



City Research Online

City, University of London Institutional Repository

Citation: Bloomfield, R. E., Bendele, M., Bishop, P. G., Stroud, R. and Tonks, S. (2016). The risk assessment of ERTMS-based railway systems from a cyber security perspective: Methodology and lessons learned. Paper presented at the First International Conference, RSSRail 2016, 28-30 Jun 2016, Paris, France.

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/15105/>

Link to published version: http://dx.doi.org/10.1007/978-3-319-33951-1_1

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

City Research Online:

<http://openaccess.city.ac.uk/>

publications@city.ac.uk

The risk assessment of ERTMS-based railway systems from a cyber security perspective: methodology and lessons learned

Robin Bloomfield¹, Marcus Bendele¹, Peter Bishop¹, Robert Stroud¹, Simon Tonks²

¹Adelard LLP, London, UK
{reb, mmb, pgb, rjs}@adelard.com
²Porterbrook Leasing Company
simon.tonks@porterbrook.co.uk

Abstract. The impact that cyber issues might have on the safety and resilience of railway systems has been studied for more than five years by industry specialists and government agencies. This paper presents some of the work done by Adelard in this area, ranging from an analysis of potential vulnerabilities in the ERTMS specifications through to a high-level cyber security risk assessment of a national ERTMS implementation and detailed analysis of particular ERTMS systems on behalf of the GB rail industry. The focus of the paper is on our overall methodology for security-informed safety and hazard analysis. Lessons learned will be presented but of course our detailed results remain proprietary or sensitive and cannot be published.

Keywords. Security assessment, Safety-critical systems, Security-informed safety, ERTMS, Railway signaling systems

1 Introduction

The European Railway Traffic Management System (ERTMS) is a major industrial project that aims to replace the many different national train control, command and signaling systems in Europe with a standardized system. In Great Britain, Network Rail are preparing to introduce ERTMS as part of the upgrade of the signaling and communications systems running on Britain's rail infrastructure. This upgrade has the potential to increase the risk of an electronic attack on the rail infrastructure, as it brings more systems under centralized control. Government and railway stakeholders identified a need to understand the security implications of the new technology more than five years ago and there have been a number of studies by industry specialists and government agencies of the impact that cyber issues might have on the safety and resilience of railway systems.

This paper presents some of the work done by Adelard in this area, ranging from an analysis of potential vulnerabilities in the ERTMS specifications through to a high-level cyber security risk assessment of a national ERTMS implementation and

detailed analysis of particular ERTMS systems on behalf of the GB rail industry. The focus of the paper is on our overall methodology for security-informed safety and hazard analysis. Lessons learned will be presented but of course our detailed results remain proprietary or sensitive and cannot be published.

2 Railway security requirements

Traditionally, computer security deals with threats to confidentiality, integrity, and availability, but here we are concerned with train movements rather than information, so our primary concern is integrity, then availability, and finally confidentiality. Loss of integrity could result in accidents or collisions, whereas loss of availability would bring the railway system to a halt. Loss of confidentiality is less of an immediate threat, but might result in the leak of sensitive operational information. Reliability is also important, since an unreliable train service will result in a loss of public confidence in the railway operators.

Thus, the hazards or potential failures or undesirable outcomes to be avoided are:

- a collision involving multiple trains;
- an accident such as derailment involving a single train;
- widespread disruption of train services over a large area;
- disruption to individual trains, or trains within a local area;
- creation of a situation that leads to panic and potential loss of life (e.g., an emergency stop and uncontrolled evacuation onto the track);
- creation of a situation that leads to passenger discomfort and dissatisfaction (e.g., stopping a train indefinitely in a tunnel);
- loss of public confidence in the railway system due to intermittent low-level problems affecting the reliability of the service;
- leak of sensitive information (e.g., movements of hazardous cargoes or VIPs).

The ERTMS safety analysis considers the effect of potentially catastrophic events on the integrity of the system. Faults that could result in an accident need to be considered in both a safety and security analysis, regardless of the underlying cause of the fault (accidental, deliberate or malicious).

3 Security analysis of ERTMS specifications

The starting point for our ERTMS work was a security analysis of the ERTMS specifications that we were commissioned to perform on behalf of key UK railway stakeholders and UK government about five years ago [1]. The aim of the study was to examine the ERTMS specifications for potential security vulnerabilities and identify systemic weaknesses in the ERTMS specifications. We were concerned with conceptual problems with the specifications rather than vulnerabilities caused by design flaws, bugs in implementations of ERTMS technology, or weaknesses in the opera-

tion or maintenance procedures for an ERTMS system. Such vulnerabilities are important but were outside the scope of our study.

