CCX TECHNOLOGIES

Securing Avionics with a Zero Trust Model

White Paper
WP-003.a

Summary

This paper explores policies and technologies that can be used to fully secure the avionics bay and cockpit from future cybersecurity threats using the Zero Trust Model.

Charles Eidsness

July 2021

Contact us at:
ccxtechnologies.com
info@ccxtechnologies.com
+1 (613) 701-6363

Table of Contents

INTRODUCTION TO ZERO TRUST

The Zero Trust Security Model is used to protect critical IT infrastructure. Unlike the Perimeter Model which implicitly trusts all traffic, devices, and users that have passed a specific gate, the Zero Trust Model assumes that no traffic, device, or user should be trusted unless it meets explicitly defined parameters. In a zero trust network levels of trust must be earned and consistently verified for all agents operating on the network.

This is similar to the use of redundant sensors in safety critical avionics systems. A pilot will not trust the output of a single sensor for safety critical operations. A set of three or more sensors are used so that matching output from at least two sensors can provide the level of trust required. If the perimeter model were used to implement the system the pilot would instead assume that a single sensor output could be trusted for an entire flight once it passed a pre-flight inspection. The pilot would assume that the sensor has not failed after the test, and assume the pre-flight check was perfect, both of which, we know from experience, cannot be safely assumed.

A Brief History of Zero Trust

The idea of zero trust networks has been around since the mid-'90s, the term and concept were reintroduced around 2010 by John Kindervag as a means to resolve unrectifiable issues with the perimeter security model.

The main issue with perimeter security is that it is not fault tolerant, it has to be perfect. Once someone breaches a perimeter they are free to move laterally throughout a system. This flaw with perimeter security has been exploited in almost every known successful cyber attack, including the Stuxnet, SolarWinds, and Colonial Pipeline attacks.

The more complex a system becomes, and the more people and devices that have access to a system, the more likely it is that someone will figure out a way to breach a perimeter. This can be seen first-hand with the increase in successful cyber attacks over the last decade, as systems have become more complex and interconnected.

A new more fault tolerant model is required for newer more complex and interconnected systems, and zero trust has emerged as the best candidate.

The zero trust model is already in use by sophisticated network operators. Following the breach of its systems by Chinese hackers, and revelations that the NSA had installed unwelcome listening posts in its network Google implemented a zero trust model on all of its internal and external infrastructure. The Department of Defense is in the process of moving towards a zero trust model and has published a Zero Trust Reference document in February of 2021. NIST has recommended a zero trust model for critical

systems and has published a guide; NIST 800-207. Further, in May of 2021 the Whitehouse issued an Executive Order on Improving the Nation's Cybersecurity with a heavy focus on Zero Trust.

Chain of Trust

Zero trust does not just apply to data or users of a network, but to all trust that is required for a specific application. For instance, in order for a pilot to trust an FMS system they must trust sensors feeding data to the FMS, they must trust the technicians that serviced the FMS, they must trust the software running on the FMS, they must trust the manufacturer of the FMS hardware, etc.

The diagram below outlines some of the agents, systems, and networks a pilot needs to trust in order to trust the GPS Coordinates displayed in the cockpit. Every one of these blocks is an opportunity to introduce something that could be used to adversely affect the operation of the system.

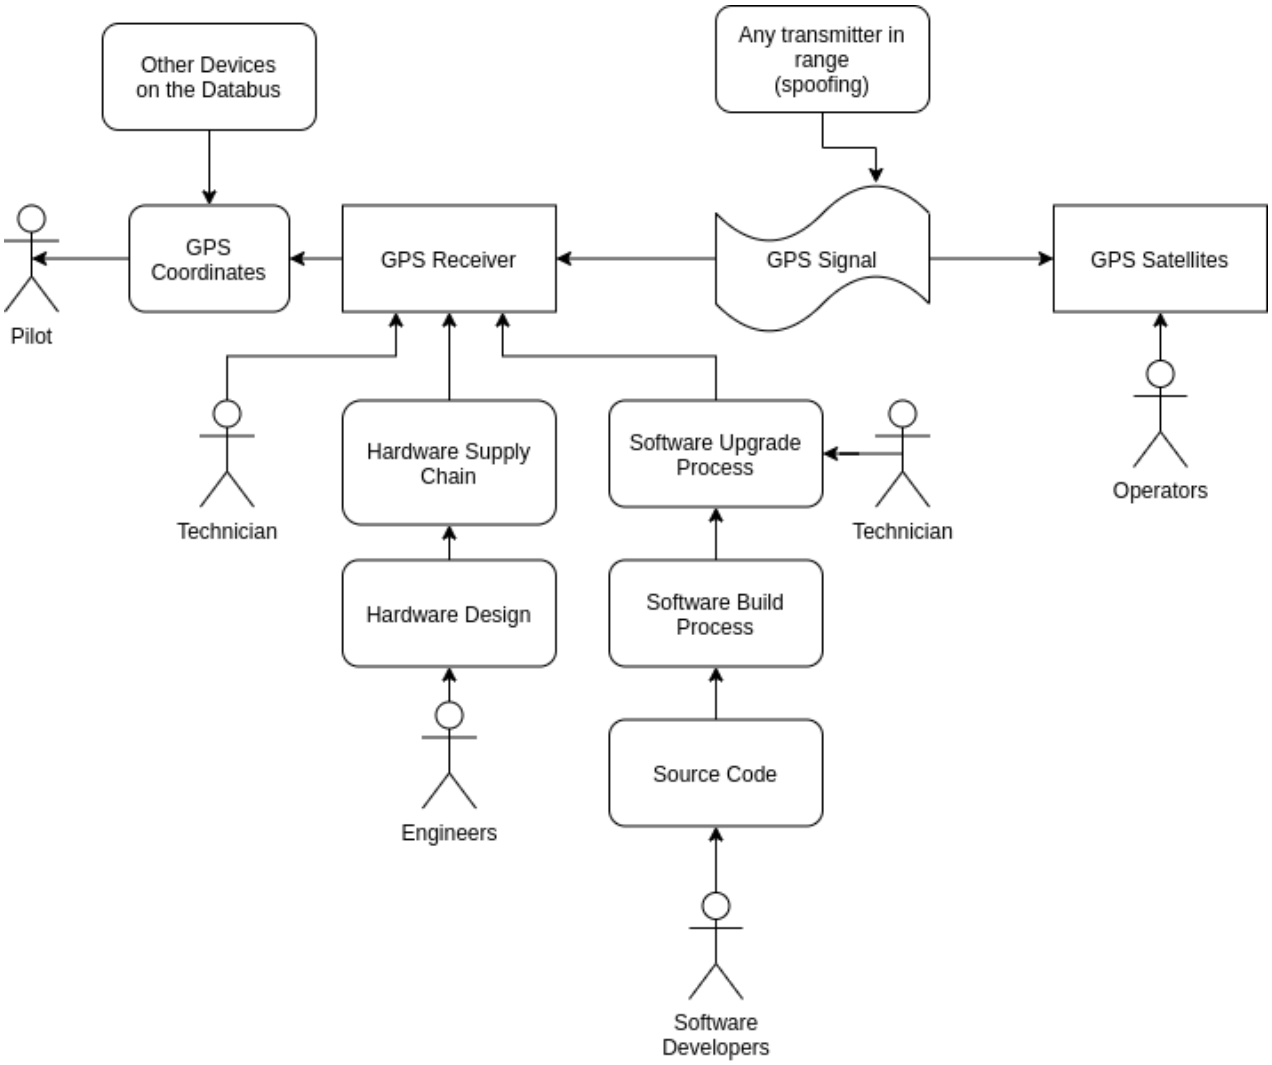

Figure 1: GPS Chain of Trust

If using the zero trust model the pilot would not implicitly trust any of these blocks unless they had some sort of verifiable process to ensure that the artifact could be trusted. For instance, it would be possible to require a login and security key in order to verify the identity of a technician and that they have been properly trained and vetted for the maintenance activity.

