

Architectures for Cyber-Security Incident Reporting in Safety-Critical Systems

Chris W. Johnson,

School of Computing Science, University of Glasgow, Glasgow, UK, G12 8RZ.

Keywords: Security, Incident Reporting, Safety-Management Systems, Security Management Systems.

Abstract

Cyber-attacks can have a devastating impact on safety-critical systems. The increasing reliance on mass market Commercial Off-The Shelf (COTS) infrastructures, including Linux and the IP stack, have created vulnerabilities in applications ranging from Air Traffic Management through to Railway signalling and Maritime surveillance. Once a system has been attacked, it is impossible to demonstrate that malware has been completely eradicated from a safety-related network. For instance, recent generations of malware use zero day exploits and process injection with command and control server architectures to circumvent existing firewalls and monitoring software. This creates enormous problems for regulators who must determine whether or not it is acceptably safe to resume operations. It is, therefore, important that we learn as much as possible from previous cyber-attacks without disclosing information that might encourage future attacks. This paper describes different architectures for encouraging the exchange of lessons learned from security incidents in safety-critical applications.

Introduction

Incident reporting has been widely recognised as a key component in many safety management systems [1]. Information about adverse events helps to warn others of potential hazards. Incident reports can also be used to disseminate the recommendations that help prevent any recurrence of previous mishaps. They also help to promote the mitigation and recovery techniques that increase our resilience to hazards that cannot be avoided. In other words, as shown in Figure 1, incident reports help to validate the likelihood and consequence assessments that drive risk analysis.

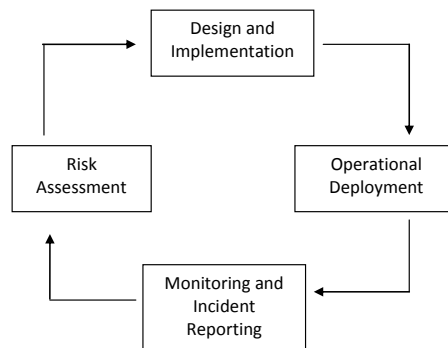


Figure 1 – Incident Reporting within the Safety Management Lifecycle

International bodies have, therefore, advocated the use of incident reporting in safety-critical applications:

"(The assembly) urges all Contracting States to ensure that their aircraft operators, providers of air navigation services and equipment, and maintenance organisations have the necessary procedures and policies for voluntary reporting of events that could affect aviation safety" (ICAO Resolution A32-15: ICAO Global Aviation Safety Plan)

International support for voluntary incident reporting systems includes near-misses. Systems that only exchange information about previous adverse events are reactive, whereas near-miss reporting schemes help to identify potential hazards before they occur:

"Companies should investigate near-misses as a regulatory requirement under the Hazardous Occurrences... Aside from the fact that near-miss reporting is a requirement, it also makes good business and economic sense because it can improve vessel and crew performance and, in many cases, reduce

costs. Investigating near-misses is an integral component of continuous improvement in safety management systems.” (International Maritime Organisation, Guidance on Near-Miss Reporting MSC-MEPC.7/Circ.7)

These initiatives have resulted in a proliferation of safety-related incident reporting tools and techniques, including but not limited to, the Australian Incident Monitoring System and Confidential Safety Reporting Information Scheme, the Canadian National Defence General Accident Information System, the European Space Agency Alert System and European Major Hazard Incidents Data Service (MHIDAS), the Japanese Maritime Incident Reporting System and Rail Accident Method, the US NTSB Aviation Safety Reporting System, FDA Adverse Event Reporting System and Manufacturer and User Facility Device Experience database (MAUDE), FRA Confidential Close Call Reporting System (C3RS), the UK Confidential Incident Reporting System (CIRS), Confidential Human Factors Incident Reporting Programme (CHIRP) and Confidential Incident Reporting and Analysis System (CIRAS). There are also generic reporting tools including Data Reporting Analysis and Corrective Action Systems (DRACAS), Failure Reporting, Analysis and Corrective Actions systems (FRACAS), Prevention and Recovery Information System for Monitoring and Analysis (PRISMA) and PRISMA-Rail, Rail-Program for Risk Informed Safety Managements, Safety Management Information System, Technique for the Retrospective and Predictive Analysis of Cognitive Errors etc.