Our analysis was holistic and considered whether a national deployment of ERTMS might introduce vulnerabilities into the national rail infrastructure. Our review focused on ERTMS Application Level 2, which made it possible to restrict attention to a number of core specifications, and ignore specifications for interacting with legacy train protection systems and trackside signaling equipment. We also considered the security of GSM-R and analyzed how GSM security impacts on GSM-R security. We were particularly interested in electronic attacks that could be launched remotely and would cause widespread disruption.

3.1 Methodology

Our approach was to consider the trust relationships between the various components of the overall architecture and analyze the consequence of a breach of trust. This enabled us to identify a set of potential weaknesses and vulnerabilities in the specifications. We then developed scenarios that showed how these weaknesses could be exploited by an attacker. These scenarios were refined and validated in discussion with railway stakeholders, and proved to be a very effective way of communicating the risks of an ERTMS implementation being compromised.

Analysis of trust relationships. ERTMS is implemented using a number of trackside and on-board sub-systems, and the ERTMS specifications describe the interfaces by which these various subsystems interact to ensure that trains move safely without exceeding their movement authority. We performed a systematic analysis of the ERTMS specifications from a security perspective by examining the on-board ETCS application, and considering its interfaces and trust relationships with other components of the ERTMS system, both trackside and on-board the train.

Development of attack scenarios. Having identified some potential vulnerabilities in the ERTMS specifications, we devised attack scenarios to explore the ways in which an attacker could exploit these potential weaknesses and vulnerabilities to achieve one of the undesirable outcomes identified in Section 2.

We devised seven attack scenarios and then analyzed each scenario in detail by considering the following questions:

- **how** is the attack performed?
- **what** vulnerabilities does the attack exploit?
- **where** can the attack be launched from?
- **what** are the possible mitigations?

We then graded each attack according to a range of criteria:

- the **type of access** required to exploit a vulnerability;
- the **level of technical sophistication** required to exploit a vulnerability;

- the **type of failure** caused by a successful attack;
- the **scale of effect** for a successful attack;
- the **scalability of the attack** from the attacker’s perspective;
- the **type of impact** caused by a successful attack;
- the **types of mitigation strategy** that are possible;
- the **level of difficulty** for implementing each mitigation.

We did not attempt to rank the various attack scenarios using a weighted average of the category scores because we believe that such a ranking would be too simplistic – the relative weighting of the various categories and the ranking of the scenarios is a matter for government and industry stakeholders. Similarly, we did not attempt to estimate the likelihood of attacks being successful because this would depend on the national implementation of ERTMS and is therefore best left to the domain experts. Instead, we used color coding (**HIGH**, **MEDIUM**, **Low**) to highlight the issues. Using this color coding, we produced a table summarizing our grading of each attack scenario under the various headings to enable the scenarios to be easily compared.

Broadly speaking, attacks that can be launched remotely do not require a high level of sophistication and are highly scalable – however, such attacks are relatively easy to mitigate. Conversely, attacks that require local access are less scalable but also more difficult to mitigate. Hence important trade-offs need to be made by the relevant decision makers and risk managers. The advantage of the analysis and grading approach presented here is that it identifies these trade-offs and helps the stakeholders to make more informed decisions.

4 Risk assessment of a national implementation of ERTMS

Following on from our initial security analysis of the ERTMS specification, we were asked to provide a risk assessment for a national implementation of ERTMS.

In Great Britain, Network Rail are planning to implement an ERTMS overlay on top of the existing signaling and control system [2]. There are also plans to introduce a new traffic management system and eliminate the need for about 800 small signal boxes by centralizing traffic management into a small number of regional control centers. This centralization will require a more network-oriented architecture with remote access to local (normally unmanned) equipment rooms via Network Rail’s fixed telecommunications network (FTN). The infrastructure is expected to evolve over time, with more equipment being centralized and the core FTN being updated to use IP-based protocols rather than dedicated voice and data channels.

Adelard were asked to determine on behalf of Government whether these changes represented a high-level risk to the national infrastructure. At this stage in the upgrade programme, the exact details of the planned infrastructure changes had not yet been defined, so we provided a high level assessment of the cyber security risks associated with a generic ERTMS-based railway infrastructure.

4.1 Approach

The first step of our risk assessment was to establish the system context and agree on the scope and motivation for the assessment with stakeholders. The major system assets and services were identified in order to ensure that the risk assessment was focused on high impact scenarios. Potential threat sources were identified and attack capabilities and impact levels were defined.

The next step was to perform a preliminary risk analysis, identifying potential hazards and consequences, and relevant vulnerabilities and causes, together with any intrinsic mitigations and controls. This analysis was then refined to identify specific attack scenarios, which were prioritized according to the capabilities required and the potential consequences of the attack.

