PREPARING FOR CYBER-ATTACKS ON AIR TRAFFIC MANAGEMENT INFRASTRUCTURES: CYBER-SAFETY SCENARIO GENERATION

Chris W. Johnson

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

Keywords: Cyber-security, safety-critical systems, transport.

Abstract

Malware poses a growing threat to a host of safety-critical systems that depend on common software components, including the Linux operating system and the Internet Protocol (IP). Threats include 'mass market' malware that is not deliberately aimed at safety-related systems. They also include more sophisticated techniques exploited by W32.Stuxnet, W32.Duqu, W32.Flame etc. Previous work in this area has focused on the consequences of a cyber-attack under 'optimal conditions'. Very little work has been done to identify more complex scenarios when malware exacerbates routine system failures that occur in all safety-critical applications. We show how Vulnerability and Violation (V2) diagrams can identify interactions between malware and degraded modes of operation. The intention is not to accurately predict future modes of attack. In contrast, the aim is to create training scenarios that test the expertise and judgement of systems engineers, operators and managers. The initial results from our work have revealed the underlying vulnerabilities that exist across safety-critical transportation infrastructures.

1 Introduction

In the past, safety-critical applications relied heavily on bespoke software. This provided a degree of protection from malware; it required specialist knowledge to devise methods of attack. In contrast, the growing use of Linux and IP infrastructures create new vulnerabilities. For example, the Voice over IP (VOIP) architectures that support communications over the public Internet are now commonly being applied in the private networks operated in air traffic management, in maritime surveillance, in fire and rescue command and control etc. Many more people have the technical knowledge that is necessary to compromise these common components; the underlying architectures are taught in most Undergraduate courses in Computing Science. The growing reliance on a small number of software architectures also creates common points of failure across different infrastructures. For the first time, we face the possibility that cyber-attacks might simultaneously affect multiple safety-critical industries.

It is difficult to balance paranoia and complacency when assessing the technical risks that malware poses to safety-critical systems. This paper focuses on scenario development for cyber-attacks on transportation infrastructures. It is hard to identify appropriate training exercises because we cannot be certain about the methods that will be used in future malware. Further problems arise because cyber-attacks are unlikely to occur under 'perfect' conditions. We must, therefore, consider the potential impact of malware on subsystems undergoing periodic maintenance or when key members of staff are unavailable.

In order to address these concerns, we have developed a graphical modelling tool, based on Vulnerability and Violation (V2) diagrams, that represents the interaction between degraded modes of operation and the impact of malware [4]. The modelling is based around known methods of attack, from previous incidents involving safety-critical systems as well as attack methods used in other domains. This encourages the development of hybrid scenarios mixing different approaches to create new attack hypotheses. The resulting models are intended to support multi-disciplinary team meetings that plan and validate cyber-exercises. The closing sections report on initial applications of this approach by European Air Navigation Service Providers (ANSPs).

2. Modelling 'Mass Market' Attacks

It is hard to assess the risks posed by malware; both in terms of the probability of an attack and the consequences for safety. There have been a small but growing number of incidents in which the defences of safety-critical applications have been penetrated, including Air Traffic Management infrastructures, Fire and Rescue dispatch systems and Maritime monitoring applications. These attacks have not yet resulted in loss of life. It can be argued that existing defences including human monitoring, firewalls, software and hardware diversity provide sufficient protection. However, it is important not to be complacent especially when there is little consensus about the nature of future threats.

Previous attacks have involved 'mass market' malware; they were not deliberately aimed at safety-related infrastructures [3]. For example, Linux.Psybot has spread between

embedded devices, including routers, running Linux MIPS (Microprocessor without Interlocked Pipeline Stages) architectures. Psybot uses relatively simple techniques, including brute-force authentication attacks on administration interfaces. These well-known attacks still cause considerable problems for safety-critical systems. Malware undermines assumptions about the processing, memory and network resources that underlie existing safety arguments. It is hard to guarantee that a safety-critical process will always have the resources that it requires if we suspect that a system has been compromised.

