

The logo for WIND, featuring the word "WIND" in a bold, white, sans-serif font with a trademark symbol, set against a black rectangular background.

WIND™

Securing the E-Enabled Aircraft

WHEN IT MATTERS, IT RUNS ON WIND RIVER

EXECUTIVE SUMMARY

E-enabled aircraft provide many benefits to operators in terms of operational efficiency, passenger comfort, and maintenance, repair, and overhaul (MRO). Systems on these aircraft transmit their data through satellite or ground-based communications networks to analytics services, which in turn give the airline operator excellent situational awareness, leading to many value-added services. Once systems are connected to a network, though, two significant challenges arise that must be addressed in order to benefit from this real-time Internet of Things (IoT) data: data bandwidth and systems security. This paper looks in more detail at the challenges of securing aircraft systems.

TABLE OF CONTENTS

- Executive Summary 2
- Aircraft Communications 3
- Security Considerations 3
- Security Scope 4
- Security Threats 4
- Security Risk Assessment 4
- Security Architecture 4
- Security Testing 5
- Summary. 5






Figure 1: Access to aircraft network domains

AIRCRAFT COMMUNICATIONS

Let's start with a high-level overview of current aircraft network domains and the security measures currently in place.

The aircraft controls domain (ACD) is the heart of the aircraft's avionics; it is where you find the flight controls, flight management, and navigation systems, among many other systems. These are typically running on a number of Integrated Modular Avionics (IMA)-based systems, and are safety-certified to the highest levels. These IMA systems are isolated from other domains, but allow read-only access from other domains through ARINC data buses to certain information—for example, altitude and heading data to show on passenger screens.

The aircraft information systems domain (AISD) contains crew systems, such as electronic flight bags, fault monitoring systems, health management, and airport ground-based communications. It provides certain data as read-only to the passenger information and entertainment systems (PIES) domain.

The PIES domain includes the in-flight entertainment (IFE) systems, cabin management systems, credit card systems, and other systems that interface with passengers.

The passenger-owned devices domain allows passenger-owned devices access to the Internet or (more often) to passenger media from the inflight entertainment system, and provides services on the passenger-owned device.

The ACD and AISD are isolated from the passenger information and entertainment domains, and of course passenger-owned devices—you certainly cannot risk aircraft safety through passengers hacking the aircraft systems.

In the future, access to higher bandwidth and advanced wireless and cell (3G/4G) connectivity will increasingly be expected not only for the PIES domain, but also the AISD, although the ACD will probably remain unconnected for the safety and security reasons already mentioned.

SECURITY CONSIDERATIONS

Data services and higher bandwidth connectivity are great, but how can you implement them in an aircraft without compromising security and, potentially, safety?

Meeting security requirements increases development complexity and, ultimately, cost. Aircraft systems are generally isolated from the Internet, and so in the past have implemented an "air gap" approach to security. But as the demand for value-added services increases, and the sophistication of security threats likewise increases, this approach needs to be updated in order to maintain aircraft security, and ultimately safety.

To handle the threat of unintentional unauthorized electronic interaction to aircraft safety, RTCA DO-326 (EUROCAE ED-202A) provides guidance for the process of aircraft certification. From DO-326A: "The purpose of the Airworthiness Security Process (AWSP) is to establish that, when subjected to unauthorized interaction, the aircraft will remain in a condition for safe operation (using the regulatory airworthiness criteria). To accomplish this purpose, the Airworthiness Security Process:

- Establishes that the security risk to the aircraft and its systems are acceptable per the criteria established by the AWSP, and
- Establishes that the Airworthiness Security Risk Assessment is complete and correct."

In other words, an appropriate level of security needs to be established—one that is relevant to aircraft safety. Securing a device is a continuous effort that spans the entire lifecycle of the aircraft, from architectural design through deployment and end-of-life. Planning and budgeting for safety and security updates throughout the entire aircraft lifecycle, along with future threat protections, are essential for any e-enabled aircraft.

SECURITY SCOPE

The first step in any security project is to define the scope of the security problem. This involves identifying the particular assets in the system, identifying the security perimeter, and documenting the security environment—in other words, a fairly straightforward review of what is in the system, where it touches the outside world, and the environment in which the system operates.

For example, assets can be broken down into hardware and software; these can be broken down further into particular data assets, such as navigation databases or software updates, and the value and impact of these assets can then be analyzed.

