Building Management Systems:

Organisational Security, Context and Convergence

Emma Boakes (*emma.boakes@port.ac.uk*)
School of Computing, University of Portsmouth

Abstract

Cyber-physical systems, such as building management systems (BMS), increase the attack surface of an organisation by introducing the possibility that a cyber vulnerability could have a physical impact and by potentially allowing new routes into an organisation's network. Responsibility for cyber-physical systems spans both physical and cyber security functions in an organisation. As a result security practitioners increasingly advocate formal collaboration, or convergence, between security functions to address the risks from this emerging cyber-physical context yet there is little guidance on how to achieve this, or evidence that it improves risk management.

This paper outlines the risks posed by cyber-physical systems. It reviews existing guidance and academic research relating to convergence, highlighting the need for further research to develop an evidence base, as well as the need to provide organisations with direction on ways to implement convergence. The paper finishes with a brief discussion of an initial interview study.

Introduction

Organisations are increasingly connecting systems to their networks to control their physical assets. These cyber-physical systems include building management systems (BMS), industrial control systems (ICS), as well as Internet of Things (IoT) devices. Connecting these previously segregated systems exposes them to cyber threats that could have a physical impact. For example, a cyber-attack targeting a BMS could disrupt heating, ventilation, or power, or could undermine physical security which increasingly relies on internet-enabled devices, such as CCTV cameras and access control. In addition, these systems are now exposed to cyber attacks that exploit their vulnerabilities to access the organisation's network. There are several instances where this has already happened (for example see ICS-CERT, U. S. Department of Homeland Security, 2016), and industry reports show that the number of groups interested in targeting such systems (Symantec, 2019) and the number of cyber incidents on physical systems (Ponemon Institute, 2018), is increasing.

There have been well documented instances of attacks on cyber-physical systems. A DDOS attack on the heating system of a residential building in Laapeenranta, Finland, rendered the building uninhabitable for several days (Janita, 2016). A spear phishing attack on a Ukraine electricity distributor allowed attackers to obtain remote access credentials to the SCADA systems which they shut down cutting power to approximately 225,000 electricity customers (ICS-CERT, U. S. Department of Homeland Security, 2016). A phishing attack on a heating, ventilation and air conditioning (HVAC) supplier to Target stores facilitated access to systems that enabled attackers to steal details from 40 million credit cards (Radichel, 2014).

Inclusion of a BMS on an organisation's network not only increases the attack surface but also places a greater load on both physical and cyber security staff, potentially leading to personnel vulnerabilities. Additionally, a BMS spans the boundaries of cyber and physical security and personnel security functions in an organisation requiring them to work together to identify, understand and mitigate vulnerabilities.

This paper argues that current guidance fails to address the challenges of integrating cyber, personnel and physical security to protect a BMS. Guidance focuses on securing the attack surface with technological solutions and by developing a 'defence in depth' approach which advocates the use of a range of security controls to deter or delay an attack. This, however, overlooks the impact on staff managing the system. Increasing system complexity puts a greater load on staff, potentially reducing their capacity, which in itself could create further vulnerabilities.

Industry practitioners have advocated formal collaboration between security teams through a converged security approach (for example, see Willison, and Sembhi, 2017) and have highlighted the benefits (for example, see Ritchey, 2018). Despite this, there is little evidence to support converged security (Kamp, 2016), and little detail about the type or level of convergence required to achieve the specified benefits. Moreover, there is no indication of how organisations might overcome the challenges of adopting such an approach (Aleem, Wakefield and Button, 2013).

This paper sets the security risks of cyber-physical systems within the broader context of organisational security. It provides a review of existing guidance and research and identifies a lack of empirical research to underpin the implementation of a converged approach to organisational security. Finally, some emerging research findings are reported along with suggestions for future research.

Organisational security & cyber-physical systems

Allowing physical systems such as a BMS to be internet-enabled increases complexity of the organisation's systems, and the attack surface, and places additional load on the staff responsible for system management. There are a growing number of groups interested in targeting cyber-physical systems (Symantec, 2019), and a growing number of incidents involving them (Symantec, 2019; Loy, K., 2018; Ponemon Institute, 2018), making securing such systems increasingly important. Organisations appear to recognise this risk and are prioritising risk reduction for these systems (Ponemon Institute, 2018), yet it is not clear how they will do this when few report having full visibility of their attack surface, and many report a lack of resource to effectively scan for vulnerabilities in their systems in a timely manner (Ponemon Institute, 2018).