Further problems arise when safety-related applications lack network monitoring tools. Critical applications often have surprisingly primitive tools for analysing network traffic. In some cases, monitoring software is deliberately removed after installation; each additional component incurs significant validation and verification overheads to ensure that they do not themselves introduce additional failure modes.

Even when network monitoring tools are available, many safety-critical organisations lack the forensic skills to diagnose the extent of an infection or to devise appropriate means of recovery. This creates particular concerns for industries that must maintain levels of service in safety-critical applications. We cannot simply hold aircraft in flight while we prepare a full forensic analysis of a Flight Data Processing System. It is, therefore, important to hold cyber-security drills and exercises that help service providers to ensure that they can respond to an attack while at the same time preserving the safety of their operations.

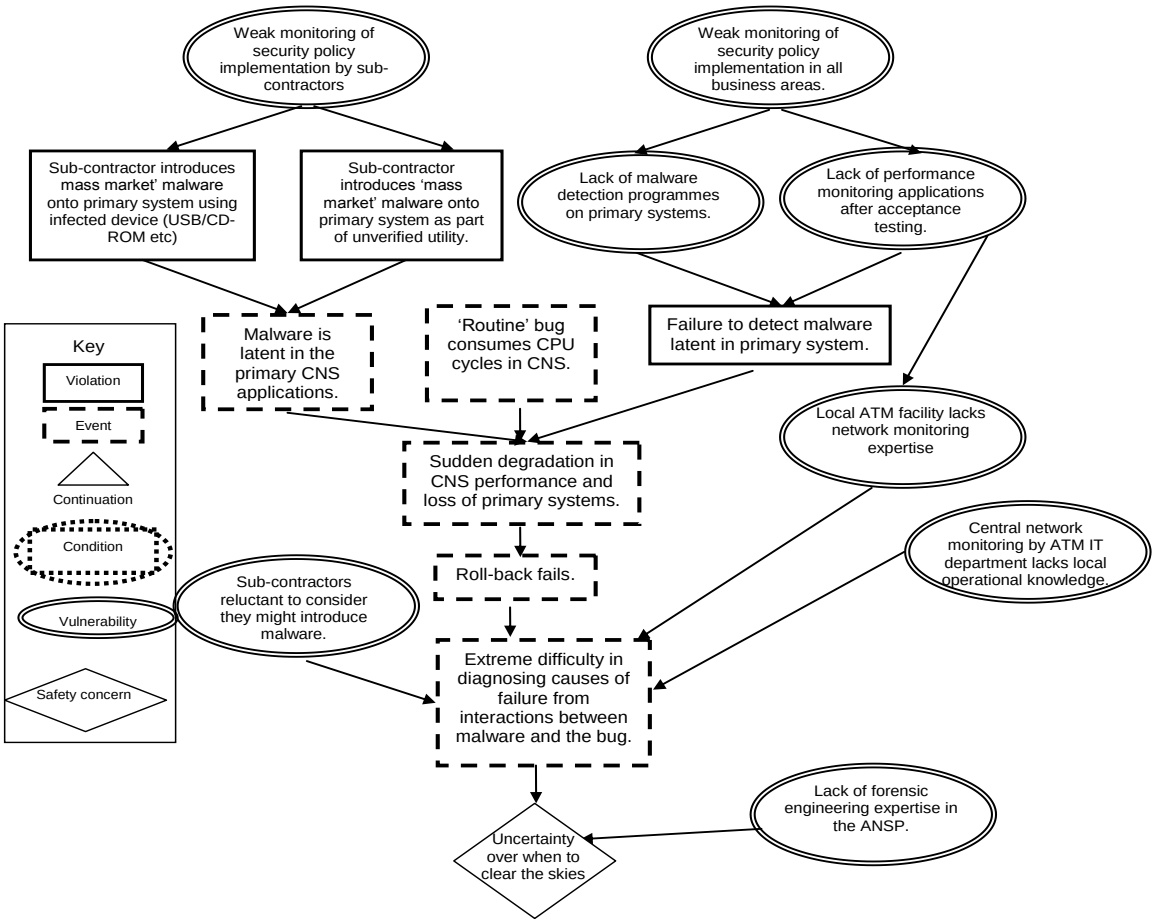

Figure 1: Violation and Vulnerability Diagram of a Cyber-Exercise Scenario Based on Two Previous ATM Incidents

