Developing cyber-security policies that penetrate Australian defence acquisitions

Stuart Fowler
Captain Craig Sweetman, Australian Army
Flight Lieutenant Sibi Ravindran, Royal Australian Air Force
Dr Keith F. Joiner, CSC, University of NSW
Dr Elena Sitnikova, Australian Cyber Security Centre

Technological edge on the battlespace has been one of Australia's military advantages within Southeast Asia. This is partly due to the relative larger size of the Australian economy and its ability to procure advanced defence systems compared to the region. However, by 2030, it is envisaged that Australia will slip to being the 23rd largest economy in the world, slightly smaller than Thailand's and similar in size to Malaysia's (Hawksworth and Chan, 2015). This will significantly erode Australia's financial advantage, as other regional nations will be able to afford advanced weapon systems, including cutting-edge cyber capabilities.

As all nations continue to develop their offensive cyber capabilities, Australian defence systems will become increasingly vulnerable. Joiner (2017) has reviewed the options of the Australian Department of Defence (Defence) for cyber-security test and evaluation and has highlighted that there are significant difficulties in attributing and predicting the final effects of cyber threats that are currently manifesting, making it very difficult for Australia to rely on the deterrent effect of a potential response from an ally like the US.

In an attempt to address these challenges, cyber security has become an increasing focus for Defence. The release of the *2016 Cyber Security Strategy* (Commonwealth of Australia, 2016) and the *2016 Defence White Paper* (Australian Government, 2016) shows that decision-makers at all levels have demonstrated an understanding of the importance of cyber security. However, this enthusiasm for protecting Defence capability from hostile cyber attacks has yet to translate into a comprehensive and coordinated set of policies at the working level (Joiner, 2017).

This is not necessarily due to ignorance or indifference. Translating the aspirations of a cyber-hardened and cyber-resilient arsenal of military systems into a workable and robust set of policies is a difficult challenge. It is amplified

by the large expenditure of public funds and the lengthy acquisition process typically associated with Defence's capability development, often cultivating a culture of risk aversion and an environment where a cyber-resiliency policy cannot simply be iteratively developed over a number of projects in any reasonable timeframe. If new Defence capability lifecycle policies are not implemented properly, cyber resiliency may be overlooked and a vulnerability may develop, or there may be significant wastage that does not materially contribute to the cyber resiliency of the capability being procured.

This article examines how Australian defence acquisition policies can adapt, mainly using US Department of Defense experiences, to deal more effectively and systematically with cyber resilience.

Project lifecycle challenges

Capability development

The capability lifecycle within Defence begins with an analysis of a desired capability concept and the development of system requirements for a solution that is ultimately procured through contracts. While this process needs to account for the entire lifecycle, there are general challenges that are faced during the early design stages that can impact the cyber resilience of major capabilities.

The capability lifecycle is based on the 'waterfall' project management model, which is suited to projects where the requirements and needs do not change over time (Bergmann, 2013). Due to the dynamic nature of the threat in a hostile cyber environment, agile iterative development practices are more suited for the cyber-security aspects of procurements wherever there are software-intensive subsystems.

Obsolescence and vulnerability management of software and hardware must be considered, and systems under evaluation must be future-proofed as much as possible. Most projects establish a technology baseline to support system integration, and test and evaluation. When this baseline is coupled with delays in acquisition, the resultant system accepted by Defence may be outdated and have known vulnerabilities, creating unnecessary risk to the capability, and potentially depriving Defence of a technological edge.

Of course, it would be misguided to see agile engineering methodologies as a panacea for the cyber-security challenges faced by Defence. On large, complex and tightly integrated systems, the level of design volatility inherent in the agile methodology can be a costly venture if contractual arrangements do not accommodate this volatility. Much of the traditional systems-engineering process is about controlling the level of design volatility to minimise risk. However, this process will need to evolve if it is to adequately address the challenge of cyber security.

An agile approach can require an increased ability to manage change, as change must be rapidly folded back into the system requirements and design to maintain the proper cadence of iterative releases. It seems clear that rising to this challenge will require a marriage between agile processes and traditional systems-engineering. This remains an ongoing research area. However, perhaps specific subsystems can follow an agile process within the confines of the requirements and constraints allocated to the subsystem, albeit with a traditional systems-engineering process at the system level.