The final step was to summarize the results of the risk analysis, identify areas of uncertainty, possible mitigations and controls, and present the results of the risk assessment in the following terms:

- a set of potential attacks on an ERTMS-based system
- the capabilities needed to implement these attacks
- the worst case impact of each attack

In order to quantify the actual risk, it would be necessary to combine these results with an intelligence assessment of the likelihood of a particular threat source having the necessary capabilities to perform each attack.

4.2 System context

ERTMS is designed to be an overlay on an existing signaling infrastructure, so it is necessary to consider the underlying railway system as part of any implementation of ERTMS. Following discussions with Network Rail, we modelled the railway system as a series of layers.

Table 1 summarizes the functionality provided by each layer and the required safety integrity level (SIL).

Table 1. Railway layers

Layer	Safety Integrity Level	Functionality
Business	SIL 0	Timetable, Train Information, Operations and Maintenance
Control	SIL 2	Traffic management, Automatic Route Setting, SCADA

Layer	Safety Integrity Level	Functionality
Safety	SIL 4	ETCS (trackside and on-board), Interlocking
Communications	SIL 0	Fixed Telecommunications Network (FTN), Radio (GSM-R)
Lineside	SIL 4	Signals, Points, Train Detection

With conventional signaling systems, the safety layer is implemented solely by trackside equipment, but the introduction of in-cab signaling and automatic train protection systems such as ERTMS means that the safety layer is now partially implemented by on-board equipment. Thus, it is important to consider both trackside and on-board equipment as part of any risk assessment.

Fig. 1 provides a high-level overview of the architecture of a national railway system implemented using ERTMS. The diagram illustrates the main interactions between the various layers and system components, and the criticality of each layer (SIL 0, SIL 2, SIL 4). Since railway signaling and control is a socio-technical system, the diagram includes people as well as equipment. The main roles considered include the controller, the driver, and the system maintainers.

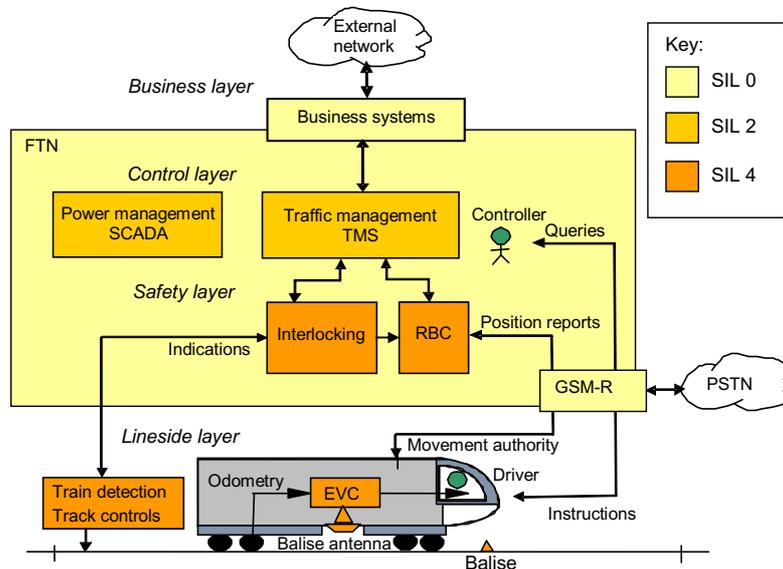


Fig. 1. Conceptual architecture of an ERTMS-enabled railway signaling system

4.3 Scope of assessment

The focus of the risk assessment was on failures of the railway signaling and control system that could have a major national impact, namely:

- attacks that result in unsafe train movements, which could cause a train accident with considerable loss of life;
- attacks that result in loss of service, which could lead to major transport disruptions.

We chose to exclude attacks that result in the theft of information because our focus was on the integrity and safety of the rail signaling and control system; loss of confidentiality is not a major concern except for some very specific attacks (e.g., on high value passengers, hazardous or high value cargoes) and the possible knock-on effect of information theft enabling future attacks on the systems.

Moreover, as this was a security risk assessment, we only considered failures resulting from the effect of deliberate attacks. We would expect failures resulting from non-malicious causes (like fallen trees, driver error, etc.) to be covered by engineering safety assessments.

4.4 Impact assessment

We assessed the impact of a successful attack on the railway system using a scale from 1 to 5, where 5 was the most serious.

Our risk assessment identified the capabilities that an attacker would need in order to achieve a high impact failure. Attacks were assessed with respect to the capability levels shown in **Table 2**.