Figure 1 illustrates the use of V2 diagrams to model cyber-attack scenarios. This diagram is based on lessons learned from two previous incidents involving European Air Navigation Service Providers (ANSPs). The V2 graphical modelling technique was initially developed to analyse physical attacks on national critical infrastructures, including the London and Madrid bombings [3]. However, it has recently been extended to model telecommunications incidents, including the Pakistan YouTube IP address spoofing [4].

The nodes at the top of Figure 2 represent a vulnerability associated with inadequate monitoring of the ways in which sub-contractors implement an organisations security policy. In the scenario shown here, the vulnerability contributes to two violations. Each of these describes a different means of 'infection' identified in the two previous incidents, mentioned above. In the first case, malware was introduced by an infected USB stick. In the second incident, security was compromised when malware was hidden in a library of modules that were not directly written by the sub-contractor

and which had not been subject to sufficient verification. In both cases, malware was brought into the primary Communication, Navigation and Surveillance (CNS) systems of European Air Traffic Management companies.

The second vulnerability, at the top of figure 1, refers to managerial and organisational weaknesses within the ANSP rather than external, sub-contractors. In one of the incidents, the ANSP had a strong central engineering department. However, the malware affected a network in a smaller regional centre. Local system engineers were responsible for their site; however, they lacked the resources available to staff in the national centre. This led to the inadequate monitoring of local networks. It was impossible to determine whether the network loading was consistent with particular operational characteristics; e.g., the number of packets generated over time by individual radar targets etc. This made it difficult to determine whether or not malware was consuming additional network resources.

Further vulnerabilities were created by the lack of anti-malware software inside primary CNS applications. The ANSP relied on a 'strong shell, soft core' security policy that was intended to ensure malware could ***never*** be introduced to a primary CNS application. Brevity prevents the inclusion of arguments for this omission within the V2 diagram. Further iterations of the analysis could include this additional information. In contrast, the remainder of this section extends the analysis of the scenario represented in Figure 1.

The lack of network monitoring tools and malware detection applications prevented the detection of malware introduced by the sub-contractors in the two incidents. There were no immediate adverse consequences because CNS applications are, typically, designed with significant additional processing and networking resources to cope with future increases in traffic levels. However, Figure 1 illustrates a pathological situation experienced by one of the ANSPs when a 'routine' software update introduced a bug into a CNS application. This is part of the everyday engineering of Air Traffic Management; the usual response is to roll-back primary systems to the previous installation. Safety is not endangered because system updates usually take place at night, during times of light traffic. In this incident, however, inadequate precautions were taken to document the configuration parameters of the CNS systems before the installation began. The roll-back could not be completed. The decision was taken to press ahead with the reinstallation of the new system and try to diagnose the source of the failure. This created huge complexity because engineers did not realise that they were trying to understand the interactions between the routine 'bug' and the effects of the malware. Making changes to the routines where the bug had been identified did not seem to remove all of the symptoms; because some were associated with the virus.

Figure 1 shows how the details from previous incidents can be used to inform the initial stages of scenario development for cyber-security training exercises. Additional details can gradually be added, reflecting the degree of difficulty that is appropriate for a particular exercise. Little is to be gained if teams of engineers are to be set a task that they have no hope of solving; equally it is important within any exercise to set challenges that will expose the lack of resources and competency in key areas such as forensic analysis.

Our analysis of previous ATM incidents has revealed a recurring problem that arises when new software upgrades are implemented at night in regional centres. It can be very difficult for local staff to obtain specialist support from central engineering teams when key individuals are not scheduled for those shifts. In many cases, central engineering teams are unaware that their regional colleagues are about to start a software upgrade. Hence, Figure 1 might be extended to describe the problems that arise when regional engineering teams try to get help from network monitoring specialists in their national engineering centres during the early hours of the morning.

