2019 NDIA GROUND VEHICLE SYSTEMS ENGINEERING AND TECHNOLOGY SYMPOSIUM
GROUND SYSTEMS CYBER ENGINEERING TECHNICAL SESSION
AUGUST 13-15, 2019 - NOVI, MICHIGAN

Holistically Increasing Cyber Resilience of Ground Vehicles

David K. Woolrich Jr

Cyber Technology, FAST Labs, BAE Systems, Merrimack, NH

ABSTRACT

This paper explores a holistic approach to increasing the cyber resiliency of Army and USMC ground vehicles. Today's current approach to securing weapon systems focuses on complying with the Risk Management Framework and applying required security controls to obtain government authority to operate (ATO). This method of securing our weapon systems is better than nothing, but runs the risk of giving us a false sense of security.

Citation: D. Woolrich, "Holistically Increasing Cyber Resilience of Ground Vehicles", In *Proceedings of the Ground Vehicle Systems Engineering and Technology Symposium* (GVSETS), NDIA, Novi, MI, Aug. 13-15, 2019.

1. INTRODUCTION

The battlefield from here on out will be contested in multiple domains: physical, the electromagnetic spectrum, and digital. Events in Syria, Ukraine and other parts of the world show that today's battlefield is now cyber contested. Through increased networking and computing capabilities, today's weapon systems are more lethal and effective than ever before, but they simultaneously grow more susceptible to cyber-attack. As this susceptibility grows, so too, does our adversaries interest in exploiting weaknesses of these systems. Researchers have shown commercial ground vehicles are susceptible to cyber-attack and weapon systems present an attractive target for U.S. adversaries. The 2018 GAO report on Weapon System Cyber Security [1] concluded that no system is impervious to cyber-attack and it is not "if" an attacker will gain access and degrade mission capabilities – it is "when."

Government and industry must develop holistic cyber *resilience-in-depth* capabilities that are affordable and scalable. How can the United States and our allies increase resilience of our legacy fleet as well as design our future platforms to be inherently more resilient by design? Further, no weapon system will ever be immune to cyber-attack. Rather than trying to make a weapon system immune to cyber-attacks, how can we remove the "low-hanging fruit," steepen the difficulty curve, detect cyber-attacks when they happen, deploy appropriate counter-measures, and if possible, recover to a known good state?

In order to secure legacy and future weapon systems against cyber-attacks, the DoD requires a holistic approach that addresses policy; informs tactics, techniques, and procedures; reduces the attack surface; applies lessons learned from enterprise; accurately detects cyber-attacks against the system, and provides an active mitigation solution to ensure the weapon system can continue

Not export controlled per ES-FL-052119-0115;
Approved for public release; Unlimited Distribution

to operate in a cyber-contested environment. Additionally, any solutions developed must be maintainable, able to move at the speed of the adversary, and scalable and affordable to meet the size of the Army and USMC ground vehicle inventory.

2. Reducing the attack surface

The first step to increasing the cyber resiliency of weapon systems is to reduce the attack surface. The attack surface represents each of the ways an adversary can gain access to a weapon system. In order to reduce the attack surface, accurate characterization is imperative. How can the Department of Defense (DoD) reduce the attack surface of their ground vehicles? This is accomplished by focusing on securing the supply chain, training the government and industry work force in the development of best practices to enable "*resilient by design"* development, leveraging commercial and enterprise best practices, and obtaining the proper accreditations.