Defence's requirements

One tempting solution for incorporating cyber resilience into Defence's capability lifecycle is simply to place cyber-resilience requirements within the functional performance specification which then becomes the basis of a prime contract with defence industry. Digging deeper, we find that it is not obvious what those requirements should be and how they should be decomposed. There is a temptation to incorporate general requirements that form the basis of security certification and accreditation, such as the *Information Security Manual* (Australian Signals Directorate, 2016) and the various emissions security standards, as part of the system specification.

However, these sorts of requirements offer thousands of controls, many of which may not be applicable to all systems, and often give little guidance on which of these controls are most important within the context of the specific capability program or known threats to the system.

Additionally, these controls focus primarily on ensuring the confidentiality of classified networks, rather than the integrity and availability of safety-critical or mission-critical systems, making them largely unsuitable for specifying cyber resilience.

More worryingly, due to their risk-based nature, it is entirely possible to meet these sorts of prescriptive requirements without providing a sufficient level of cyber resilience or even cyber-defence testing. The *Information Security Manual* is a strong resource when applied to strategic or corporate ICT systems but a reliance on this document to ensure cyber resilience of deployed systems is a flawed model.

One way to overcome this would be to specify the need for cyber resilience in a general sense, where a desired output is specified—a system which is cyber resilient—rather than specific input-driven requirements. The major downside to this approach is that this sort of requirement is not particularly measurable or quantifiable, making it difficult to formally verify and validate. To examine this issue more thoroughly, we can look at how similar problems have been addressed within broader Defence acquisitions; and the two parallels drawn here are with system safety and system survivability.

System safety shares the problem of being difficult to specify through requirements, and having a high-level desired output that can be difficult to measure, that is, that a system must be safe. It also operates at the intersection of defined engineering standards and risk assessments, much like cyber resilience. System safety is typically handled as part of the contractual 'statement of work', rather than as a technical system requirement, and this tends to achieve a more holistic outcome. The safety risk assessment, review and acceptance process itself is mandated contractually, and culminates in a safety case report that summarises the safety analysis performed, the state of any safety controls and the residual risk levels.

By necessity, system-safety processes take a holistic approach that, when done properly, helps transcend contractual boundaries. Taking a similar approach to incorporating cyber resilience into Defence's capability lifecycle has some merit, as it would allow for a more consistent approach across projects, encourage a collaborative approach between contractual parties, and provide a standard framework for which cyber-related risk is captured.

The US Department of Defense has recently incorporated into its standard acquisition policies the concept of a Program Protection Plan (Reed, 2015; Brown *et al,* 2015), which is a customer-owned document that captures cyber-related issues, plans and resolutions. The program protection plan is a consistent and comprehensive plan that covers a wide range of project-specific cyber-security information, from classification guidance to threat intelligence, supply chain, architecture and risk assessment, and can be viewed as analogous to a safety case but in the cyber-security domain.

If a program protection plan could be developed in parallel with an operational concept document and functional performance specification, the maintenance of the program protection plan, along with other process-based cyber-security issues such as risk assessments, certification and accreditation, could be mandated within the contractual statement of work for the prime contractor. This would free up the functional performance specification to contain specific functional requirements that support the cyber resilience of the system, while ensuring the foundational 'cyber-hygiene needs' are still flowed down to the contractor.

System survivability, the ability of a system to survive within a contested environment, also shares some challenges with cyber resilience. Both require continued operation within an unpredictable situation where a hostile entity is actively trying to degrade or destroy the system by evolving the threat. Again, this can be a difficult attribute to specify in a requirement that does not atrophy quickly, because the concept of setting a minimum acceptable floor for system survivability is difficult to rationalise, given the ambiguous constraints of active hostile action.

Any honest measure of survivability would shift over time as hostile capabilities increased, and it is an attribute that needs to be maximised within other constraints, rather than simply meeting an arbitrary minimum level. The standard approach to system survivability is to treat it as a key performance parameter, or a measure

of effectiveness, within the system itself. Again, this is often treated separately from functional requirements.