Levels of Trust

Like risk, the level of trust is not absolute. A zero trust model compares the verifiable level of trust of a specific agent vs. a policy to determine if an agent should be permitted to perform a task. This is similar to how a safety critical system requires a higher design assurance level than an in flight entertainment (IFE) system.

Verification of Trust Management Systems

The avionics industry is already familiar with the concept of third-party audits and approvals, these are also critical for a zero trust network. Every step in the trust system needs to be auditable in a way that will identify any issues with the assignment of trust. This includes monitoring network traffic for abnormalities, auditing login logs on systems, auditing software build and code check in procedures, etc.

The model assumes that it's possible to breach the system at any point, and this includes any systems that have been put in place to implement the model itself.

Additional Benefits of Zero Trust

In addition to increasing system security, implementing zero trust in avionics systems does provide other benefits. Zero trust adds tools that can be used to trust configuration and control data coming from existing and new devices. These tools can be used to facilitate the addition of new devices and features in a safe manner.

If using a zero trust network it would be possible to limit all data accepted by a device to data generated by other devices that are RTCA DO-178 Design Assurance Level (DAL) A certified. It would also be possible to filter data based on DAL levels at the receiver, for instance, the FMS could ensure that GPS data is sourced from a DAL A device, but flight plan data may be sourced from a lower DAL level Electronic Flight Bag (EFB). With these policies enforced in the FMS itself it would increase the safety of the system and make it possible to add new, innovative features to a system with less risk and uncertainty.

By implementing zero trust the industry would be improving the overall safety of avionics systems, and providing a less costly and time consuming path to adding future innovations, like EFBs, AI based tools, remote assistance tools, etc.

THE CURRENT STATE OF AVIONICS CYBERSECURITY

The primary goal of improving the cybersecurity posture of a system is to increase the cost of developing, weaponizing, and/or using a vulnerability to affect the operation of the system to the point that it no longer makes economic sense for an attacker to attempt.

A secondary goal is to provide data and tools that can be used to discover an attack in progress (or after it has been carried out), limiting the damage and reducing the cost and down-time associated with the attack.

In order to achieve these goals a comprehensive security plan for an entire system is required, but there currently is no comprehensive cybersecurity plan for avionics systems. There are a handful of established and upcoming standards like ARINC-842, RTCA DO-326, and ED-202 that provide guidance on how to secure portions of the system, but nothing comprehensive that covers everything in the chain of trust from the supply-chain to technician access, to real-time operations.

The concept of cybersecurity for avionics is fairly new, partly because cybersecurity is a relatively new field, and partly because the avionics industry tends to be very conservative (for valid reasons) and can be slow to adapt to new things.

Fortunately some of the features designed into avionics systems to ensure safety also provide a level of cyber security. Avionics systems tend to be isolated, are kept as simple as possible, and are designed in a way that they can survive a failure to one or more parts of the system. As a result avionics systems are starting at a much more secure level than other industrial control systems.

Even though some existing features provide a solid level of security, cybersecurity is still a unique problem compared to traditional engineering problems and requires a unique solution. Avionics engineering focuses on removing design and operational risk that result from unintended mistakes, cybersecurity engineering focuses on removing design and operational risk from intended attacks on a system by a third-party.

Avionics engineering focuses on a static system design, a design can be completely tested, refined, and evolved to maximize utility and minimize risk. Cybersecurity engineering focuses on a dynamic system design where new threats can be discovered at any time, and new defenses must be created for them.

The Avionics Engineering Process

Most avionics systems are developed using the same traditional engineering process.

The process starts with a finite set of technical requirements derived from business and regulatory requirements.

A system is then designed that fulfills these technical requirements, and a set of tests are performed to verify that the system achieves all of the technical requirements while not interfering with the operation of any other equipment on the aircraft.

The entire process of requirements, design, and verification is tracked and documented. Any updates to the design after the release of the original product are documented and reviewed using a similar process.

There is an element of risk in this process, especially as interconnected systems become more complex and behave in unanticipated ways, but this risk can be reduced by following well understood processes like DO-160 testing, FMEA analysis, DO-254, DO-178, etc..

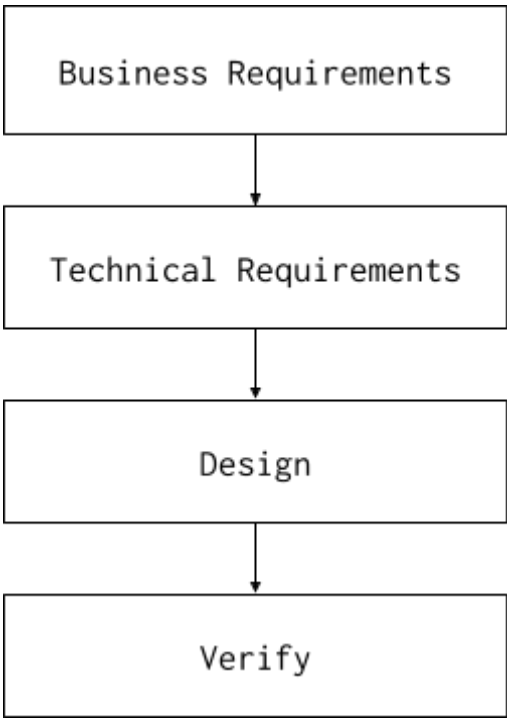

Figure 2: Traditional Engineering Process

How Cybersecurity is Different

When developing solutions for cybersecurity problems it isn't possible to follow a clean deterministic path. We aren't currently aware of vulnerabilities that have not yet been discovered (or have not yet been disclosed) so it's not possible to define requirements for, and design defenses for these unknown vulnerabilities. Cybersecurity is an ongoing risk mitigation activity with traditional engineering at the beginning to resolve known vulnerabilities.

What Makes Cybersecurity Difficult

Over time all systems, if left unchanged, become less secure.

The offensive side of the cybersecurity domain is a resources over time problem. The more people and equipment that are committed to developing and weaponizing vulnerabilities on a system the sooner a vulnerability will be discovered and made available for use.

The only way to secure a system is to continually invest resources in defense of that system, it simply isn't possible to have a secure system with no ongoing investment.

New tools are continually being developed that can be repurposed as offensive cyber-weapons, like low-cost software defined radios, fuzzers, software decompilers, information leaks, etc. We can't assume that just because an attack vector is prohibitively expensive today it will not be a risk at some point in the lifetime of a product, especially when the operational life of the product can be decades.

IMPLEMENTING A ZERO TRUST MODEL FOR AVIONICS

Fortunately the avionics industry is already familiar with a lot of the tools required to secure a system using the zero trust model.

The concept of trust is not new to the avionics industry, we already have processes like DO-254, DO-178, and DO-160 that require audited and witnessed verification that a piece of safety critical hardware or software will perform all specified tasks and not interfere with the operation of any other safety critical systems.

All of these systems, policies, and procedures are currently only intended to identify design defects and component failures that may result in increased safety risk but they can also be extended to prevent intended malicious activity.

