other information that is linked or linkable to an individual, such as medical, educational, financial, and employment information.

Currently, the FTI contract does not have the Privacy AMS clauses applied. FAA will have to re-assess the requirements for the FTI-2 vendor to possess and therefore protect PII data for federal employees and their contractors according to regulations.

5.4 PREDOMINANT AND EMERGING THREAT VECTORS

The one constant attribute associated with threat vectors against FTI is that they are ever evolving. As the architectural landscape for FTI expands to include increased use of shared infrastructures, e.g. FAA cloud services or Carrier Ethernet, and expand of variety of end points devices to delivery of new and emerging services (IP voice and data), the security architecture and its ability to protect the FTI-2 network will likewise need to evolve. The proverbial "Cat and Mouse" game will continue to be played throughout the term of the FTI-2 contract. To prepare for that game and to ready ourselves for the dynamic, adaptive and resourceful tactics adversaries will levy, we need to understand what the predominant and emerging threat vectors will be during the FTI-2 period of performance.

Many industry reports, include McAfee Labs Threat Predictions 2016 report and Cisco's 2015 Annual Security Report among others present the threat landscape for the cyber battle space for the next several years that include some of these more notable predictions:

- **Insider Threats** — Malicious or not, insider threats will be a dominant threat vector for FTI-2 or for any critical network for that matter. Ensuring people, processes and systems that lie inside the FTI-2 protective boundary will act appropriately and abide by policy will become increasingly important. Access and Authorization control mechanisms will remain a vital component of the security approach on FTI-2. Regular and random audits and the conduct of comprehensive insider threat awareness training will also serve to counter this vector.
- **System firmware-based attacks, especially as the FAA migrates in to cloud-based infrastructures** — This requires that FTI-2 be abundantly aware of its system components below the operation system (OS) level and how those components can be exploited or leveraged for attack. To that end, Compliance will become increasingly paramount to assure FTI-2's configuration and vulnerability baselines are up-to-date and patched to current releases. Investment in compliance assessment and response tools will be essential.
- Application vulnerabilities will continue to be an ongoing problem as evidenced by the proliferation of Zero-Day attacks especially derived from open source software. Increasingly, embedded systems, the Internet of Things, and infrastructure software will become the targets for advanced threats and these attacks.
- **Cloud service vulnerability** — The advantages of cloud-based infrastructure and service is highly attractive in offering a cost-efficient, highly available, flexible and

scalable architecture to host FAA systems and services. They are also attractive to cybercriminals who wish to steal poorly protected information. Customers of these services are at the mercy of security controls at the hosting service and often have little insight into the service provider's security posture. Policy assurance and visibility into a cloud-service provider's security posture and Incident Response (IR) approach will be vital.

Internet Protocol version six (IPv6) is the most fundamental change in the Internet since its inception and will eventually touch every network attached device, Operations, Administration and Management (OA&M) systems, monitoring gear and higher layer applications whose data is transported across the network. All current major operating systems, Microsoft, Apple OS and most Linux versions have IPv6 enabled by default and in some cases, (specifically Microsoft) cannot be turned off at the end user device as it leaves the system in an untested and unstable mode with systems failing.

There are no more Internet Protocol version four (IPv4) only networks as tunneling is enabled automatically (though it can be safely disabled) and a number of network administrator and engineers are ignorant of the latent IPv6 threat on their networks. It is the recommendation of security subject matter experts (SMEs) that native IPv6 traffic and IPv6 tunnels over IPv4 be dropped by edge network security devices until such time as Agencies have management, security and operational policies and processes in place.

IPv4 was designed in an era of developing computer to computer communications with a small group of trusted machines while IPv6 was developed in an increasingly hostile environment. IPv6 to assist in securing the protocol, but these mechanisms have to be employed and operational to be effective. IPv6 is not necessarily more secure than IPv4; it depends on how effectively and efficiently network and security engineers configure the communications systems.