Figure 2 – Incident Reporting within the ENISA Key Security Governance Processes [2]

Given the proliferation of incident reporting within safety-critical applications, it is no surprise that organisations, including the US Department of Homeland Security as well as the European Network and Information Security Agency (ENISA) have promoted similar schemes to track cyber-security concerns. For example, Figure 2 shows how ENISA include incident reporting within their key processes for the governance of security concerns. Risk assessments help to identify the potential targets of an attack and to determine whether there are known vulnerabilities. Security measures are then taken to protect those targets and to ensure that measures continue to be implemented over time. Collecting incident reports helps to understand “weaknesses in security measures and to evaluate and validate the risk assessment”. This triangle is, typically, supervised by a government agency, such as a regulator, or by an industry association, including groups of professional auditors [2]. As mentioned, the US DHS advocates similar arrangements through the National Cyber security and Communications Integration Center (NCCIC). The NCCIC coordinates the information collected through incident reporting to improve situation awareness for cyber communities across government and the private sector. Similarly, the US Computer Emergency Response Team (US-CERT) provides direct operational advice on the reporting of security incidents, based on the NIST guidelines for incident reporting [3]. These initiatives have been supported by a growing number of tools and techniques that provide means of reporting security incidents. These include commercial and open source tools such as AbuseHelper; Application for Incident Response Teams (AIRT); Assuria Auditor and Request Tracker for Incident Response (RTIR). Most national Computer Emergency Response Teams have reporting applications, including those of the US and UK CERTs. Professional and industry groups have also coordinated security incident reporting, these include the Forum of Incident Response and Security Teams (FIRST) as well as the industry bodies supported within the UK Centre for the Protection of National Infrastructures Warning, Advice and Reporting Points (WARPs) programme. There are semi-automated security incident reporting tools, such as the US Air Force Automated Security Incident Measurement (ASIM) infrastructure and the new Einstein programme for collecting, analysing, and sharing computer security information across the US Federal Civilian Government. Other reporting systems support particular sectors, such as the US Federal Communications Commission’s Disaster Information Reporting System (DIRS) and Network Outage Reporting System (NORS).

Unfortunately, there have been very few attempts to integrate the reporting of safety and security incidents even though it is clear that cyber-attacks can have a profound impact on the safety of most complex systems [4, 5]. The increasing reliance on mass market Commercial Off-The Shelf (COTS) infrastructures, including Linux and the IP stack, have created vulnerabilities in applications ranging from Air Traffic Management through to

Railway signalling and Maritime surveillance. Once malware has infected a complex system, it is impossible to predict the potential impact on safety requirements. We cannot assume that the development techniques used in creating a virus or Trojan would meet the strict requirements of an industry regulator within safety-critical industries! Although many papers have been written about the impact that COTS software might have on meeting Safety Integrity Levels or Software Assurance Levels, very few have considered the implications of malware. It would be difficult to guarantee that critical processes continue to receive necessary network or processing resources without a sustained forensic analysis of the malware. This creates further problems given the length of time required to conduct such studies. In the short term, we cannot keep aircraft circling while we determine whether or not an Air Traffic system can safely be used to guide their descent. Beyond that, it is difficult to contemplate the business consequences of closing air space for the length of time it might take to convince a regulator that an infrastructure is safe to resume operations. It is impossible to demonstrate that malware has been completely eradicated from a safety-related network. For instance, recent generations of attack use zero day exploits and process injection with command and control server architectures to circumvent existing firewalls and monitoring software. This creates enormous problems for regulators. It is, therefore, important that we learn as much as possible from previous cyber-attacks without disclosing information that might encourage future attacks. The following pages describe integrated architectures for encouraging the exchange of lessons learned from security incidents in safety-critical applications.

Internal Reporting Architectures

Figure 3 represents one of the simplest architectures for an incident reporting system; in this case the focus is on a safety-related application. A contributor submits a report based on the occurrence that they have witnessed or are concerned about. This submission process can be implemented using printed forms, by telephone calls, or increasingly using computer-based techniques. In some cases, automated systems can detect adverse or near miss events that may subsequently prompt further investigation – for instance, a Short Term Conflict Alert (STCA) in Air Traffic Management. An investigator is then required to gather further evidence, including system logs and witness statements. These are then used to map out the events leading to an incident. The reconstruction supports more detailed studies of the causal and contributory factors. From there it is possible to identify those actions, which are intended to reduce the likelihood or mitigate the consequences of any recurrence. The previous stages of the analysis are then documented and distributed to stakeholders so that corrective actions can be implemented.

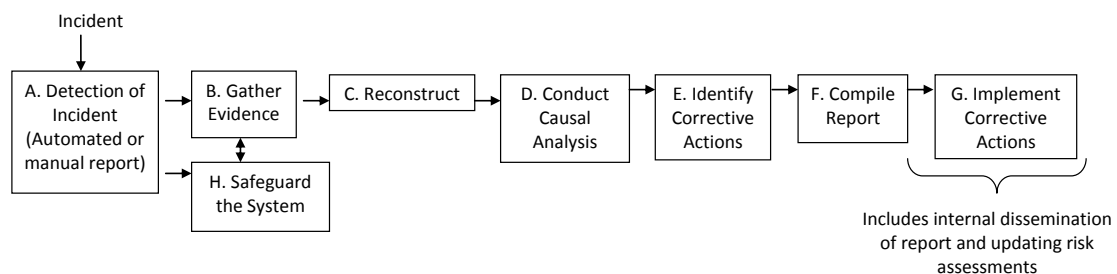


Figure 3: An Internal Incident Reporting Architecture

This process is a simplification because it assumes that all incidents will be analysed to the same level of detail. In practice, there is typically a preliminary risk assessment after initial evidence has been obtained. Most reporting systems lack the resources to conduct detailed causal analyses for all adverse events and near miss incidents. In consequence, only those events with a higher risk of recurrence will go through all of the stages illustrated in Figure 3.

