Providing Cyber Situational Awareness on Defense Platform Networks

Patrick M. Hayden
David K. Woolrich
Katherine D. Sobolewski

ABSTRACT

Modern defense platforms are at increasing risk of cyber-attack from sophisticated adversaries. These platforms do not currently provide the situational awareness necessary to identify when they are under cyber-attack, nor to detect that a constituent subsystem may be in a compromised state. Long-term improvements can be made to the security posture of these platforms by iterative application of cyber risk assessments and subsystem hardening, but this is a time-consuming and costly task. Monitoring platform communication networks for malicious activity is an attractive solution for achieving improved cyber security on defense platforms in the near term. The MIL-STD-1553 bus is central to the operation of a broad range of defense platforms, making 1553 security solutions generally applicable. This article presents our research into the susceptibility of modern defense platforms to cyber-attack. We discuss risk factors contributing to cyber access, and command and control channels. We then describe a range of platform cyberattack classes, while considering the observables and indicators present on the 1553 bus. Finally, we examine factors and considerations relating to implementation of a "Cyber Warning Receiver" solution approach for detection of such attacks.

THE THREAT IS REAL

For as long as weapons system platforms have been called upon to perform missions in contested spaces, the military has sought to protect the warfighter by equipping these platforms with survivability equipment. This equipment detects threats from across the various domains in which the platform operates, and alerts operators while taking appropriate response measures. As technology and connectivity of these platforms evolves, and increasing sophistication is realized through automation, a new threat domain has emerged. This threat lurks in the dark, escaping detection by human eyes

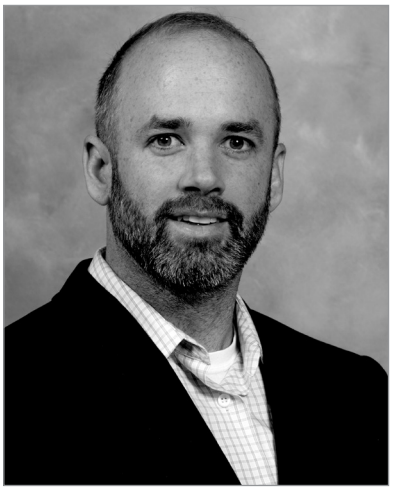

Patrick M. Hayden is the Chief Engineer for Cyber Electronic Protection programs at BAE Systems. Patrick has 10 years of experience in the cyber security field, performing systems assessment, software reverse engineering, and vulnerability research covering both offensive and defensive perspectives across a wide range of targets and applications.

and ears, yet it has a clear potential for harm to the warfighter and the mission. This is the cyber threat, and it is real.

Cyberattacks become a credible threat if there is a reasonable expectation that a malicious actor could gain access to a defense platform, achieve a persistent malware presence, and subsequently trigger this malware to impart a damaging effect. While there is a lack of openly documented cyberattacks against DoD platforms, published examples against similar systems in other industries provide a compelling case for the feasibility of such attacks.

Unlike traditional kinetic attacks, cyberattacks are not limited in range. In cyberspace, there are no concrete boundaries or borders. A malicious actor in a faraway land can achieve the same reach as someone attacking a target from the same city. Cyberattacks also have greater flexibility in their timing than most traditional attack types. A complete cyberattack may begin well in advance of the realization of any ultimate effect. Attackers can leverage a latent presence at a critical moment in the future to achieve their end goals. This may occur at a predetermined time, or when a predetermined condition is met and may affect a single platform or an entire compromised squadron simultaneously.

Our platforms are at risk regardless of their location, from the battlefield to their home base. Despite these realities, many weapons system platforms operate without sufficient means of providing detailed situational awareness into their cybersecurity state.

LESSONS FROM INDUSTRY

Throughout industry and academia, we hear more and more about attacks against embedded systems and other smart devices. Attacks originate from threats that range from individual troublemakers

David K. Woolrich is a program manager in BAE Systems–Survivability, Targeting, and Sensing Solutions business area. His background is information security, information assurance, and cyber electronic warfare. He is focused on increasing awareness and survivability of DoD assets from cyber attacks.

to state-sponsored hacking groups. These attacks can be foul-mouthed hackers yelling at children via smart baby monitors [1], using SmartTVs as entrance points to home networks [2], entire automobiles being taken over remotely [3], or debilitating modification of industrial control processes [4].

