CyberSafety: On the Interactions between CyberSecurity and the Software Engineering of Safety-Critical Systems

C.W. Johnson,

Department of Computing Science, University of Glasgow, Glasgow, Scotland, UK, G12 8RZ.
johnson@dcs.gla.ac.uk, http://www.dcs.gla.ac.uk/~johnson

Keywords: Cyber Security, Cyber Safety, incident reporting.

Abstract

A range of common software components are gradually being integrated into the infrastructures that support safety-critical systems. These include network management tools, operating systems-especially Linux, Voice Over IP (VOIP) communications technologies, Satellite Based Augmentation Systems for navigation/timing data etc. The increasing use of these common components creates concerns that bugs might affect multiple systems across many different safety-related industries. It also raises significant security concerns. Malware has been detected in power distribution, healthcare, military and transportation infrastructures. Most previous attacks do not seem to have deliberately targeted critical applications. However, there is no room for complacency in the face of increasing vulnerability to cyber attacks on safety-related systems. This paper illustrates the threat to Air Traffic Management infrastructures and goes on to present a roadmap to increase our resilience to future CyberSafety attacks. Some components of this proposal are familiar concepts from Security Management Systems (SecMS); including a focus on incident reporting and the need for improved risk assessment tools. Other components of the roadmap focus on structural and organizational problems that have limited the effectiveness of existing SecMS; in particular there is a need to raise awareness amongst regulators and senior management who often lack the technical and engineering background to understand the nature of the threats to safety-critical software.

1 Introduction

In the past, the specialized nature of infrastructure engineering has limited the failure modes that could cross multiple systems boundaries. The future looks very different. The increasing integration of critical infrastructures creates new opportunities, for instance through the development of Smart Grids for the generation and distribution of electricity or through the use of EGNOS satellite based timing and location services for railway signaling (Pederson et al, 2006). However, this integration creates new vulnerabilities. Safety-critical applications increasingly rely on a small number of common operating systems and network protocols. Very similar algorithms are being used by the same suppliers across both primary and secondary systems. For example, Voice Over IP (VOIP) communication technologies are being used for backup and principle systems in areas ranging from Air Traffic Management to emergency response.

The increasing use of common software components across many different safety-critical industries creates concerns that the consequences of any bugs might extend across multiple applications. For instance, previous work has shown that design flaws have been carried between the GPS applications that were initial developed for aviation and then were subsequently integrated into maritime bridge information systems (Johnson, Shea and Holloway, 2008). There are also significant security concerns where safety-critical systems rely on Commercial Off The Shelf (COTS) operating systems and network infrastructures. This brings significant savings to the developers and operators of safety-critical systems. However, they also attract a host of 'mass market' viruses. For instance, Linux variants are increasingly being used in Air Traffic Management. There have been several recent cases where engineering teams have discovered malware affecting these systems; typically introduced by contractors using infected USB sticks. These incidents illustrate the problems that Air Navigation Service Providers (ANSPs) face in implementing their existing security policies. The Linux attacks have not yet had significant safety implications; Air Traffic Control Officers (ATCOs) maintain sufficient situation awareness to continue service provision even when primary systems have been compromised. However, with increasing traffic loads and greater systems integration planned by the US NextGen and European SESAR programmes, there is no scope for complacency over the consequences of future security threats (Johnson, 2011).

2 Assessing Security Threats to ATM Safety

The increasing reliance on common software components leaves us unprepared for the consequences of coordinated attacks on safety-critical infrastructures. These vulnerabilities are compounded by the problems that arise when determining who is responsible for meeting the costs associated with national resilience. Government security agencies rely on support from industry. However, few

companies can afford to meet the costs of design diversity and redundancy that provide higher levels of assurance. Further problems arise when commercial organizations fail to monitor the effectiveness of their security management systems. Many policies and procedures only exist on paper and are never used in daily operations. In consequence, many safety-related applications remain highly vulnerable to a wide range of cyber security and cyber defense threats.

