Taylor & Francis
Taylor & Francis Group

Australian Journal of Multi-Disciplinary Engineering

ISSN: 1448-8388 (Print) 2204-2180 (Online) Journal homepage: https://engineersaustralia.tandfonline.com/loi/tmul20

Architectural choices for cyber resilience

Geoffrey Brennan, Keith Joiner & Elena Sitnikova

To cite this article: Geoffrey Brennan, Keith Joiner & Elena Sitnikova (2019) Architectural choices for cyber resilience, Australian Journal of Multi-Disciplinary Engineering, 15:1, 68-74, DOI: 10.1080/14488388.2019.1664210

To link to this article: https://doi.org/10.1080/14488388.2019.1664210

Published online: 06 Sep 2019.

AUSTRALIAN JOURNAL OF MULTI-DISCIPLINARY ENGINEERING
2019, VOL. 15, NO. 1, 68–74
https://doi.org/10.1080/14488388.2019.1664210

Architectural choices for cyber resilience

Geoffrey Brennan [a], Keith Joiner [b] and Elena Sitnikova [c]

[a]School of Engineering and Information Technology, UNSW, Canberra, Australia; [b]Capability Systems Centre, UNSW, Canberra, Australia; [c]UNSW Canberra Cyber, UNSW, Canberra, Australia

ABSTRACT

Open system architectures have been proposed as a solution to addressing issues surrounding the acquisition of mission-critical platforms and critical infrastructure. These issues relate to the extended timeframes associated with introducing advanced and complex systems at a rate commensurate with technological and threat advancement. Adopting an open system architecture helps to address these issues but comes with considerations that influence the capacity for those systems to maintain cyber resilience. Generally, these considerations relate to the use of openly available interface standards, the capability for dynamic network management, maturity of processes and, the ability for platforms to evolve in response to emerging cyber threats. This paper discusses advantages and disadvantages of a bespoke or open architecture and highlights the impact of each approach on mission-critical platforms. The resulting recommendations focus on using a project's projected lifecycle to indicate the most suited architecture and therefore leverage the relative benefits of each methodology.

ARTICLE HISTORY
Received 11 June 2019
Accepted 28 August 2019

KEYWORDS
Open system architecture; cyber resilience; system engineering

1. Introduction

As described in Moore's law, every year communications and information systems continue to undergo exponential growth in capability (Mack 2015). Concurrently the cycles of adaptation by cyber-attackers are also increasing, placing a difficult challenge on Managed Security Service Providers (MSSPs) (Chan 2018), defended mission-critical systems and the many legacy systems without cybersecurity evaluation or defence (Joiner and Tutty 2018). This presents a dilemma for traditional project management and acquisitions processes whereby the speed in technological advancement and cyber-threat outpaces the capacity to introduce contemporary systems into service. One proposed solution to this issue has been to adopt open architectures in major systems procurement and allow for modular system components (Serbu 2013). An open architecture is 'a technical architecture that adopts open standards supporting a modular, loosely coupled and highly cohesive system structure that includes publishing of key interfaces within the system and full design disclosure.' (Department of Defense 2013) This concept supports a faster acquisition cycle by allowing for components to be upgraded and replaced as technology evolves without having to re-engineer the entire system, but it is predicated on robust and well-defined interface standards (Rose et al. 2014).

This concept has been cited as a possible solution to issues surrounding the procurement of mission-critical military platforms as a means of addressing shortfalls in the system update cycle (Serbu 2013). One of the early adopters was the US Navy with their use of the Open Architecture framework for the delivery of major platforms which, in turn enabled key components such as submarine sonar arrays to maintain an upgrade cycle commensurate with technological advances (Stevens 2008; Castelle, Dean, and Daniels 2019). Following this success, a number of other programmes and standards have been introduced, including the Department of Defence Modular Open Systems Approach (MOSA), detailing the methodology and processes to be followed during systems procurement, an approach which is also now incorporated into the Defence Acquisition Guidebook (Department of Defense 2013) and National Defence Authorisation Act (NDAA) of 2017 (Schwartz and Peters 2018). Other military programmes to build on the open architecture concept include the US Air Force's Mission Systems Open Architecture Science and Technology (MOAST) programme (Littlejohn et al. 2017), the US Army's Common Operating Environment (COE) programme (Heininger 2016), the Future Airborne Capability Environment (FACE) Technical Standard (Koontz and Johnson 2015) as well as the UK Government's Open Standards Principles (Pearson et al. 2015). The utility of this approach, particularly when considering the additional complexity of coalition operations, has also resulted in the development of NATO Standardisation Agreement (or STANAG) 4754 describing a NATO Generic Vehicle Architecture (Dawson 2018). The approach is perhaps best exemplified in the United States