In 2015, security researchers Dr. Charlie Miller and Christ Valasek were able to remotely access an unaltered SUV, controlling everything from the volume of the radio, to the transmission and steering of the Jeep. Initially, takeover of the SUV required access to the USB connection on the automobile, which is normally reserved for vehicle maintenance. With time, however, Dr. Miller and Mr. Valasek were able to gain access to the SUV through its onboard cellular network, traverse multiple Jeep subsystems, and ultimately control physical aspects of the SUV from their hotel room while the Jeep was traveling on a highway.

In 2017, security consultants ARS were able to demonstrate the insertion of malicious code over a broadcasted TV signal. This malicious code was transmitted via the digital video broadcasting—terrestrial signal and once executed allowed full remote control of the TV with no physical access required. The transmitted code was able to exploit a vulnerability in the smart TV's web browser enabling root access for the attacker. If a broadcast station were compromised, this attack could be delivered to any vulnerable TV within the broadcast towers' range.

As systems become more complex and gain more parts, supply chains for devices and systems become more spread out and global. This creates difficultly in validating the pedigree of 100% of the components on any one system. A single system could be comprised of hundreds or thousands of components. Without rigorous vetting of all parts, it is possible that compromised or counterfeit parts could be introduced into the system. This fear was realized by the DoD

Katie D. Sobolewski is a Technology Development Manager for Cyber Electronic Protection at BAE Systems with experience in cyber defense, cyber electronic warfare, and platform protection. Katie has a background in algorithm development, signal processing, and optimization for increased system performance.

when foreign chip manufacturer Lenovo was suspected of introducing phone-home capabilities into their chipsets [5], sparking fear within the US government that their systems could be compromised. A 2017 Defense Science Board Task force on Cyber Supply Chain confirms the supply chain to be a real risk to DoD assets.

INCREASINGLY CONNECTED PLATFORMS

The examples above represent three distinct attack access vectors against embedded systems: supply chain compromise (microprocessor compromise), maintenance pathways (vehicle USB), and compromising data links (broadcasted malware in TV signal). Current trends in weapons system platform modernization suggest that these same vectors are also applicable to defense platforms.

Most platforms are comprised of a diverse mix of commercial off-the-shelf, government off-the-shelf and custom hardware and software. Components have been developed over multiple iterations and many years. These components are sourced from a wide array of providers, each with different security practices. They leverage different processor types, operating systems, and source codes. Although this diversity may help improve the security of the system to prevent an attack from spreading [6], it also provides a large surface area for attackers to address, increasing the risk that they could establish at least a single point of presence via supply chain compromise.

Platforms also employ a range of data products throughout the course of their lifecycle to accomplish their mission. Flight-line maintenance activities, mission preparation, and post-mission analysis activities all involve connecting platforms to a variety of support equipment. These numerous pathways each create new opportunities for an attacker to gain presence or provide control.

Also like their commercial counterparts, platforms are increasingly interconnected via data links and tactical networks during mission execution. Connectivity via these links provides pathways that could extend attack impact beyond a single infected platform, by which sophisticated malware could propagate from one platform to another, or by which attackers could exert control over their payloads.

PARALLEL SECURITY APPROACHES

The trends of increasing computer automation and platform interconnectivity are here to stay, as they enable distinct tactical advantages. Platform security must improve to address these trends head on.

The two complementary approaches are common when it comes to traditional IT security measures. These apply in the world of defense platforms as well. The first is host-based security, where the security of the individual boxes on a network are improved to achieve increased security for the system overall. The second is network-based-security, where communications between hosts on a network are monitored to detect and potentially intercept malicious activity.

Build Secure

Improving the security of each subsystem on a platform is a great option and a necessary step in securing future platforms, but it's time-consuming and costly. There is certainly much to be gained through a thorough security review of each subsystem on a network, along with the implementation of bug fixes, configuration hardening, host-based security state monitoring, and other general security improvements. In many cases though, platform subsystems are not actively involved in current upgrades. Given the range of implementations present across all the subsystems on a given platform, there is no single silver bullet solution for host-based protection, such as a "platform antivirus" or the like. Instead, platform stakeholders should consider incorporating cybersecurity hardening requirements during subsystem upgrades, as informed by the outcomes of cyber risk assessments against their platform.