A BMS spans the boundaries of cyber and physical security functions within an organisation, so cyber, personnel and physical security teams need to work together to identify, understand and mitigate vulnerabilities. 57% of security professionals think that combining security activities is an important or useful response to cyber-physical threats (Dorey, Willison and Sembhi, 2012). Security industry publications and practitioner interviews document the benefits of more formal integration of previously siloed security teams. Benefits include improved visibility and alignment with the business with a single Chief Security Officer (CSO) (Ritchey, 2018; Slater, 2005), improved cross training opportunities and communications which facilitate business continuity (Slater, 2005), improved synergy of mission, efficiency and costs (Ritchey, 2018), improved prioritisation of vulnerabilities, and a reduction in meetings (Willison, Sembhi and Kloet, 2012; PWC, 2010).

Expert opinion provides a good overview of security convergence, and is not to be dismissed, yet its limitations need to be recognised, for example, there is little indication of the type or level of convergence necessary to achieve the benefits specified. Where a specific aspect of convergence is associated with a benefit there is still a lack of detail, for example adopting a CSO may improve visibility, but it is unclear what activities the CSO would need to do to achieve the benefit, and exactly what that benefit would look like. Without understanding the specific nature of convergence and any supporting actions that have taken place by the organisation, it is not possible to determine how reported benefits materialise, and this makes it difficult to establish the veracity of any claims. Consequently, specific convergence measures cannot be replicated elsewhere to either test their effectiveness or to achieve the same beneficial results. In addition, it could also be argued that some benefits are not solely achieved through convergence, and similarly, it is questionable whether convergence by itself could bring about

some of these benefits. For example, does converging security functions alone improve communication?

Some organisations have already taken a converged approach with a single head of security or security executive (Gill and Howell, 2016; Dorey, Willison and Sembhi, 2012). This indicates the integration of security teams through organisational structure. It is not clear, however, how organisations have embarked on this, what sources they have used to facilitate their decision or implementation of such integration, how the security functions work together in practice, or what type or level of convergence is required to achieve any benefit. As Tyson (2007, p4) notes, '*security convergence can be as much or as little as is useful to an organization*' and Dorey, Willison and Sembhi (2012) indicate that organisations are not operating convergence in a standard way, with some organisations only implementing a partial integration, whilst others report complete convergence. This demonstrates that there is potentially an array of factors that impact on convergence, and that some organisations may be more predisposed to implement converged security functions than others.

Guidance on securing cyber-physical systems.

Government and industry bodies help organisations improve the security of cyber-physical systems through guidance. Examples of such guidance include: the Industrial Control Systems Cyber Emergency Response Team (US Department of Homeland Security, 2016), US Department of Defense (Dalton, 2016), Centre of Protection for National Infrastructure (CPNI) and the British Standards Institute (BSI, 2015) and CPNI and the Institute of Engineering and Technology (Boyes, 2013).

These guidance documents appear to be based on good practice, as perceived by the authors and contributors, rather than being based on empirical evidence. It is difficult to ascertain the validity of the suggestions, or on what basis they were included. The expertise involved in their generation is not always clear with contributions limited to a list of organisations (BSI, 2015; Boyes, 2013), or a mention of technical committees (BSI, 2015), rather than detailing the expertise of those contributing, such as ICS cyber security specialists (US Department of Homeland Security, 2016). Without a broad range of expertise feeding into the development of guidance there is a missed opportunity to explore additional and alternative ways of ensuring security. There is also a risk that such guidance is self-perpetuating, where guidance documents are written by subject matter experts, fed into organisational policy and system design, which in turn is used in generation of future guidance. This perhaps explains why guidance focuses primarily on technological solutions, which is a narrow response to the underlying factors behind a cyber-physical breach and the protective measures that enhance security. The

challenge of how security teams should work together to understand the threats they face, and decide on the measures they can put in place to mitigate them, are not addressed. There may be many safeguards that could be put in place for any single vulnerability, and this is not adequately addressed in the guidance.