CONTACT Geoffrey Brennan geoffrey.brennan@student.unsw.edu.au

Navy's 'Compile to Combat in 24 Hours (C2C24)' framework, which has four principles (U.S. DoD 2018). The first principle is to use a common data standard, such as the Extensible Mark-up Language (XML) for security to data elements. The second principle is using shared architecture to reduce the attack surface. Third, is to simplify risk management by modularising applications and only certifying new modules in any patching, while fourth is to use commercial cloud-based services so as to access advanced data analytics and align with patching services.

Wider adoption of this approach has also been considered for use within Australian Defence Force capability acquisition programs (Toohey 2017) beginning with the Generic Vehicle Architecture (Dawson 2018), however, adopting an open architecture does have a number of considerations for the system's cyber-resilience that must be understood to allow for an informed risk-based decision to be made.

When considering the relative cyber-resiliency benefits for a purely open or bespoke system architecture, this paper uses the MITRE cyber-resiliency goals of anticipate, withstand, recover and evolve (Bodeau, Graubart, and Laderman 2014) as the framework for analysing the relative merits of bespoke and then open architectures before discussing the architectural overheads and future solutions. The paper concludes with recommendations and future research.

2. Bespoke architectures

A bespoke architecture is, by its nature inherently integrated, which can present fewer system gaps to be exploited, enabling it to better withstand an active threat. This enables a coherent approach to technical 'defence in depth' with fewer system inconsistencies (Bodeau, Graubart, and Laderman 2014). It also allows for the system design to be hardened from the outset, provisioning only essential connectivity requirements and limiting those that can exploited (Littlejohn et al. 2017). While the bespoke integration concept requires more time and resources upfront, it also supports the strict control of system baselines (Pearson et al. 2015). Knowing when threat actors are present on a system and what they are attempting to achieve requires a thorough understanding of normal operations to enable the detection of the abnormal or 'absence of the normal.' The relatively static system design resulting from a bespoke architecture enables the establishment of this baseline. Conversely, the dynamic nature of changing components and functionalities in open architecture can result in unpredicted consequences thus increasing vulnerabilities and attacks vectors which therefore require constant investment in evaluation activities to identify these issues and mitigate them.

Many bespoke systems are designed without security built-in and have unique or proprietary designs, necessitating the use of contractual support to conduct detailed system monitoring (Stevens 2008) as well as provide the function of sensor fusion and analysis (Bodeau, Graubart, and Laderman 2014). This limits the ability for system managers and operators to adequately understand the network's performance and anticipate threats, effectively outsourcing the capacity for dynamic mapping and profiling to commercial vendors. While outsourcing can sometimes support cyber-resiliency by offloading the management of non-essential functions to more capable providers (Bodeau, Graubart, and Laderman 2014), the mission criticality of major military platforms and the limitations in getting contracted support into theatres of conflict often precludes this from being viable option within the military context.

One key cyber security advantage drawn from bespoke architectures is the ability to prevent or avoid threats based on the increased requirement for a nefarious actor to conduct detailed reconnaissance in support of their attacks. The unique configuration and protocols in use require a greater effort from threat actors to understand the system and design vectors to achieve their malicious intent within it. In a limited sense, this contributes to the technique of creating unpredictability as a means of supporting cyber resilience by ensuring that systems are not constrained to standard designs with their inherent limitations (Bodeau, Graubart, and Laderman 2014). However, this implies a reliance on secretive design to support the overall cyber security strategy which makes bespoke architectures particularly vulnerable to insider threats or intellectual property theft (Sledge 2015). Once the system details are known, there is likely to be limited capacity to adopt a dynamic reconfiguration technique for cyber resiliency and negate or delay the attack in progress (Bodeau, Graubart, and Laderman 2014).