The final node in Figure 1 captures the safety concerns that arise from the cyber-attack scenario. At low levels of traffic, safety can be assured using procedural Air Traffic Management techniques. The longer that any failure persists then the harder it is to maintain levels of service and safety. The decision to close air space involves detailed consultations between systems engineers, managers, operational staff etc. This is a key stage in any cyber-exercise. As before, the V2 diagram illustrated in Figure 1 can be extended to represent additional safety concerns that might be introduced into a cyber-exercise scenario. For example, the decision to close air space can suddenly increase the workload for neighbouring centres who must now handle the traffic from the failing centre in addition to their own traffic. It is still rare for European ANSPs to conduct training exercises that involve staff from multiple centres, although this may arguably become more common through the development of Functional Airspace Blocks across national borders. In our example, we might extend the scenario to address the safety concerns that might when neighbouring centres receive additional traffic even when they might also be infected by the same malware. This is plausible given the increasing network integration envisaged under the SWIM components of the Single European Skies programme, known as SESAR.

This section has mapped out the interactions between 'routine' system failures, including software bugs, and mass-market malware, such as Linux.Psybot variants. Most of these attacks are easily identified and contained in non-critical systems using commercial anti-virus software. This is more difficult in safety-related applications. It is important to continually update anti-virus software so that it maintains definitions of the most recent threats. However, this creates problems when engineers must demonstrate that updates to anti-viral software do not introduce new failure modes into safety-critical systems. It is, therefore, particularly important that cyber-attack exercises are conducted to increase confidence in the ability of key staff to respond to the threat posed by malware. This section has also shown how V2

diagrams can be used to map out training scenarios. They can be used by planning teams to ensure the technical validity and veracity of an exercise.

3. Modelling Sophisticated Attacks

The involvement of state agencies in the development of malware has, arguably, resulted in a more complex range of threats illustrated by W32.Stuxnet, W32.Duqu and W32.Flamer. At the time of writing this paper, the author is unaware of these forms of attack having been detected by ANSPs. However, W32.Duqu has been detected inside the firewalls of several European safety-related industries. It is, therefore, critical that tools for the generation of cyber-security scenarios can consider the impact of more sophisticated attack vectors on public safety.

One of the most innovative aspects has been the integration of social media with command and control servers. The use of remote hosts to update versions of malware and to forward sensitive information from an infected machine is well established. However, the development of these technologies has been sustained by the growth of Internet based services where accounts can be created in minutes with minimal authentication. Similarly, anonymity servers have provided proxy services and levels of indirection for new generations of malware. These innovations enable more advanced forms of malware to cross firewalls that safety-critical organiations previously considered to be impervious from most forms of attack. For example, Stuxnet and Duqu used command and control servers to collate information about anti-viral software on an infected host. Additional modules could then be downloaded by the malware or the attack could be suspended if new versions of protective software were detected.

More advanced malware architectures resemble those of the protective systems that companies deploy against them. Many viruses resemble loaders that are able to change the nature of their attack by downloading different definitions and component structures. This helps to avoid detection. New malware can also be downloaded to gather different information during an attack. These innovations have already been used in the examples, cited above. It remains to be seen whether the operators of safety-critical systems, including ANSPs would have the skills that are required to detect and contain these more sophisticated threats.

Malicious software has also been detected in system components that would not normally be considered by many anti-viral systems, including graphics cards. Malware has also exploited multiple transmission vectors across local and wide area networks. There have also been innovations in transmission through different media, including USB sticks, SD memory cards etc. Stuxnet limited the number of infections that could occur from an individual device; deleting itself after the malware had been transmitted to five new hosts. This made it difficult for administrators to trace the course of any infection. Many safety-critical organisations continue to assume that they are protected from malware because they are isolated from the public Internet. The use of these devices provides a direct mechanism for transmission between training and development systems that are connected to the public Internet and those machines that are supposedly isolated from infection.