The architecture illustrated in Figure 3 can also provide a template for security incident reporting systems. Many of the concerns in the implementation of such architectures are the same as they would be in safety-related industries. For instance, if the definition of a reportable incident is set too low then scarce resources will be wasted as analysts investigate thousands of false positives. For example, supervisors' time can be wasted by STCA alerts if the system is configured to generate alarms that are within the bounds of normal, safe operation. Similarly, automated network monitoring systems will detect adverse events even when there is no threat to security if they have not been correctly configured for normal traffic patterns.

A number of problems complicate the use of this simplified safety incident reporting architecture for cyber-security concerns. Firstly, it is often more difficult to detect cyber-attacks than it is to identify safety-related incidents. Many security threats take elaborate measures to hide within a network. This is another reason for

the integration of security and safety reporting systems, given that malware will often show the same symptoms as more routine bugs or system failures. For example, engineering teams often become suspicious when network monitoring tools identify unexpected transmissions or when memory/processing resources seem to be compromised. These concerns could be triggered by malware or by routine configuration problems and it is often impossible to know the cause without more sustained analysis.

Thirdly, stages B to E are characterised as ‘forensic analyses within security management systems [6, 7]. This raises a host of concerns that are not, typically, considered within safety-related incident reporting systems. For example, the systems and networks that are affected by a suspected cyber-attack can be considered a crime scene and evidence must be preserved according to legal principles and guidelines. It will be necessary to uncover normal, hidden, deleted, encrypted and password-protected files to gain as much information as possible about the nature and scope of any attack. Further problems arise because the tools and techniques that support the causal analysis of cyber-attacks lags many years behind those available to accident and incident investigators in other domains. Further work is required to determine whether the existing application of root cause analysis techniques, including those using counter-factual reason and systemic models, can be extended to support the reporting of cyber-attacks [1].

Further differences arise at the end of the reporting chain with the drafting and dissemination of lessons learned. In safety-related systems there is usually a presumption that as many stakeholders as possible should be informed of any lessons learned; even if the presentation of those lessons may be tailored to particular audiences using newsletters, technical reports, daily briefing documents etc. In the aftermath of security related incidents, there is a concern that any subsequent dissemination should not undermine the future security of an application process. In some cases, disclosing that an attack has been identified will itself provide adversaries with important information on ways to refine future cyber threats.

There are a number of limitations with the simple reporting architecture shown in Figure 3. In particular, there are no guarantees that a company will take any corrective actions or that the actions implemented in stage G will address the underlying causes of an incident. Similarly, there is a danger that different organisations will respond in different ways to similar incidents across the same industry. This inconsistency creates the opportunity for future failures if an organisation fails to correctly safeguard the system. Similarly, in security reporting systems there is a concern that a known vulnerability would only be patched by the company suffering an attack and that any other critical infrastructures would remain exposed. A further problem is that there is no external validation of an incident report. This creates concerns that a lack of technical expertise or problems of political bias might undermine the response to previous incidents. These limitations are addressed by an increasing role for external agencies, including professional bodies and industry associations, in the following reporting architectures.

Gatekeeper Architecture

Figure 4 illustrates a more elaborate architecture for reporting adverse events. This model explicitly represents different agents within the scheme. As can be seen, reports are generated by a host of sources from both inside or outside an organization. These are forwarded to a supervisor or ‘Gatekeeper’ who gathers the initial evidence in the aftermath of an adverse event. They will then conduct an initial risk assessment to determine whether the mishap warrants a full investigation. The term ‘gatekeeper’ is used because this individual plays a key role in determining the focus of subsequent investigations. If an incident is identified as a high risk for any recurrence then stages F through to J follow those in the simplified architecture of Figure 3. Otherwise, only a summary report is developed, however, the supervisor may also be required to explicitly document the reasons why it was NOT investigated. Several of these low risk incident summaries can be compiled, for instance every six months. The collated documents can then be analysed for underlying safety or security trends. In other words, several low risk or near miss incidents might collectively justify a more sustained analysis than any individual incident.

Figure 4 further extends the simplified architecture of Figure 3 by considering the reporting chain for adverse events within an organisation. The supervisor or gatekeeper controls the day to day running of the system. However, an internal security or safety management group provides strategic and tactical oversight. In the case of a high-risk incident, they are immediately informed and may, in turn, choose to notify external agencies of a significant threat to safety or security. In addition, they are responsible for monitoring the implementation of corrective actions taken both in the aftermath of high risk incidents and also to resolve common concerns amongst the periodic reviews of low risk or near miss events. In other systems, the management group may

have a more direct role in approving or rejecting recommendations – when, for instance, recommendations require major sustained investments or management commitment at more senior levels than the supervisor/gatekeeper.