The US has recently incorporated cyber resilience and cyber survivability into its standard system survivability key performance parameter (Office of the Secretary of Defense, 2014). Contractors are required to show their performance against these parameters at project milestones and design reviews, and are required to meet thresholds for these parameters or risk their projects being defunded. While this may not be directly applicable to the Australian context, given different funding arrangements, there is a strong argument for including a cyber-survivability or cyber-resilience key performance parameter in acquisition contracts, perhaps with incentive payments tied to the level at which the key performance parameter is met.

Architecture and design

A striking difference between the Australian and US acquisition contexts is that Australian projects are more focused on the integration of off-the-shelf components or the acquisition of existing platforms from international partners (Joiner, Sitnikova and Tutty, 2016, p2). The reluctance, both of prime contractors and the government, to carry the risk of significant development activities in a climate of off-the-shelf acquisitions can place a constraint on the cyber-security design and development of a cyber-resilient architecture. This approach drives architectures during the tendering phases of an acquisition project to focus on off-the-shelf solutions that provide lower costs while meeting the minimum functional requirements.

The trade-off made in this methodology is the loss of flexibility in dealing with traditional non-functional quality attributes, such as cyber security and system safety, relying instead on test and evaluation to identify the risks being accepted. This approach can also increase the integration and interoperability risk for Defence if the off-the-shelf system uses proprietary protocols which can be difficult to integrate into wider Defence systems. Limited integration of systems may reduce the ADF's effectiveness in the battlespace, as an integrated force is a force multiplier and a cornerstone of the ADF's strategy and doctrine.

Strong and comprehensive cyber-security requirements need to be reinforced during tender evaluations by treating cyber resiliency as an integral capability discriminator. A tender reviewer would not compromise on the ability of a fighter plane or naval ship to dispense chaff or electronic countermeasures—and they need to take the same mindset when reviewing the ability of the system to resist cyber threats, as they can be just as damaging to mission effectiveness, safety and human life as kinetic threats.

Another challenge presented by the off-the-shelf approach is that there is very little ability to alter the native capabilities of the products, and many of the commercially available security appliances are poorly suited for defence systems; such as often assuming the presence of a stable internet connection or high data rate connectivity to a central server. A potential solution to this issue may be to start incremental risk reduction activities to look at this issue in advance of major capability projects.

An example of small-scale cyber-resilience projects is the High-Assurance Cyber Military Systems (HACMS) program developed by the US Defense Advanced Research Projects Agency. It is software designed to 'thwart cyber-attacks once deployed in any context, like a defence system or an unmanned aerial vehicle' (Paganini, 2014). The goal of the program is to develop technology for high-assurance cyber physical systems using 'clean slate, formal methods based approach to enable semi-automated code synthesis from executable, formal specifications' (Richards, 2016). In a demonstration using an unmanned helicopter, the program was able to prevent hackers from disrupting critical systems, despite the hackers having 'unfettered access' to the aircraft's computer (Slezak, 2015).

Similarly, the development of the Resilient Hull, Mechanical, and Electrical Securities (RHIMES) system by the US Navy was in response to the threat of an attack on control systems. RHIMES is a cyber-protection system with the design purpose of increasing the cyber resilience of shipboard mechanical and electrical control systems. The resilience is provided by preventing an attacker from taking control of programmable logic controllers that manage such functions as damage control, firefighting, anchoring, climate

control, electrical power, hydraulics, steering and engine systems. In the words of Rear Admiral Mat Winter, Chief of US Naval Research, 'the purpose of RHIMES is to enable us to fight through a cyber-attack' (Freeman, 2015).

Incremental risk reduction activities could focus on the development of a cyber-resiliency toolkit; a series of modules, components or products that are specifically crafted to work within the context of deployed systems that could be incorporated into a larger architecture in an off-the-shelf fashion. These modules would encompass traditional cyber-defence strategies, such as boundary protection, system management, intrusion detection, cyber-defence execution and cyber-state visualisation, as well as more complex cyber-resilience principles, such as adaptiveness, deception, diversity, unpredictability and reconstitution. The consistent implementation of these modules across projects would prevent the implementation of bespoke systems with a different look-and-feel on different platforms, and lessen the training burden of developing a new generation of cyber-smart military personnel.

