

Conceptual Model for Monitoring NextGen FAA Services

Michael Olsen
Andrew Parcel
Harris Corporation



Conceptual Model for Monitoring NextGen FAA Services

ABSTRACT

Challenges exist today with monitoring Federal Aviation Administration (FAA) services, including: understanding how FAA services impact each other, determining how data flows from program to program within the FAA, knowing which FAA applications use other FAA applications, and, when problems arise, quickly determining root causes. Early FAA services began with collections of simple point-to-point serial connections evolving into complex webs of point-to-cloud Internet Protocol (IP) services to net-centric application (Layer 7) services. What is needed is a single conceptual model that is capable of describing the many types of FAA services, data flows, and applications into a common language. This paper proposes a common conceptual model that can enhance the FAA's ability to monitor its infrastructure, services, and NextGen applications to provide an enhanced situational awareness of the National Airspace System (NAS)

INTRODUCTION

As NextGen introduces new technologies across the NAS, new tools and techniques will be required to monitor the health and status of FAA services. Transitioning the NAS to take advantage of these new technologies is important; however, transitioning technical support processes is critical to ensuring these technologies will be successfully employed in the long run.

Proper monitoring of the NAS means alerting the designated individuals when problems arise and enabling them to quickly pinpoint an issue's root cause, which, in turn, allows for the timely restoration of operations. In order to identify the root cause of faults or problems, an understanding of what hardware, networks, systems, and applications impact other hardware, networks, systems, and applications is required. An understanding of these details is also essential for determining impacts of non-corrective, maintenance actions. Without knowledge of the FAA Enterprise Architecture (i.e., a complete view of the hardware, network, systems, and applications, along with their interdependencies), a technician performing maintenance on a sensor at an airport has no idea that commercial airlines consuming Airport Surface Detection Equipment (ASDE-X) data from the System Wide Information Management (SWIM) NAS Enterprise Messaging Service (NEMS) are missing their updates on their situational displays to manage planes on the tarmac. When users see their situational displays go blank, they will likely report the problem to a help desk to determine the root cause. A lot of investigative time will be spent trying to isolate whether the issue is a software bug, a network issue, or a security failure. The help desk has to spend time researching the associated systems and applications to further isolate the issue, until it is discovered that it was just a scheduled maintenance activity of an airport sensor.

Each hardware component, network service, system, and application that makes up the NAS can be thought of as a puzzle piece. Puzzle pieces have varying degrees of complexity, corresponding to the numbers of rounded tabs and blanks on each. Pieces interlock with one or more other pieces, just as NAS components are connected with one another. While every puzzle piece contains a small part of a picture on it, only when all of the pieces are connected does one have a complete picture. Today, there is no system that presents a clear, complete picture of a puzzle as large and elaborate as the NAS. As new technology is

introduced to the NAS and the number and complexity of these components increases, the challenge of tracking the many interrelationships between systems, and disseminating this information in meaningful ways, increases. The detailed architectural information associated with each piece of the NAS puzzle needs to be captured in a common conceptual model that everyone can use to enhance situational awareness of the FAA Enterprise Architecture, as a whole.

THE CHALLENGES OF MONITORING FAA SERVICES

The challenge of monitoring the FAA Enterprise Architecture is that the many systems and applications are all different with unique alerts and interfaces. There is no one central monitoring application or person today that can understand how each FAA system works and interacts with other systems.

The NAS communications architecture is built from a foundation of infrastructure consisting of routers, A/B switches, multiplexers, and switches. These core components are put together to build services and networks. These services and networks are combined to form FAA applications and systems. Systems exist today to monitor individual network services; however, they may not be able to quickly determine how the infrastructure and network services combine to build FAA systems and applications. Having this information readily available at users' fingertips is critical to the future of NextGen operations and safety.

Some schools of thought suggest that if all the network information and alerts for all the hardware and services are obtained, a complete view of the FAA Enterprise Architecture can be created. Unfortunately, this view of the FAA Enterprise Architecture cannot be built by merely capturing all the network traffic or alert information; it must be built by aggregating all the knowledge in each area of the FAA enterprise into a single, cohesive, conceptual model.