2.1. Characterizing the attack surface

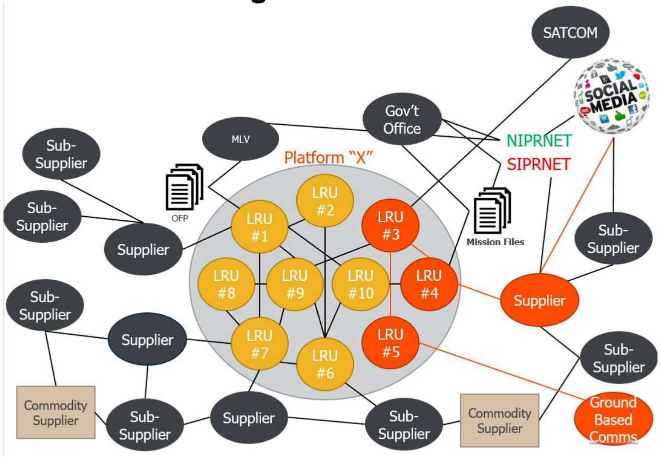

Figure 1- a representative system's highly complex interconnectivity

The attack surface of a ground vehicle is defined as the set of all distinct points where an attacker can enter or extract data from the vehicle. It is important to develop an accurate picture of potential vectors an adversary can use to negatively affect a system. With external vectors defined, we must understand how each system interacts with each other. Modern systems rely on numerous subsystems that are inter-connected via networks, avionics buses, or other means. In mapping the internal connections of subsystems we are able to discover subsystems that lack direct external connectivity but are actually networked to the outside world. Sophisticated adversaries can target these external facing systems, establish a foothold on the system and then move laterally to affect connected subsystems.

2.2. Securing the supply chain

A single ground vehicle uses hundreds of different suppliers such as the original equipment manufacturer (OEM), line replaceable unit (LRU) providers, individual component providers, and test equipment providers. At any point, this supply chain can be compromised and the entire chain's integrity jeopardized.

The commercial world provides an example of a supply chain compromise resulting in damage. The 2013 cyber-attack against Target department stores resulted in the theft of 40 million people's information and cost Target over $200M in legal fees and settlements. When criminals attacked Target, they did not initially attack Target directly. Rather, their initial attack vector was through a HVAC company who had access to Target's supplier network. This HVAC company was targeted via phishing emails and legitimate login credentials were stolen. With the login credentials, attackers were able to gain access to Target's system, move laterally through Target's internal systems, and discover a vulnerability which enabled them to install malware on their points of sale stations siphoning the credit card information of 40M+ consumers. [2]

Despite rigorous DoD cyber-hardening initiatives and security education requirements, a single compromised component supplier can create a cascade of effects felt by the warfighter. For example, consider the following hypothetical scenario. A supplier of ruggedized single board computers (SBC) is targeted and successfully attacked via phishing email. The attacker leverages their position within the supplier to compromise firmware that is loaded onto every single SBC. These SBCs are shipped to multiple 2nd tier vendors who develop individual LRUs for a DoD combat system. These systems are now installed onto a ground vehicle with the compromised firmware.

2.3. Resilient by design development

As systems are designed or upgraded, small, smart changes can provide large gains in providing cyber resilience. If developers within the DoD and industry are taught secure coding and best practices, they can increase the cyber resilience of their components with minimal additional cost and lower future risk of expensive patching. Establishing bounds checking to mitigate buffer overflows is one example. Using cryptographically secure hashes rather than simple checksums is another.

Beyond faults, holistic resilient design examines the system's security assumptions. For example, current ground platforms often implicitly trust the internal subsystems. Incorporating authentication requirements and establishing formal trust relationships would increase system resilience by preventing unauthorized subsystem access and exposing authorized but malicious components. These same cryptographic functions can be applied to lockdown the test and update ports on the system, which prevents the introduction of malware from these highly privileged, but poorly protected ports.

2.4. Leveraging best practices

The DoD can adopt additional best practices from the commercial industry. Efficient patching, while mundane, is critical to providing cyber resilience. While 0-day exploits traditionally dominate news headlines and DEFCON presentations, it is often known vulnerabilities that are exploited. The May 2019 cyber-attack against the city of Baltimore highlights this. Reports indicate the ETERNALBLUE exploit was used to attack and ultimately encrypt Baltimore's municipal computers via ransomware. A patch for this vulnerability was released two years prior in 2017. Following patching best practices would have saved Baltimore time, money, and credibility.

2.5. Proper accreditations

Having an accreditation is the minimum threshold for business. Following the Risk Management Framework (RMF) along with obtaining and maintaining system accreditation is a start and applying RMF controls provides an initial level of protection. However, RMF alone is incapable of providing the cyber *resilience-in-depth* needed for platforms operating within a cyber-contested battlefield. This is because RMF (and related controls outlined in NIST 800-53) were developed to protect traditional IT infrastructure; not weapon systems and their embedded subsystems. NIST 800-53 [3] security controls fail to address specific attack vectors nor account for the dynamic, real-time threats of operational environments. To address these challenges, BAE Systems is developing a RMF overlay for embedded systems. BAE Systems' overlay seeks to account for the embedded systems and incorporates offensive architecting into the accreditation process.