A byproduct of the avionics industry's acute focus on safety is that avionics systems are starting at a higher level of security vs. other industrial control systems. A lot of the tools, processes, and organizations are in place that could be used to fully secure avionics systems. A comprehensive plan, some new pieces of technology, and a will to implement them is currently missing.

What is Required to Implement Zero Trust

Fundamentally the zero trust model requires a means to assign a trust level to every user, device, and network in a system along with a policy that either allows or rejects access to parts of the system based on the assigned level of trust.

On an IT system this is typically implemented by separating the control plane from the data plane and adding a separate trust management system which is used to provide access to the system. BeyondCorp from Google is an example of an implementation like this.

Unfortunately a centralized system like this will not work for an avionics system. For the centralized system to be able to mitigate a simple Denial of Service (DoS) attack, the system must reject all traffic if the centralized system fails. This is acceptable for an IT system but would create an unacceptable single point of failure in an avionics system.

A failure resilient distributed system is required for avionics, but the concepts and policies used in a more general purpose IT system can still be leveraged.

Notes on Trust Level and Access Policies

Every entity in the system is required to have a trust score and access to parts of the system are restricted based on that score, but that does not mean that an entity can't have a score of zero and still be able to access parts of the system.

For instance, since there is currently no way to verify the source of ADS-B traffic it would have a trust score of zero. It could be used in a system as long as the system was designed in a way that took this into account.

It is also possible to increase the trust score of ADS-B data by cross-referencing it to other data with a higher trust score, or by using a receiver that can identify different RF sources for the data, or by implementing software that verifies that the data is tracking a physically possible vector, which in most spoofing cases it will not.

The trust score for ADS-B data can not achieve the highest trust level since there is no absolute way of defining the source (without the addition of a crypto signed signature field) but it can be increased by adding features to the system.

Cryptographic Signing of Data

Cryptographic signing of data is used extensively within a zero trust network. Cryptographic signing uses public/private key pairs and mathematical algorithms that are easy to perform operations in one direction, but difficult to perform operations on in the other. For instance, addition would be a terrible algorithm to use since subtraction is just as easy as addition, multiplication would be better if no one had a calculator because long-division is more difficult than multiplication for most people.

Signatures like these are currently used in email systems, and to sign and verify the certificates that are used by browsers to secure and verify connections to secure web pages, like a bank.

Only data of a specific length can be signed using public key cryptography so typically a summary of the data is first created using a hashing algorithm like SHA256. The hash is then signed using a private key. Any agent in possession of the public key can calculate the hash of the data and then verify that the signature provided by the sender was created with the corresponding private key.

Signing of data allows the receiver to elevate the level of trust of a data element to the level of trust of the public and private key, which can usually be protected to a level higher than the data itself.

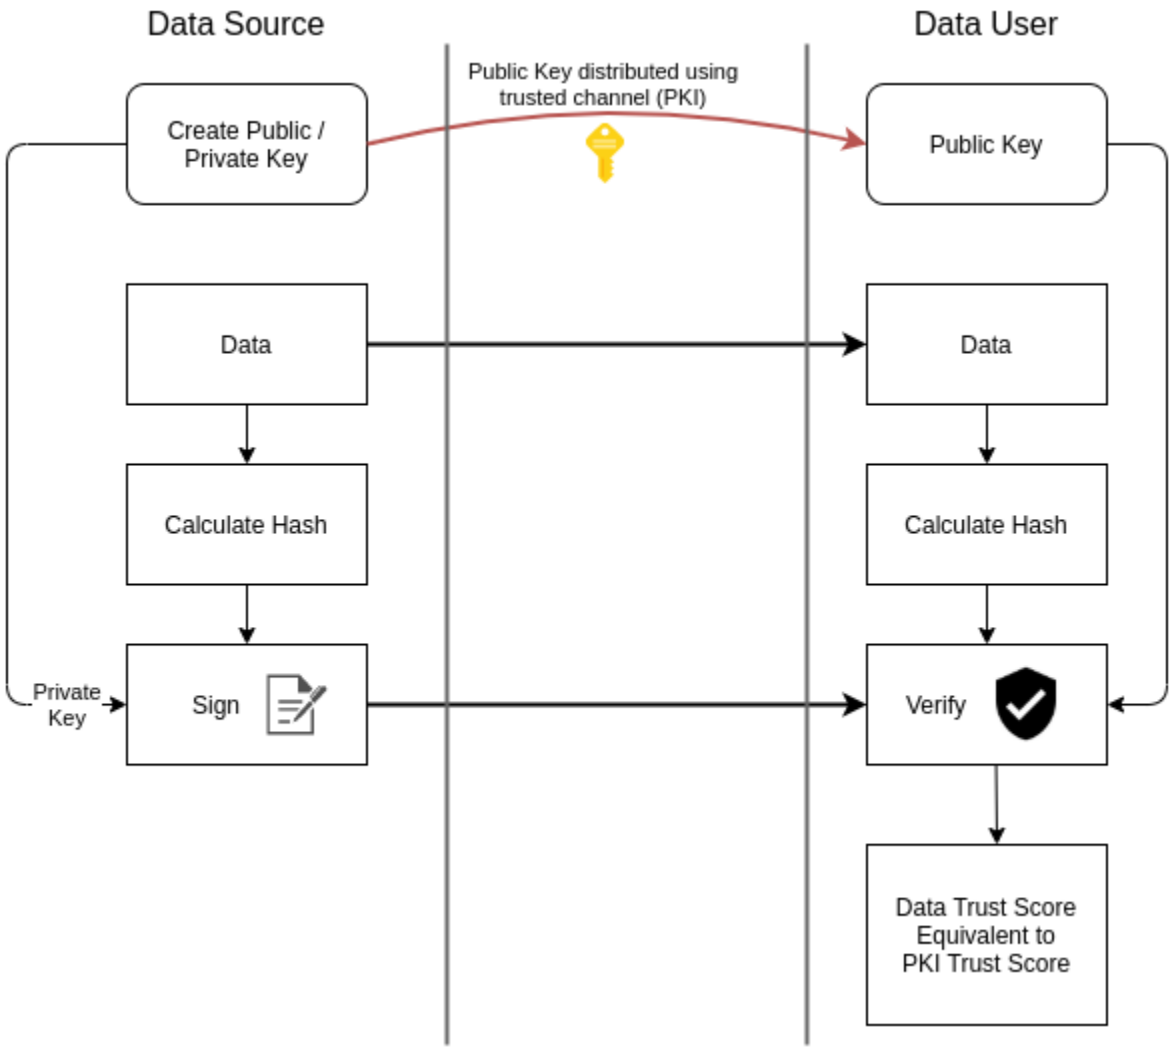

Figure 3: Cryptographic Signing of Data

Trusting a System

Referring back to Figure 1: GPS Chain of Trust we can see that there are different types of entities that we have to establish a trust score for:

- Design Intellectual Property (source code, schematics, etc)
- Data inputs (RF inputs like GPS, local databus traffic)
- Hardware (supply chain, post installation tampering)
- Processes (software build process, software loading process)
- People (technicians, engineers, operators)

Most of these elements will apply to most avionics systems, there may be additional types of entities required for different systems.

Trusting a Design

Before starting development of a new avionics device it is important to understand the expected level of trust of the design in a complete system. A trust score can be assigned to a design based on the trust assurance features incorporated into the design, any known vulnerabilities in the design, including the likelihood that untrusted entities have had access to the design, and any ongoing cybersecurity research on the design.