These arguments can be illustrated by the General Accounting Office (1998) review of CyberSecurity in US Air Traffic Management. They found that the FAA was "ineffective in all critical areas including -operational systems information security, future systems modernization security, and management structure and policy implementation". They further concluded that the "FAA is similarly ineffective in managing systems security for its operational systems and is in violation of its own policy". They had "performed the necessary analysis to determine system threats, vulnerabilities, and safeguards for only 3 of 90 operational ATC computer systems, or less than 4 percent"; only one of the nine telecommunications networks had been analysed for security vulnerabilities. The GAO also found that the FAA "does not consistently include well formulated security requirements in specifications for all new ATC modernization systems, as required by FAA policy".

Many of the same concerns were again raised by the US Department of Transport (DoT, 2009). This identified problems in both corporate information systems and operational infrastructures. In February 2009, an intrusion compromised the social security details of more than 48,000 staff held on FAA servers. In other attacks, the administrators' passwords were obtained for FAA networks in Oklahoma and Alaska. These intrusions focused on web-based information systems but the interconnected nature of FAA operations created significant concerns for the operational networks where surveillance, communications and flight information is processed. It was hard for the systems engineers to guarantee that these attacks could not have any impact on service provision. The Department of Transport report argued that "In our opinion, unless effective action is taken quickly, it is likely to be a matter of when, not if, ATC systems encounter attacks that do serious harm to ATC operations."

The (DoT) report went on to reiterate many of the arguments that have been made in the opening sections of this paper. The introduction of commercial software and Internet Protocol technologies has provided significant cost savings to FAA modernization initiatives. However, they also introduce greater security risk compared to previous generations of proprietary software; "Now, attackers can take advantage of software vulnerabilities in commercial IP products to exploit ATC systems, which is especially worrisome at a time when the Nation is facing increased threats from sophisticated nation-state-sponsored cyber attacks" (DOT, 2009). Their concern was exacerbated by the inadequate response to previous incidents. In 2008, there were more than 870 cyber incident alerts. By the end of the year, 17% (N=150) of these had not been resolved; "including critical incidents in which hackers may have taken over control" of the computational infrastructure for ATM operations.

Europe lags behind the United States in this area. There are no surveys of ATM security practices, which might be compared to those produced by the GAO and DoT. We lack detailed evidence about the extent of CyberSecurity vulnerabilities across member states. As part of the preparation for this paper, the author visited engineering teams in more than a dozen ANSPs across Europe. Every centre reported having experienced problems from malware. In several cases, these intrusions had forced them to rely on secondary communication systems or Flight Data Processing systems.

The vulnerabilities are likely to increase. For example, EUROCONTROL's CASCADE programme has considered a range of security concerns associated with the unauthorized use of Automatic Dependent Surveillance - Broadcast (ADS-B) information. ADS-B relies on aircraft transmitting their identity, position etc, to support ground surveillance by ANSPs. The data can also be used by on-board avionics to improve the situation awareness of other aircraft. However, the increasing use of this technology creates the potential to deliberately introduce false targets into the system or to use aircraft identity and position information for malicious purposes.

There is an urgent need for more information about existing vulnerabilities and future threats. It is, therefore, important for the European Commission to review the CyberSecurity of Air Traffic Management across member states. This should be completed before the SESAR programme for the modernization of ATM infrastructures further increases our reliance on software systems.

3 A Roadmap for CyberSafety Engineering

Previous sections have described the security vulnerabilities that permeate software infrastructures in some safety-critical applications. We have also argued that these vulnerabilities will increase as common software components, including network management tools and operating systems, are used across multiple systems in many different industries. These concerns have motivated the development of a roadmap to increase resilience against cyber-attacks on safety-critical systems (Johnson and Holloway, 2011). Figure 1 (overleaf) provides an overview of the approach advocated in the remainder of the paper.