Such projects could be completed through a partnership between the ADF's single-Services, strategic security agencies, defence industry, the Defence Science and Technology Group, universities and relevant security stakeholders, and could include research into self-hardening networks, tailored network-monitoring solutions or utilise modern virtualisation and container technologies to build on existing research into self-reconfiguring networks (Beraud *et al*, 2011) and self-reconstituting networks (Ramuhalli *et al*, 2013).

Test and evaluation

Cyber-security test and evaluation is another area where Defence's acquisition policy needs to be improved (Joiner, 2017). The vast majority of cyber-security test and evaluation within Defence's capability lifecycle is focused on ensuring confidentiality of information on classified networks though the certification and accreditation process. It is worth noting that a highly-critical system, such as a platform control system or navigation system, will typically not be assessed with as much rigour from a cyber-security perspective as a system that processes classified data. Cyber resiliency, however, needs to focus primarily on these critical systems to ensure they can continue to operate in a cyber-threat environment.

Joiner (2017) argues for cyber-security test and evaluation to become a larger part of the broader test and evaluation component of the capability lifecycle. This is an important step that needs to be taken, because it is the one point in the lifecycle where the cyber resilience of the system can be measured and reviewed. A robust cyber-security test and evaluation program can identify the residual risks in the system and validate the risk assessments that have been performed during early design activities.

The conduct of vulnerability and exploit assessments during detailed design reviews and developmental test and evaluation will increase the likelihood that any vulnerabilities are discovered and can be resolved before the design process is finalised. Potential vulnerabilities are cheaper to rectify early in the design process than after production, where significant engineering changes may be required. This process mitigates the risk that new cyber-security vulnerabilities are discovered as part of the acceptance test and evaluation, and operational test and evaluation, allowing these activities to focus on the analysis of, and where necessary acceptance of, residual risk.

While the conduct of cyber-security test and evaluation will not prevent poor security by design, real-world experience would be gained by the testers, and this could drive the specification of more relevant and specific cyber-security requirements for future capabilities. The collected lessons learnt from cyber-security test and evaluation would flow into existing capabilities, and into the initial design of future capabilities, ensuring security and resilience is built from the ground up.

This test and evaluation of cyber resilience requires that defence industry is able and appropriate to test that resilience against Defence's cyber threats. If the US Defense acquisition is any guide, such test and evaluation, once it begins in earnest, is likely to be Commonwealth-led with restricted access for many industry system developers, providing mainly recommended changes to software architectures without

necessarily explaining why, at least at a threat-level. Such Commonwealth-led cyber resilience evaluation may need to adjust Defence's preferred contracting models in Australia and some of the cultural preferences in Australian Defence acquisition of who should own what aspects of the designs at various stages of the system life-cycle.

Change management

General change-management policies also need to be reassessed to understand their impact on cyber security. One area where this is particularly important is in the application of security patches and upgrades to software in critical systems. Any change to a system that controls critical systems, such as weaponry or platform functionality, will trigger a recertification and functional regression testing of the system. This is often very costly and can act as a barrier to change within the system, with the resultant effect of leaving these critical systems running on unsupported or unpatched versions of software.

This regression test is required for good reasons; ordnance control and personnel safety rightly outweighs cyber-security concerns, so these tests should not be discarded. However, there needs to be a reasonable compromise between maintaining certification boundaries in critical systems and applying security patches, and any unnecessary process overheads must be streamlined.

Dedicated test environments where these security updates can be quickly tested must be maintained as part of platform sustainment, and it is recommended that the inability of this environment to implement security updates in a timely fashion be treated as a security risk in and of itself (McGinn, 2015). Automation of regression testing is the one feasible way to achieve the level of agility required to implement patches and security updates on mission and safety critical systems, and this must be part of any change-management policy that addresses cyber security.

Organisational challenges

Defence industry

Currently, Australia's defence industry sector is not usually the pioneer of cutting-edge technology (Ercis *et al,* 2015). Innovative industries such as start-up companies and small businesses sometimes cannot afford, nor survive, the

timelines and costs associated with Defence's bureaucracy (Müller and Fischer, 2014).