Most safety-related organisations have security policies that place tight controls on the use of portable memory devices. However, these policies are often violated and many devices are found during security audits. Sub-contractors often do not understand the reasons for restrictions on their use. Even if they understand the concerns over USB sticks then they often fail to realise that many of the same risks also apply to SD cards, CD ROMS etc.

Recent incidents have shown that more advanced forms of malware target specific systems. They use virulent transmission mechanisms to compromise a large number of hosts but they only damage a smaller number of intended targets [3]. The malware is intended to have a minimal effect on most of the machines that are infected, if they are not one of the small group of hosts that are the focus of the attack. They can be thought of as guided weapons. Like conventional munitions, however, there remains a considerable risk of collateral damage to safety-critical systems. We cannot assume that the inadvertent infection of safety-related applications would have minimal consequences even if they were not the intended target of the malware. The additional resources consumed by malware can invalidate safety arguments. There is also the possibility that hostile states may in the future use these methods to deliberately target safety-related applications beyond those associated with the Iranian nuclear programme.

A further area of innovation has been the integration of state machines. These carefully control the impact of malware on the host that is under attack. If software causes predictable damage – in terms of CPU, network bandwidth, memory consumption etc, then it is relatively easy for systems administrators to track the source of an infection. It is also possible for engineers to make assumptions about the likely behaviour of the system under attack. In contrast, Stuxnet used a state machine to disguise the impact of the malware. It was deliberately designed to frustrate detection and diagnosis. After the virus infected a targeted machine, the malware did nothing for several weeks. It then used an internal clock to vary the impact of the damage over several weeks. It would then return to a state in which it had no observable effects. Such variations create huge problems for the teams that must detect any infection. It is even more difficult for safety-critical industries that have out-sourced the provision of application software and network infrastructures to external companies. The diagnosis of any attack requires interaction and coordination between IT service providers; who often have a minimal understanding of the safety-critical nature of particular operations. These problems are likely to become far worse if there is any expansion in the relatively small

number of safety-related service providers that have begun to exploit Cloud architectures.

The use of state machines to frustrate detection and diagnosis has led to 'land-mine' attacks. In conventional warfare, mines are not intended to kill the victim but to ensure that scarce resources are used to care for the injured. By extension, more recent attacks have not been intended to halt the operation of a critical system but to consume scarce engineering resources that might otherwise sustain operations. Staff become engaged in a fruitless search for the source of what appears to be a 'normal' bug within application software. This further emphasises the importance of cyber-attack exercises so that safety-critical engineers are encouraged to consider that adverse behaviour might possibly be the result of malware rather than a more routine coding or configuration error.