3. Increasing cyber situational awareness

In commercial and DoD environments, cyber-situational awareness is a necessity. In the commercial world, an advanced adversary averages 146 days [4] before being detected on a victim's network. Worse, on a weapon system, an adversary may entirely escape discovery without advanced situational awareness capabilities. A recent GAO report concluded that testers were "able to take control of these systems and largely operate undetected." [1] Without the ability to detect a cyber-attack, no informed defensive actions can be taken.

Increasing cyber situational awareness is more than a single vehicle knowing it is under attack – it involves the ability for a battlefield commander to gain an understanding of the cyber domain as much as the physical domain. By providing a commander with situational awareness across a battalion or battlefield, they can potentially detect trends against their systems and make informed command decisions.

To respond to a cyber-attack against a vehicle, we must first be able to accurately detect the cyber-attack and properly attribute the attack source and victim system.

3.1. Detecting cyber-attacks

Cyber attack detection fundamentally relies on the ability to observe a resulting change in system behavior. Examples can be in the form of malformed or mistimed messages, corrupted messages, changes in expected behavior, or changes in physical characteristics of the system.

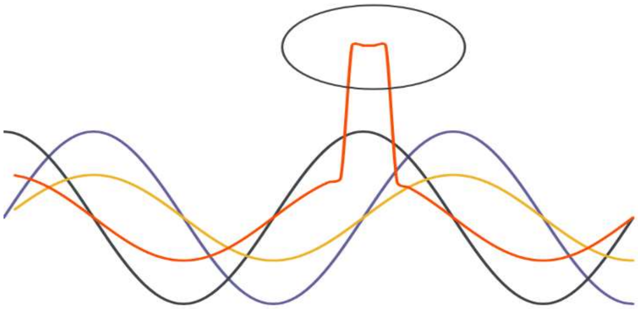

Figure 2- A deviation from normal can indicate a potential cyber attack

Unlike traditional intrusion detection systems, a ground vehicle system would need to operate on non-traditional bus networks like MIL-STD-1553 and Control Area Network (CAN).

To detect these changes, we can install systems that are capable of monitoring the behaviors of the system and comparing them against a previously determined model of "good" behavior. This monitoring can be bus-traffic message monitoring, enabling the monitoring and detection of cyber-attacks for systems connected to the bus.

We can also monitor the physical characteristics of the systems to detect correlated changes in secondary behaviors indicative of a potential system compromise. Physical characteristics can include thermal load, the physical characteristics of the signals (i.e. voltage, amplitude, and phase), changes in the electromagnetic spectrum surrounding the systems, and internal system resources such as memory usage. A complete monitoring system would be capable of fusing all of these data sources into a single coherent platform picture, detecting cyber-attacks with a high degree of confidence.

3.2. Determining impact

Anomalies alone do not indicate a cyber-attack against a weapon system, as anomalies occur normally throughout the operation of the vehicle. An effective detection system must be able to determine that 1) an anomaly has occurred, 2) the anomaly is a cyber-attack, and 3) impacts the mission.

Impact to the mission should be determined by working closely with platform and mission SMEs. All three should occur prior to alerting an operator. Anomalies that occur and are deemed to be a cyber-impact, but have no mission impact should be logged for future analysis.

3.3. When to alert the operator

Based on interviews and discussions with ground and air platform operators, system operators are not looking for more information to flood their already over tasked mental load. Additionally, if there's no impact to the mission or a critical system, it may not be worth displaying to an operator. Here, a system should log all relevant information and store it for post-mission forensic

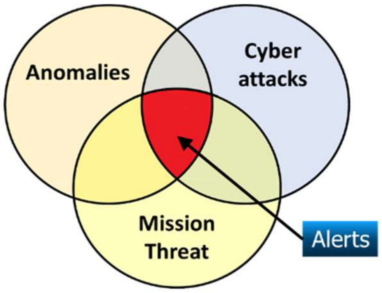

Figure 3- By combining multiple indicators, an operator is only notified when there is a mission impact

analysis. Any alert displayed to an operator must be high in confidence and must provide relevant information about the attack (without providing too much distracting detail), and any recommended countermeasure.

3.4. Proper attribution

Accurate attribution of a cyber-attack – that is the identification of the target affected system and the compromised attacking system – helps drive response and recovery options. If an operator or response mechanism can successfully locate the rogue system and isolate it, they may be able to reduce the impact of the system. If a critical system is compromised, an operator can take steps to minimize or mitigate mission degradation.