PROPOSED CONCEPTUAL MODEL FOR MONITORING THE FAA ENTERPRISE ARCHITECTURE

The conceptual model is based on a simplified Open Systems Interconnection (OSI) model, reduced to three layers. The infrastructure, hardware, and cables reside on the physical layer (OSI Layer 1). This layer needs to document the physical hardware configuration, including what cable plugs into what port on what piece of hardware. The network layer (OSI Layer 3) segments are comprised of point-to-point serial services, point-to-cloud IP services, and voice services that travel from a source to a destination over multiple pieces of infrastructure. Many different network services can travel over the same cable or piece of hardware on the physical layer. The network layer needs to have a mapping of each piece of infrastructure that a network service rides over. The application layer (OSI Layer 7) consists of applications/systems that are made up of multiple network segments. Many different applications/systems can ride over common network services. The applications and systems also need to map to the network services, which are used for communication. Without the mapping between the layers, if issues on the physical layer occur, then impacts to higher layers cannot be determined. The conceptual model is represented in Figure 1.

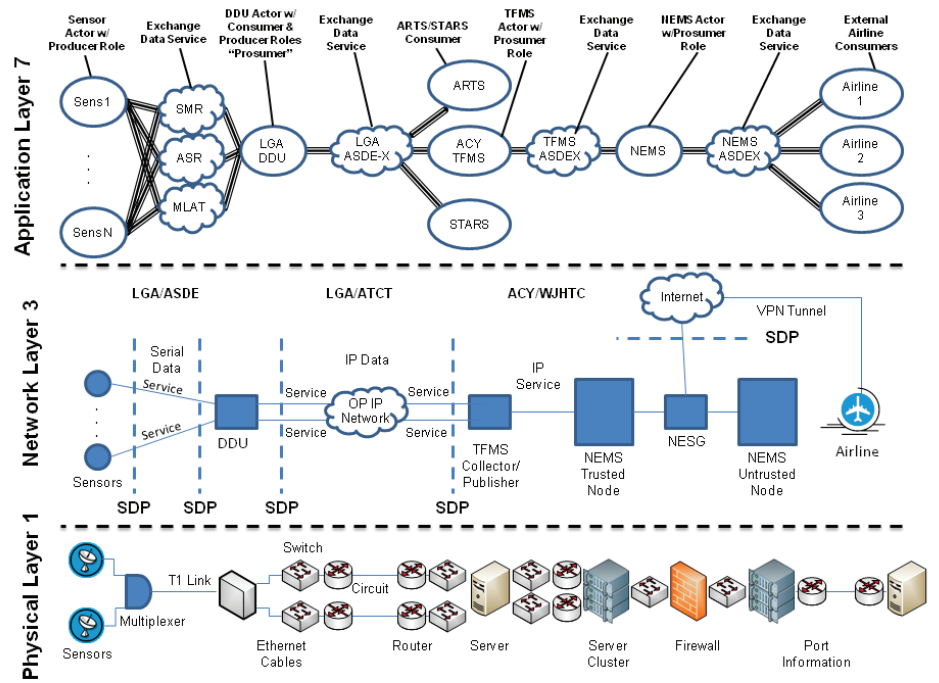


Figure 1: Conceptual Model for Monitoring NextGen FAA Services

Figure 1 describes an example end-to-end data flow from the source sensor data at LaGuardia (LGA) Airport to situational displays at commercial airlines. Three types of sensor data are transmitted to the LaGuardia Data Distribution Unit (DDU): Surface Movement Radar (SMR), Airport Surveillance Radar (ASR), and Multilateration (MLAT). The LaGuardia DDU provides ASDE-X data to government users (ARTS, STARS, and federal law enforcement) and non-government users (air carriers, air freight operators, and airports). The Traffic Flow Management System (TFMS) collects and performs filtering of raw ASDE-X sensor data and produces ASDE-X data in a SWIM-compliant format to the NEMS. NEMS receives the ASDE-X data and delivers it to external commercial airline consumers for processing on their situational displays.