To identify the security perimeter, all touch points of the system to the outside world must be considered, including maintenance interfaces, digital interfaces such as passenger devices, crew systems, and connections between avionics systems; and existing security mechanisms within systems must also be identified.

Finally, the environment needs to be defined and analyzed, including other systems that the system under evaluation may come into contact with, such as air traffic systems or passenger booking systems. These outside systems must be identified, and the security threat analysis must cover potential threats from these sources.

In addition, this is a continually evolving system, so although the security scope may be identified at the start of the project, provision must also be made to update it throughout the lifecycle of the system. This could involve introduction of new technology such as 4G or cloud computing, or new systems within the environment.

SECURITY THREATS

After the security scope of the system is identified, the threats to that system need to be considered, and the conditions under which these threats might emerge identified. For example, a

passenger may be allowed to connect his or her device to the IFE in order to stream information, but this condition could now lead to threat injection into the IFE system.

Part of this analysis includes documenting security requirements for outside sources and identifying whom you can trust and to what level. RTCA DO-356 (EUROCAE ED-203, out for consultation at time of print) defines “trust levels” you can use to measure these outside environments. These trust levels are mapped directly to DO-178C safety levels such as “No Safety Effect” and “Catastrophic,” so level twE is “Not trustworthy to use or manage assets with any safety impact above No Effect” and level twA is “Trustworthy to use and manage assets of Catastrophic safety impact.”

SECURITY RISK ASSESSMENT

Once the security scope and threat analysis are in place, a security risk assessment can be performed that maps threat scenarios onto the security system to identify potential vulnerabilities that may need to be mitigated. This assessment is used to map these vulnerabilities onto failure conditions as defined in CFR 25.1209 and EASA CS-25 35.1309, allowing definition of the impact in well understood terms, such as those already used in safety analysis, from “No Effect” through to “Catastrophic.”

This analysis also identifies the risk associated with each threat identified, so that a value judgment can be made on the protection required. This is the same analysis as currently performed in safety systems engineering.

SECURITY ARCHITECTURE

The security architecture can now be implemented to mitigate the risks identified and to protect the assets within the security scope. Concepts such as defense in depth and layered assurance ensure that any particular threat has to penetrate multiple security measures in order to succeed, and this “chain of protection” provides greater security than just a single protection mechanism.

This layered protection should, for each system, cover system design, boot, run time (data in transit), and power down (data at rest). For each of these, the systems should align the security architecture to the identified threats that need to be mitigated.

As with DO-178C, design involves the process through which the

code itself is developed. As with safety, the higher the required protection, the more investment needed in the code development process. One way of reducing risk for code development is to use commercial off-the-shelf (COTS) components wherever it makes sense; so, for example, the operating system itself would be a COTS component. To support this decision, Wind River® supplies its VxWorks® and Wind River Linux operating environments with full security capabilities defined in security profiles to use in developing the system security architecture.

In order to implement layered protection, security measures must start as soon as the hardware powers up. In IT environments, attacks at the early boot and initialization stages are the most difficult attacks to remove (rootkits for example). As the hardware executes its firmware, the system needs to make sure the firmware has not been compromised, and that it then boots the expected (and secure) environment. As this is tied to the hardware system, this involves customization as well as COTS secure boot technology.

Using a COTS operating system that supports safety and security mechanisms also makes sense for the run-time architecture. The security threat analysis needs to go deep enough into the system architecture to make sure that the required OS features are enabled and configured accordingly—for example, if password protection is needed, or if any user interaction must be eliminated or limited.

Data-at-rest protection can be very simple, in the form of encrypted storage, or more sophisticated, using Anti-Tamper technology implemented both in hardware and software.

SECURITY TESTING

Security testing needs to look for specific vulnerabilities both in OS code and in any middleware such as networking code, as well as in the applications themselves. This testing covers many aspects of security vulnerabilities such as confidentiality, integrity, authentication, availability, authorization, and non-repudiation.

This security testing is over and above any functional testing to verify that the system does what it is supposed to do. The level of testing will be as identified in the security scope and threats, and must also include a plan for further testing across the lifecycle of the device, both in terms of testing the existing solution and in terms of testing for new threats as they become apparent.

SUMMARY

The move toward e-enabled aircraft is inevitable, and will provide many benefits for operators, manufacturers, and passengers. But these benefits will only be realized if the additional services can be made secure, without compromising the industry's impressive safety record.

