Trends in Information Security

Topic Description: Security of Safety-Critical Systems

Prof. Chris Johnson,
School of Computing, University of Glasgow.
Johnson@dcs.gla.ac.uk
http://www.dcs.gla.ac.uk

Introduction

There are strong similarities between safety and security – they are both regarded as non-functional requirements. In other words, you cannot devise exhaustive tests to show that a system is entirely safe nor is it possible to demonstrate that a system is totally secure – when the nature of potential attack methods may change over time. Many techniques that were originally developed to support the design of safety-critical systems have been extended to support the design of secure systems. For example, risk assessment techniques focus on the probability and consequence of different hazards. Tools that support this form of hazard analysis in safety-related systems, including HAZOPS and FMECA, have been applied to identify security risks. Similarly, safety management systems that link risk assessment to incident reporting have also been extended to support security management systems. However, there are important differences. For instance, many of the mathematical techniques used to analyse the probability of safety hazards cannot be used in security critical systems. In safety related applications we can assume independence between the probabiility of hardware failures. In contrast, in security applications the fact that one subsystem has been attacked will massively increase the probability of an attack on other systems given that attacks are coordinated and directed by attackers.

Different Perspectives

Some researchers have argued that safety and security are orthogonal. In other words, they are entirely separate concerns. For example, a system may be safe but not secure – such as a medical information system that enables a doctor to directly access patient records without entering a password. Alternatively, a secure system may not be safe – for instance, if it takes the clinician so long to enter a password that the patient dies before they can access their records.

Others have argued that these two concerns are so closely related that you must consider them as part of a combined approach to design. In other words, an insecure system will never be safe. Conversely, a safe system must always have been subjected to a security assessment that is continually maintained to identify changing threats over time.

A range of European and US government agencies have begun to draft policies and guidelines that are intended to ensure the security of critical information infrastructures. However, a host of research questions remain to be addressed.

Open Research Questions

Previous sections have argued that there are important differences between safety and security that force changes to be made when re-using existing techniques from one area so that they can support the other. At the moment, there are very few approaches that can be used to support the development of both safe AND secure systems. In particular, it is hard to identify interactions – where increased security might undermine safety and vice versa. For instance, intrusion detection systems can be used to monitor unusual behaviour through the (ab)use of computational resources. However, in order to be approved for use in safety-related applications code has to go through a range of detailed verification and validation activities that take considerable time to complete. This means in many cases that there would be a significant delay between the identification of a new attack pattern and the introduction of intrusion detection systems. At present it is unclear how we can update malware monitoring systems without running the risk that a bug in the protection software might not cause the loss of a safety critical system.

Similar research questions focus on forensics – if malware is detected it may not be possible to isolate and power down a safety-critical system without increasing the risks to potential users. For example, we cannot simply switch off an air traffic control system with aircraft still in the sky. Other questions can be summarised as follows:

- Is it technically possible to increase security through the logical or physical isolation of safety-critical systems from the Internet?
- How do we increase security without undermining safety or imposing constraints that undermine the economic justification for many complex applications?
- How do we secure safety-critical infrastructures with an operational life of more than 50 years, such as the UK next generation of nuclear reactors?
- How can companies that operate safety-critical systems provide security against state sponsored attacks and who should pay for these protection measures?
- Can we secure safety-critical systems that rely on equipment and software provided by companies from other countries?

References

C.W. Johnson (2012), Preparing for Cyber-Attacks on Air Traffic Management Infrastructures: Cyber-Safety Scenario Generation. In Proceedings of the 7th IET Conference on Systems Safety and Cyber-Security, Edinburgh, Scotland, 15-18 October 2012, IET, Savoy Place, London, 2012.

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

U.S. Nuclear Regulatory Commission (2010), Cyber Security Programs for Nuclear Facilities, Office Of Nuclear Regulatory Research, Regulatory Guide 5.71, January, 2010.
http://pbadupws.nrc.gov/docs/ML0903/ML090340159.pdf