A Cyber Warning Receiver, designed to look for malicious activity on the 1553 bus can provide the broadly applicable solution necessary to achieve near-term game-changing platform security enhancement.

Network Lockdown

The actions necessary to conduct a cyber-attack, and the effects will, in the majority of cases, be observable via the data networks used to communicate commands, status, and data between systems on a platform. Although a compromised box could affect its function without leveraging any network communications, attacks against other system components will involve the use of platform networks. With this in mind, monitoring these networks for malicious activity can provide the situational awareness necessary to detect an attack and inform an appropriate response.

A common set of networks covers the vast majority of communications occurring on these platforms. In particular, the U.S. Army's Common Avionics Architecture System (CAAS) depicted in Figure 1 relies heavily on Ethernet and MIL-STD-1553 (or fiber optic 1773) networks, and also includes support for RS-232, RS-422, Arinc 429, analog and discrete signals [7].

Within a broad range of platforms employed by the Army and other services, 1553 networks form the backbone for communications between platform subsystems. They provide the critical link between pilot interface equipment like displays and keypads, and the endpoint devices that actually implement critical control or measurement capabilities. Monitoring the 1553 bus would provide a high degree of visibility into cyberattacks. A Cyber Warning Receiver, designed to look specifically for malicious activity on the 1553 bus can provide the general broadly applicable solution necessary to achieve near-term game-changing platform security enhancement. This device can be rapidly adapted to fit a range of platforms and provide immense benefit to the cyber security posture of the overall fleet.

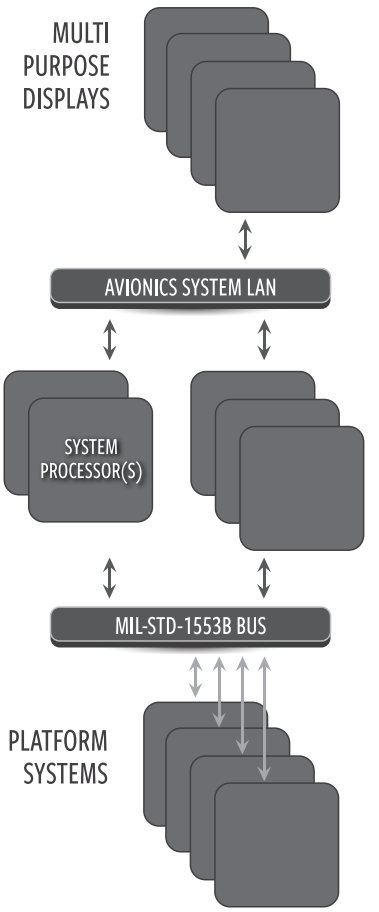

Figure 1: The Common Avionics Architecture System

THE MIL-STD-1553 NETWORK

MIL-STD-1553 is a serial messaging interface that prescribes a physical layer and data link protocol for exchange of data between a set of terminals residing on a bus. The physical network topology is flat, with all remote terminals (RTs) connected and listening to the same bus signal [8].

BC to Specific RT(s)	BC to All RT (Broadcast)
1. Controller to RT Transfer	1. Controller to RT(s) Transfer
2. RT to Controller Transfer	2. RT to RT(s) Transfers
3. RT to RT Transfers	3. Mode Command Without Data Word (Broadcast)
4. Mode Command Without Data Word	4. Mode Command With Data Word (Broadcast)
5. Mode Command With Data Word (Transmit)	
6. Mode Command With Data Word (Receive)	

Table 1: MIL-STD-1553 Message Types

All communications are facilitated by a single terminal designated as the Bus Controller. The Bus Controller implements a schedule on which it sends and receives information to and from the other terminals, or instructs them to pass messages between one another. Each message in the schedule is repeated at a prescribed rate, typically ranging from 50 times per second to once every two seconds. The bus also supports asynchronous messaging and supports polling for RTs that need to send an extra message on a given cycle. The 1553 bus is designed for determinism, reliability and redundancy, and comprises at least two redundant busses, and two redundant bus controllers (a primary and a backup) to enable failover in the event of a single failure conditions.