Agility can be highly beneficial in many aspects of defence acquisition, and small, agile industry participants do have an important role to play in ensuring a cyber resilient defence force. Through defence innovation hubs, they can tackle the incremental risk reduction activities advocated for previously to develop solutions and components which can be incorporated into major system acquisition projects.

However, defence acquisition for major projects is typically risk averse, and defence organisations would often rather pay a premium to engage with an established business within the defence industry. The start-up mantra of 'fail fast, fail early, fail often' does not hold much credence when spending billions of dollars in public money, where politically, defence organisations cannot afford such a hit-and-miss approach. Consequently, for major acquisitions, they usually look toward the recognised businesses which can be held responsible if something goes wrong, and which are established enough that they will absorb the costs to resolve an issue rather than simply declaring bankruptcy. This mitigates the risk to Defence by shifting some of it to industry.

The downside to the current approach to industry engagement is that defence-related acquisition and development projects tend to focus on industry's ability to satisfy the specific functional requirements to which they are to be contracted, with less consideration given to the attributes of the system to which they are not directly measured yet, such as cyber resilience and as argued earlier, system survivability more generally.

Due to the evolving nature of the cyber threat, security issues must often be addressed during the in-service phase of the capability lifecycle, and if contracting models do not adequately account for this eventuality, addressing these threats may be considered out-of-scope for the industry participants. If any additional costs are beyond the project's budget and Defence cannot afford the solutions, it must accept the risk of the vulnerabilities, which can lead to systems with significant residual risks being accepted. Consequently, if sustainment contracts are not cognisant of this need for addressing emerging cyber vulnerabilities, the management of cyber-security vulnerabilities that are detected during the in-service phase of the capability risks consuming a higher proportion of the system sustainment budget than should be the case.

It is therefore vital that industries that specialise in identifying and quantifying cyber vulnerabilities must be engaged throughout the lifecycle, from requirements through acquisition to disposal, to ensure mission assurance and effectiveness for Defence in cyber-intensive environments. To achieve this, Defence must work collaboratively with prime contractors to ensure this cyber-vulnerability support is part of the engineering design and development processes. Therein, the program protection plan referred to earlier is key to such industry support being appropriately scoped and engaged.

One way to implement this would be as an extension to the Infosec Registered Assessors Program, currently administered by the Australian Signals Directorate. This program brings together cyber-security experts who can be engaged by government agencies and their contractors to audit and assess ICT systems which process Australian government information. A complementary program could be implemented that would provide a group of independent experts who could be engaged by prime defence contractors, or by Defence itself, to assess designs and systems for cyber resilience and cyber worthiness on behalf of Defence.

Defence's cyber systems

The cyber-security landscape in a software intensive mission system or platform is inherently complex and dynamic. Most Defence systems are acquired through a project led by the Capability Acquisition and Sustainment Group, which is then delivered to the responsible capability manager within Defence. The personnel who acquire the system may not be involved during operations and sustainment, and Capability Acquisition and Sustainment Group is often measured against the timeliness and cost of a capability, which risks incentivising short-sighted thinking to get a project to its final operating concept milestone rather than the necessary deeper analysis to address complicated cyber-resiliency challenges. The problems of incentives in US Defense acquisition are covered extensively in a meta-analysis by Smith *et al* (2016).

From a Defence perspective, the cyber-security landscape is further complicated with different organisations responsible for the provision of ICT systems at the strategic, operational and tactical levels. This risks dilution of accountability and control of system cyber security due to the various organisations involved with overlapping roles. It also risks diminished accountability to record, control and monitor cyber-security issues. These risks are concerning for a tightly integrated small military force like that of Australian Defence, as a highly effective network-centric warfare system is paramount to maintain dominance over numerically superior forces, and a lack of organisational coordination can have a negative impact within this paradigm.

A single organisation should be responsible for managing Defence's cyber-security risk, responsible to the Chief of Joint Capabilities as the capability manager for joint C4ISR, space and cyber. As the challenges in cyber security are highly technical, the organisation must contain both technical and operational expertise on Defence's various software-intensive mission systems and platforms. Furthermore, the organisation should contain a mixture of people from the different Services, Reserve forces and external contractors, who are specialists in the cyber domain.