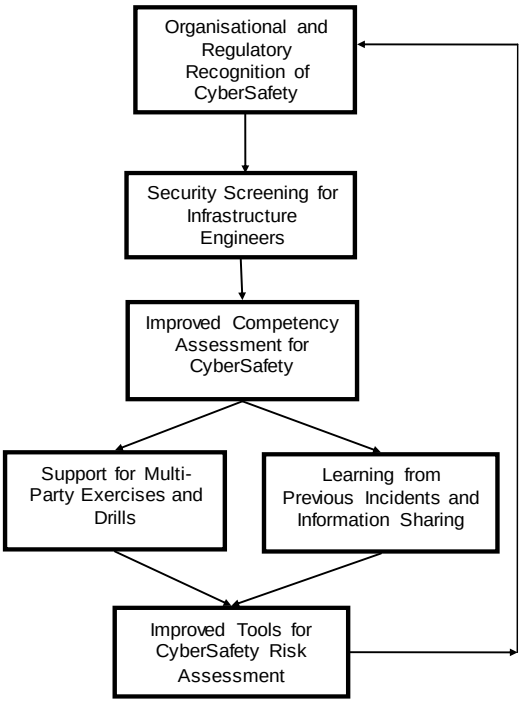

Figure 1: A Roadmap for CyberSafety Engineering

3.1 Addressing Managerial & Regulatory Complacency

In the past, operational staff could intervene if they had concerns about problems in the underlying software infrastructures. Air Traffic Control Officers (ATCOs) retain the right to 'close the skies' or adjust the amount of traffic in response to periodic malfunctions or if there is evidence of an intrusion in networked systems. However, this position cannot be sustained. The next generation of automated systems will stretch the ability of operational staff to directly intervene without decision support tools. It is also becoming increasingly difficult for ATCOs to distinguish malicious behaviours from more benign bugs or from 'normal operation' in complex, integrated systems. This significantly increases the consequences of any breach in the security of underlying software infrastructures by delaying the time before an intrusion is detected.

These problems are exacerbated by a lack of strategic leadership in CyberSecurity. Most regulatory organisations do not understand the security threats that are posed to software architectures. Very few regulators have expertise in software engineering or the integration of complex distributed systems; there are few incentives for leading technical staff to join regulatory bodies where the salaries and career prospects may be less attractive than in many companies. These problems are compounded by the under-representation of engineers at higher levels of management in safety-critical industries. The appointment of operational and financial experts deprives engineering teams of the strategic guidance needed to address CyberSafety concerns. This creates a situation in which many governments across Europe and North America have identified the potential threat and taken action to create specialist agencies, such as the UK's Centre for the Protection of National Infrastructures. However, their warnings about the vulnerabilities of technologies such as the Internet Protocol for safety-critical applications have not been acted upon. These criticisms extend well beyond the field of Air Traffic Management.

3.2 Security Screening for Infrastructure Engineers

Our roadmap includes improvements to the security screening of technical staff across safety-critical industries. Most attention has focused on external cyber threats. The 'insider threat' has been ignored. In consequence, Air Traffic Management Security Management Systems often do not require background checks on the staff employed by sub-contractors. There is minimal vetting for the engineers and technicians who work on critical infrastructures in Europe and North America. These limitations are compounded by a lack of guidance on how to mitigating vulnerabilities to insider attacks within existing security standards, such as ISO 17799 (Theoharidoua et al, 2005).

It is important to stress that these are not hypothetical concerns. NIST's Industrial Control System Security project report how a disaffected worker used their knowledge of the underlying software infrastructures to attack SCADA controlled sewage equipment (Abrams and Weiss, 2008). On at least 46 occasions, he issued radio commands that caused 800,000 litres of raw sewage to contaminate local parks and rivers. When he was arrested, he was found to possess a laptop that ran a version of the sewerage control application. This was connected to a Motorola M120 two-way radio; the same device that was used to send 'legitimate' commands to the equipment. The serial numbers showed that the radio had been ordered by the operating company. He also had a PDS Compact 500 computer control device with an address that mimicked a spoof pumping station enabling him to test out the impact of his commands.

The perpetrator of this attack was initially employed by an IT subcontractor to the sewage company. His everyday work provided him with a good understanding of the underlying software infrastructures. This has strong parallels in Air Traffic Management. Many ANSPs have outsourced IT infrastructure provision. In such cases, it can be particularly difficult for operators to identify and diagnose potential attacks. In this incident, the perpetrator took steps to disguise his actions, interspersing them between less malign system failures. It was difficult to distinguish the impact of the malicious attacks from design flaws and malfunctions.