An additional advantage of using a zero trust model and applying a trust score to each design is that it makes it possible to quickly determine the system level impact for any newly discovered vulnerabilities. This can help in making operational decisions based on newly discovered vulnerabilities, devices with known vulnerabilities don't necessarily need to result in aircraft grounding for instance, if an updated new trust score doesn't have a detrimental impact on the system level trust score.

Trust Assurance Features

In order to achieve the highest possible trust score a design requires features that provide access controls on any data inputs and requires features that provide a means to identify that it is the source of any data outputs. This includes configuration and control data, operational data, and software updates.

Due to system level restrictions this may not always be possible, if using unverifiable ARINC-429 data from a legacy device for instance. In this case the trust score for any outputs derived from the untrusted data source should be lowered.

These features would typically utilize:

- Public Key Infrastructure (PKI) to sign and possibly encrypt outgoing data and verify incoming data, authentication and authorization systems to ensure that only trusted personnel can adjust configuration settings on a design
- Secure boot and encrypted data at rest to ensure that software and configuration setting can't be changed by someone with physical access to the device
- Signature verification systems for software updates.

These features are all discussed in more detail in upcoming sections.

Cybersecurity Risk Assessment

It is important to assess and track any known vulnerabilities in a design that may adversely affect the trust score of the design. This should start with a risk assessment before a design is released, any vulnerabilities that may have an impact on the device's level of trust must be assessed, and the trust score adjusted until the vulnerability can be addressed.

Traditionally, a risk assessment would be performed on a design, as opposed to a trust assessment. For the purposes of this assessment risk can be thought of as the inverse of trust. Risk is the likelihood that something bad will happen compared to how bad the event is, ie. likelihood times impact.

Assigning values to likelihood and impact can be useful for comparing the risk of multiple vulnerabilities and making decisions of what vulnerabilities need to be addressed. When using risk assessments in this way it is important to take into account that likelihood and impact estimates can vary greatly depending on the experience and viewpoint of the person making the estimates.

A system like the Common Vulnerability Scoring System (CVSS) can be used to reduce risk assessment variability. These systems are highly recommended to reduce risk estimate variation, we use a version of CVSS for all of our avionics vulnerability assessments.

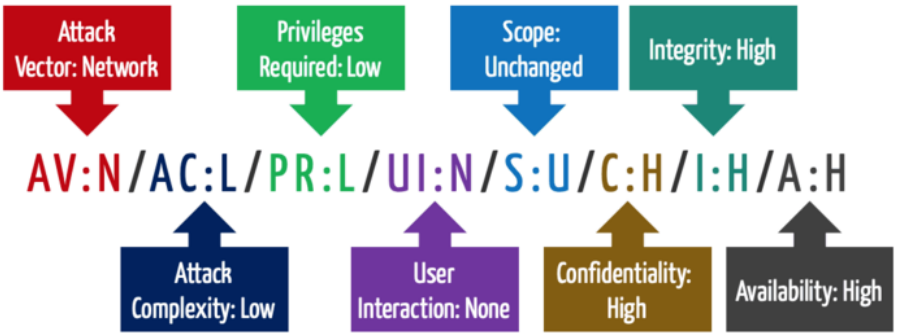

Figure 4: CVSS Element Breakdown

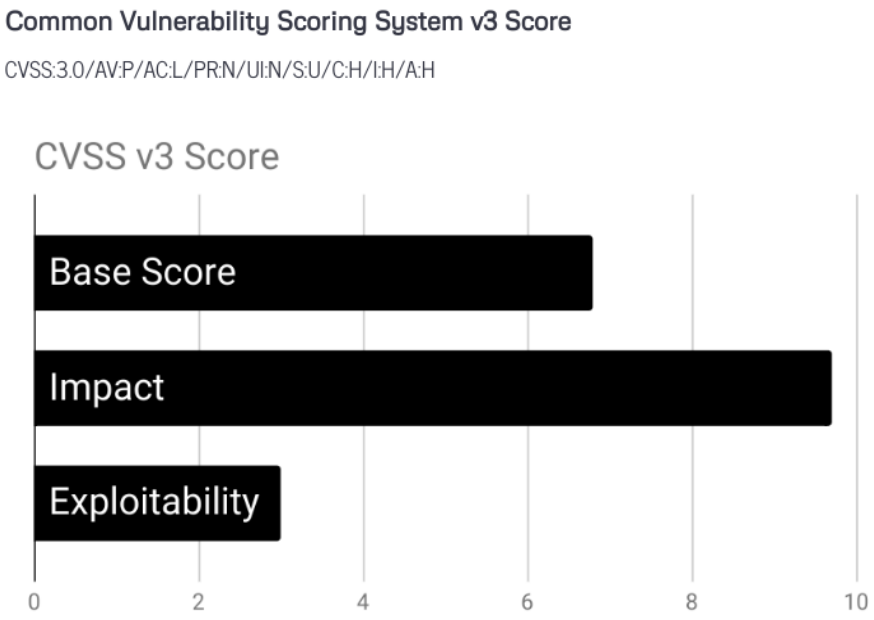

Figure 5: Example CVSS Score from CCX Technologies Tools

The aviation industry is in the process of implementing the DO-326/ED-202 standards which can be used to perform an initial risk assessment. This is a great starting point, but unfortunately these standards are already dated. These standards rely on the concept of a security perimeter, which has been proven to be a failed model and is no longer used when architecting secure IT systems. These standards also use dated definitions of likelihood and risk which can result in large variance in assessments depending on the experience level and opinions of the assessor. These new standards also rely on a database of known avionics vulnerabilities, which does not currently exist.

Automated CVE Analysis and a Software Manifest

We currently require Certificate of Conformance tracking for all hardware components in an avionics system but we have no such requirement for software. If all equipment had either a software manifest, or if we were to require continual automated checks for software dependencies against known CVE databases it would provide a starting basis for assigning trust levels to a software design.

At CCX Technologies we automatically run these checks for all of our dependencies on every major release.

It is important to note that just because a design uses software that has a known vulnerability it does not mean that that design is vulnerable itself. It may not use that specific feature, or it may include additional features that ensure that a vulnerability is not exposed. Automated CVE checks are not a panacea but they are easy to run and do provide a good starting baseline for generating a trust score for a software artifact.

Ongoing Cybersecurity Research

In order to secure a design it is imperative that there is ongoing research to identify unknown vulnerabilities. Without this research it is impossible to ensure that the trust score for a design hasn't been compromised due to a new innovation. The longer a design is in service, the more likely that there are new vulnerabilities that weren't considered in the initial assessment.

Though the avionics industry does benefit from general cybersecurity research, like the discovery of the Spectre and Meltdown vulnerabilities, compared to other industries there is very little avionics specific research being performed at this time.

This is primarily due to the lack of access to equipment that can be safely used for research. At CCX Technologies we have our own lab that we can use to simulate active flights on actual avionics equipment. We use this lab to develop and test our products, but there are very few of these labs in existence; the USAF has one as well, as do a handful of other organizations. Broader access to these labs is critical in ensuring that avionics systems remain safe and secure.

Another significant obstacle in performing responsible cybersecurity research on avionics systems is lack of access to documentation and software; the ARINC and RTCA standards are readily available for a reasonable fee, but avionics manufacturer documentation is essentially impossible to obtain through legal channels unless the equipment is being directly installed in a specific aircraft (ie. you will need a tail-number). It's unfortunate because the offensive side does have access to all of these documentation sets through less-legal channels. If there was a document repository where registered and responsible researchers could get access to documentation and software updates it would greatly benefit the industry as a whole.