CYBER ATTACKS AND 1553

The breadth of published work on 1553 attacks is small in comparison to research for similar consumer, commercial, and industrial networks. Such networks are more openly accessible to security researchers for characterization. In particular, security research in the field of Industrial Control Systems (ICS) and Supervisory Control and Data Acquisition (SCADA) systems has illustrated the potential vulnerability of similar serial messaging interfaces. The MODBUS serial protocol, which has several features in common with 1553, has been the subject of extensive cyber security research. Huitsing, Chandia et. Al., in their paper describing attack taxonomies for Modbus Protocols [9], propose 15 distinct attacks across five classes for the Modbus serial protocol. Such findings are a useful starting point when considering the cyber security of 1553.

Through internal investment, we've adapted existing platform System Integration Labs to create a 1553 cyber security test bed. Using this as a research tool, we have begun to explore and characterize the space of 1553 attacks, considering attacks that directly target, exploit, or misuse 1553 functionality, and also attacks for which 1553 networks are involved, but not directly targeted. Our ongoing research has shown that many of the attack types conceived for other network types are also applicable to the 1553 network. The standard does not provide any security features, such as authentication or encryption that would mitigate such misuse.

The attack types available to an attacker exploiting the 1553 network depend on the specific foothold they achieve on a platform. In general, there are several positions an attacker might hold on a platform with respect to the 1553 system:

1. Attacker presence on systems outside the 1553 network that leverage data sent or received via the 1553 network;
2. Presence on a Remote Terminal connected the 1553 network;
3. Presence on a Bus Controller for the 1553 network; and
4. Multiple points of presence creating a combination of these states

Given this set of states, some of the attack types we've described and characterized are:

- Methods by which a compromised bus controller could impact the system. A compromised bus controller enables a high degree of control. It enables an attacker to initiate new messages, remove existing messages, or intercept and modify data in transit between remote terminals.
- Methods by which a compromised Remote Terminal could initiate new messages on the 1553 bus without coordination with the bus controller, impersonate a different Remote Terminal, or even attempt to become the bus controller.
- Methods by which any compromised host on the 1553 network could deny messaging between other remote terminals.
- Attacks in which basic rules and conventions of the 1553 standard, or the application layer data they contain, are violated.
- Attacks where a compromised host deliberately sends incorrect data to another host as part of the normal data exchange cycle. This could include measurement data, control commands, system status or other types of information.

Each of the attack types above have been hypothesized along with specific details relating to their realization on the 1553 bus. Some have been tested in practice. Discussion of these specific implementation details are beyond the scope of this article. Consideration of possible attack types and characterization of their effects helps inform a robust design for a platform security detection system like a Cyber Warning Receiver.

ATTACK OBSERVABLES

As the attacks described above take place on a 1553 network, they produce side effects that are observable to a high-fidelity bus monitor. For the purpose of organizing these observable side effects, the 1553 network can be considered as being comprised of several network layers, as depicted in Figure 2.

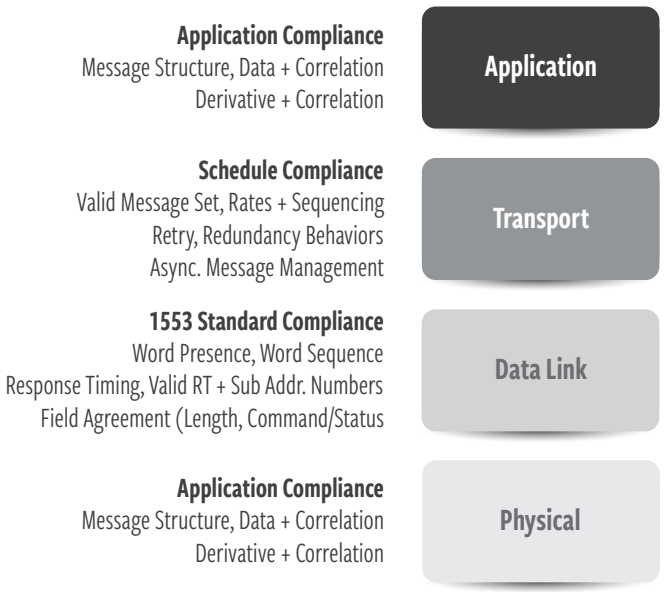

Figure 2: 1553 Network Layers and Observables

The bottom layer is the physical layer, which contains observables relating to the fundamental electrical environment necessary for proper operation of 1553. Certain attacks can cause disturbances at this level, especially in cases where misuse of the 1553 bus causes message collisions.