4. Responding and recovering from cyber-attack

Once a cyber-attack is detected, what should be done? Through a combination of existing technology, new technology, and modifications to operator and system behavior, we can increase the cyber resilience of ground vehicles. We next discuss how after successfully detecting a cyber-attack we can use these existing modes to mitigate the effects of a cyber-attack, enabling cyber resilience. Lastly, we will explore active response methods and deployment CONOPS where, through in-line hardware, we can identify malicious behaviors and defend core system functionality.

4.1. Developing new TTPs

Tactics, Techniques, and Procedures (TTP) should be adapted and updated to account for cyber-attacks experienced on the battlefield. New TTPs should be developed to address situations which current TTPs are insufficient. Current TTPs do not account for what actions an operator must take when under cyber-attack. Working with operators, researchers, and labs, new TTPs can be established to increase the resilience of ground combat vehicles.

4.2. Leveraging existing modes

Existing systems and subsystems often have fault tolerant or degraded modes. These existing modes represent a simple, pre-existing mitigation to some effects of a cyber-attack.

For example, consider an attacker that has compromised a GPS unit and is causing bad data to be distributed on the network. If detected, an operator can failover to an alternative position source (such as inertial navigation) and ignore the external GPS data: effectively mitigating the GPS attack. While this is only a localized solution, it provides a method of defense that already exists within the system in a feature intended for fault tolerance and redundancy.

Modifying existing TTPs to account for cyber-attacks would enable operators to mitigate cyber effects with the systems they have today.

4.3. Active measures

Modifying existing TTPs, resilient by design development, and detection of cyber-attacks will not be enough to defeat advanced adversaries. Active mitigation solutions must also be developed and deployed. This section explores a few areas of interest but is not all-encompassing of potential solutions.

Data traveling over CAN and 1553 is currently observable by any system on the bus. With the interface control document, it is possible to understand the contents of the data also. This means that with physical access, an adversary could send erroneous messages or intercept messages on the bus and in real time modify them causing unexpected behavior. Encrypting bus traffic can provide dynamic protection of bus traffic, preventing interception and manipulation of bus traffic providing a level of resilience. Encrypting traffic on the bus would defeat any man-in-the-middle or malicious hardware implant types of attacks. Encryption is only one part of resilience-in-depth. If the host system were compromised, it would still be vulnerable to attack – as the compromised system would have the encryption key.

A constantly evolving and dynamic attack surface dramatically increases the difficulty of an attacker to gain a persistent presence on our combat vehicles. Advanced persistent threats thrive in static and homogenous environments – where every system is the same and remains the same throughout the life of the system. Introducing a heterogeneous environment creates an environment where a single exploit cannot impact an entire fleet. It also reduces the amount of time a threat can survive on a system, as the instance will be cycled and thrown away, requiring an adversary to reestablish their footprint on the system. Amazon Web Services currently employs a similar approach, where they cycle and discard cloud instances. Their approach is transparent to users, who do not notice this occurring, but increases resilience of their system by increasing the difficulty level needed for a threat to remain persistent. Introducing this to DoD combat vehicles presents challenges and needs to be researched more, but it would create an environment that would significantly increase the cyber resilience of our combat vehicles. A dynamic attack surface would create an environment where a single exploit cannot disable all Bradley fighting vehicles (as an example) worldwide.

4.4. Filter, firewalls

Through hardware shims and modified connectors, the DoD could install physical filters and firewalls on each line replaceable unit connection interface to observe and defend all bus traffic. As cyber threats or attacks are detected, these filters and firewalls could automatically drop and/or disable a malicious traffic from the system: interrupting the attack.

However, there are risks associated with LRU-based filter and firewall use. How does this impact Size Weight And Power (SWAP)? How does it impact the qualification of the platform? How do we prevent malicious misuse of such a response to amplify the effects of a cyber-attack? How do we ensure critical systems are not blocked? Research and development with government, industry, and academia is required.

4.5. Deployment CONOPS

When applying new features to a combat system, the deployment strategy presents are a very real challenge. Introducing new hardware and software has physical SWAP constraints and also procedural impact. The qualification and requalification of systems can be an onerous process. Does every single system need to be protected? Or can we take the same approach as body armor where we protect the vital organs and manage the risk for exposed extremities? If we can identify critical systems and subsystems, we can provide enough protection to enable the system to operate in a "safemode" and get back to base if needed. This could also reduce the amount of system qualifications needed.