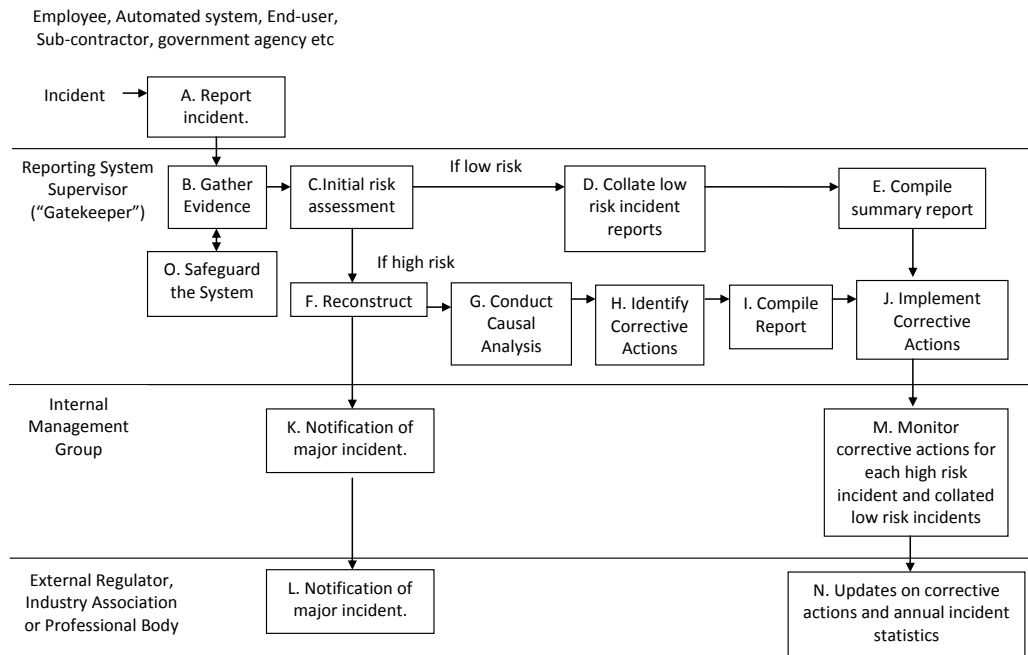


Figure 4: Gatekeeper Reporting Architecture

This revised architecture also shows how safety or security management groups provide an interface with external agencies including regulators, professional bodies or industry associations. In safety-critical applications, this provides important mechanisms for ensuring that incidents are not repeated across an industry. Even in safety-related areas, this raises a host of concerns. Some companies are reluctant to provide their competitors with a commercial advantage by providing information about lessons learned; even though this is indicative of a poor safety culture. Other organisations can be more worried about a potential loss of reputation or of market confidence. There are also concerns that information about previous failures, for instance involving management decisions or employee actions, might provide the basis for subsequent litigation. In consequence, many reporting systems support the submission of confidential and anonymous reports. This can reduce the utility of the lessons that are provided when readers cannot obtain details about the context in which an incident or near miss occurred. These concerns are exacerbated when incidents relate to the security of complex systems. There is often a reluctance to provide information outside of the immediate organisation suffering the attack. Concerns focus on the third party release of details that might further undermine security or public confidence in the aftermath of a cyber-incident. Without legal guarantees, most commercial organisations will only provide anonymous summaries of minor adverse events to these agencies – in some cases; even this would be refused unless there are reciprocal benefits from the further exchange of security information by competitors.

Further problems affect the extension of this approach to capture both safety and security related incidents. In particular, the success or failure of the system depends on the skills and expertise of the supervisor or gatekeeper. If they decide that an incident does not merit further analysis then the management committee will only see a summary report. Even if the committee then decided that the adverse event/near miss required further investigation, many organisations would only be able to retrieve part of the necessary forensic evidence. The focus on the gatekeeper is even more critical because there is little practical guidance about how best to assess the risks of potential cyber-threats. To illustrate the dilemmas, a European Air Traffic Management service provider recently detected a problem with the local area network that integrated radar and flight plan data. The problem caused an intermittent degradation in quality of service. The systems engineering supervisor took the decision to investigate the incident further but could not find the cause. The symptoms then disappeared. At this stage, the supervisor has to make a decision – it might have been a non-malicious bug in the network management system or an intermittent hardware fault or a result of interactions between the thousands of applications that exchanged data over the infrastructures. Alternatively, the loss of service might have been the first symptom of a cyber-attack. The supervisor had limited resources and was in the middle of a periodic

software upgrade on another application and decided not to investigate any further. Some weeks later, the symptoms recurred and the causes were traced to a keystroke logger running on a Linux installation that was not specifically focussed on the ATM service provider. It is possible to criticise the supervisor’s decision, however, the site had no specialist expertise in cyber-security and they had never before experience malware in an operational system. As mentioned before, they also lacked any formal tools to help them decide whether or not the initial symptoms should have triggered a deeper investigation, given that the malware had several weeks to operate without being detected inside the organisation’s firewalls.