A bespoke system architecture is also limited in its ability to anticipate and address emerging cyber threats by being largely restricted to the threat analysis done during its initial system design (Fahey 2015). Initial designs are usually tailored around the known and persistent threats of the time, however, the requirement to commit to the complete system '*as designed*' makes it relatively incapable of, or too cost prohibitive to, support an ongoing threat analysis feedback cycle in the long term. In this way, the primary disadvantage of a bespoke system is usually the limited capacity to rapidly recover and evolve from a threat. Once a threat has been identified, the bespoke architecture means that minor upgrades often require major engineering changes, usually contractually constrained to a single vendor (Pearson et al. 2015). This means that systems often have prolonged periods of exposure to already identified

threats and remediation efforts become particularly expensive. The relatively restrictive system design can also limit the overall capacity to make changes without requiring a completely new system.

By committing to a complete system design, there is also a certain commitment to maintaining the technology for the full life of the product which induces additional risks on the through-life supply-chain guarantee and service (Alberts et al. 2017). The later supply-chain arrangements are a factor recently termed as cyber-provenance, due to the cyber security advantages of having all aspects of critical systems locally sourced (or close by alliance) (Joiner 2018) or trusted (Hakan and Cagal 2016). This can become increasingly difficult for bespoke designs with long periods in-service. Techniques to perform such risk assessments early are readily available (U.S. NIST 2015) including in a systems engineering context (Nejib, Beyer, and Yakabovicz 2017) and Defence context (U.S. DoD 2015); albeit these are not yet in widespread use in Australia (Joiner, et al., 2018). Towards the end of the platform's lifecycle, the system is also required to maintain operations using potentially outdated software and hardware that can be more expensive to operate than more readily available and technologically superior products (Pearson et al. 2015). In addition to the increased supply chain costs of maintaining obsolete and proprietary components, there is also an increased cost associated with engineering effort to address known exploits (Stevens 2008). The capacity to develop new cyber-defences on outdated systems is limited as it is constrained by the existing hardware. When coupled with the bespoke system design used, cyber-threat mitigations often require expensive and novel mitigations that take significantly more time and resources to develop.

3. Open architectures

Castelle, Dean, and Daniels (2019) recently found that an open-systems approach to ship development '*offers numerous benefits including increased return on investment (ROI) over the service life and total lifecycle, parallel development, reduction of development costs and duration, rapid prototyping of payloads, improved flexibility through standard interfaces and equipment, efficient technology refresh, and the ability to postpone decision-making under uncertainty*'; all of which assists cyber-resilience. Their research examined challenges in introducing agile development, especially in software and IT, into the longer and somewhat conservative ship development problems. The hierarchical and siloed organisational structure in Defence departments and the bureaucratic processes of policies, directives, orders, committees and commands, can make the agile practices called for by Bishop et al. (2017), Schreider (2017), and Castelle, Dean, and Daniels (2019) difficult, which in turn limits the cooperation, collaboration, information sharing, and synchronised action necessary for cyber operations and resilience. Industry is less constrained in implementing reform such as the 'DevSecOps' movement outlined by Matteson (2017), whereby there is a blending of traditional IT security approaches with agile collaboration approaches.

With an open architecture comes known interface standards that enable the monitoring and auditing of system data in a manner that cannot typically be achieved in solely proprietary systems by focusing on the network interfaces between components (Littlejohn et al. 2017). The well-defined and understood interface standard allows for higher fidelity system analysis that itself can be considered a system component and upgraded at a faster rate (Rose et al. 2014). This empowers the system managers and operators to identify threats and respond appropriately without a reliance on contracted support, minimising the overall impact of a cyber-threat. This directly enables a vast array of the cyber resiliency objectives by catering to the 'understand' function and providing accurate threat-based intelligence to better tailor responses and security controls such as network anomaly detectors and cyber-dashboards (Bodeau, Graubart, and Laderman 2014).