The next layer up is the data link layer, which covers low-level implementation details of the 1553 protocol. At this level, we can detect that only valid hosts and sub-addresses are present, and also that the expected message structure is intact, including the allowed message types and expected word sequences. Some attack types can cause changes to this ordering or produce multiple repeated copies of certain message words. The typical request and response timings for 1553 transactions can also be monitored at this level.

The next level up is a transport layer, in which platform specific attributes relating to the use of 1553 are defined. Messages that occur on 1553 can be uniquely identified by attributes including their type, source, destination and length. At this layer, we can verify that the system is using the set of messages expected to occur as part of the schedule, with the appropriate sequence and timing. Monitoring systems must account for changes to this schedule that may result from different operating modes for the platform. At this level, it's also possible to enforce that retransmit or redundancy features spreading messages across multiple busses are performing as expected without misuse.

The top level is the application layer. Details at the application layer are specific to the individual systems on the bus and their implementations. A navigation device may transmit one type of data using message formats and data representations established by its developers, while a threat warning system may use a completely different representation

for its data. Detection of a valid structure is one useful observable. Where data fields are specified or can be otherwise identified, a set of normal behaviors can be observed based on their values. For example, data may be known to have a limited range of values, to exhibit a known distribution, or to have a limited rate at which it can change. In other cases, multiple data fields might exhibit correlations, such as always moving together, or negating one another. Performance outside of these norms could be indicators for a cyber attack.

DETECTING ANOMALIES

A Cyber Warning Receiver operates by monitoring traffic and discovering anomalies in the behavior of these observations and measurements. The normal set of behaviors for each of the measurements must be characterized before deployment based on the 1553 specifications and specific inputs for the platform to be protected. Examples of specific input may include valid RT and sub-address ranges in use, and message schedule in different operating modes, and observations from collections of real world data.

In general, the higher the layer at which observation and characterization are required, the more specified a solution is to a particular attack, and the more data will be required to establish normal behavior and detect anomalies. Leveraging observable side effects that are agnostic to specific attack implementation details enables detection of attacks that have not before been observed in the wild, or preconceived by defenders. For lower layers, the number of possible attack approaches is limited, making it tractable for subject matter experts (SMEs) to explicitly define a spanning set of detectors. At these lower levels, detectors are also more portable than for higher levels. This simplifies the task of implementing cyber threat detection across platforms.

Although there are many advantages to monitoring the 1553 bus at lower levels, observations derived from these layers are not sufficient by themselves. There are important classes of cyber-attack that do not produce observable impacts at these layers. For example, manipulation of data from a given device would only be observable by changes in the platform-specific messages that exist in the application layer, as would violation of application layer message formatting. To characterize these forms of attack via the application layer, and detect them on-the-fly, more sophisticated anomaly detectors are required.

Creating anomaly detectors to operate at the application layer introduces several practical challenges:

1. Scalability to address the sheer volume of data relationships that would exist for all systems and messages across a complete defense platform. Do all of these relationships need to be enumerated by hand?
2. Managing the specifics of the application layer message formats and field locations for dozens of devices and hundreds of unique messages. Do these formats need to be manually specified to enable a practical system?

3. Discovery of subtle or secondary correlations that might escape the intuitions of human cyber defense experts and therefore remain open to exploitation by malicious parties.

These limitations suggest the use of more automated techniques for anomaly detector creation.

MACHINE LEARNING AS A KEY ENABLER

Given the typical platform, which contains multiple busses, each with multiple communicating 1553 devices sending multiple messages between one another, how can we equip detection systems with the ability to detect attacks occurring at the application layer? In short, a Cyber Warning Receiver must be programmed or trained to recognize how a system should behave under normal operating conditions, and how this behavior would manifest in the various observable measurements described above.

Advances in machine learning provide this capability. Machine learning also innately addresses the three challenges identified at the end of the previous section. Powerful parameter estimation and model structure detection techniques from machine learning are beneficial for system identification [10]. These capabilities help address the breadth of anomaly detection instances required to form a robust monitoring solution. Multiple examples of using observations to establish normal behavior models for complex systems exist [11]. Activity outside that expected by the normal behavior models is thus anomalous and becomes a data point for cyberattack investigation.