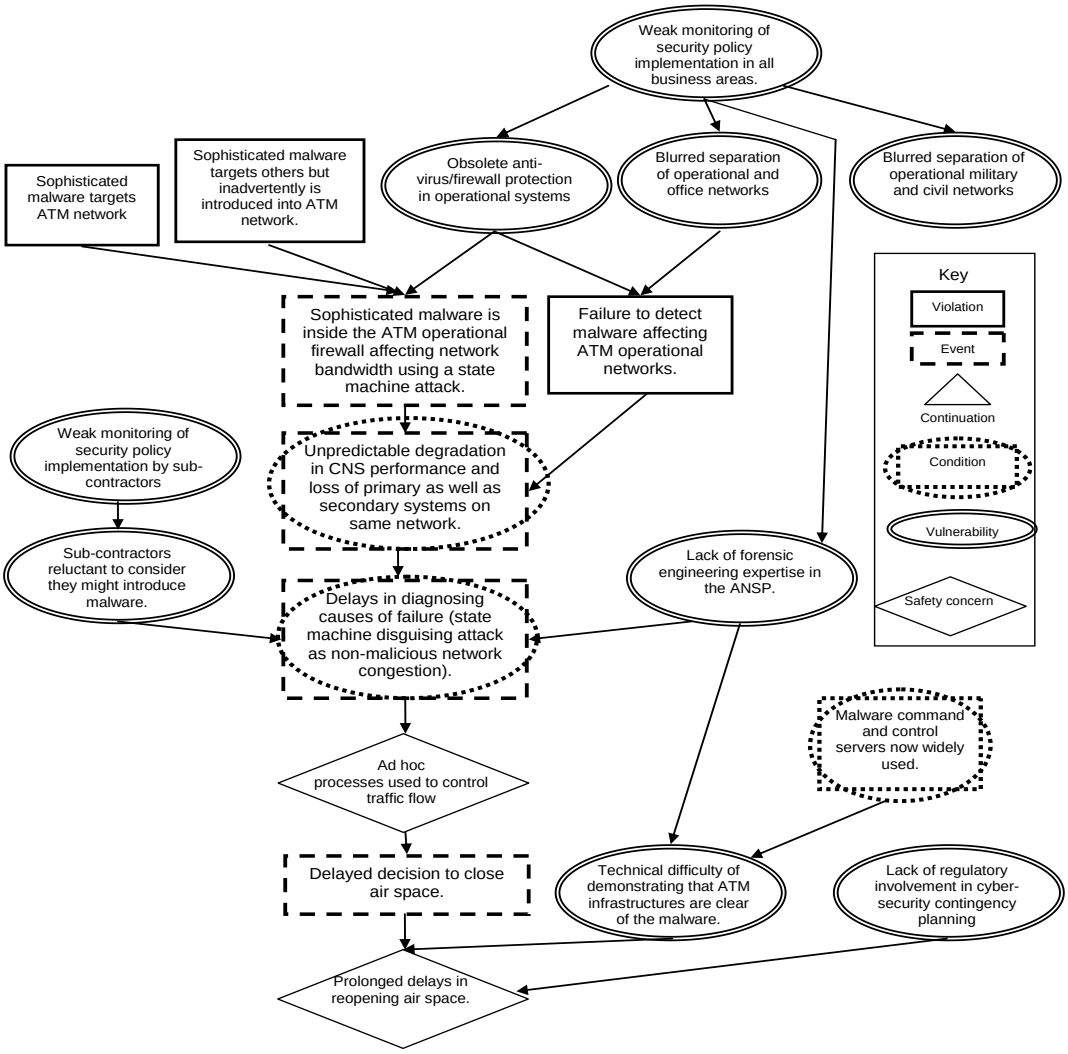

Figure 2: V2 Diagram of a More Complex Cyber-Attack Scenario

Regulators are important stakeholders in cyber-security exercises. In preparation for this paper, the author visited more than twenty European states to assess the state of regulatory support in this area. The results are depressing. At present most state agencies are under-prepared for the possible consequences of a cyber-attack. They lack personnel with forensic training. They do not participate in the (limited) cyber-security training that is provided by many service providers. In consequence, many companies will not inform their regulators about an attack. Reports are drafted to show that an infection only affected office systems even though these applications directly supplied data to operational systems. In other cases, pressure is put on engineering teams not to delay the recovery from an incident by engaging with the regulator. Some of these problems arise from structural problems at a higher-level in Europe and North America. We have created agencies whose primary focus is on safety, such as the European Aviation Safety Agency, and others that deal with cyber-security, such as the European Network and Information Security Agency. They produce documents and guidance that reveals little understanding of the concerns that influence other agencies. Safety-related risk assessment techniques often cannot be applied to security threats; which are not characterised by normal probability distributions. Security related guidance often assumes that it is possible to immediately isolate a system, which would directly threaten

public safety in safety-critical applications. We argues that there is an urgent need for these agencies to work together before we suffer a significant attack on safety-critical infrastructures.

Many of the security concerns described in this section affect all safety-critical systems. Cyber-exercises can also identify particular threats to specific industries. For instance, some European states have a close integration between the computational infrastructures that support military and civil air traffic management. In some cases, this extends to the use of shared network components and routers. Although there systems are, typically, well protected, it is essential to test the effectiveness of these defences which are not always as strong as their proponents might claim.