Most ANSPs have not considered the impact of a similar 'insider' attack on the systems that they operate. The consequences are difficult to exaggerate. For most service providers, the threat of a deliberately introduced bug would be enough to halt service provision with no immediate way to trace whether safety requirements continue to be met across many millions of lines of code (Haley et al, 2008). In Europe, many military air traffic management systems share

the same machine rooms as their civil counterparts. In some states, they share a common network infrastructure.

3.3 Competency Assessment for Cyber Security

The roadmap in Figure 1 also includes the need to improve the competency of engineering teams in dealing with cyber threats. For instance, very few technical staff are aware of recent developments such as Maddalon and Miner's (2004) work on intrusion tolerance for computational systems in Air Traffic Management.

There has been some progress in considering security requirements within the modernisation programmes mentioned in previous sections. For instance, the FAA's System Wide Information Management programme distinguishes between security mechanisms to be implemented at the application layer and those that are implemented at lower levels within a prototype Secure Service Gateway. These plans will have little impact if existing engineers lack the skills to configure and maintain the new security infrastructures. Today, many Linux installations are not protected by anti-virus programs because technical staff remain unaware of the potential threats. Operational pressures have 'forced' staff to cut corners by installing updates using unverified media. Security patches are not always introduced in a timely manner across the many different systems that support ATM services.

The lack of CyberSecurity competency contrasts with the detailed training requirements that have been developed for operational staff, for example in the European Manual of Personnel Licensing - Air Traffic Controllers (EUROCONTROL, June 2004). The site visits that motivated this paper revealed that none of the engineers in the twelve European countries had any formal training in computational forensics. When malware was detected, it was often through chance rather than the result of coherent security monitoring techniques. For instance, previous attacks have been detected as an indirect consequence of efficiency concerns by ANSP engineers trying to understand where additional CPU cycles were being lost.

3.4 Supporting Multi-Party Exercises and Drills

Figure 1 identifies three further requirements in the roadmap for CyberSafety. The first advocates multi-party exercises and drills to assess the potential consequences of a cyber-attack. The Idaho National Laboratory launched a well-publicised attack on a SCADA system to show that this could result in physical damage to electricity infrastructures. Questions have since been raised about the veracity of these tests. Other simulated attacks have identified vulnerabilities in the SmartGrids proposals for energy distribution in Europe and North America. There have also been attempts to rehearse large scale attacks across infrastructures, these are illustrated by the US Cyber Storm and CyberEurope exercises. Such initiatives have had little impact on ANSPs; there has been minimal involvement in previous studies.

Significant advances have been made in helping promote cooperation between operational users – for instance through the use of simulation tools and training in the management of critical events. However, these techniques are not used to prepare engineers for the demands that are placed on them during cyber-attacks. Very few training exercises or drills have been held so that technical staff can practice their response to malware or an intrusion into the underlying systems and networks that support Air Traffic Management. It is, therefore, unsurprising that systems engineers often find it difficult to coordinate their response when a threat is discovered.

It is particularly important to rehearse the response to a potential attack because cyber-threats affect many different stakeholders. Engineering teams must integrate their response with the ATCOs and managers who are responsible for maintaining service provision. Additional help may be required from external experts, for example in computational forensics. In some circumstances, the operational consequences of an attack must be communicated to neighbouring states. It is also important to involve a range of national bodies, including regulators as well as police and investigatory agencies. Very few ANSPs have integrated their plans with the support provided by national Computer Emergency Readiness Teams, including US-CERT.

In Air Traffic Management, the sophistication of software infrastructures has stretched the engineering capabilities of most service providers. In consequence, many states rely on external support for the operation of their network infrastructures. This creates a situation in which ANSPs depend upon the security management practices of a small number of contractors. In an ideal situation this might increase resilience by providing a pool of committed engineers with significant expertise in detecting and mitigating threats. In reality, the quality of external personnel and their ability to implement the principles of security management is variable. Sub-contractors often share the lack of interest that many ANSPs show in Cyber Security until they become victims. In contrast, our roadmap argues for increased investment in multi-party exercises and drills that help develop team resource management skills in response to a range of simulated attacks.

3.5 Sharing Lessons Learned

Site visits to ECAC ANSPs revealed considerable reluctance to let anyone outside of their own company know that they had been the victims of an attack. In some cases this information was not passed to senior management or to the national regulator. None of the incidents were passed to the European agencies responsible for exchanging 'lessons learned' across member states.