Modern machine learning approaches incorporate feature engineering and credit assignment as key elements. Deep machine learning techniques, for example, combine input observations (e.g., values in each 1553 message data field) into more abstract aggregate features that, while no longer representing actual physical measurements, provide an excellent basis for making decisions (i.e., normal behavior or not) [12]. Machine learning automatically selects which learned features contribute to making such decisions and which are essentially irrelevant—they assign credit to the various features. Over and above increasing the predictive power of the learned normalcy models, these characteristics of appropriate machine learning approaches obviate the challenge of identifying the most important data fields within the 1553 application layer. This is a huge benefit over the alternative of manual specification of data fields and their relative importance. Manual specification is cumbersome, especially considering that application layer message definitions may not exist in one place, but may be scattered across multiple disparate interface description documents, each utilizing different formats which makes them poorly suited to automated parsing.

Machine learning enables reasoning over much larger volumes of data than would be possible for human experts alone. Anomaly detectors increase the visible range of subtle

interactions and mutual patterns of behavior exhibited by disparate elements on the 1553 bus. These patterns may seem innocuous to cyber defense experts trying to envision attack vectors. However, these are exactly the oversights that inevitably get exploited. Finding instances of such subtle relationships has enhanced situational awareness in other domains [13]. Interestingly, insight into such patterns may also prove advantageous in system evaluation and trouble-shooting when non-attack anomalies surface.

Leveraging observable side effects that are agnostic to the specific attack implementation details enables detection of attacks that haven't before been observed in the wild, or preconceived by defenders.

By addressing the three challenges outlined above for reasoning about platform security using deep inspection of data at the application layer, machine learning is a key enabler for cyber situational awareness. Use of machine learning is not exclusive to the application layer, however, and is useful at the lower protocol layers as well. For example, machine learning algorithms can learn the normal message schedule for the platform as a function of the different operating modes, and/or establish normal electrical signal levels at the physical layer. Moreover, these adaptive algorithms can help eliminate the need for tuning and tailoring of detection systems for each instance of the protected platform. Instead, they enable deployment of solutions applicable across an entire platform fleet.

TRAINING FOR CONTINUED SUCCESS

With machine learning comes a need for algorithm training, the process by which machine learning algorithms ingest relevant data, extract features, and build their representations of expected behavior. For a practical defense system, this training should not impose intensive requirements for data collection. Suitable machine learning algorithms operate initially with bus data recorded during field trials and qualification testing and improve their performance upon acquisition of additional data.

One avenue for the collection of additional training data involves incorporation into the mission cycle for a given platform. Bus-recordings collected post-mission would support incremental updates to training sets and learned behavior models. Distributing new models across platform instances at regular intervals enables all protected platforms to benefit continuously from learning over collective data. With more data and collective knowledge, the performance of these machine learning based systems would continue to improve, providing a defense system that evolves with new threats, and adapts to defeat them.

MEASURING MALICE

Not every anomaly means the platform is under attack. Systems are regularly entering and exiting new states and scenarios and experiencing abnormal conditions resulting from a range of incidental activities or failure modes. The key distinction between system glitches and cyberattacks are the correlations that exist between observations, and the story they tell.

Any single cyberattack step would generate a set of measurable side effects and artifacts. Multiple steps in sequence begin to form a picture of the current attacker presence and their objectives in an attack.

A data fusion system is the key element required to put these pieces together. Data fusion formulates the best possible estimate of the underlying system state based on observations, then determines the likelihood that anomalies are caused by an underlying failure, engagement in a scenario or operating mode not previously characterized, or a cyberattack.

BUILDING A COMPLETE PICTURE

A final consideration in defining a Cyber Warning Receiver capability is the question of appropriate output format. The output should never distract a pilot or other key mission personnel unless the findings suggest an imminent survivability threat. Coordinated cyber and kinetic attacks in a combat situation would need to be prioritized to ensure a manageable feed of critical information to the operator.

There is still work to be done to establish the exact manner in which a platform and its occupants should respond in the face of a cyber threat. To follow a general model, this would mean informing or alerting operators given the high probability of compromise for a mission critical system, or if the attack trajectory suggests movement in that direction. Providing too much information, or generating excessive nuisance false alarms might be cause for an operator to disable a system, eliminating the protection and defeating the purpose.

The key distinction between system glitches and cyberattacks are the correlations that exist between observations, and the story they tell.