The gatekeeper architecture also suffers from increased complexity. Individuals and teams will only remain motivated to contribute information about safety or security incidents if they feel that their concerns are being addressed. It can be hard to them to follow the progress of a particular incident report through the various stages of causal and forensic analysis. Similarly, they may be frustrated if and their concerns are classified as ‘low risk’ and do not trigger more detailed investigations. Some companies have addressed these issues through the introduction of incident tracking systems so that reporters can trace each action being taken in response to the safety or security issues that they raise. This also enables the internal management to review any open corrective actions that have still to be implemented following a major incident. Several of these systems were listed in the opening sections of this paper. Very few of them have been extended to support security management, hence when individuals do report concerns over violations of security policy they often report that little seems to have changed [8]. This undermines both the long term future of the reporting system and the utility of any immediate lessons that might have been drawn from a particular concern.

Active External Monitoring Architectures

The reporting architectures illustrated in this paper are deliberately intended to reflect different levels of safety or security maturity. The simple system in Figure 3 focuses on the internal dissemination of information. The more elaborate gatekeeper architecture assumes that the reporting organisation has sufficient confidence and legal protection to share lessons learned with external regulators, industry associations or other professional bodies. However, figure 4 assumes a relatively passive role for external oversight. In contrast, active monitoring architectures provide for additional support from external agencies in the investigation of adverse events. As can be seen in Figure 5, external bodies are notified after a high risk incident has been detected. However, in contrast to the earlier models regulators, industry associations or other professional bodies offer assistance both in safeguarding the system and in investigating the incident. As can be seen, it is assumed that this more active support would only be appropriate for high risk incidents or in special circumstances, for example where a company lacked specialist expertise in forensic analysis. As might be expected, it takes a higher level of trust and safety/security maturity to encourage this level of participation from external agencies – or in other cases, this level of involvement may be the consequence of specific legislation to ensure the protection of critical infrastructures. As before, there may be considerable concern to ensure that any industry feedback in stage M does not compromise commercially sensitive information or the public reputation of a company participating in the scheme. Great care must also be taken to ensure that other companies can use any lessons learned in a report while at the time protecting the future security of safety-related applications.

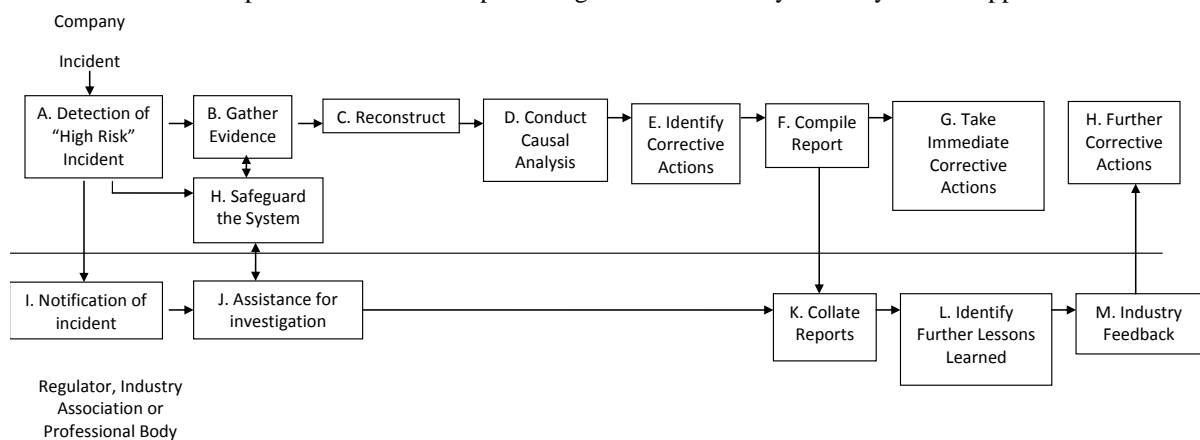


Figure 5: Active External Monitoring Architecture

As with previous architectures for incident reporting, a number of factors complicate the use of active external monitoring as an integrated approach to safety and security management. At present, many industries suffer from artificial barriers or silos between the different external agencies that address the problems of safety and