It should be noted however that a detailed understanding of the prescribed interface does not constitute a complete understanding of the system. The existence of proprietary sub-systems representing a '*black box capability*' from the platform perspective means that this monitoring and damage assessment is limited to the interface layer (Rose et al. 2014). Threats can still persist within sub-systems and would require the same vendor support required in a bespoke architecture albeit, at a reduced scale. Furthermore, the constantly changing system baseline that comes from rapidly changing components requires regular system analysis to occur in order to ensure that a normal state can be established for the system despite significant changes (Pearson et al. 2015). Without this, all forms of analytic monitoring, including risk monitoring through Cyber Table-Topping (CTT), are significantly hindered, further emphasising the requirement for deliberate test and evaluation stages as part of the open architecture model (Christensen 2017). Cyber Table Topping builds on regular cooperative vulnerability and penetration testing by systematically assessing how new threats are targeting a network or system, the likelihood of exploitation, and the consequences so as to risk manage until the next test period (Christensen 2017). Such a management strategy aligns with how configuration control boards manage obsolescence and mission creep.

An open architecture design has a potentially contradictory influence with relation to the '*withstand*' goal of building cyber-resilience. The same detailed and published standards that allow for modularity in open architecture systems, thereby supporting the '*design once, use many times*' concept (Stevens 2008), also

provides valuable system information enabling a threat actor's reconnaissance phase. In order to enable vendor competition, these standards must be well defined and openly available within the community-of-interest so as to enable sub-system integration, but doing so potentially makes this information available for threat actors to understand system limitations and therefore develop exploits (Littlejohn et al. 2017). Chan (2018) overviews the phenomenon of several open source tools being used by cyber-attackers as '*accelerants*', before then proposing artificial intelligence as a countermeasure for such accelerants (see also Klemas & Chan):

> 'Cyber attackers are becoming increasingly adept. Just as MSPs [managed service providers] and MSSPs [managed security service providers] are leveraging early warning indicators, such as the National Vulnerability Database (NVD) and Sentient Hyper Optimized Data Access Network (SHODAN), cyber attackers are also leveraging these assets for exploitation opportunities and as attack accelerants.'

The presence of a core network within the system that provides the component interface means that malware also has the potential to move laterally though the system relatively easily and attack other components or a conduct a coordinated system wide attack (Littlejohn et al. 2017). This makes the prospect of isolating a threat considerably more difficult. Any inherent flaws in the design of the architecture's interface can pose a risk to all the modular components utilising it and is therefore a risk that must be owned by the client mandating the standards (Pearson et al. 2015). In this way, when considering attack vectors, the components themselves can act as hard surfaces, with the common interface becoming exploitable system gaps. By design, these interface standards are relatively rigid, and the discovery of any significant exploits may require platform-level changes in a similar manner to that required in a bespoke system architecture.

With regards to system recovery, provided a threat is contained to a specific or limited set of system components, the open architecture design allows for rapid system recovery by using the cyber-resiliency technique of dynamic reconfiguration to negate or curtail adversary attacks (Bodeau, Graubart, and Laderman 2014). This concept allows for individual components to be rapidly replaced, giving the system owner options to restore mission-critical function as soon as possible with a lower requirement for vendor support or the need for a full system redesign (Littlejohn et al. 2017). Through the better known and well-defined open interface standards, this process can be intelligence-led, supporting adaptive management in the face of identified threats using the higher fidelity network monitoring and damage assessment (Bodeau, Graubart, and Laderman 2014).