Table 2. Attack capability levels

Capability level	Interpretation for railway systems
E	An expert in security engineering who can: <ul style="list-style-type: none">• use tools specific to the domain, which may be customized for the attacks;• develop novel equipment and tools specific to the attack;• use publicly available and proprietary information on how the system works and what mitigations are in place against attacks;• develop large test beds and trials for the attack;• coordinate timing of several attacks;• influence expert insiders.

Capability level	Interpretation for railway systems
D	<p>An expert in security engineering who can:</p> <ul style="list-style-type: none"> • use tools specific to the domain, which may be customized for the attacks; • access equipment for trials and attack development; • use publicly available and proprietary information on how the system works and what mitigations are in place against attacks; • influence knowledgeable insiders.
C	<p>Someone with a basic understanding of security engineering who can:</p> <ul style="list-style-type: none"> • use tools specific to the domain but without customization; • use publicly available information on how the system works and what mitigations are in place against attacks; • influence insiders (but at routine skill level).
B	<p>Someone with physical access to the system, for example:</p> <ul style="list-style-type: none"> • an engineer who is able to plug a maintenance console into the equipment but has no specific training or authorization to access the system in this way; • an unwitting participant, using a compromised machine or device.
A	<p>Someone without access to the system, for example:</p> <ul style="list-style-type: none"> • unskilled individuals using scripts or programs developed by others to attack computer systems and networks; • someone who has been co-opted into scaling a distributed denial of service attack; • an enterprise IT user.

Although our risk assessment was mainly concerned with cyber attacks, we also considered the effect of physical attacks on cyber assets because the infrastructure is geographically distributed and is therefore more open to such attacks. We used a similar set of criteria (skills, resources, equipment, etc.) to grade the capability needed for physical attacks on cyber equipment.

Evaluation of the likely attack frequencies and capabilities of specific threat sources is outside the assessment scope and would normally be undertaken by government agencies.

4.5 Risk analysis

In this section we describe each step of our risk analysis, which considered possible attack scenarios that could compromise railway assets to cause either:

- unsafe movements;
- no movement when it is safe to proceed.

Preliminary fault tree analysis. The initial stage of risk analysis was to construct fault trees in order to identify possible attacks on operational assets that could lead to the top events (unsafe movements and no movement). The fault trees systematically considered:

- attacks on messages sent between systems, typically by:
 - blocking transmission;
 - modifying / inserting messages;
- attacks on the systems themselves, typically via compromises of:
 - system firmware;
 - system configuration data.

The fault trees considered the effect of application-level attacks and only dealt with the consequence of these attacks, not their technical difficulty or potential impact.

Attack vectors and capabilities. The next stage of analysis was to consider what capabilities (as defined in **Table 2**) were needed to implement each attack scenario. The scale ranges from A (little skill required) to E (capabilities usually possessed only by nation states).

The preliminary risk analysis identified a number of possible attack vectors, so attack capabilities were estimated for each of these attack vectors. The primary attack vectors considered were:

- physical attacks;
- cyber intrusion;
- data preparation / installation;
- software maintenance;
- network attacks.

The estimated attack capabilities took account of the safety integrity level (SIL) of the system being attacked because we would expect the vulnerabilities and defenses to differ between SIL 0 and SIL 4 systems. However, because our analysis was based on a generic system architecture for a national implementation of an ERTMS-based signaling system, our estimates of attack capability were necessarily quite broad. For a more precise assessment, we would need to have detailed knowledge of the actual system architecture.

Attack scenarios. Using the fault trees and attack capabilities required for each attack vector, we developed a series of potential attack scenarios. Each scenario identifies the target asset, the potential attack vectors, and the capability required for the attack. These capability estimates were fairly broad to accommodate uncertainties in the security features present in the systems and the maintenance processes.

We also considered the immediate effect and the potential scale of each attack, which we used to inform the impact assessment.

Impact assessment. Our criticality scale distinguishes between loss of service and loss of life, so we make this distinction in our impact assessment.

Loss of life. It is credible that an attack that resulted in “unsafe movement” could cause an accident with 100 or more deaths in the worst case. The Eschede [3] and Amagasaki [4] train accidents exceeded this level while the Santiago de Compostela [5] accident was just below it. One could envisage multiple attacks causing multiple accidents and several hundred deaths, but it is likely that rail operators would respond to multiple accidents by shutting down the network.

However, we also need to consider the associated disruption. For a physical attack, we estimate that the disruption would be localized to a particular part of the network and would last for about a week until the physical repairs were completed. In contrast, if the accident was shown to be due to a systemic cyber security problem within the safety, communications, or lineside layer, the disruption could be far greater. To respond to a systemic cyber problem:

- all assets of the same type within the rail infrastructure would need to be assessed in order to determine if they were vulnerable to the same attack;
- operational changes would need to be put in place to minimize the risk. This would imply degraded service levels for all vulnerable parts of the network;
- systems will need to be updated and validated before normal service can be restored.