Snyder *et al* (2015) assert that managing cyber-security risk has three components, namely, minimising vulnerabilities to systems; understanding the threat to those systems; and minimising the impact to operational missions. The recommended organisation must cater for these components and implement a comprehensive cyber-risk management system for all Defence mission systems, which would require it to be involved in all lifecycle states of software-intensive mission systems. There are numerous cyber-risk management methodologies that are currently available and could be leveraged (such as Cherdantseva *et al,* 2016). Another source of guidance is the risk-management strategies used by allied forces for their cyber-defence risk quantification.

During the lifecycle of a defence-related mission system, the recommended organisation needs to:

- Influence decisions that shape the software architecture, underpinned by rigorous and timely test and evaluation, so as to ease cyber-security management and troubleshooting of the system in the later lifecycle stages. Cyber-security issues identified during design may result in significant cost-savings;
- Ensure that cyber security of the system is part of the design across the entire lifecycle of the system by developing and championing specific project protection plans;
- During the procurement, developmental and operational testing phases, ensure that the system complies to a cyber-security risk management framework;
- Ensure consistent and timely change management of the software-intensive systems across platforms including but not limited to upgrades, patches, modifications, technology refreshes etc;
- Continually monitor the cyber-security landscape, the changes in military concepts of operations and its effect on the software systems to continuously ensure that these changes do not introduce new vulnerabilities; and
- Ensure that the disposal of software-intensive military systems is performed without compromising other military systems by revealing vulnerabilities.

The recommended organisation must also have rapid prototyping capabilities and laboratory facilities that can simulate Defence's mission-system environments in a controlled environment (Van Antwerpen, 2015). This would also give Defence the ability to experiment and learn about the current mission systems, including their vulnerabilities, and provide an opportunity to experiment with emerging technologies and understand their limitations in a tactical environment. The knowledge gained from these rapid prototyping facilities would also allow Defence to be a 'smart buyer' as detailed in the recent first principles review of defence acquisition and sustainment.

The involvement of this recommended organisation during the pre-acquisition phases of a project is also likely to have the effect of encouraging defence industry to act on these sorts of concerns by making it clear that this capability is a discriminator in the selection process. An organisation such as the one described could become a strong influence for cyber security across Defence, provided its powers of review

and oversight were enabled within acquisition contracts, much like they currently are for safety regulators. There may be a temptation to merge this organisation with existing security regulators, such as the Australian Signals Directorate or the Chief Information Officers Group. However, this article would argue that the impact of cyber vulnerabilities should be managed by the capability manager for joint C4ISR, space and cyber.

Conclusion

Within Defence, cyber resiliency must be incorporated across the entire capability lifecycle, and the associated policy development needs to be cognisant of this holistic need. The specification of cyber-resiliency requirements needs to be improved, further importance needs to be placed on cyber security in the architecture and design phases, and cyber-security test and evaluation needs to be an integral part of overall test activities. To achieve this, a new Defence organisation has been proposed and its broad responsibilities outlined. The burning question that remains is what sequence of actions do we need to take to organise the Defence enterprise to adapt to the rapidly changing cyber environment?

The first step should be to implement program protection plans across major acquisition projects, move foundational cyber-hygiene processes into the statement of work, rather than the functional performance specification, and reserve those functional cyber security requirements for only those needs specific to the capability being acquired.

The second step would be to build the cyber-security test expertise and infrastructure through targeted trials to create a body of knowledge that can be leveraged by newer systems to enable more resilient design. The third step should be the joint development of a cyber-resiliency toolkit through a series of risk reduction activities, which would provide off-the-shelf modules that can be incorporated into acquisition projects. The final step should be the adoption of a key performance parameter for cyber resilience, and articulated outcomes for cyber-security test and evaluation on future projects.

This sequence would allow for all stakeholders to develop their cyber-resiliency capabilities over time, and the shared responsibility would drive a collegial partnership to ensure the outcome of a more cyber-hardened and cyber-resilient Defence. Industry, academia and the government in Australia would all need to work together to achieve this vision. The free flow of threat information, security solutions and best practice between these participants would also be essential in providing a cyber-secure future for Defence.