To illustrate, the example given in Boyes (2013) highlights the vulnerability of a CCTV to wifi jamming, which poses the risk that intruders will not be recorded, with the suggested safeguard to install wired CCTV at points of entry and exit. This would have been an opportunity to explore broader issues such as discoverability of the CCTV equipment wifi broadcast, technical options to detect or prevent interference, the ability of physical security officers to detect hostile intent or the information about the set up discoverable on suppliers' websites. After all, wired CCTV will not stop an intruder, it will simply record them for investigative purposes. Additionally, a piecemeal approach will not uncover systemic factors that might undermine the security put in place. For example, the need for additional resources to monitor and update increasingly complex systems and security controls, without which security is undermined.

To help organisations make decisions about implementing security controls, guidance advocates a risk assessment process to prioritise systems that would be most impactful should they be exposed to a cyber security attack (Dalton, 2016; US Department of Homeland Security, 2016; BSI, 2015; Boyes, 2013). This may be problematic as the potential impact of a system may change depending on the threat landscape and connectivity. It also does not help organisations think about their systems as a whole, recognising that vulnerabilities, and the risks they pose, may exist due to many different factors; for example low level technical design issues, such as not forcing a default password reset on device set up, to broader issues, such as system details being posted online by the supplier. This is particularly important for cyber-physical systems, such as BMS, where the vulnerabilities might be in either domain.

Guidance needs to address this complexity and provide information at an appropriate level for its audience so they can make decisions about implementing security controls. For a BMS, stakeholders might include physical security, facilities management, external contracting organisations as well as cyber security. Yet guidance does not always specify the target audience (US Department of Homeland Security, 2016) or relies on broad descriptions, for example 'asset owners' (BSI, 2015). This makes it difficult to establish whether the information provided is at the appropriate level, and whether the guidance prompts engagement with the right stakeholders to help in the decision-making process. Additionally, the level of detail across the guidance varies considerably, often leaving room for interpretation. This again places the emphasis on the organisational decision-making process. For example,

> *'The contracts and site operating procedures should define responsibilities and acceptable practice to address the risks associated with the frequency of personnel changes'* (Boyes, 2013, p.31)

This gives no indication of the *'responsibilities'* to be defined, for who, or what *'acceptable practice'* looks like. In contrast, the US Department of Defense (Dalton, 2016) guidance is more detailed, prescribing requirements for selection and use during the design process of a facility building control system. Even here, there are still a number of decisions that could impact which security controls are identified, such as the precise user actions that need authentication.

> *'The organization identifies and defines organization defined user actions that can be performed on the information system without identification or authentication consistent with organizational missions/business functions'* (Dalton, 2016, p.135)

Unlike other guidance, the US Department of Defense guidance is directed at designers of facility building control systems, a specific target audience. While this guidance outlines who is responsible for security controls, however, it is not always clear who has responsibility for, or should be involved in, the decisions about them.

The focus on technological solutions across the guidance (Dalton, 2016; Boyes, 2013) overlooks the impact the solutions themselves may have on staff managing the increasingly complex system. As system complexity increases, there is more opportunity for vulnerabilities to creep into the system, resulting from insufficient prioritisation of business critical systems, inappropriate decision making about controls to be implemented and the systems they should be implemented on, and the maintenance of an increasing number of those controls. These are very human considerations, and again indicate that adopting technical solutions is insufficient without considering the support offered by the broader system to facilitate the implementation and maintenance of an increasing number of technical solutions.

Without recognising these broader contextual factors the effectiveness of any technological controls could be undermined. This is exemplified by the Target hack (Radichel, 2014), where systems-monitoring software had been introduced and standards had been adhered to, but these were ineffective due to system issues around staffing levels and training which meant unusual activity went unnoticed. Current guidance does not recognise the broader context, and does not address the collaboration required to identify, understand and mitigate potential threats. In short, there is a gap between the documented guidance and the practicality of implementing it in a real organisation.

An Overview of Convergence Research

Formal collaboration between security resources has been labelled '*convergence*' (Tyson, 2007). Broadly there is agreement that convergence aims to bring together security professionals to provide a more holistic view of organisational security (Kamp, 2016; Aleem, Wakefield and Button, 2013), by removing organisational silos to encourage information exchange (Rahman & Donahue, 2010), and through the use of processes (Aleem, Wakefield and Button, 2013).