In the worst-case scenario, the resulting disruption could be nationwide and last several weeks.

Loss of service. There are many attacks that could result in a wide-scale loss of service, particularly at the business, control and communications layers.

Cyber attacks on the business and control layers (for example, attacks on the timetable or traffic management system) would be a cause for concern, but it might be easier to accept systemic vulnerabilities in these layers if the attacks could be detected and rapidly corrected (e.g., by restoring systems from secure backup storage). Given rapid system restoration, a recovery to normal service might take 1 or 2 days. However, the impact might be increased by repeat attacks if the vulnerability could not easily be addressed.

Successful physical attacks could also have a widespread effect at the business and control layers but again recovery would be fairly rapid (a few days) unless there were repeated attacks.

Loss of service could also be achieved by physical attacks on the safety and lineside layers but the effect would be localized and physical repairs would only take a few days, so the impact would be low. Repetition of attacks is possible but the impact would still be fairly low.

In practice, it is difficult to be too specific about the impact from loss of service as this depends on the resilience built into the system architecture. In particular, the impact of a cyber attack depends on the recovery process and could be reduced by switching to a fallback mechanism.

Impact vs capability summary. We combined the capabilities needed for the attacks on specific layers to obtain an overall capability range and assigned a worst-case impact based on the rationale outlined in the previous section. We then summarized our results graphically, as shown in **Fig. 2**. The lines plot the range of impact and likelihood for the different layers, attacks and impacts. The figure identifies the highest impact and lowest capability attack for each layer and shows the scope for driving the risks down by reducing the impact or increasing the capability for each attack.

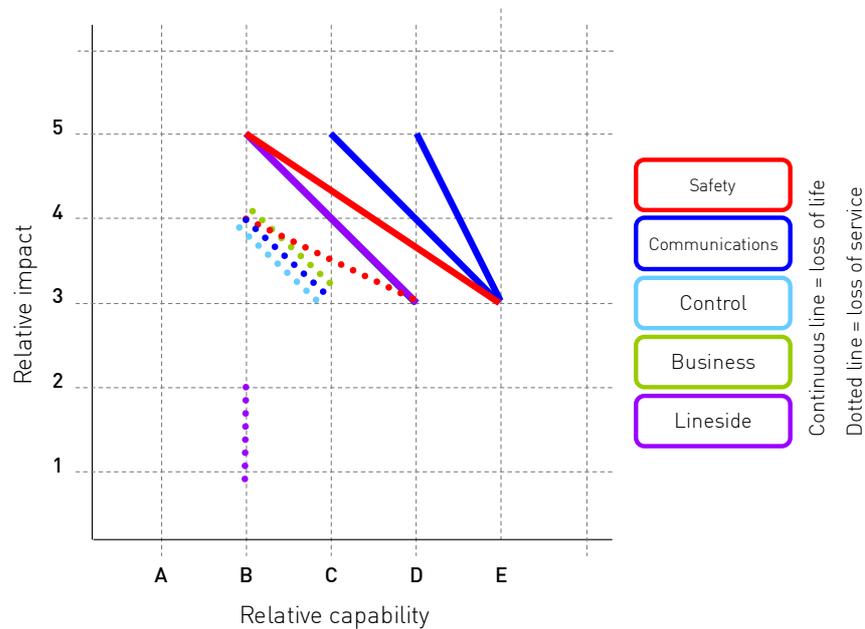


Fig. 2. Impact vs. capability diagram

Further information about the implementation would enable us to develop more precise capability estimates. Similarly, we could reduce our impact estimates if the implementation included features to limit the level of disruption caused by a successful attack.

Although our analysis identified cases where relatively low capability attacks could have a high impact, this is partly due to our uncertainty about the actual capability needed to perform cyber attacks.

The capability required for physical attacks is easier to assess and relatively modest capabilities can have quite significant effects.

For cyber attacks on the network, the capability needed at the communications layer to cause loss of life depends critically on the protection provided at the endpoint subsystems in the safety and lineside layers, which in turn depends on whether the network is considered to be open or closed. A cyber attack on the connection between the interlocking and lineside equipment is currently difficult. However, this may change as newer technology (like IP) is introduced. If the communications layer is always regarded as untrusted and the endpoints are protected, the capability needed for a successful cyber attack rises from C-E to D-E.

The other low capability-high criticality attacks relate to attacks on the data used to configure SIL 4 systems in the safety layer. Our capability B assessment is at the low end of the capability range and might be overly pessimistic.