Another key feature of a Cyber Warning Receiver is the operator interface, which allows operators to explore underlying system security state, and examine the evidence supporting those assumptions. Such data could be analyzed outside of critical moments to enable early detection of malicious actors, or activities relating to the initial establishment of cyber presence on the platform.

Finally, a Cyber Warning Receiver can provide the capability to perform post-mission forensic analysis of anomalous data, in order to provide better threat insights and

preparedness for future engagements. This is enabled by capturing and recording the raw data that is deemed anomalous.

ACTIVE DEFENSE

A major decision to be made with respect to Cyber Warning Receiver technology is the location and configuration of the unit within the system. Two possibilities exist: active or passive.

Cyber warning capabilities form a key addition to the suite of platform survivability equipment, providing visibility into the cyber domain and keeping the warfighter safe in the face of this emerging advanced threat.

The first possibility is to configure a cyber-warning receiver as a passive device, monitoring the system for malicious activity and alerting operators of anything suspicious, but never actively interacting with the network. In this case, the device would need to be positioned within a system to enable monitoring of all applicable buses. This is analogous to the Intrusion Detection System concept from traditional IT security. This option provides a degree of safety from a regression test standpoint, and the likelihood of any performance impact of a Cyber Warning Receiver on critical mission activities is minimized.

Alternatively, a Cyber Warning Receiver could be positioned in line with critical 1553 bus subsystems, prepared to take rapid and decisive action to stop cyber-attacks in their tracks. Given that cyber-attacks can happen in the blink of an eye, active defense may in some cases be the only reasonable way to stop an attack from occurring. The risk with an inline device is that it could be tricked by attackers into providing an inappropriate response, in effect becoming a part of the attack itself. Design precautions would be necessary to ensure that attack suppression actions delivered by an inline Cyber Warning Receiver could not create consequences beyond what the original attack would have achieved by itself.

Given its role, and especially when considered as part of an active defense configuration, a Cyber Warning Receiver as envisioned might itself become an attractive target for adversaries. As the core of cyber security operations on a platform, attackers may make it a priority to disable or interfere with this system to enable their other objectives. As such, any Cyber Warning Receiver would have to be built with the utmost secure design in mind. This could include applying provable security approaches, or leveraging security hardened hardware and software through an active security development lifecycle that includes regular software patching.

CONCLUSION

Modern weapons platforms continue to reach new heights of interconnectivity and software-defined automation. With these enhancements comes the need to address the increasing cyber security risks. Evidence from the commercial and industrial sectors suggests that many of the access vectors and attack methods observed there also apply to DoD platforms, with consequences that are potentially much more severe. Despite this reality, many modern weapons system platforms currently operate without any means of providing detailed situational awareness into their cyber security state.

Platform stakeholders should consider a two-pronged approach to improving platform cyber security posture. This approach begins with implementing survivability equipment that can monitor platform networks for malicious activity. Network monitoring enables near-term capability to detect or prevent cyber-attacks that are a very real threat today. The second facet of this approach involves making ongoing security improvements to individual subsystems, which will help reduce the overall platform attack surface over time.

The MIL-STD-1553 bus is identified as a prime location for observing cyberattacks in progress. This bus is pervasive across both modern and legacy defense platforms and forms the backbone for an exchange of commands, status, and data between operators and the critical subsystems essential to the function of a platform. Cyber Warning Receiver technology can monitor this bus for a range of malicious activities and attack types. This includes attacks that are being carried out to exploit the 1553 bus itself and also attacks that cause deviation from established system behavior norms for data traversing this bus.

Through continuing research, we have characterized a wide range of 1553 network-based attacks and established a corresponding set of observables. A Cyber Warning Receiver measures these observables over time and identifies anomalous or malicious activity. It implements detectors from two categories: explicit detection rules defined by subject matter experts, and system behavior models derived using machine learning. Use of explicit detection rules enables monitoring of the 1553 physical and data link layers for anomalous activity that violates the 1553 standard or does not agree with basic attributes of the known system configuration. The use of learned system behaviors enables deep inspection of messages traversing the 1553 interface to verify they are operating on schedule, that the expected correlations exist between various data fields, and that data ranges and rates of change are within their expected values.

When a cyberattack occurs, the observations and anomalies that result are collected and examined using a data fusion process. This process estimates the underlying security state of the platform and tracks attacker actions. When critical systems are involved, or a survivability risk is identified, a Cyber Warning Receiver can alert operators. Cyber warning capabilities form a key addition to the suite of platform survivability equipment, providing visibility into the cyber domain and keeping the warfighter safe in the face of this emerging advanced threat.