Stuart Fowler is a Systems Engineer working in the Australian defence industry. He has worked on major naval acquisition projects, with a primary focus on information assurance and cyber security. He holds a Bachelor degree in Mechanical and Mechatronic Engineering, and a Masters degree in Engineering Management.

Captain Craig Sweetman, Royal Australian Corps of Signals, has served in a variety of regimental and staff appointments. He has a Bachelor degree in Science (Computer Science and Information Systems), and Masters degrees in Science (Information Technology), Project Management, and Cyber Security Operations. He is currently attending the Capability and Technology Management College at the Australian Defence College.

Flight Lieutenant Sibi Ravindran graduated from Monash University with a Bachelor (Honours) of Engineering (Mechatronics). His postings have included 1 Combat Communication Squadron, Project Zodiac (assessing operational and deployed communication systems) and the Land Network Integration Centre. Flight Lieutenant Ravindran has a Masters degree in Cyber Security and is currently working in Electronic Warfare Operations at the Air Operations Centre, Air Command.

Dr Keith Joiner joined the Royal Australian Air Force in 1985 and, over a 30-year career, became an aeronautical engineer, project manager and teacher. From 2010 to 2014, Group Captain Joiner was the Director-General of Test and Evaluation for the ADF. In 2015, he joined the University of NSW as a senior lecturer in test and evaluation. Dr Joiner is a Certified Practising Engineer and a Certified Practising Project Director.

Dr Elena Sitnikova is a researcher and academic within the Australian Centre for Cyber Security at the University of NSW, teaching inter alia cyber governance. Her main research interests are in critical infrastructure protection and cyber security, quality assurance and enterprise process capability improvement. Elena currently leads the critical infrastructure area, carrying out research projects in cyber security and process control systems with industry, and state and Federal government partners in Australia.

References

Australian Government, *2016 Defence White Paper*, Department of Defence: Canberra, 2016.

Australian Signals Directorate, 'Information Security Manual', *Australian Signals Directorate* [website], available at <https://www.asd.gov.au/infosec/ism/> accessed 31 March 2017.

Beraud, P., A. Cruz, S. Hassell and S. Meadows, 'Using cyber maneuver to improve network resiliency', paper presented at Military Communications Conference, Baltimore, 8-10 November 2011, pp. 1121-6, available at <http://ieeexplore.ieee.org/document/6127449/> accessed 23 June 2017.

Bergmann, K., 'Capability development reform: a work in progress', *Asia-Pacific Defence Reporter*, Vol. 39, Issue 10, 2002, p. 34 (subscription only).

Brown, C., P. Christensen, J. McNeil and L. Messerschmidt, 'Using the developmental evaluation framework to right size cyber T&E test data and infrastructure requirements', *ITEA [International Test and Evaluation] Journal,* Vol. 36, Issue 1, 2015, pp. 26-34 (subscription only).

Cherdantseva, Y., P. Burnap, A. Blyth, P. Eden, K. Jones, H. Soulsby and K. Stoddart, 'A review of cyber security risk assessment methods for SCADA systems', *Computers & Security,* Issue 56, 2016, pp.1-27 (subscription only).

Commonwealth of Australia, *Australia's Cyber Security Strategy: enabling innovation, growth & prosperity*, Department of the Prime Minister and Cabinet: Canberra, 2016.

Ercis, A. and M. Unalan, 'Analysis of the world's most innovative companies on the basic of industry: 2005-2014', *Procedia-Social and Behavioral Sciences*, Issue 195, 2015, pp. 1081-6, available at <https://openaccess.firat.edu.tr/xmlui/bitstream/handle/11508/8385/Analysis%20of%20The%20World%27s%20Most%20Innovative%20Companies%20on%20The%20Basic%20Of%20Industry%3A%202005-2014.pdf> accessed 30 May 2017.

Freeman, B., 'A new defence for Navy ships: protection from cyber attacks', *Office of Naval Research* [website], 17 September 2015, available at <https://www.onr.navy.mil/en/Media-Center/Press-Releases/2015/RHIMES-Cyber-Attack-Protection> accessed 20 May 2016.

Hawksworth, J., and D. Chan, 'The world in 2050: will the shift in global economic power continue?', *PricewaterhouseCoopers* [website], February 2015, available at <http://www.pwc.com/gx/en/issues/the-economy/assets/world-in-2050-february-2015.pdf> accessed 15 July 2016.