The application layer is broken up into actors and exchange data services. An actor is an application/system that can assume roles of a producer, consumer, or both. Actors with a producer role send data to actors with a consumer role via exchange data services. Exchange data services describe the information sent from a producer to a consumer. In the example shown in Figure 1, SMR, ASR, MLAT, and ASDE-X are all exchange data services. Actors with both a producer and consumer role are "prosumers" and link FAA applications/systems together to form a chain of applications/systems from the original source of data to the end user destination at the application layer. The definition of exchange data services (e.g., how many exchange data services exist, what the exchange data services consist of, etc.) is left to the application domain experts for that particular area. The model is flexible enough to allow for exchange services to be as granular or as broad as needed. The exchange data services are linked to the lower layer network services. This model at the application layer is adaptable so that no matter the differences of the applications/systems, they can still be broken down into data flows being produced from one system to being consumed on the other. The FAA Enterprise Architecture can be broken down into these chains of applications/systems, linked to network layer services, linked to the physical layer of infrastructure.

Impact and Root Cause Analysis

With this conceptual model, operations staff can quickly determine the impacts to end users at the application layer, proactively forecast maintenance impacts, and narrow-down the root cause of an issue faster than before. Figure 2 depicts a multiplexer at LaGuardia Airport having a failure, shown in red. This physical layer failure of the multiplexer causes the serial network services from the sensors to the DDU on the network layer to go to outage. This, in turn, affects the SMR, ASR, and MLAT exchange data services, causing them to go to outage (depicted in red). Without these exchange data services, the LaGuardia DDU cannot feed ASDE-X data to the TFMS Collector at Atlantic City (ACY). The ACY TFMS collector can still, however, receive ASDE-X data from other airports, such as Newark (EWR) and John F. Kennedy (JFK). Since the TFMS collector is receiving some but not all data, it is in an impaired state (depicted in orange). This impaired state impacts the NEMS and impairs the situational displays of commercial airlines. Operations staff can traverse the model to see the source of the problem is at the multiplexer and then deploy technicians to replace the unit.

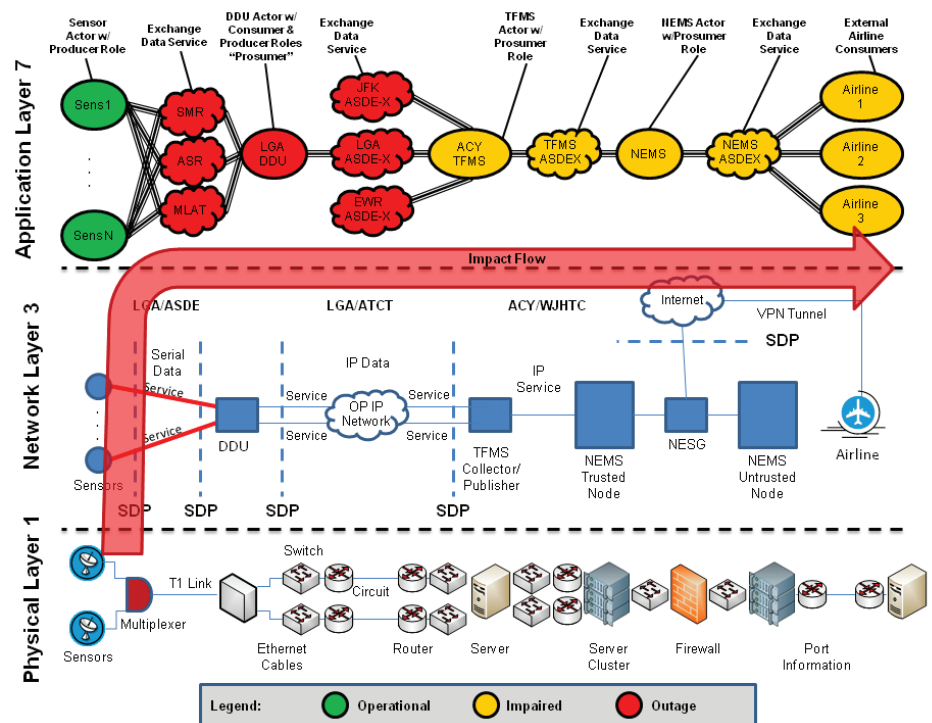


Figure 2: Depiction of Impact Analysis on the Conceptual

Impacts flow from the bottom layers to the top layers and from producers of data on the left to consumers of data on the right, as shown by the transparent red arrow in Figure 2. Impacts at the higher layers do not affect the lower layers. If there is an issue at the application layer, it does not impact the network or hardware underneath.