5 Cyber security risk assessments of ETCS on-board systems

As part of the ERTMS upgrade programme, the companies that own the trains (Rolling Stock Operating Companies or ROSCOs) are in the process of tendering for ‘first-in-class’ fitments of ETCS on-board systems for each class of locomotive that will be used on ERTMS-enabled infrastructure. In the light of concerns about the security of ERTMS, Adelard and MWR InfoSecurity were commissioned by Porterbrook on behalf of the National Joint ROSCO Programme (NJRP) to provide advice and guidance on any additional security requirements that might need to be included in the contract. Adelard have expertise in risk assessment whilst MWR InfoSecurity have expertise in security testing.

Each risk assessment was informed by our generic research into ERTMS security issues, the results of a security-focused Hazop workshop that was held with the suppliers, subsequent analysis of the system by Adelard, and the results of security testing performed by MWR InfoSecurity at each supplier’s test facility.

5.1 Security-informed Hazop methodology

A series of workshops was held to study the security risks associated with each system. The workshops took the form of security-informed Hazard and Operability (Hazop) studies, and were attended by experts from each supplier.

A Hazop study is a structured approach to the identification of potential hazards and deviations from design and operating intention. The technique is qualitative, and aims to stimulate the imagination of participants to identify potential hazards and operability problems.

The study is based on the architecture of the system and involves a multi-disciplinary team of experts. Each element of the system is reviewed systematically, using a set of guidewords to prompt the experts to identify potential hazards. The experts are asked to identify

- causes of a potential malfunction
- potential consequences of the malfunction

- any system features that can detect or mitigate the malfunction
- any follow-up activities

Each study was based on a simplified architecture diagram that was intended to capture the most relevant components and interfaces of the ETCS on-board system from a cyber security perspective. Adelard created this diagram after reviewing the various documents provided by the supplier.

The goal of each Hazop study was to identify potential attacks on the ETCS on-board system that could be investigated further during the security testing, and to suggest some additional controls and assurance activities that would provide confidence that the system was protected against such attacks.

The workshops also provided an opportunity to clarify the system architecture and the test environment available for the security testing, and identify particular areas of concern to be the focus of the security testing.

The findings of each workshop were systematically recorded as a series of Hazop tables, and recommendations were categorized and numbered to ensure consistency. Each study resulted in a detailed analysis of possible attack scenarios, potential hazards, existing protections, and recommendations for additional security controls.

5.2 Security testing

In this section, we describe our general approach to security testing and the specific objectives of the security testing that was performed on each supplier's system by a team of experienced penetration testers from MWR InfoSecurity.

General approach to security testing. An ETCS on-board system can be attacked externally via interfaces that are required for ERTMS interoperability or internally via interfaces that are proprietary to the system. Attacks can be at the application level, network level, or platform level. In particular, the underlying platform might be built using commercial off-the-shelf components that contain security vulnerabilities or expose additional services that are not required for the application.

At the application level, security weaknesses in the ERTMS specifications allow a variety of attacks that are not described here for obvious reasons.

At the network level, security testing should include robustness testing of all the major interfaces, both external and internal, in order to probe whether the system is robust against deliberately crafted messages that pass the integrity checks but are invalid at the application level.

Testing should also challenge closed network assumptions. This requires investigating the security of the network used to connect together components of the ETCS on-board system and assessing the damage that could be done to the system by an attacker with access to these networks.

Security testing objectives. The overall goal of the security testing is to explore this range of attack vectors and determine whether any of the attacks are feasible. In prac-

tice, depending on the test environment, it may not be possible to perform the full range of tests, so the aim is to achieve broad coverage of the possible attack vectors.

More specifically, the security testing objectives can be broken down as follows:

- explore the feasibility of attacks allowed by the ERTMS specifications and discover whether the driver receives any notification if something unexpected happens;
- determine whether the ETCS implementation is robust against malformed messages or whether it is possible to crash the system or cause it to behave in an arbitrary way;
- investigate whether the closed network assumption is valid and determine what damage could be done by an attacker with access to the network and some inside knowledge;
- perform a security audit of the underlying platform and any third-party components.

Some of the test results exposed anomalies or ambiguities in the ERTMS specifications. Although these anomalies do not raise any safety or security concerns, it is important to resolve any ambiguities in the ERTMS specifications in order to remove the potential for an attacker to exploit differences in behavior between implementations.

5.3 Recommendations

Our final set of recommendations were divided into four categories:

1. technical or procedural controls that would improve the security of the ETCS on-board system;
2. assurance activities to improve confidence in the security of the ETCS on-board system;
3. recommendations for the national implementation of ERTMS;
4. suggested changes to the ERTMS specifications.