5.4.1 IPv6 and the US Government

US Government representatives began working with Internet Engineering Taskforce (IETF) and other Internet stakeholders in the mid-nineties on IPng, (Next Generation Internet Protocol) that in early 2000 would be relabeled IPv6. In 2001 the Department of Defense (DOD) Defense Research and Engineering Network (DREN) brought up the first IPv4/IPv6 dual stacked network (dual stack referring to having both an IPv4 and an IPv6 protocol stack operating at the same time). By 2009 the suite of IPv6 protocols had been well defined and packet processing was moved into silicon.

In 2003 the DOD put out the first Memo on adopting IPv6 within the DOD community; it has been delayed several times since then but is in full effect today. In 2005 the Office of Management and Budget (OMB) released its first adoption memo and required Agencies to pass IPv6 traffic over their IP backbones by 2008. A majority of Agencies then turned it off except for the National Library of Medicine/National Institute of Health/Health and Human Services, which has been fully dual stacked since 2009.

In 2009 Federal Acquisition Regulations (FAR) were amended to require IPv6 capability in all IT and other procurements that had network interfaces that utilized IPv6 protocol. Guidance was given to Agencies to use the Capital Planning and Investment Control (CPIC) budgeting process to request funds for IPv6 adoption within the Agency with the proviso that Tech refresh cycles and funds were to be used to replace non-conforming equipment, services and applications at the appropriate times.

In 2010, IT/Networking Leaders in Agencies, DOD and OMB ascertained that no real movement in adopting IPv6 had occurred after the 2008 guidance; a 2010 memo was released to set two new goals for 2012 and 2014. As the number of mobile devices attached to the Internet went up drastically, wireless carriers began deploying hundreds of millions of devices with IPv6-only addresses. The administrations goals were equal citizen access to services over either IPv4 or IPv6 and to maintain a leadership position in the Next Generation Internet that was only going to be IPv6 addressed.

The 2010 goals were for citizen services such as web sites, e-mail and the Domain Name System (DNS) to be available on native IPv6 which could coexist over the same infrastructure as IPv4. (Note: Native IPv6 means that the IP packets are not encapsulated over IPv4 or translated; in other words that there be end to end IPv6 protocol).

The 2014 guidance went further into the Agencies networks to require that Government workers or contractors that had to access the web could do so over native IPv6 networks with all of the devices, network links and services used in those IPv6 capable networks.

It is the stated intent of the Federal Chief Information Officer (CIO) that all mission critical network infrastructure only use native IPv6 protocols, as soon as feasible while supporting mission requirements. All Federal Agencies have been given guidance to inventory all networks and attached devices, services and applications, ascertain which systems are to be end-of-life' d either because there is no longer a need for the system or the technology and applications are obsolete and are too expensive to maintain or can no longer be maintained.

Agency CIO shops are to develop plans based on these inventories and Senior Agency Leadership decisions on which system to transition and which to retire with appropriate milestones and resource requirements. Dual stacking is a step toward an IPv6-only network and requires that both IPv4 and IPv6 protocol stacks be resident on the gateway or other system as each protocol will not rely on any aspects of the other protocol. Dual stacking as a phase in IPv6 adoption is to be as short as possible while supporting the networking requirements of the mission

6.0 SECURITY IMPLICATIONS

There are people, processes and equipment/services that make up the engineering, operations, management parts of the network. If we define a secured network as one that is both operational and protected which requires that management, operations and security use best practices because if there is a broken network, it does not matter if it was sloppy operations or a security event that took it down as it is still broken and not mission capable.

Successful management of a dual stacked or IPv6-only network, as well as larger IPv4 networks, requires a discipline of personnel, the tools, systems and processes they utilize. If the technical sophistication of Cloud and SDN are super imposed on the network infrastructure, the level of complexity grows dramatically.

6.1. SPECIFIC IMPACTS

6.1.1 Network Engineering

Older network engineers will remember multi-protocol networks with Digital Equipment Corporation Network (DECnet), Novell, Microsoft and Transmission Control Protocol/Internet Protocol (TCP/IP) as very prominent protocols. Each protocol has its own set of rules as well as the layer 2 links the protocols are running on. Original Ethernet had a max number of devices in a subnetwork in the 250 range.