Through cooperation and collaboration it is possible for the industry to establish cybersecurity testing labs, access to equipment and documentation, and create a central database of known vulnerabilities. Without this it will be very difficult to completely secure avionics systems.

Restricted Access to Design Files

In order for a design to be trusted all of the trusted elements in the design, and the people adding content to the trusted elements of a design need to be trusted.

This does not necessarily mean all elements of a design need to be trusted, but if an element has to be trusted it's important that there are access restrictions for that trusted element.

Tools currently exist that can be used to limit and log write access to design files, at CCX Technologies we use GitHub for all design files.

It is also possible to assess dependencies and ensure that dependencies are being built from assessed source code. At CCX Technologies we use buildroot, which will verify source code against a precalculated sha256 checksum before building. Other tools are available that can be integrated into any design process.

Trusting Data Input

In a zero trust system each consumer of data needs to be able to determine the level of trust of any input data before use. This allows the consumer to either reject data with too low of a trust score, or limit the usage of this data; similar to how the degraded accuracy tag can be used in an ARINC-429 system.

There are multiple ways to establish a trust score for incoming data, if possible we would recommend implementing all of them.

Cryptographic Signatures (the best way)

Verification of source can be accomplished using cryptographic signatures. A hash of the data can be signed before (or even after) it is sent using a private key and the signature verified by any consumers using a trusted public key. This requires Public Key Infrastructure (PKI) so that each consumer can be supplied with trusted keys, either before each flight, or during maintenance activities. These keys are similar to those described in ARINC-842 and used by airlines to load firmware and operations data.

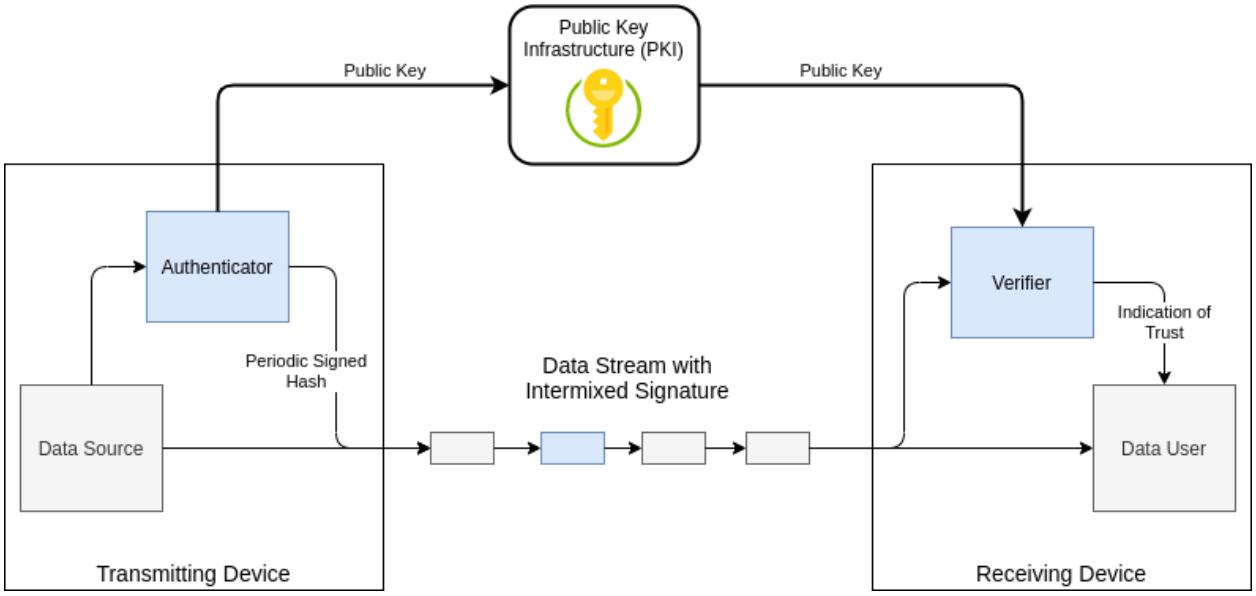

Figure 4: Signing Data

IT networks that utilize the zero trust model will typically use an overlying system to verify signatures and apply trust scores. That topology isn't realistic for an avionics system because it would create a single point of failure. A distributed model like the one pictured above is much more robust for safety critical systems.

Legacy Avionics Data Protocols

A significant barrier to adding data verification features to avionics systems is the almost universal use of legacy protocols that don't include verification features. The most popular avionics networks, like ARINC-429, ARINC-717, ARINC-664, AFDX, CAN Bus, MIL-1553, ADS-B, and ACARS do not provide the required feature sets. Any data received on these interfaces will be treated as legitimate data, regardless of the actual sender.

In addition, most of these interfaces are high-impedance when there is no active driver. This makes it possible to install a data injector on the bus and transmit illegitimate data when the bus is inactive. The RF based systems like ADS-B are even easier to inject data into, since no physical connection is required.

It is possible to add data-signing features to legacy protocols in a backwards compatible way; a solution for ARINC-429 is discussed in another CCX Technologies white paper, WP-002. This would require software updates to all of the equipment on the bus, or the addition of hardware at the wiring harness connectors.

The only viable long-term solution to this problem will require the avionics industry to come together and adopt common data authentication methods that can be integrated into equipment from all manufacturers. It is technically feasible to create a zero-trust network using legacy avionics networks, but it will require an industry commitment to implement it.

Redundant Data Sources (a distant second)

Another method that can be used to increase the trust score for incoming data is to compare it to the same or similar data from alternative sources. Since it's theoretically possible to affect the operation of every sensor in a system this method could never be used to score the highest possible level of trust but it can be used to increase the trustworthiness of data as it would greatly increase the cost of an attack.

As an example, velocity can be calculated using the GPS output and compared to velocity measured by on-board accelerometers. If they are equal then the velocity data would have a higher trust score.