Unfortunately, we cannot publish any of our recommendations here because they implicitly identify potential vulnerabilities in the systems.

6 Discussion / lessons learned

6.1 Context

There is a growing awareness that safety and security can no longer be considered in isolation and that a system cannot be considered to be safe unless it has also been shown to be secure. However, there is currently a lack of underpinning analysis to demonstrate how and whether cyber security issues can be integrated in to hazard and risk analyses, and hence a lack of consensus about the best way to integrate safety and security. In particular, there are no clear guidelines about methodology, and standards

in this area are still evolving. As a result of the work on security-informed safety that Adelard and others have been doing in the railway industry, this situation is changing within the UK. The Department of Transport has recently published guidance on Rail Cyber Security [6] and commissioned work to develop a code of practice for the railway industry on how best to develop security-informed safety cases. Adelard has been active in this area and worked with the Railway Safety and Standards Board (RSSB) to develop a security-informed safety case as an exemplar for the railway industry. However, security-informed safety is not just a concern for the railway industry – Adelard was a partner in the SESAMO project [7], which was concerned with security and safety modelling for embedded systems across a wide range of industrial domains, including avionics, automotive, industrial control, medical, smart grid as well as rail. There is now a much greater awareness of the need to consider cyber security in the design of safety-critical systems, and the focus has shifted from raising awareness to developing guidance, standards and worked examples.

6.2 Strategy

Adopting a phased approach towards cyber security assessment has proved to be an effective strategy. We started by performing a security audit of the ERTMS specifications, which enabled us to identify a number of systemic vulnerabilities in the specification and potential areas of concern. These were refined by developing specific attack scenarios, which proved to be an excellent way of communicating and engaging with railway stakeholders because the attacks became real rather than theoretical and abstract.

The next stage was to conduct a high-level risk assessment of a national implementation of ERTMS, which was used to inform the national risk register. Focusing on the potential risks at a national level gave our risk assessment a sense of proportionality and perspective.

In practice, the worst-case impact in terms of loss of life or loss of service depends on many implementation factors (including provisions for resilience) that have not yet determined at this stage in the upgrade programme. Thus, our assessment will need to be revisited as the upgrade progresses and more operational experience is gained. However, we believe that our main findings are robust.

Our risk assessment of a national ERTMS implementation was based on a generic system architecture with little specific information about the vulnerabilities and defenses that might exist in the actual system. In contrast, our risk assessment of ETCS on-board systems from each of the major suppliers looked at real systems in detail and took into account the results of security testing and vulnerabilities discovered in the configuration of each system. These assessments were performed on behalf of the rolling stock operating companies (ROSCOs), who wished to purchase ETCS on-board systems for new and existing trains and needed to have some reassurance that their assets would be robust against cyber attack. The results of the assessments were used to inform the procurement process for the ‘first-in-class’ fitment programme to install ETCS on each class of locomotive, and the recommendations from each assessment were written into the contract with each supplier. The assess-

ments were beneficial to both the purchasers and the suppliers because they enabled the purchasers to reduce their risk whilst providing guidance to the suppliers on how to improve the security of their products.

6.3 Where next?

Over the last few years, government and industry have been mobilizing and commissioning research and support for developing cyber security strategies and guidance and there is now a plethora of groups working in this area. It is important to develop a coherent strategy that clearly identifies roles and responsibilities at different levels of governance (project, industry, government) and identifies gaps where further research and development of standards and guidance is necessary.

Railway-specific issues. In the railway context, management of cyber risk is complicated by the divided responsibilities for maintaining safety in an ERTMS-based signaling system. Responsibility for the safety layer is split between the trackside and the train, which are owned and managed by different organizations. Security needs to be embedded in the processes used by all stakeholders in order to maintain the overall safety and integrity of the signaling system.

Another complicating factor is the widespread use of legacy systems that were designed in a different age to protect against different threats. Closed network assumptions are no longer valid but it is not always possible to add security features to legacy systems, so alternative approaches are needed.

At a more general level, we need to consider if there is adequate oversight for the introduction and operation of new technology like ERTMS and whether there are sufficient technical resources available to the regulator.

Incident reporting. It is important to ensure that we can learn from incidents, so that safety issues with the new technology can be identified and rectified. Ideally, incident reporting should be undertaken by all ERTMS users and suppliers. We recommend the introduction of policies for the collection, analysis and sharing of cyber incident information, even when such incidents have no safety impact.

Resilience requirements. There is currently a lack of any clear definition of resilience requirements from a policy perspective. While safety is governed by existing legislation, there do not appear to be any system level resilience requirements. Governance and business models should be established to ensure that sufficient resilience is provided by the system as a whole. Incentives may need to be provided for diversity that is justified from wider societal considerations rather than from an infrastructure owner's business case.