For devices in an IP network to communicate they have to be part of or addressed in the same subnetwork. A base IP address and a mask size identify a subnetwork and delineate the total number of addresses in the subnetwork that can be used for device interface identifiers. Current IPv4 addressing utilizes Classless Inter-Domain Routing (CIDR), which is a variable sized subnetwork field out of the IP address for interface identifiers, and for example, here is a list of the total number of addresses in several IPv4 address ranges. Two of the IPv4 addresses are reserved, zero and broadcast, which leaves the number of usable device addresses are:

- 192.168.10.0/30 4 -2 = 2
- 192.168.10.0/29 8 -2 = 6
- 192.168.10.0/28 16 -2 = 14
- 192.168.10.0/24 256 -2 = 254

A network engineer will first:

- Identify which super block they will need to get a sub block of IPv4 address allocated from and how many they need
- Identify which sub block has sufficient size to meet their requirements
- Allocate that block and mark them used

IPv6 networks have defined a fixed size subnetwork of a /64 which gives 18,446,744,073,709,551,616 number of usable device interfaces so network engineers no longer need to worry about the number of devices on any one subnetwork.

When American Registry for Internet Numbers (ARIN) allocates IPv6 blocks for federal agencies they do so based on a formula which takes into account the number of sites an Agency has and the number of aggregation points (locations) the Agency maintains. While ARIN allocates a /48 to a site giving 65,536 /64s, FAA may elect as part of their IPv6 address planning to have smaller locations get a /56 which would give 256 /64s. For very limited communication locations a /60 would give 16 /64s in the case additional equipment were ever needed at those locations.

At certain FAA campus locations it may make sense to have a collection of sites for technical or administrative requirements and have either multiple /48s or mixed /48s and /56s aggregating up.

The FAA has multiple routing domains and in ARIN lexicon is a Local Internet Registry or LIR. The minimum allocation for an LIR is a /32 though this may or may not be enough for the FAA without further analysis which is beyond the scope of this section. There are multiple routing domains with the FAA purvey including NAS, FAA mission space and potentially other partners that may require IP address space from FAA blocks.

FAA IP addressing plans and allocation strategies are significant areas for security review and policy enforcement to maintain network security.

IPv4 layer 2 ARP has been replaced with Neighbor Discovery Protocol (NDP) which introduces a layer 3-like protocol as a link layer protocol in IPv6. Internet Control Message Protocol (ICMP) in IPv4 was rather limited and in IPv6 several sets of ICMPv6 messages pass different network boundaries.

6.1.1.2 Training and Experience

All members of the team from Program Manager to the Technician or Administrative assistant need IPv6 training. There will be at least four or five different levels of training as well as a hands-on lab that mimics the real network. A major risk in IPv6 networks is the newness and lack of familiarity of engineering, operations and security personnel for either dual stacked or IPv6-only networks. The Federal IPv6 Task Force/Working Group was constituted to allow sharing of experiences, lessons learned and best practices. Standard Operating Procedures (SOP) need to be developed and tailored to Agency (FAA) needs and validated.

There is no substitute for turning up a test network to get familiar with both normal and abnormal operation of the network, as well as to exercise the test network to gain additional insight and understanding.

A next step is using a section of the operational network that does not have the same HA level and run controlled tests to verify configurations and to note any deviations between the test and operational network. The FAA already has duplicate networks for testing which allow IPv6 capability to use the same validation methods as other tech refresh cycles.

6.1.1.3 Security Engineering

The reader is recommended to review the IPv6 Security Architecture section of the Architectures White Paper for a thorough treatise on IPv6 security. Some areas will be addressed here.

A number of tools and techniques are directly transferable to securing the IPv6 protocol utilizing IPv6 addresses, of course. Access Control Lists (ACL), Firewall/Intrusion Prevention System (IPS)/Intrusion Detection System (IDS) rules and policies apply to IPv6 packet streams as well as IPv4. For Dual Stack there will be an IPv4 ACL, an IPv6 ACL (that will not be the same ACLs and IPv4 with only the addresses changed to IPv6) and potential ACLs because of a potential IPv4/IPv6 interaction.