security. In the United States, the Government Accountability Office has issued a series of reports over the last twelve months with titles that include: “A Better Defined and Implemented National Strategy Is Needed to Address Persistent Challenges” [9] and “Cyber security: National Strategy, Roles, and Responsibilities Need to Be Better Defined and More Effectively Implemented” [10]. Part of the confusion arises because safety has, traditionally, been devolved to organisations with a specific focus on particular safety-critical industries – these include the FDA, FRA, FAA with cross-sector organisations looking after more general forms of occupational health and safety, through OSHA. However, cyber-security has been seen as a cross-cutting concern that requires specific expertise. Governments have, therefore, created distinct agencies to deal with these threats. In particular, the US Federal Information Security Management Act (2002) provides the wider context for this paper. FISMA requires that Federal agencies have “procedures for detecting, reporting, and responding to security incidents”. The US National Institute of Standards and Technology (NIST) coordinate the technical implementation of FISMA, with operational leadership from the Department of Homeland Security and the US CERT, mentioned in previous sections. These divisions create a dangerous situation where cyber-security agencies have almost no understanding of the impact that malware could have on the technical operation of safety-critical systems. Conversely, the regulatory agencies established to monitor the implementation of safety standards have almost no expertise in cyber-security; many suffer from a long legacy of physical security specialists whose talents provide little help in mitigating new generations of advanced persistent threats.

In Europe, the legislative context is set by EU Directive 2009/140/EC. A key element within the directive has become known as ‘Article 13a’ on the security and integrity of public communication networks. Paragraphs 1 and 2 of Article 13a require service providers to ensure the security and integrity of their networks and to ensure continuity of service. Paragraph 3 requires that service providers report significant security breaches and losses of integrity to national regulatory agencies. They must then forward summaries to the European Network and Information Security Agency (ENISA). EU Directive 2009/140/EC focusses on the resilience of operators irrespective of whether their services are being used in safety-critical infrastructures or in mass market applications. However, the European Commission has recently proposed the extension of obligatory reporting requirements as part of the European Union’s 2013 Cyber-Security Strategy. Again significant work remains to be done before common incident reporting structures can be established for safety concerns. For example, the European Aviation Safety Agency has traditionally avoided any consideration of cyber-threats, even though they can have a considerable impact on the operation of complex, critical infrastructures. ENISA lacks specific expertise in the aviation domain. This creates a situation where companies lack clear guidance – for example on how best to respond using remaining airborne and ground systems when malware is detected within the Flight Data Processing or Surveillance systems of an Air Navigation Service Provider [4].

One partial solution is to draft letters of agreement between safety and security regulators to clarify responsibilities and establish an agenda for future cooperation within national and international programmes for critical infrastructure protection. Alternatively, professional bodies and industry associated can support the integrated reporting of safety and security concerns. This raises questions about whether regulatory agencies would interpret the use of these systems as acceptable means of compliance with legal reporting requirements across different industries. Other questions relate to the funding of incident reporting systems through professional organisations. In some of the smaller European member states, one or two companies compete in a limited market. In can be difficult to justify funding more complex monitoring mechanisms. In other countries, existing industry associations lack the organisational and technical expertise to support such an enterprise.

It seems likely that these organisational barriers will be resolved through political and organisational changes over the next decade. It remains to be seen whether the necessary changes can be completed before lives are lost across national critical infrastructures. In the meantime companies face a host of practical challenges in integrating a unified approach to incident reporting using the active monitoring architectures illustrated in Figure 5. Previous sections have described an incident in which Air Traffic Management engineers were initially alerted to potential malware through intermittent delays in data passed across a local area network. It took several days to determine whether this was due to a conventional safety-related concern, covered by reporting infrastructures under EASA, or whether there was a potential cyber-attack, reportable under the separate extensions to Article 13a cited above. Without a more unified approach, it is very difficult for companies to know how to obtain the “assistance for investigation” envisaged in stage J of the active monitoring architecture.

Joint Public-Private Architecture for Cyber-Safety

Figure 6 illustrates an integrated architecture for the reporting of potential cyber-incidents in safety-critical infrastructures. It focuses on a joint public-private approach based on cooperation between industry and

government with implicit mechanisms for cost sharing. As can be seen, it builds on the previous architectures. However, it also assumes the creation of a Joint Monitoring Group for cyber-security incidents. This is intended to represent a wide range of stakeholders but, in particular, safety regulators from a range of different industries as well as various government security agencies, including national CERTs. The Joint Monitoring Group should also include companies with the technical expertise required to both provide advice and help disseminate the lessons learned from cyber-incidents in national critical infrastructures. Given the increasing range of novel threats to the security of complex systems, the monitoring group would help to focus the more general assistance that is available through CERTs and other government bodies, which today have little experience or knowledge of safety-critical software engineering standards. This is reflected in the active role that the monitoring group plays in stage Q ‘Safeguarding the System’. It is intended that a company would only notify the joint monitoring group if they had initial evidence that a safety-related application might have been the target of a cyber-attack.