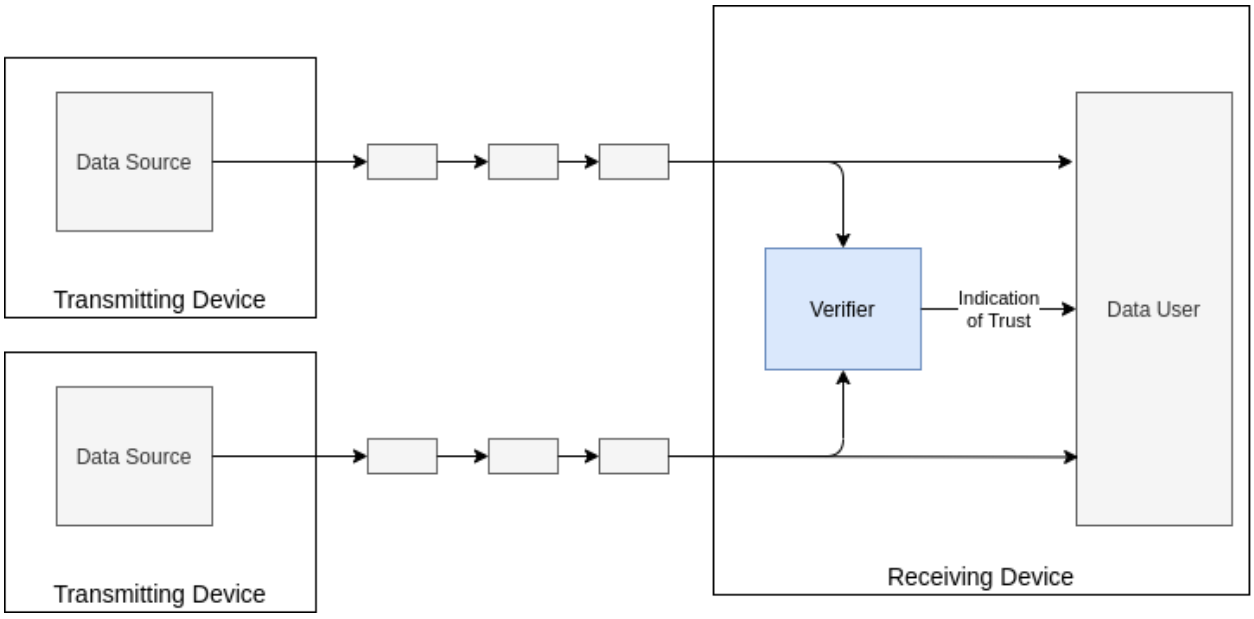

Figure 5: Comparing Data Sources

Comparison based trust scores can be added to existing systems without affecting the operation of the system. CCX Technologies Avionics IDS is an example of such a system, it can collect data from

multiple data-buses and compare data from different data sources. If an anomaly is discovered an alert is raised which indicates that the system itself is operating at a lower trust level than desired.

Unexplainable Changes in Data (the last resort)

The least attractive but still marginally better than nothing method to increase the trustworthiness of data is to compare the current data samples to previous samples to identify unrealistic rates of change.

This method can be added to an existing system, or even a single device.

For instance, some airports have added a system that will monitor ADS-B transponder outputs for sudden location changes to identify possible spoofing attacks. It is also possible to implement this method using the CCX Technologies Avionics IDS.

Continual Monitoring and Logging

Continual monitoring and logging of avionics data provides an opportunity for ongoing cybersecurity research, which can be used to increase trust in the system design and operational policies, and it also can be used to reduce the down-time experienced when there is a cyber related event.

Without continual monitoring and logging there is no way of knowing if a system has been compromised, it is a fundamental component to any cybersecurity solution. It would be similar to a tight-rope walker working without a net, it may look more exciting but perhaps excitement shouldn't be a goal of cybersecurity engineering.

CCX Technologies has developed a complete system that can be used to collect and monitor avionics data for cybersecurity purposes. Data can be collected from any avionics network and can be pushed into an Intrusion Detection System (IDS) that can identify different types of attacks.

The collected data can also be used post-flight to identify unusual patterns and events, and can be used during an investigation after an event. This is very similar to, and can overlap with a FOQA program..

Trusting Hardware

In order to trust the software running in a system, and the data generated by that software it is paramount that the hardware that comprises a system is also trusted.

If it is possible for anyone who comes into contact with the hardware to alter the operation of the system then in order to create a zero trust network the system's hardware would have to be locked in some way that only provides access to trusted individuals. This would require the addition of physical security devices, adding weight to the aircraft, and would potentially affect the design of the aircraft itself.

Instead of physical security controls it is also possible to utilize secure boot and encrypted data at rest technology to achieve the same goals.

Secure Boot / Encrypted Data at Rest

Most avionics systems don't currently support secure boot and encrypted data at rest. Without implementing these features with physical access to equipment it is straight-forward to update configuration, update software, install malware, extract private keys, etc.

The lack of secure boot and encrypted data at rest has multiple implications:

- It turns physical security into a single perimeter defense
- It provides a means to extract keys and software from one system that can be used to access another system of the same type.
- Software can be extracted and decompiled, decompiling software from binaries is a very common way to identify software vulnerabilities.
- It provides a straight-forward means to extract data from the equipment, like aircraft tracking data, engine performance data, and flight operations data

Secure boot provides a mechanism that ensures that only signed software can boot on a specific system. Without secure boot it is possible to boot an alternative bootloader and operating system which can be used to affect the operation of the system and extract data from the system. Since secure boot starts with secure boot ROM firmware some hardware features are required to implement it.

If data is not encrypted at rest it is possible to extract that data from any storage device, regardless of what type of device it is. This data can include firmware, software, private keys, operational data, etc.. The only way to protect this data is by encrypting it with a protected key.

The primary problem encrypting data at rest on an embedded platform is protecting the key used to decrypt the data. On a laptop, PC, etc. a user can be queried for a password but that isn't possible on an embedded system, the key must be stored in a secure manner on the system itself. Special hardware is required that can decrypt a protected key after a system has booted from verified secure boot software.