Increasing the frequency of penetration (pen) testing is a best practice until all of the new security mechanisms have been thoroughly tested and there have been no changes after a period of time. Pen testing should continue to be done when there is any change in server, end user or network device or security device configuration change. (You want the good guys to find the hole and not be on CNN.)

A number of Network Security Devices and services are fully IPv6 capable. It is in Agencies best interest to get Supplier's Declaration of Conformance (SDOC) on the vendors IPv6 capabilities as well as test the capabilities in their Test networks to validate that the required functionality is there.

6.1.1.4 Operations and Secured Network Services

IP address management, DNS and Dynamic Host Configuration Protocol (DHCP) services can be the Achilles heel of IP networks so they have to be as automated and integrated as possible as well as secured. The FAA High Activity (HA) requirements cannot tolerate the error rates in manual DDI ((DNS, DHCP &I IPAM)) systems. Unique IP addresses have to be assigned to thousands of devices with hundreds of those devices needing DNS records with those addresses and unique DHCP scopes for all of the end user systems that are not being statically addressed. IP address spread sheets have normal error rates of 10 to 20 percent and non-integrated network services can still achieve a 5% to 7% error rate with manual rekeying and cut and paste errors. Review and validation of data is required to scrub the data but that process cannot introduce more errors in the data than what it is fixing.

6.1.1.5 Secure DNS

DNS has become the second most used attack vector behind Hypertext Transfer Protocol (HTTP), as firewalls, IDS/IPS systems have to let port 53 traffic (DNS resolution request and response) through or the network does not function. Next generation firewalls are used to block any protocol piggybacking on port 53 (are the packets destined for port 53 actual DNS packets, or are they, for example, port 80 HTTP traffic with a forged port number?).

There are three sets of DNS security features that should be implemented in Agency IP networks and they are RPZ (Response Policy Zone), DNS DoS mitigation and blocking of data exfiltration using DNS packet streams. RPZ technology is based on a collection of known bad actors URLs that are determined by companies that provide threat data feeds. When end user systems send out resolution requests (for example I want to know the address of www.faa.gov and they get back a 2001:db8::0123:faa as the answer.) if those URL are identified as malware sites the RPZ response can be to ignore the request or send back an address for the honey pot system. For DNS Denial of Service mitigation, DNS servers have additional hardware that parses the DNS resolution requests in real time, discards mal formed packets and allows correctly formatted packets through.

Correctly formed DNS request and response packets can have long valid text fields that can be leveraged by specially constructed DNS servers to bypass firewalls and other security devices to exfiltrate sensitive data out from behind security perimeters. These publicly available servers are used to bypass ISP access charges in hotels and other locations.

The way DNS servers in the Internet are constructed there are very definite patterns for normal DNS traffic and multiple requests and responses to the same server standout in higher level data analytics as being abnormal and can then be interdicted, which requires highly automated systems to manage effectively.

6.1.1.6 Advent of IPV6

As IPv4 matured and networks grew in size, additional network services, such as DNS, DHCP, IP address management (IPAM), and Network Time Protocol (NTP) were developed deployed and became required for the correct operation of the network. Each device attached to the network required a unique 32-bit IP address for correct operation and packet routing across the network, as the management, allocation and remote discovery of that set of IP addresses is key.

Initially this management was done manually and DHCP assisted in the registration of the growing number of end user devices attached to the network. DNS allowed the segmentation of the network into more manageable domains which consisted on automated protocols to discover the relevant IP address but not to generate the records in the first place.

Due to the lack of usable and routable address space in IPv4, Network Address Translation (NAT) was introduced to extend the life of the protocol by sharing globally routable IP space among a set of users behind a NAT device. Management of NAT space was delegated to the local support staff as the NAT addresses only had local significance. In IPv6 networks the addresses have global significance and an enterprise or Agency wide address authority is required. This is the same address management concepts that Service Providers have utilized for years.

IPv6 is the latest layer 3 Internet Protocol and is a fundamental protocol for internetworking which needs to be secured. Communications protocols are designed to transport data between computer systems and there is a dichotomy and a balance between securing and communicating. As a layer 3 protocol, IPv6 will very efficiently carry malware and botnet command and control traffic at higher levels in the same fashion that IPv4 carries the bad traffic today. A number of current practices for securing IPv4 networks can be applied to IPv6 networks such as Access Control Lists (ACLs), next generation firewalls, IPS and IDS systems.