Building and Maintaining the FAA Enterprise Architecture

This conceptual model of the FAA Enterprise Architecture must be built from different sources of knowledge in each area at design time. When the applications are being designed, the domain knowledge is still present to be able to describe the links to other applications and adjacent layers in the model. Figure 3 shows the scope of operational visibility that exists for each area and layer of the model. A person with operational visibility for the LaGuardia DDUs, for example, has no visibility of how the NEMS works further downstream. A network engineer may have little knowledge of the applications riding above. However, the knowledge from each domain can be combined into a common conceptual model, which will allow operations staff supporting the end users and services to find the necessary information in a timely manner.

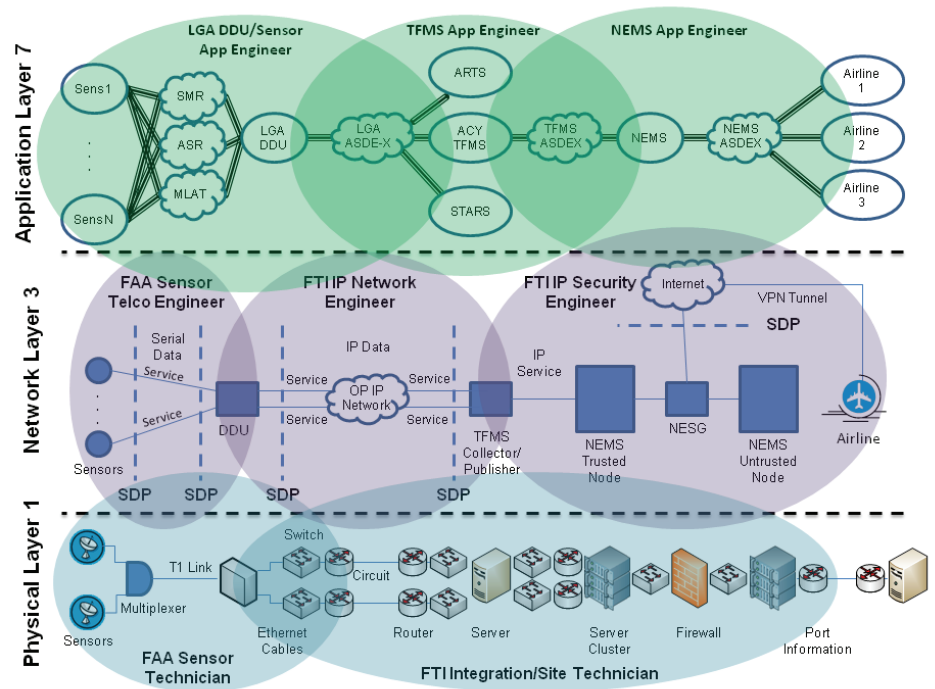


Figure 3: Scope of Operational Visibility in the FAA Enterprise Architecture

As FAA systems and applications change over time and the number of end users grows, the model must be updated as the architecture is re-designed while the knowledge exists and is readily available. When services are ordered and provisioned, the required information must be entered into the model in order for it to remain current. Without up-to-date information, the model ceases to be valuable to operations staff, engineers, and end users.

CONCLUSION

In an enterprise as large and complex as the NAS, it is unreasonable to expect any one person to fully comprehend the complexities associated with every FAA system at the application layer, let alone how each system relies on lower layer network services and infrastructure. The example highlighted in this paper (the multi-layer impacts associated with a multiplexer failure at LGA) is only one of thousands of chains that exist today in the NAS. As NextGen introduces new technologies to the NAS, the capabilities, as well as the complexity, of the NAS will continue to grow. Linking the various chains of applications, systems, network services, and physical infrastructure into a common conceptual model to document the FAA's Enterprise Architecture is needed in order to effectively transition, maintain, support, and grow NextGen systems and applications. If a common conceptual model is developed, interdependencies in the NAS can be captured, allowing for reduced labor costs/duplicated efforts, proactive maintenance impact analysis with less expended effort, and even quicker resolution of operational issues. As the FAA develops requirements for the next generation of the NAS, a common conceptual model should be built into the framework to create a central FAA Enterprise Architecture.



Government Communications Systems
P.O. Box 37
Melbourne, FL USA 32902-0037
1-800-4-HARRIS (1-800-442-7747)
www.harris.com/atc



Non-Export-Controlled Information

harris.com

Harris is a registered trademark of Harris Corporation.
Trademarks and tradenames are the property of their respective companies.
© 2014 Harris Corporation 08/14 523243 ECL d0742