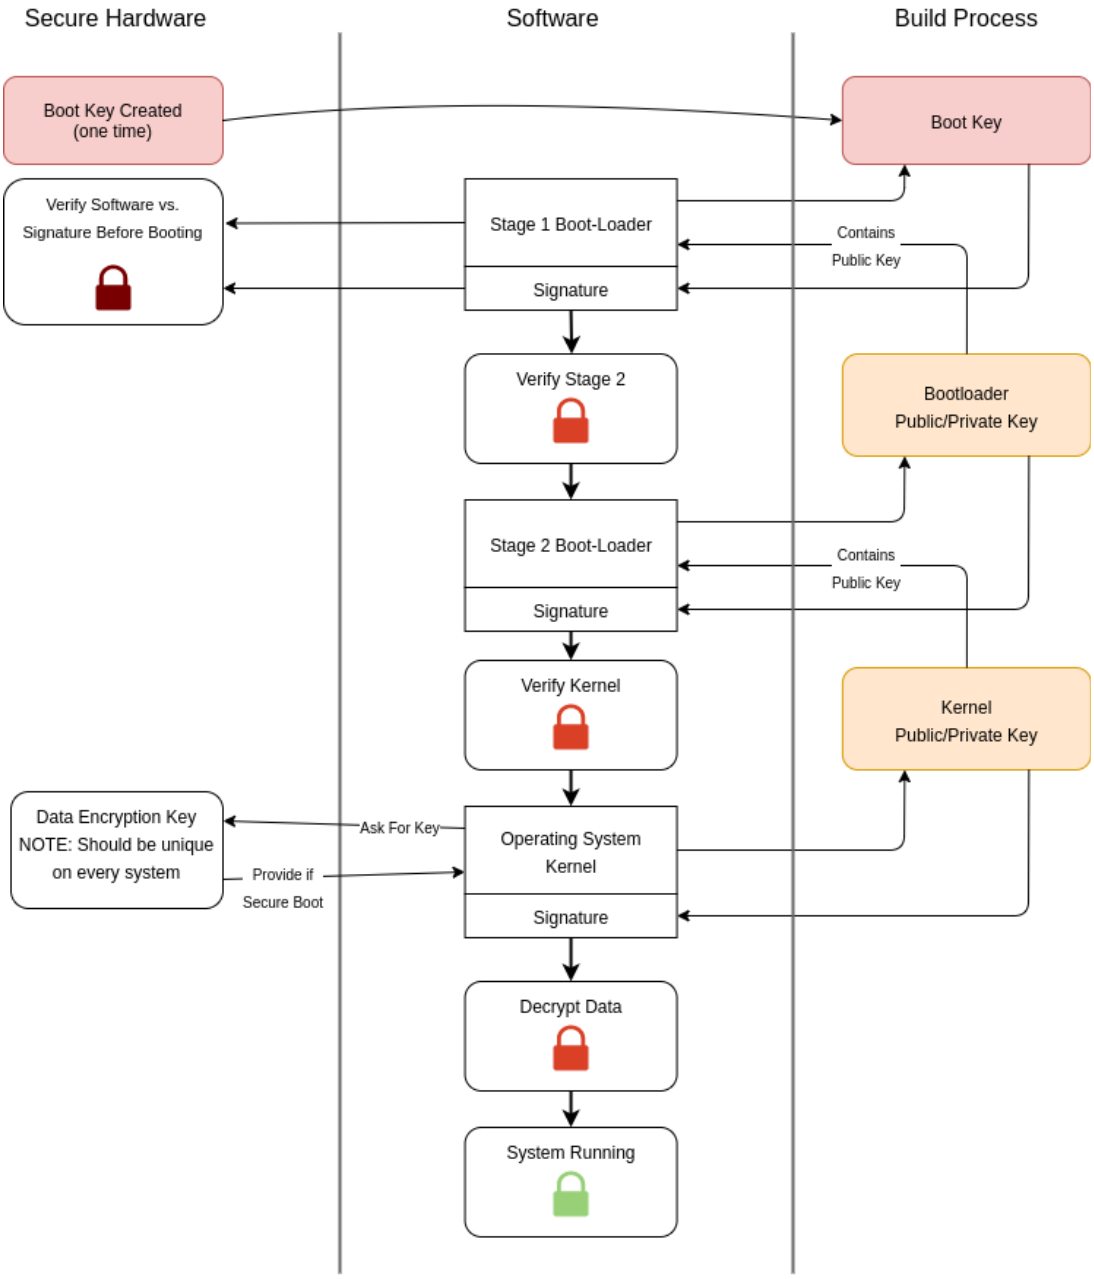

Figure 6: Secure Boot Process

Supply Chain

Trust in a hardware component can only be as high as the trust in the supply chain used to build or acquire that component. Though it would be very expensive; potentially impossible, to completely trust a supply chain and thereby a hardware component it is possible to assign a high level of trust to a supply chain by implementing some protections.

Trickle-down requirements, similar to how the DoD will be requiring CMMC certification on all contractors, sub-contractors, is the best, proven, way to ensure and enforce supply chain security. A system that includes individual requirements for security and tracking for all subcontractors that can be audited and certified would greatly increase supply chain trust. A system like this would take buy-in and organisation from either a large industry body, or a regulatory agency.

There are other things that individual manufacturers can do today to help secure their supply chain.

There are some steps in most supply chains that need to be explicitly protected, burning secure boot keys for instance. Specific policies and procedures are required for these steps, and if any of these steps are performed by a subcontractor, audits and proof of conformance is absolutely required.

Certificate of Conformance documentation can be requested and archived for all system components.

It is important to use trusted manufacturers for all at-risk components, for instance, the manufacturer of a key component like a processor should have a high trust score and should be able to provide documentation and information about processes that secure the manufacturing of the device. Other suppliers, like a machine shop producing a machined aluminium part may not require a high trust score, or possibly any trust score at all.

Trusting Processes

There are several ongoing design processes required to maintain an avionics system that need to be trusted in order to trust the system itself. These include things such as building software, loading software, releasing service bulletins, etc.

Software Build Process

The SolarWinds attack has demonstrated that a system is only as trustworthy as its least trustworthy component, and in a lot of cases the least trustworthy component can be the software build process.

A trusted software build process requires a means to collect and verify trusted source code, build it into a package with a set of trusted compilers, on a trusted system, and then sign the resulting output files so that they can be verified to be trusted by end users.

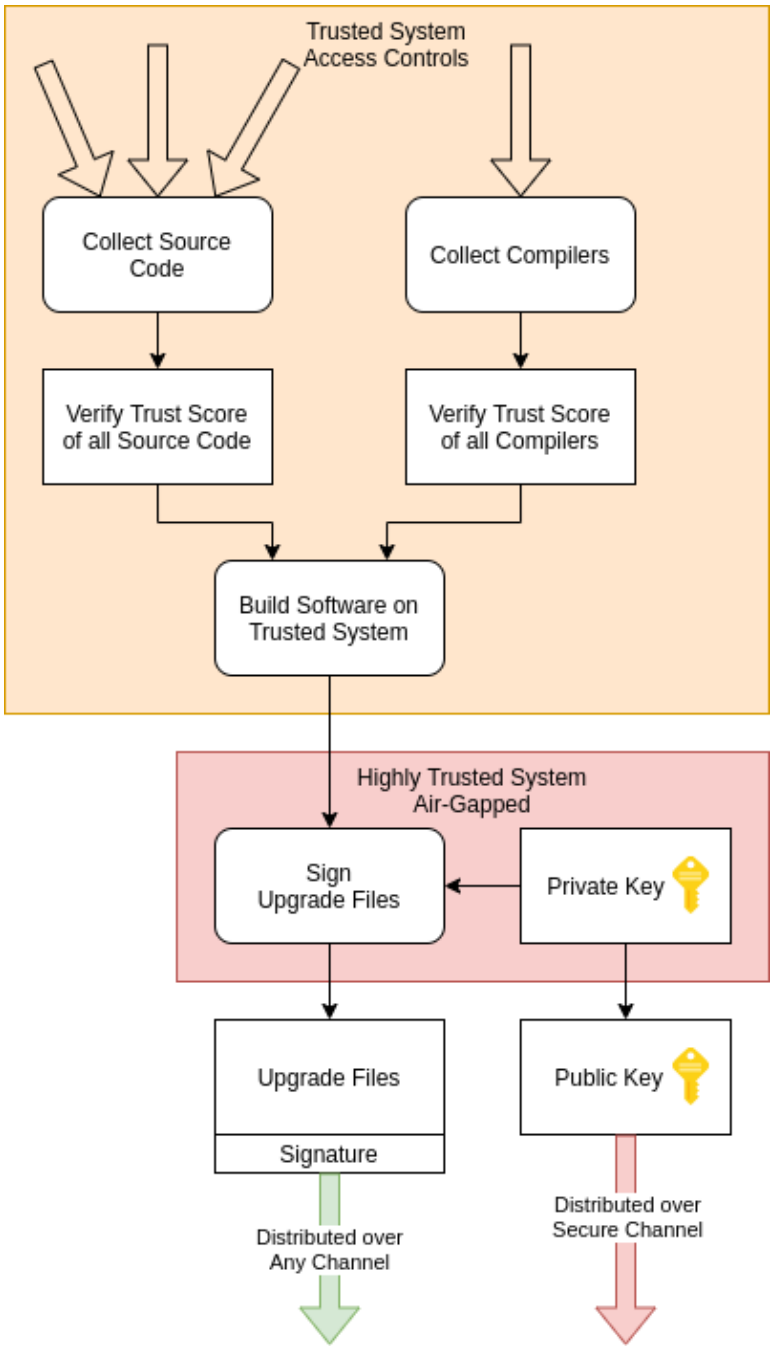

Figure 7: Trusted Build System

Software Loading Process

A trusted software loading process requires a means to verify that the software being loaded is from a trusted source and trusted build process before it is loaded. This can be implemented using public key infrastructure and cryptographic signing. The system described in ARINC-842 can be used to implement this.

Trusting People

In order for a system to be trusted, people who have the agency to affect the operation of the system must also be trusted. For an aircraft this may include pilots, technicians, flight crew, baggage handlers, etc.