Secure by design. There is also clearly a need for industry guidance on methodology and guidance for developing and assessing systems that are intended to be both safe

and secure. Suppliers need guidance on how to build security into their products, and purchasers need to be informed about cyber security and be given tools to help them assess whether a product is adequately secure for its intended use. This is particularly important during the procurement phase of a large railway project, where there is an opportunity to influence both the generic product and a specific application of the product to the GB context.

There are already a number of sources of guidance available, including:

- 20 Critical Controls for Effective Cyber Security Defense;
- DHS Cyber Security Procurement Language;
- Trustworthy Software Initiative;
- Cyber Essentials;
- BSIMM, OWASP, Microsoft SDL, SafeCode;
- Common Criteria.

These need to be customized and adapted for the railway sector.

Standards and legislation. In Europe, any significant changes to a mainline railway system must be assessed in accordance with the Common Safety Method. This is a legal requirement. Similarly, the ERTMS specifications form part of the Technical Specification for Interoperability, which is mandated by European law. This makes it difficult for GB concerns about cyber security within the railway industry to be addressed at a national level, and makes it necessary to engage at a European level to influence the development of these standards to ensure that they include adequate provision and protection against cyber attacks.

Risk and uncertainty. A risk assessment should be a living document and needs to be revisited periodically during the system life cycle. Risks can change during the development of the system and also during its operation, so it is important to understand the risks and the mitigations in place at every stage of the life cycle. This is particularly true for risks arising from cyber security threats – security decays faster than safety.

Our risk assessments of ETCS on-board systems were assessments of real systems and were performed with the benefit of detailed design documents, access to system experts, and the opportunity to perform security testing on the actual system to determine whether potential vulnerabilities existed in reality and could be exploited. The systems were still under development but the manufacturers were receptive to our recommendations and willing to incorporate changes into the design of their systems to make them more robust and resilient against cyber attack.

In contrast, our risk assessment of a national ERTMS implementation was performed at an early stage in the upgrade programme, and was therefore based on a generic system architecture. As a result, there is significant uncertainty in the results and it is therefore important to revisit the assessment as more implementation detail is provided and more operational experience is gained. An updated risk assessment would need to address:

- the impact of the differing responsibilities of the multiple stakeholders (operators, leasing companies and the supply chain) for safety management and hence cyber risk;
- the susceptibility of data preparation and maintenance processes to cyber attack;
- the extent to which the overall system architecture is designed to limit cyber attack as the system evolves (e.g., when there are changes in network technology);
- the resilience and recovery from cyber attack provided by fallback options (both in fixed infrastructure and on board the train);
- the co-operation and security culture of the stakeholders.

7 Conclusions

The next generation of railway signaling and control systems will potentially have more risk and less resilience than the current generation of systems due to security vulnerabilities and increased connectivity. However, this increased connectivity means that the new systems could potentially be engineered with stronger controls, greater defense in depth, and improved recovery mechanisms, thus eventually presenting less risk overall and providing greater resilience. The risk assessments presented in this paper are one contribution to ensuring that this is the case.

Acknowledgements. We are grateful to our sponsors for their permission to publish this summary of our work over the last five years. We would also like to acknowledge the contribution of Richard Bloomfield and Ilir Gashi to our initial analysis of the ERTMS specifications.

References

1. Bloomfield, R., Bloomfield, R., Gashi, I., Stroud, R., How secure is ERTMS? In: Ortmeier, F., Daniel, P. (eds.) Computer Safety, Reliability, and Security Proc. SAFECOMP 2012 Workshops: Sassur, ASCoMS, DESEC4LCCI, ERCIM/EWICS, IWDE, Magdeburg, Germany, September 25-28, 2012. LNCS, vol. 7613, pp. 247-258. Springer-Verlag Berlin Heidelberg (2012)
2. Network Rail, Strategic Business plan for 2014/2019, January 2013.
3. Wikipedia, Eschede train disaster. http://en.wikipedia.org/wiki/Eschede_train_disaster
4. Wikipedia, Amagasaki rail crash. http://en.wikipedia.org/wiki/Amagasaki_rail_crash
5. Wikipedia, Santiago de Compostela derailment. http://en.wikipedia.org/wiki/Santiago_de_Compostela_derailment
6. Department for Transport, Rail Cyber Security, Guidance to Industry, February 2016. <http://www.rssb.co.uk/Library/improving-industry-performance/2016-02-cyber-security-rail-cyber-security-guidance-to-industry.pdf>
7. SESAMO – Security and Safety Modelling, ARTEMIS Embedded Computing Systems Initiative 2011, Project Number 295354, May 2012.