Convergence is seen as a way of mitigating vulnerabilities that might be introduced as the boundary between physical and cyber becomes increasingly blurred (Kamp, 2016), for example, by networked technology solutions (Rahman and Donahue, 2010). Research that specifically looks at drivers for convergence is lacking and whether convergence is actually adopted to mitigate the potential for a blended threat has not been explored. Equally, the lack of research in this area misses an opportunity to establish why organisations may not be adopting convergence and whether there are any barriers to adoption. Understanding the drivers for adopting security convergence is important as it may have an impact on how an organisation approaches their convergence efforts. For example, an organisation's motivation may impact the extent to which they try to converge security, where they converge security, how they converge security, and the effort in making convergence a success. It could be argued that organisations that are driven to adopt convergence to provide improved, more holistic, security may be focused on a different type of convergence than those using convergence as a cost saving measure.

There is also little detail across the literature about how convergence should happen. There is no roadmap organisations can use to help them start implementing convergence, no tools or techniques that might facilitate the process in whatever form it might take for them, and no measures to assess whether they have been effective. The academic literature that has sought to address *how* organisations converge security resources is limited, and has focused primarily on the development of models and frameworks (Aleem, Wakefield and Button, 2013; Rahman and Donahue, 2010).

The approaches put forward in the academic literature (Aleem, Wakefield and Button, 2013; Rahman and Donahue, 2010) offer a high-level guide within which organisations have the flexibility to achieve a form of convergence that is suitable for them. Rahman and Donahue (2010) put forward three alternative converged solutions based on changes to organisational structure: alignment of security teams under one manager, with either the different security functions integrated or kept separate, or conversely, the use of a risk council to deal with security issues. Rahman and Donahue appear to offer no obvious evidence base, however, and it

is not clear whether they are reporting ways that organisations have adopted convergence, or strategies that organisations might like to take. In contrast, Aleem, Wakefield and Button (2013) propose a more procedural approach that doesn't require organisational restructuring, something proposed as a way of smoothing the convergence process by Rahman and Donahue (2010). Aleem, Wakefield and Button (2013) developed their model from a case study, where processes across security functions were mapped and aligned. Their approach encourages security functions to collaborate on certain procedures, such as risk mapping and business impact assessment; crisis and disaster management and business continuity; and awareness and training. These models and frameworks offer a high level guide to convergence, yet the level of detail is superficial and it is doubtful an organisation could use the literature to embark on a successful converged approach. In addition, they do not provide any indication of how an organisation's security functions might actually collaborate and communicate once they are converged to ensure success.

The convergence literature recognises the potential challenges organisations might face when adopting convergence, although these do not appear to be grounded in research. Challenges to convergence include: cultural differences between cyber security and physical security that arise from their traditionally diverse backgrounds in information technology and law enforcement (Rahman and Donahue, 2010); and the disparity in salary between traditionally higher paid cyber security staff and the lower paid physical security staff (Rahman and Donahue, 2010). Such challenges need to be addressed if collaboration between the security functions is to occur. Any measures that generally facilitate convergence are seemingly untested, and are deficient in the level of detail required to be effectively implemented in an organisation, for example, achieving support from the board (Rahman and Donahue, 2010; Aleem, Wakefield, and Button, 2013), and undertaking cross training between security functions (Aleem, Wakefield and Button, 2013; Rahman and Donahue, 2010).

In short, it is unclear how an organisation should adopt security convergence. In addition, the models and frameworks put forward in the academic literature do not appear to have been tested and it is not possible to ascertain how generalisable or effective they are, or their applicability or usability in real life settings. It is no surprise that interview studies have noted that organisations '*struggle with the implementation of a converged security approach'* (Kamp, 2016, p47). Overall, the evidence base for convergence is inadequate, and of limited use to organisations making the decision to move to converged security, and deciding how, exactly, to adopt the approach. This gap is acknowledged by the academic literature (Kamp, 2016; Aleem, Wakefield and Button, 2013).

It is clear that cyber-physical systems, such as BMS, introduce new security challenges in organisations because they increase the attack surface and add complexity to the risk profile. Such challenges are amplified as responsibility for cyber-physical systems spans the boundary of security functions, and places additional load on often already overloaded security practitioners. Convergence has been advocated within the security industry as a way of identifying and mitigating risks from such systems, yet there is little guidance on how organisations might implement the approach or evidence that it is an effective solution. Research is needed to build an evidence-base for convergence in order to help organisations decide whether it is a good option and to develop guidance that can be applied in real-world settings.

The remainder of the paper discusses some preliminary findings from an initial interview study to identify commonalities in the methods used by organisations to converge security resources. A brief outline of future research is also introduced.