5. Challenges faced

Any solutions adopted by the DoD need to be scalable to account for the sheer amount of platforms within each of the services. Solutions cannot be boutique, expensive ones specifically tailored for each individual system. Detection systems must be able to account for changes in system environment, LRUs that are interchanged, damage while on the battle field, and other "abnormal normal situations."

An "abnormal normal situation" is a situation that during routine operations does not occur, but external events trigger abnormal conditions that are now normal and critical to the system. An example of this would be a vehicle being engaged in combat, a missile being fired at the system, and the missile detection system declaring a threat is incoming. In 99.99% of the vehicle's operating life cycle a missile is not being fired at the vehicle, but in that moment, while abnormal, it is very real. Any cyber-attack detection and mitigation solutions must be able to account for this and not create a situation where it claims this is an abnormal event and a cyber-attack.

Any solutions introduced must account for the real-time processing requirements of the embedded systems. Encrypting and keying all systems can potentially introduce latency to the system and needs to be explored and resolved prior to implementation.

Patch management, while important, is difficult. It is sufficiently challenging within traditional IT equipment, let alone world-wide deployed combat systems. Following current DoD testing and qualification guidelines, a single patch could take more than 12 months to be properly vetted, much longer than the recommended quarterly patching cycle. Research is required on how to patch a system and provably show no adverse effects to the immediate system as well as second and third order effects downstream.

6. The BAE Systems approach

BAE Systems has developed a platform cyber-attack framework and platform cyber defense framework to document the current landscape and challenges faced. We continue to work internally to develop an RMF overlay for embedded systems. In partnership with the DoD, academia and other industry partners we continue to develop tools that augment the manual characterization of the attack surface, to the ability to automatically reverse binaries and discovery hidden vulnerabilities, to the ability to detect and defeat cyber-attacks in real time to increase vehicle cyber resilience.

BAE Systems has developed and continues to mature cyber-attack detection capabilities for weapon systems. Our Cyber Warning Receiver (CWR) has built on a legacy of anomaly detection capabilities and expanded that capability to detect certain classes of cyber anomalies on CAN and MIL-STD-1553. CWR is the first step in providing cyber situational awareness for military vehicles. CWR is a passive system, capable of monitoring traffic on vehicle buses for anomalous behavior and notifying an operator of the event. Our CAN variant is currently TRL-6 and is expected to be TRL-8 by the end of 2019. CWR provides a logging capability that allows for post-mission cyber forensic analysis, a capability currently lacking on DoD weapon systems. The solution is available as either as a software load to an existing system or as a standalone hardware capability.

7. Conclusions

Hope is not lost. Working with platform primes, sub-system developers, users, requirements generators, and acquisition professionals, we continue to explore response mechanisms that are practical, implementable, effective, and affordable in increasing the cyber *resilience-in-depth* of DoD vehicles.

While there are challenges ahead of us, they are not insurmountable. Weapon systems and vehicles today are susceptible to cyber-attack, but the government and industry are working to rise to the challenge. There is no single “magic bullet” or fix that will solve the DoD’s vehicle cyber security challenges. We must shift our focus away from defense-in-depth to holistic *Resilience-in-Depth.* Our adversaries are capable and will find a way to impact out systems. Expanding beyond RMF and shifting from defense to resilience accounts for this, acknowledging an adversary may be successful, but providing a way to continue to fight and win in a cyber-contested environment.

8. References

[1] United States Government Accountability Office, "WEAPON SYSTEMS CYBERSECURITY - DOD Just Beginning to Grapple with Scale of Vulnerabilities," GAO, 2018.

[2] T. Radichel, "Case Study: Critical Controls that Could Have Prevented Target Breach," the SANS Institute, 2014.

[3] National Information Technology Laboratory, "NIST Special Publication (SP) 800-53," 2013. [Online]. Available: https://csrc.nist.gov/publications/detail/sp/800-53/rev-4/final.

[4] FIREEYE, "Cyber Threats: A perfect storm about to hit Europe?," 2017.

[5] Director, Operational Test and Evaluation, "FY 2018 Annual Report," 2018.