Political, economic and regulatory concerns limit the extent to which security information is passed within industries. This forms a strong contrast with the growing use of reporting sites for the exchange of information about safety concerns and

operational problems. In consequence, many companies remain completely ignorant about the attacks the have been suffered by their neighbours and colleagues (DoT, 2009).

The need to share lessons from previous attacks is apparent when discussing cyber threats with engineers and senior management across many European states. None of the sites visited in the review phase of this project had read the studies of US security in Air Traffic Management. This had important implications. For example, the DoT (2009) review identified that intrusion detection systems had only been deployed in eleven out of three hundred ATM facilities. If these reports were more widely disseminated then service providers might be encouraged to consider the possible consequences of CyberAttacks on their infrastructures.

3.6 Tools for CyberSafety Risk Assessment

The final element of the roadmap uses the insights from drills and exercises together with the lessons learned from previous incidents to revise security risk assessments in safety-related applications. This is non-trivial. In previous work we have considered the potential consequences that a CyberAttack might have on service provision (Johnson and Atencia Yepez, 2011). This has demonstrated that most of the evidence gathered to support safety cases can be undermined by the detection of malware or unauthorised access. For instance, it can be difficult to guarantee response times given the potential impact on processor and memory resources.

Existing software architectures lack many of the security features being considered, for instance within the FAA's SWIM program. Intrusion detection programmes and access control techniques have been installed in a piecemeal fashion. In consequence, it is very difficult to create convincing arguments that safety can still be maintained once an attack has been detected. Typically, the affected systems are shut down and engineering teams assume that secondary applications are unaffected. This assumption can be difficult to support when common infrastructure components are increasingly being used across multiple redundant systems. In some cases, ECAC states have created secondary systems that rely on identical software to their primary applications. This reduces the costs associated with design diversity and helps to mitigate the risks of introducing errors by trying to maintain two different applications. However, this approach creates enormous concerns when security vulnerabilities are common to primary and secondary systems. Risk assessment tools help to identify the likelihood and consequences of attacks that exploit the vulnerabilities created by common software components.

The final stage in the roadmap communicates the insights from revised risk assessments back to senior management and regulators. The lack of strategic leadership can only be addressed if information is provided to senior management about future forms of attack.

5. Conclusions

Common software components are gradually being integrated across many safety-critical infrastructures. These include network management tools, operating systems such as Linux, Voice Over IP (VOIP) communications techniques, Satellite Based Augmentation Systems for navigation/timing data etc. The increasing use of these common components creates concerns that bugs might affect multiple systems across many industries. It also raises significant security concerns. Malware has been detected on a range of safety-related systems in the power distribution, healthcare, military and transport industries (Anderson, 2008). Previous attacks have not been targeted on critical applications. However, there is no room for complacency in the face of increasing vulnerability to cyber attacks on safety-related systems. In order to illustrate the vulnerabilities, we have identified security threats to Air Traffic Management infrastructures.

This paper has presented a roadmap for increasing resilience to future CyberSafety attacks. We must raise awareness about the potential threats to safety-related systems amongst regulators and senior management. Without greater strategic leadership there is a danger that ANSPs will continue to respond to security breaches in a piecemeal way that leaves major vulnerabilities in our underlying infrastructures.

A second element of the roadmap focuses on improved screening, competency assessment and training for engineering staff. Most ANSPs continue to ignore the 'insider threat' and lack the expertise either to diagnose or resolve potential attacks. It is for this reason that organizations including the FAA and EUROCONTROL should promote and provide training in anticipation of future incidents.

Other areas for action include the use of drills and exercises to support team resource management in the aftermath of an attack. These exercises can be tailored to scenarios derived from previous incidents; this depends upon initiatives to exchange information about those attacks that have already occurred. At present, there is no forum for lessons of this nature. The final element of the roadmap proposes a new generation of tools that use lessons learned from previous attacks together with the insights from drills and exercises to assess the risks of future cyber attacks. For example, safety cases can be used to map the impact of a potential threat in terms of the arguments that are undermined by an attack (Johnson and Atencia Yepez, 2011). Other techniques provide support for more formal reasoning about the safety consequences of malware and intrusions into critical infrastructures (Johnson, 2011).