Most airlines and airports already have controls that limit and record access to aircraft, and equipment. These systems can be used in establishing zero trust avionics systems but since they are not necessarily intended for this specific use some additional processes or technology may be required.

Most business aviation and general aviation operators and airports have limited or no access controls for their aircraft, ideally the industry could come up with a complete solution that would benefit all operators.

Configuration and Diagnostics Access

Technicians require access to the configuration database and diagnostic tools in order to service an aircraft. Access is typically provided through a maintenance port, through a browser-based GUI, or through offline tools and a configuration loading procedure. Access to these systems is typically provided by shared credentials, like a shared password or a shared key file.

Due to the transient nature of aircraft these systems need to be serviceable at multiple geographic locations, this makes a shared password system convenient.

The are three main issues with a shared password system:

1. Anyone who has ever serviced that equipment will have access to it at any future time, and can provide access to anyone else. This greatly expands the network of trust that a password affords, potentially pushing it into untrusted parties, or even publicly available.
2. Shared passwords are very difficult to change if there is a breach as all parties using these credentials will have to be informed of the change. This tends to limit the amount of times the passwords are updated, if ever.
3. Shared passwords make it impossible to attribute system access to individuals, which is critical information when investigating any potential cyber events.

In order to responsibly provide equipment access in a zero trust manner a system that includes a means to identify and log individuals, and providing and revoking access to specific aircraft and systems is required. The tools for such a system exist today and are in use by most large organizations to provide access to IT infrastructure but in order to use these tools on a avionics system there would need to be an industry consensus specific tools and tool management.

IMPLEMENTING ZERO TRUST ON AN EXISTING SYSTEM

It's possible to evolve an existing system closer to a zero trust model in increments. Equipment and software can be added to a system to monitor the current trust levels without the system necessarily being able to automatically react to changes. The trust score can instead be presented to an operator (either the aircrew, or in an operations center on the ground) who can make operational decisions and in turn affect the operation of the system.

For instance, it's possible to passively monitor all of the traffic in a network and identify specific events that would indicate that the system is operating at a lower trust level. It's possible to collect information on all people that have had access to a system, which can be referred to at a later date if one of the people has been deemed to be less trustworthy. It's possible to add supply chain monitoring and verification services in place independently of the operation of the system. All without making major changes to the existing system.

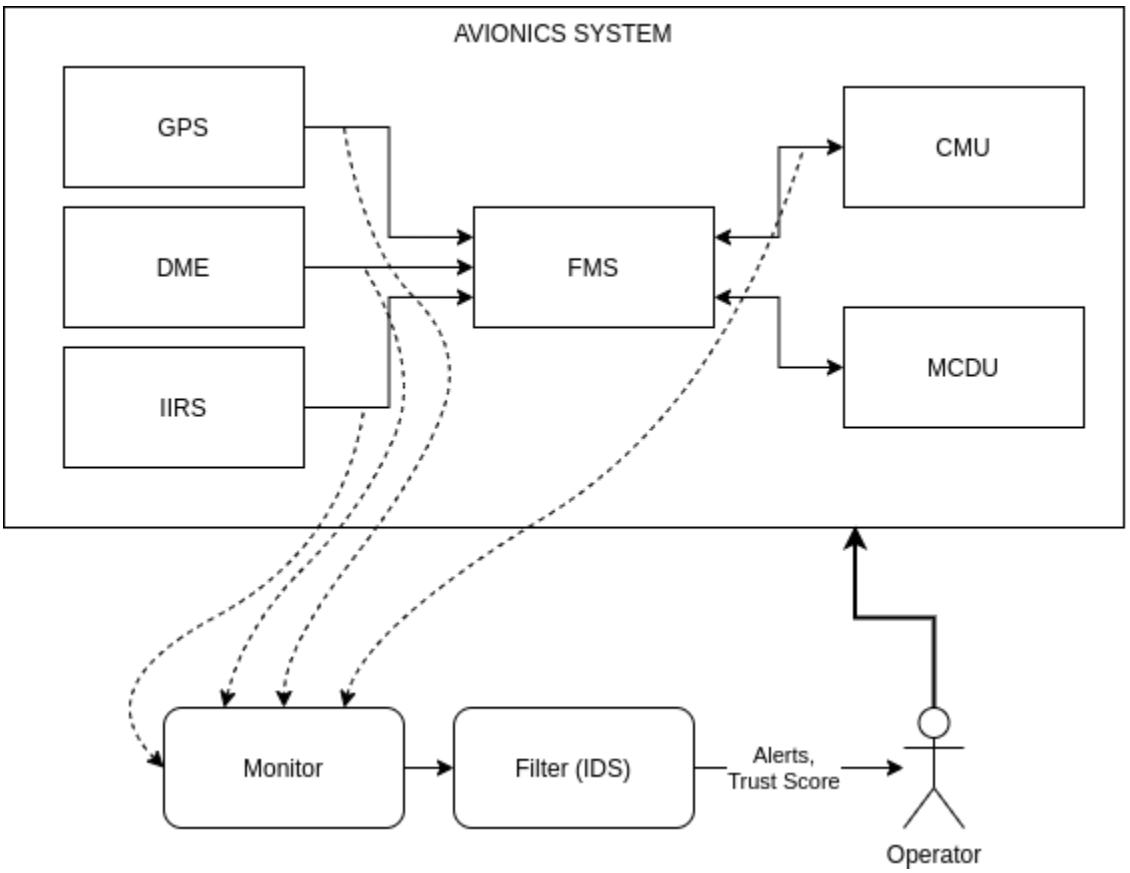

Figure 8: Overlay Network Monitoring for System Trust

As these technologies are added, the outputs of these technologies can be integrated into the operation of any newly designed equipment that is added to the system, to improve responsiveness (and security) and reduce operator load.

A Complete and Comprehensive Plan

In order to fully implement a zero trust model on an avionics system we need to start with a comprehensive plan that provides a means to establish trust in every part of the system and then introduces technology to track and verify trust while the system is in use.

Once this plan has been created parts of the plan can be implemented before others, any new security feature, if properly implemented, will increase the security of the overall system. The ultimate goal of addressing cybersecurity vulnerabilities is to increase the cost of any attack to the system. This isn't an all or nothing process, there is a lot of value in starting anywhere and building the system from there.

Without a comprehensive plan though, it will be difficult to identify the best place to start, and it will reduce the possibility of inadvertently adding technology that provides no net benefit.

CONCLUSIONS

This paper introduced the concept of the Zero Trust Model and proposed processes and technologies that could be used to develop a complete Zero Trust Avionics Network.

In order to develop a complete solution new regulations and guides are required, which include descriptions of solutions to measuring and verifying trust in all aspects of the system.

In each system a complete chain of trust needs to be documented, established and audited as the system is designed, maintained, and in operation.

It is possible to introduce some zero trust technologies before implementing a complete solution, any increase in security will benefit the system as a whole.

Recommendations from this paper include:

- Document a complete chain of trust for all devices and systems.
- Create a trust policy defining minimum trust levels required for all elements in the chain of trust
- Perform a cybersecurity vulnerability assessments before releasing any new designs
- Ensure that all design source code and files are protected during the design process
- Perform ongoing cybersecurity research on all avionics systems to identify unknown vulnerabilities
- Provide tools that can be used to verify the source of all data in a network
- Provide tools to continually monitor and log the state of a system in operation
- Provide secure boot and encrypted data in rest features for all new avionics devices
- Add auditable controls to ensure supply chain integrity
- Ensure all software build processes are secure and verifiable
- Provide a means to control system access to individual people