NOTES

1. Doug Gross, 2013, Foul-mouthed hacker hijacks baby's monitor, August 14, http://www.cnn.com/2013/08/14/tech/web/hacked-baby-monitor, (accessed April 25, 2017).

2. Dan Goodin, 2017, Smart TV hack embeds attack code into broadcast signal—no access required. March 31, https://arstechnica.com/security/2017/03/smart-tv-hack-embeds-attack-code-into-broadcast-signal-no-access-required/(accessed April 25, 2017).

3. Charlie Miller and Chris Valasek, 2015, "Remote exploitation of an unaltered passenger vehicle." Black Hat USA 2015.

4. Nicolas Falliere, Liam O Murchu, and Eric Chien, 2011, W32.Stuxnet Dossier Version 1.4. Malware Analysis, Symantec.

5. Bill Gertz, 2016. The military is concerned that Chinese computer products could pose a cyber security threat, October 24, http://www.businessinsider.com/pentagon-chinese-computer-products-cyber-security-threat-2016-10, (accessed April 25, 2017).

6. Per Larsen, Stefan Brunthaler, and Michael Franz, 2014, "Security through diversity: Are we there yet?" IEEE Security & Privacy 12, no. 2 28-35.

7. Paul Clements and John K. Bergey, 2005, The U.S. Army's Common Avionics Architecture System (CAAS) Product Line: A Case Study. Technical Report, Pittsburgh, PA: Carnegie Mellon Software Engineering Instutute.

8. Wikipedia Contributors. n.d., "MIL-STD-1553," Wikipedia, The Free Encyclopedia, https://en.wikipedia.org/w/index.php?title=MIL-STD-1553&oldid=776707274, (accessed April 25, 2017).

9. Peter Huitsing, Rodrigo Chandia, Mauricio Papa, and Sujeet Shenoi, "Attack taxonomies for the Modbus protocols.," International Journal of Critical Infrastructure Protection 1 (2008): 37-44, doi:10.1016/j.ijcip.2008.08.003.

10. Lennart Ljung, Hakan Hjalmarsson and Henrik Ohlsson, 2011, Four encounters with system identification. European Journal of Control, 5-6, 449-471; Pillonetto, Gianluigi, 2016, The interplay between system identification and machine learning, arXiv:1612.09158v1.

11. B.J. Rhodes, N.A. Bomberger, M. Zandipour, D. Garagic, L.H. Stolzar, J.R. Dankert, A.M. Waxman, & M. Seibert, 2009, Automated activity pattern learning and monitoring provide decision support to supervisors of busy environments, Intelligent Decision Technologies, 3, 59–74; B.J. Rhodes, N.A. Bomberger, M. Zandipour, L.H. Stolzar, D. Garagic, J.R. Dankert, & M. Seibert, 2009, Anomaly detection & behavior prediction: Higher-level fusion based on computational neuroscientific principles. In N. Milisavljević (Ed.), Sensor and Data Fusion, 323–336, Croatia: In-Teh.

12. Y. Bengio, A. Courville, and P. Vincent, 2013, Representation learning: A review and new perspectives. IEEE Trans, PAMI (Special issue: Learning Deep Architectures), 35, 1798–1828, doi:10.1109/tpami.2013.50; Hinton, G. E.; Salakhutdinov, R. R. (2006), Reducing the dimensionality of data with neural networks. Science, 313 (5786), 504–507, doi:Y. 10.1126/science, 1127647.

13. M. Zandipour, B.J. Rhodes, and N.A. Bomberger, 2009, Probabilistic prediction of vessel motion multiple spatial scales of maritime situation awareness, In Proceedings of the 10th International Conference on Information Fusion, Cologne, Germany, June 30–July 3, 2008; J.R. Dankert, M. Zandipour, N. Pioch, B. Biehl, R. Bussjager, C.Y. Chong, M. Schneider, M. Seibert, S. Zheng, & B.J. Rhodes, (2010), MIFFSSA: A multi-INT fusion and discovery approach for Counter-Space Situational Awareness, In Proceedings of 2010 Space Control Conference (SCC), Lexington, MA, USA, May 1–3, 2010.