Figure 2 presents a V2 diagram for a more sophisticated cyber-attack scenario than that illustrated in Figure 1. The top level vulnerability again describes the weak monitoring of security policies. In this case, concerns focus on the blurred separation of operational and office systems. The US Department of Transport found numerous interdependencies between operational and business networks in the US National Air Space [5]. Further vulnerabilities arise from the use of obsolete protection given the difficulty of verifying the safety of anti-viral software updates, noted in previous sections. Figure 2 also includes a node to indicate the vulnerability of military and civil systems; this is not expanded here for obvious reasons. This second V2 diagram focuses on the consequences of a deliberate attack on ATM networks, for instance by causing a machine to broadcast spurious packets onto a local area network in an attempt to deliberately undermine quality of service for legitimate network traffic. This scenario is based on previous byzantine failures caused by hardware problems that led to the closure of European air space [2]. Subsequent events in the V2 diagram illustrate the introduction of malware using a state machine attack. This complicates diagnosis and response during a cyber-exercise. Figure 2 introduces additional challenges to the participants; engineering teams must coordinate with sub-contractors; managers must consult regulators. This scenario raises significant safety concerns in the recovery phase, as we have seen the development of malware loaders frustrates attempts to ensure that any threat has been removed from a network. Recent contingency events in European aviation, including the Eyjafjallajokull eruption, have shown the difficulty of sustaining technical arguments in the face of political and public pressure to open airspace [1]. As with figure 1, each of these elements in the V2 diagram form the focus for subsequent exercises that are intended to build the skills that needed by teams of different stakeholders, including but not limited to sub-contractors, engineering staff, operations, management and regulators.

6. Conclusions and Further Work

Malware poses a growing threat to safety-critical systems. The growing reliance on common software components, including the Linux operating system and the Internet Protocol (IP) have increased exposure to ‘mass market’ malware that is not deliberately aimed at safety-related systems. There is also an emerging threat from more sophisticated attacks, illustrated by W32.Stuxnet, W32.Duqu, W32.flame etc. It is, therefore, important to ensure that engineers, managers, operational staff and regulators consider what might happen if safety-critical systems were to be compromised by a cyber-attack. This paper has used Vulnerability and Violation (V2) diagrams to map out the safety concerns that arise from a range of different scenarios in air traffic management. The intention is to create multi-agency exercises that test the expertise and judgement of regulators, systems engineers, operators and managers as they work to ensure safe and successful operation. The relevance of our work has been demonstrated by the inclusion of information from previous incidents and also by the use of these scenarios in training staff from a number of European ANSPs.

This paper is a first step – many other tools and techniques support scenario development. For instance, a related project is making extensive use of Excel worksheets to develop detailed scripts for participants in a NATO exercise. V2 diagrams arguably provide a better overview of the safety concerns during initial scenario development. Further work is required to determine whether we can integrate both approaches to support initial planning and more detailed script development for cyber-security exercises in safety-critical applications.

References

[1] A. Jeunemaitre and C.W. Johnson, Future Directions for Contingency Planning in European Air Traffic Management: A Response to the 2010 Eyjafjallajokull Volcano Eruption. In A. Alemanno (ed.), Proceedings of the 1st HEC Paris Workshop on Regulation; Emergency Regulation under the Threat of Catastrophy - the Volcanic Ash Crisis, Jouy en Josas, France, 2010.

[2] Irish Aviation Authority, Report Of The Irish Aviation Authority Into The ATM System Malfunction At Dublin Airport, Dublin, Ireland, September 2008.

[3] C.W. Johnson, *CyberSafety: On the Interactions Between CyberSecurity and the Software Engineering of Safety-Critical Systems*. In C. Dale and T. Anderson (eds.), Achieving System Safety, 85-96, Springer Verlag, ISBN 978-1-4471-2493-1, 2012.

[4] C.W. Johnson, Lessons from Major Incidents Influencing and Influenced by Telecoms Failures. In P. Theron and S. Bologna (eds.), Critical Information Infrastructure Protection and Resilience in the ICT Sector, IGI Global, Pennsylvania, USA, 2012, in press.

[5] 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.