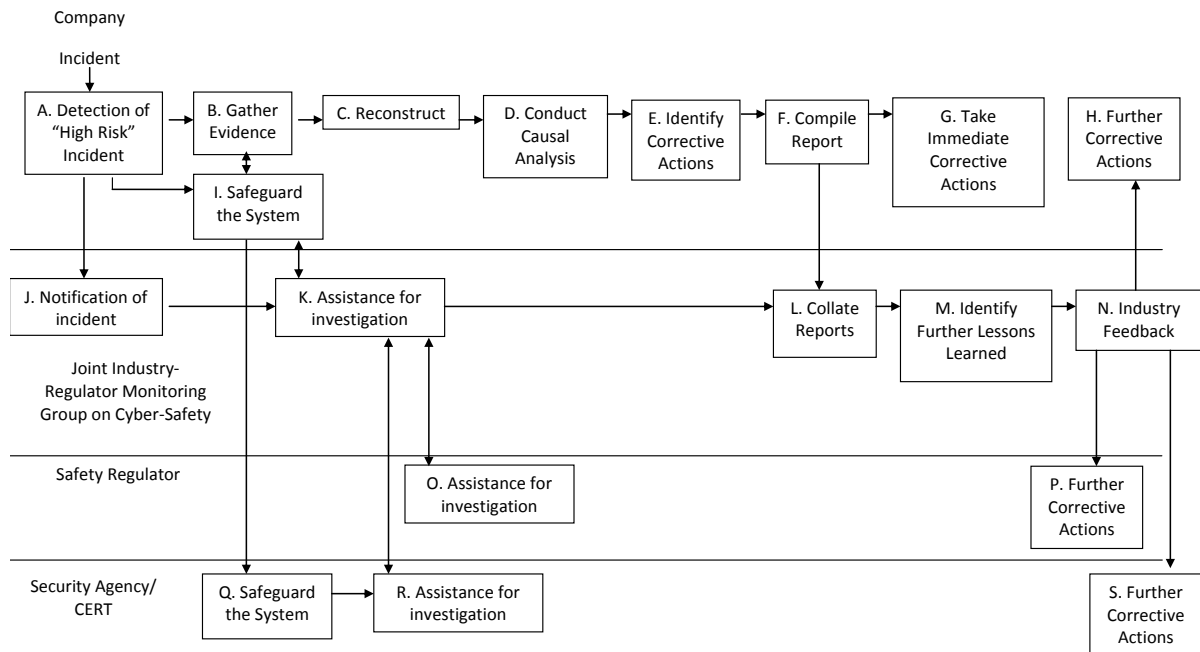


Figure 6: Joint Public-Private Coordination of Cyber-Safety Reporting Systems

All of the reporting architectures illustrated in this paper represent compromises. Internal reporting systems based on the Figure 3 architecture provide simple, low cost approaches to ensure that lessons are disseminated inside a company. However, they provide limited support across an industry or in alerting other critical infrastructures to the potential threat from future cyber-attacks. At the other end of the complexity spectrum, Figure 6 represents an elaborate approach in which multiple public bodies and private companies work together to increase resilience by sharing investigatory resources and lessons learned across critical infrastructures. They assume significant input from safety regulators and a continuing commitment from commercial participants to invest resources of time and expertise to support other companies when attacks occur.

Conclusions and Further Work

This paper presents a number of architectures that can be used to support the development of incident reporting systems in safety-critical industries. The aim has been to identify ways of integrating the response to both safety hazards and cyber-attacks. For example, we have identified the stages that are typically involved in the development of simply internal reporting systems used within an individual organisation. It is necessary to secure sufficient evidence to reconstruct an adverse event, conduct a causal analysis, identify corrective actions etc. Many of these stages that were originally identified within safety-critical systems have their parallels in the forensic investigations of cyber-incident reporting applications. However, malware also poses unique challenges – in particular, how to safeguard an application and the public in the immediate aftermath of an attack. For instance, how do we land the aircraft in flight when we fear that an Air Traffic Management infrastructure is compromised?

Subsequent sections presented a more elaborate Gatekeeper architecture involving cooperation with internal company management and with external organisations, including regulators, industry associations or

professional bodies. The intention was to extend a purely internal reporting system so that other companies in the same industry or across other critical infrastructures might be alerted to potential future attacks. The system supervisor or gatekeeper must determine which events are passed onto the company management; hence they play a key role in filtering the information that is eventually passed to external organisations. A key recommendation is that these individuals urgently require tools that help them to assess the safety consequences of potential cyber-incidents to ensure the coherence and consistency of their decision making.

Active external monitoring architectures build on the gatekeeper approaches but also assume that regulators, industry associations and professional bodies will become more closely involved in assisting cyber-incident investigations. This is important when most safety-critical companies lack any in-house forensic expertise. These systems also assume a greater degree of both safety and security maturity given that companies must be willing to accept support and guidance from external organisations. It seems unlikely that these approaches will succeed without additional legislative and regulatory protection for the companies that participate in the programme. However, they may be necessary to ensure adequate protection for national critical infrastructures; without resource pooling we can have little confidence in the investigation of more complex cyber-attacks across many of our industries.