6.2 Remote Connections VPN/FRAC for MDTS and Use of Bring Your Own Everything (BYOX) Devices

FTI-2, like the predecessor FTI contract, will continue to rely on the use of SSL VPN connections to remote connect to the Mission Support (MS) administrative network to access FAA email systems and FTI administrative tools (e.g. NMO – Network Management and Operations, IBS – Integrated Business System, and FRAC – FTI Remote Access Capability). What has not been permitted and should not be unless careful and deliberate understandings of risk and acceptance of risk has been made is the concept of remote connections in the FTI network elements. To date, the risks associated with these connections regardless of security mechanisms used to make them have far outweighed the convenience or improvement they might have to FTI service availability (e.g. improved outage restoration time).

The question of FAA's continued use of MDTs that rely on remote SSL VPN connection in to certain NAS systems by selected FAA technicians should continue to be evaluated as to its security level and need.

6.3 DOD Classifications (Intra-Agency Collaboration)

Currently, there is no DOD classification requirement for the FTI program. Recommend FAA consider including requirements that the FTI-2 network provider establish a group of key staff that have at minimum Secret, preferably Top-Secret/ Sensitive Compartmented Information clearances to promote readiness and the ability to collaborate on potentially sensitive, DOD classified subjects pertaining to security events, threats and intelligence collection.

7.0 RECOMMENDATIONS

The FTI-2 security approach, like that on FTI, should be an integrated and internal component of the program performed embedded seamlessly within the deployment, operations, engineering and logistics functions performed. The degree of verification, compliance, audit and assurance needed to uphold and preserve acceptable operational and program risks cannot be achieved otherwise.

From a Governance perspective, the FAA should ensure that security-related policies are kept relevant and current to the latest technologies and processes; ideally being revised every 3-5 years. Additionally, the FTI-2 program should include requirements to comply with the latest FISMA, NIST and FIPS regulations and continue to follow its standing policies to conduct continuous monitoring and oversight of the security implementation.

FTI-2 will experience many technology transitions over its lifecycle; including the sunsetting of analog TDM and the proliferation of cloud-based and wireless service providers for transport and last mile delivery. For all of these transitions, Security focal points should be involved from the beginning to guide developers to the optimal approach that satisfies both operational and security requirements. There is no 100% secure solution; only one that provides an acceptable risk position and balance of investment. Each should be embraced and fully understood to ensure the FTI mission is achieved.

There will be many emerging threats facing FTI-2; some that have not been conceived or yet considered. To adequately and proactively address them, there should continue to be open and collaborative intelligence collection and sharing across the federal government and industries that support its critical infrastructure, including FTI. FAA may wish to consider introducing DOD clearance requirements to certain FTI security personnel to facilitate this measure. Transparency of security process, controls, and reporting are essential elements of any MOUs undertaken between the agency and a CSP.

As with any emerging technology or new program, the FAA and its FTI-2 vendor should be an active and regular participate in discussions, decisions and actions taken to evaluate and develop emerging capabilities and requirements at the earliest point possible to ensure they are developed with the appropriate security measures to protect the agency from threats and at the very least unknowingly accepting any undue risk.

Overall, the security approach for FTI should be adaptable, scalable, and holistic. It should provide a defense-in-depth architecture and have ability to measure, monitor and control all traffic flows that traverse within its boundary. Redundant and resilient gateways should be defined and staunchly controlled at key boundary locations to provide authorized users access to only the voice and data they are authorized to produce and/or consume. The posture of the network should be scanned and remediated to ensure continuous configuration and vulnerability compliance. Reporting of unusual behavior should be prompt and effective in reaching key decision makers within the FAA (e.g. NCO and CSMC) all in an effort to keep the NAS, the FAA and the flying public not only safe but secure.

Authors and Affiliations

Larry Nace, Harris Corporation

Michelle Head, Harris Corporation

John Lee, Infoblox

Andy Chen, Cisco Systems, Inc.