A series of site visits to European ANSPs informed and motivated this work. These revealed the need to provide engineering teams with a forum for sharing common concerns and proposed solutions to Cyber Security threats. For example, many of the groups stressed the difficulty in

distinguishing whether abnormal behaviors are due to a security breach or to everyday bugs. The visits also demonstrated the need for a more systematic overview of the vulnerabilities across member states; the European Commission should replicate the studies prepared by the US GAO and DoT.

We are now beginning to witness the first wave of successful attacks. Viruses are being detected on primary systems in safety-critical applications. There is, therefore, an urgent need to act without any further delay. We also need to think strategically and plan ahead for future threats. Several ANSPs have begun to consider ways in which their operations might exploit Cloud-based infrastructures. This opens up powerful and cost-effective tools for operational and engineering staff; it also raises a host of security concerns about the threats to the next generation of safety-critical systems (Mather, Kumaraswamy and Latif, 2009).

References

M. Abrams and J. Weiss, Malicious Control System Cyber Security Attack Case Study–Maroochy Water Services, Australia, NIST/Mitre Corporation, NIST Industrial Control System Security Project , http://csrc.nist.gov/sec-cert/ics/index.html 2008.

R. Anderson, Security Engineering: A Guide to Building Dependable Distributed Systems, Wiley, Indianapolis, USA, 2008.

C. Haley, R. Laney, J. Moffett and B. Nuseibeh, Security Requirements Engineering: A Framework for Representation and Analysis, IEEE Transactions on Software Engineering, (34)1:133-153, 2008.

C.W. Johnson, Using Assurance Cases and Boolean Logic Driven Markov Processes to Formalise Cyber Security Concerns for Safety-Critical Interaction with Global Navigation Satellite Systems. In J. Bowen and S. Reeves (eds.), Proceedings of the 4th Formal Methods for Interactive Systems Workshop 2011, Limerick, Ireland, 2011.

C.W. Johnson and A. Atencia Yepez, Mapping the Impact of Security Threats on Safety-Critical Global Navigation Satellite Systems. In Proceedings of the 29th International Systems Safety Society, Las Vegas, USA 2011, International Systems Safety Society, Unionville, VA, USA, 2011.

C.W. Johnson, C. Shea and C.M. Holloway, The Role of Trust and Interaction in GPS Related Accidents: A Human Factors Safety Assessment of the Global Positioning System (GPS). In R.J. Simmons, D.J. Mohan and M. Mullane (eds.), Proceedings of the 26th International Conference on Systems Safety, Vancouver, Canada 2008, International Systems Safety Society, Unionville, VA, USA, 2008.

C.W. Johnson and C.M. Holloway, A Roadmap for Safer Systems Engineering, IET Systems Safety Conference, The IET, Savoy Place, London, 2011.

J.M. Maddalon and P.S. Miner, An Architectural Concept for Intrusion Tolerance in Air Traffic Networks, NASA Langley Technical Report, Integrated Communication Navigation and Surveillance (ICNS), Annapolis, Maryland, 2003.

T. Mather, S. Kumaraswamy and S. Latif, Cloud Security and Privacy: An Enterprise Perspective on Risks and Compliance. O'Reilly Media, California, USA.

P. Pederson, D. Dudenhoeffer, S. Hartley and M. Permann, Critical Infrastructure Interdependency Modelling: A Survey of U.S. and International Research, Technical Report INL/EXT-06-11464, Idaho National Laboratory, US Department of Energy, August 2006.

M. Theoharidoua, S. Kokolakisb, M. Karydaa and E. Kiountouzisa, The Insider Threat to Information Systems and the Effectiveness of ISO17799, Computers & Security, (24)6:472-484, 2005.

US General Accounting Office, Air Traffic Control: Weak Computer Security Practices Jeopardize Flight Safety (Letter Report, 05/18/98, GAO/AIMD-98-155), 1998.

US Department of Transport, Report on Review of Web Applications Security and Intrusion Detection in Air Traffic Control Systems, FAA Report Number FI-2009-049, Washington DC, USA, May 2009.