Joiner, K., 'How Australia can catch up to US cyber resilience by understanding that cyber survivability test and evaluation drives defense investment', *Information Security Journal*, Vol. 26, Issue 2, 2017, abstract available at <http://www.tandfonline.com/doi/abs/10.1080/19393555.2017.1293198> accessed 30 May 2017.

Joiner, K., E. Sitnikova and M. Tutty, 'Structuring defence cyber-survivability T&E to research best practice in cyber-resilient systems', paper presented at Systems Engineering Test and Evaluation Conference, Melbourne, 2016, available at <http://search.informit.org/documentSummary;dn=255735974744392;res=IELENG> accessed 23 June 2017.

McGinn, K., 'A centralized holistic approach to information assurance testing', *ITEA Journal*, Vol. 36, No. 3, September 2015, pp. 208-10 (subscription only).

Müller, J.P., and K. Fischer, 'Application impact of multi-agent systems and technologies: a survey', in Onn Shehory and Arnon Sturm (eds.), *Agent-oriented software engineering: reflections on architectures, methodologies, languages and frameworks*, Springer Berlin Heidelberg: Heidelberg, 2014, pp. 27-53, abstract available at <https://link.springer.com/chapter/10.1007%2F978-3-642-54432-3_3> accessed 30 May 2017.

Office of the Secretary of Defense, 'Procedures for operational test and evaluation of cybersecurity in acquisition programs', memorandum dated 1 August 2014, available at <http://www.dote.osd.mil/pub/policies/2014/8-1-14_Procs_for_OTE_of_Cybersec_in_Acq_Progs(7994).pdf> accessed 30 May 2017.

Paganini, P., 'Hack-proof drones possible with HACMS technology', *InfoSec Institute* [website], 2014, available at <http://resources.infosecinstitute.com/hack-proof-drones-possible-hacms-technology/> accessed 23 May 2016

Ramuhalli, P., M. Halapanavar, J. Coble and M. Dixit, 'Towards a theory of autonomous reconstitution of compromised cyber-systems', paper presented at 2013 IEEE International Conference on Technologies for Homeland Security, 2013, subsequently published in *Homeland Security Affairs,* Supplement 6, April 2014, available at <http://cryptome.org/2014/05/cybersys-reconstitution.pdf> accessed 30 May 2017.

Reed, M., 'System security engineering for program protection and cybersecurity', paper presented at 18th Annual NDIA Systems Engineering Conference, Springfield, Virginia, 29 October 2014, available at <http://www.acq.osd.mil/se/briefs/16994-2014_10_29-NDIA-SEC-Reed-SSE-PP-vF.pdf> accessed 30 May 2017.

Richards, R., 'High assurance cyber military systems', *Defense Advanced Research Project Agency* [website], 2016, available at <http://www.darpa.mil/program/high-assurance-cyber-military-systems> accessed 20 May 2016.

Slezak, M., 'Unhackable kernel could keep all computers safe from cyber attack', *New Scientist*, 16 September 2015, available at <https://www.newscientist.com/article/mg22730392-600-unhackable-kernel-could-keep-all-computers-safe-from-cyberattack-2> accessed 21 May 2016.

Smith, N.C.; E.D. White, J.D. Ritschel and A.E. Thal, 'Counteracting harmful incentives in DoD acquisition through test and evaluation and oversight', *ITEA Journal*, Issue 37, 2016, pp. 218-26 (subscription only).

Snyder, Don, James D. Powers, Elizabeth Bodine-Baron, Bernard Fox, Lauren Kendrick and Michael Powell, 'Improving the cybersecurity of US Air Force military systems throughout their life cycles', *RAND Corporation* [website], 2015, available at <http://www.rand.org/pubs/research_reports/RR1007.html> accessed 30 May 2017.

Van Antwerpen, C., and D.K. Bowley, 'An Australian approach to concept development and experimentation: linking strategy to capability', *Journal of the Operational Research Society*, Vol. 63, Issue 2, pp. 278-92, abstract available at <https://link.springer.com/article/10.1057/jors.2011.28> accessed 30 May 2017.