Initial interviews

A set of five pilot interviews have been carried out to explore how organisations converge security resources, how different security disciplines collaborate, and the barriers and facilitators experienced. The semi-structured interviews were with senior security staff at organisations from five different sectors where a converged security approach had been adopted. Interviews lasted between 40 minutes to 1 hour 40 minutes, and were carried out face to face or over the telephone. Interviews were recorded using a digital voice recorder, transcribed by a transcription service, then anonymised by the researcher. At the time of writing, interviewees are reviewing and redacting any sensitive or identifying information from the transcripts. Once this is completed, a thematic analysis will highlight key themes using the 6-stage process suggested by Braun and Clarke (2006). Familiarisation with the data however has already highlighted some initial preliminary findings.

Preliminary Findings

All organisations that participated had a high level security function that overarches several security areas, either embedded in the organisational structure or facilitated through a cross cutting committee. Convergence was operationalised through forums, formal and informal meetings across the individual security areas, and staff having cross cutting responsibilities that traverse the different security areas. This indicates that convergence in these organisations is more than a change to organisational structure, and has required specific activities and job design to enable convergence on the ground. The interviews also highlighted the breadth of

organisational security, encompassing areas such as business continuity, organisational resilience, incident management, education and culture.

A champion for convergence appears to be an important component in initiating and facilitating the approach; a person who instigates convergence and who drives the security agenda and culture within the company. More generally, interviewees voiced that a security function needs people who can communicate effectively and build relationships across the business, to provide support and security advice. This aligns with the idea that security should be seen as a function that supports the business and adds value and as such has to understand, align, and provide a service to the business; this is akin to a security consultancy, where the security function is directly responding to the needs of the business and providing security specific expertise.

Consequently, security staff have a role in helping the business to understand the concept, principles and objectives of security, and therefore they have to engage with people who do not have, and are not expected to have, a security background. This makes security an increasingly outward facing function that needs to have an understanding of the different business areas they engage with.

Further analysis of the dataset is ongoing, however, it is of note that interviewees were self-selecting, and it is unlikely that organisations where convergence is more problematic would come forward therefore there is a bias in the dataset. Additionally, organisations not operating converged security were excluded, therefore no comparison can be made between converged and non-converged security functions. The aim is to address this in future research.

Future research.

The review of practitioner guidance and academic research, along with the initial interviews that have been carried out, indicates that there is further scope to research the area of convergence. Further research in this area will adopt a case study design. Organisational security is a complex phenomenon, and case studies allow for an in-depth exploration of convergence in the real world context in which it occurs (Boblin, Ireland, Kirkpatrick & Robertson, 2013). This will enable the exploration of why and how organisations have adopted convergence.

The case studies will be structured using an evidence-based practice approach, which looks at the evidence used in the decision-making process, and, importantly, the evaluation of that evidence (Briner, 2019). Case study design and evidenced-based practice complement each other in terms of the range of evidence sources used. Using case studies will not only allow an exploration of the evidence used by organisations, but will also, in turn, help to build the

evidence base for convergence. Future research will also aim to identify best practice and method(s) that facilitate the adoption of converged security.

Conclusion

This research responds to the need for organisations to enhance their security posture through collaboration across their security teams. Attacks that exploit the gaps between security domains demonstrate the need for organisations to identify and mitigate vulnerabilities that cross cyber, physical and personnel security. Converged security is widely advocated to address this, yet there is limited evidence and little information to assist organisations who want to adopt this approach.

This research seeks to understand the barriers and enablers to converged security through an in-depth exploration of operational converged security functions. Preliminary findings from an initial set of five interviews have illustrated how convergence works operationally. These findings have begun to indicate factors of converged security operations that are thought to enhance the role of security within the organisation, such as security as a consultancy. Further research will look at exploring and building the evidence for converged security through a case study methodology using an evidence-based practice framework.

References

Aleem, A., Wakefield, A., & Button, M. (2013). Addressing the weakest link: Implementing converged security. *Security Journal*, *26*(3), 236-248. https://doi.org/10.1057/sj.2013.14

Boblin, S. L., Ireland, S., Kirkpatrick, H., & Robertson, K. (2013). Using Stake's qualitative case study approach to explore implementation of evidence-based practice. *Qualitative health research*, *23*(9), 1267-1275. https://doi.org/10.1177/1049732313502128

Boyes, H. (2013). *Resilience and cyber security of technology in the built environment*. Published by the Institute of Engineering and Technology, London, UK.

Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3, 77-101. https://doi.org/10.1191/1478088706qp063oa

Briner, R. (2019). The Basics of Evidence-Based Practice. *HR People and Strategy*. Advanced retrieved from https://www.cebma.org/wp-content/uploads/Briner-The-Basics-of-Evidence-Based-Practice.pdf

BSI, S. (2015). PAS 1192-5: 2015: Specification for security-minded building information modelling, digital built environments and smart asset management.

Dalton, J. (2016). Unified Facilities Criteria (UFC): Cybersecurity of Facility-Related Control Systems. U. S. Department of Defense, Washington DC United States. Retrieved from https://www.wbdg.org/FFC/DOD/UFC/ufc_4_010_06_2016_c1.pdf

Dorey, P., Willison, J., and Sembhi, S. (2012). *Converged Security Management Survey 2012*. Retrieved from http://personal.rhul.ac.uk/vsai/149/Convergence%20Survey%20Final%20070312_V6.pdf

Gill, M., and Howell, C. (2016). *Tackling Cyber Crime: The Role of Private Security*. Retrieved from https://perpetuityresearch.com/wp-content/uploads/2016/09/SRI-Report-2016.pdf

ICS-CERT, U. S. Department of Homeland Security (2016). Cyber-attack against Ukrainian critical infrastructure. Alert (IR-ALERT-H-16-056-01). Retrieved from: https://www.us-cert.gov/ics/alerts/IR-ALERT-H-16-056-01

Janita (2016, November 7). DDoS attack halts heating in Finland amidst winter. *Metropolitan.fi.* Retrieved from http://metropolitan.fi/entry/ddos-attack-halts-heating-in-finland-amidst-winter

Kamp, G (2016). Security Convergence in a Critical Infrastructure Framework and Enablers for Successful Implementation. Masters Thesis retrieved from: https://openaccess.leidenuniv.nl/bitstream/handle/1887/53731/2016_Kamp_CSM.pdf?sequence=1

Loy, K. (2018). Implementing Cybersecurity Best Practices in Five Steps.Retrieved from https://www.securityindustry.org/2018/09/14/implementing-cybersecurity-best-practices-in-five-steps/

Ponemon Institute (2018). Measuring & Managing the Cyber Risks to Business Operations. Retrieved from https://static.tenable.com/marketing/research-reports/Research-Report-Ponemon-Institute-Measuring_and_Managing_the_Cyber_Risks_to_Business_Operations.pdf

PWC (2010) *Convergence of Security Risks: Addressing the security dilemma in today's age of blended threats*. Retrieved from the ASIS website: http://www.asis.org.uk/articles/Security%20Risk%20Convergence.pdf

Radichel, T. (2014). Case study: Critical controls that could have prevented Target breach. SANS Institute InfoSec Reading Room.

Rahman, M., & Donahue, S. E. (2010). Convergence of corporate and information security, *International Journal of Computer Science and Information Security*, *7*, (1), 63-68. Retrieved from https://arxiv.org/ftp/arxiv/papers/1002/1002.1950.pdf

Ritchey, D (2018). The Unstoppable Convergence Between Physical and Cybersecurity. Retrieved from: https://www.securitymagazine.com/articles/88847-the-unstoppable-convergence-between-physical-and-cybersecurity

Slater, D (2005). Physical and IT Security Convergence: The Basics. Retrieved from: https://www.csoonline.com/article/2117824/physical-and-it-security-convergence--the-basics.html

Symantec (2019). Internet Security Threat Report, Volume 24, February 2019 . Retrieved from https://img03.en25.com/Web/Symantec/%7B1a7cfc98-319b-4b97-88a7-1306a3539445%7D_ISTR_24_2019_en.pdf

Tyson, D. (2007). *Security convergence: Managing enterprise security risk*. Elsevier.

US Department of Homeland Security (2016). Recommended Practice: Improving Industrial Control Systems Cybersecurity with Defense-in-Depth Strategies. US-CERT Defense In Depth.

Retrieved from: https://www.us-cert.gov/sites/default/files/recommended_practices/NCCIC_ICS-CERT_Defense_in_Depth_2016_S508C.pdf

Willison, J., Sembhi, S., and Kloet, F. (2012) Security Convergence and FMs: the Learning Curve. Retrieved from https://www.ifsecglobal.com/uncategorized/security-convergence-and-fms-the-learning-curve/