The closing sections presented a more complex approach based on the creation of public-private partnerships. Joint working groups consist of companies across safety-critical industries, of national safety regulators and of existing security agencies. The intention is to eliminate the silos that have arisen when each industry has its own safety regulator which are separate from state cyber-security agencies. In consequence, very few safety regulators have any cyber-security expertise. Conversely, most CERTs lack any specific understanding of safety-related development practices. Companies are encouraged to share incident information through their representatives on the joint working group and to seek support when needed.

The architectures in this paper are drawn from a number of international projects to establish incident reporting systems across Europe and North America. Previous work has shown that there is no 'ideal approach'. In the past, reporting systems have failed because their proponents over-estimated the maturity of the host organisation. Elaborate reporting architectures seldom succeed if industry employees are worried about retribution or prosecution when they report an incident. Many European reporting systems have died as soon as public funding is reduced. In such circumstances, it may be better to encourage simplified systems with the exchange of anonymous summaries each year. This helps to establish the reporting culture that is a prerequisite for more elaborate schemes [1]. Irrespective of the architecture that is used, the underlying argument behind this paper is that we must act now to integrate reporting mechanisms for cyber-attacks on safety-critical, national infrastructures. At present, companies do not know where to report their concerns in consequence many attacks are treated as isolated incidents. There is a reluctance to tell safety-regulators because of the consequences for the certification and approval of underlying software/hardware infrastructures. There is also a clear lack of regulatory guidance on the tools and techniques that can be used to assess and mitigate the safety hazards of cyber-attacks. All of these concerns create an urgent need to coordinate government and commercial action before public safety is placed at greater risk.

References

1. C.W. Johnson, Failure in Safety-Critical Systems: A Handbook of Accident and Incident Reporting, University of Glasgow Press, Glasgow, Scotland, 2003.
2. European Network and Information Security Agency, Critical Cloud Computing: A Critical Information Infrastructure Protection Perspective On Cloud Computing Services, Version 1.0, Heraklion, Greece, December 2012.
3. NIST, Computer Security Incident Handling Guide: Recommendations of the National Institute of Standards and Technology, NIST Special Publication 800-61 Revision 2, August 2012.
4. 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, London, UK, Paper to accompany a keynote address, 20th Annual Conference of the UK Safety-Critical Systems Club, ISBN 978-1-4471-2493-1, 2012.

5. C.W. Johnson, The Telecoms Inclusion Principle: The Missing Link between Critical Infrastructure Protection and Critical Information Infrastructure Protection. In P. Theron and S. Bologna (eds.), Critical Information Infrastructure Protection and Resilience in the ICT Sector, IGI Global, Pennsylvania, USA, 2013.
6. U.S. National Institute of Standards and Technology (NIST) (2006), Guide to Integrating Forensic Techniques into Incident Response, Special Publication 800-86, Gaithersburg, Maryland, 2006. <http://csrc.nist.gov/publications/nistpubs/800-86/SP800-86.pdf>
7. U.S. National Institute of Standards and Technology (NIST) (2012), Computer Security Incident Handling Guide (Draft), Special Publication 800-61 Revision 2 (Draft), Gaithersburg, Maryland, 2012. <http://csrc.nist.gov/publications/drafts/800-61-rev2/draft-sp800-61rev2.pdf>
8. J. Wiik, J.J. Gonzalez and K-P. Kossakowski, Limits to Effectiveness in Computer Security Incident Response Teams. In Twenty Third International Conference of the System Dynamics Society. The System Dynamics Society, Boston, MA, July 17-21, 2005. <http://www.cert.org/archive/pdf/Limits-to-CSIRT-Effectiveness.pdf>
9. US Government Accountability Office, A Better Defined and Implemented National Strategy Is Needed to Address Persistent Challenges, GAO-13-462T, Washington, DC, USA. March 7, 2013
10. US Government Accountability Office, Cybersecurity: National Strategy, Roles, and Responsibilities Need to Be Better Defined and More Effectively Implemented, Washington, DC, USA, GAO-13-187, February 14, 2013

Biography

Chris.W. Johnson, DPhil, MA, MSc, FBCS, CEng, CITP,
School of Computing Science, Univ. of Glasgow, Glasgow, G12 8RZ, Scotland, UK.
Tel +44(141)3306053, Fax +44(141)3304913, Johnson@dcs.gla.ac.uk, <http://www.dcs.gla.ac.uk/~johnson>

Chris Johnson is Professor of Computing Science at the University of Glasgow in Scotland. He heads a small research group devoted to improving the reporting and analysis of incidents and accidents across safety-critical domains ranging from healthcare, to the military to aviation and rail.