The greatest advantage in using open system architectures to support cyber resilience lies in the capacity to evolve and address system vulnerabilities with rapid component upgrades. If a specific component is identified to be a vulnerability, that individual component can be redesigned and replaced without having to redesign the system as a whole (Pearson et al. 2015). The lifecycle of components can also therefore be reduced to shorter periods, allowing them to keep pace with technological and threat advancement (Stevens 2008). Rather than requiring novel threat mitigations, an open architecture system can simply upgrade the hardware and leverage industry-known mitigations for discovered exploits, reducing the cost and time required to mitigate future attacks (Rose et al. 2014). The vendor competition that will result from having their component rapidly replaced following an untreated cyber vulnerability also promotes security as an essential service, to be incorporated by default into the initial design phases.

4. Architectural overheads and future solutions

Joiner et al., 2018a outline in some depth the ICT governance challenge of developing and maintaining cyber-resilient systems, especially in the necessary software test infrastructure and skills which are receiving renewed focus (Knyish 2019; Pittet 2018; Dougall 2018; Pearson 2015; Vassilev and Celi 2014). While they outline that cybersecurity testing can leverage extant ICT or software testing like usability testing and performance testing to some extent, there are also unique cybersecurity test infrastructure demands and skills which are in short supply. This practical research largely assumes elements of the ICT infrastructure, even if only at the aggregated system-of-system level, are bespoke and that elements of the threat emulation and vulnerability are either classified, or restricted in some way as to limit commercial outsourcing; for example the need for timely monitoring, in-house risk management (i.e. cyber table-topping) and response (i.e. unable to be delegated). Put another way, this research suggests that an organisation must grow in-house capability for cyber-resilience, especially testing (Fowler et al. 2017) and what Australian Defence now terms cyberworthiness (Australian Senate 2018a). That assumed obligation is magnified when systems are bespoke, since the community that would enable secure outsourcing and have the necessary tools to enable and automate it are less likely to exist commercially.

The advantage of open architectures in outsourcing security and keeping pace with cybersecurity is outlined by Chan (2018) as follows, similar to the USN C2C24 Principle 4 outlined earlier:

> 'MSP responsibilities are increasingly shifting from repairs, patches, delivery of new software, and incorporation of cloud services to that of data-related security services. According to Gartner, a new class of MSP, the Managed Security Service Provider (MSSP), has emerged to provide outsourced monitoring and management of security devices and systems. Prototypical managed services now include, among others, managed firewall, virtual private network, vulnerability scanning, anti-viral services, and intrusion detection. Outsourcing to MSSPs has typically improved the client ability to deter cyber-threats … . MSSPs have burgeoned not only in industries that have experienced massive compromises in recent times (e.g. healthcare), but also in areas that are at unprecedented levels of risk (e.g. energy sector).'

An example of this trend in the Australian public sector is the recent hearing in the Australian Senate that focused a lengthy line of questioning on why a cloud provider had achieved high-assurance accreditation for the Government in a comparatively shorter timeframe and into an area hitherto restricted to boutique Australian providers (Australian Senate 2018b). The short answer to this high assurance approach was that it was required for the purposes of necessity, risk management and expertise (see also Joiner et al. (2018), for a coverage of high assurance challenges).

There does appear to be a further solution developing in the U.S. to cybersecurity of systems, whether they are open architecture or bespoke, or more precisely where they lie on that spectrum. Developed as part of the U.S. efforts on autonomy, it is described as the concept of cyber-sidecars (U.S. DSB, 2016), where artificially intelligent monitors are placed onto systems to provide oversight, take critical actions if necessary and provide a degree of evidence for attribution as well as improve resilience afterwards. These systems would of course be controlled and highly secure; ideally with the adaptability to 'deep learn' bespoke system behaviours as much as the more readily available open architectural behaviours. Such artificially intelligent overwatch capabilities are already in use in networking on many government departments and critical industries but they focus at present on common networking traffic areas and require high levels of personnel support and expertise to run in each context. The research by Chan (2018) illustrates an open-source networking security stack that is somewhat analogous to such current controlled and artificially intelligent cybersecurity systems. The cyber-sidecar context is envisaged to be much more universal and suited to implant to more standalone systems, much like controlled cryptographic systems operate. The spectrum of open to bespoke architecture may in the future simply dictate the cost and extent of deep learning the cyber-sidecar(s) must undertake.

5. Conclusion

The decision to adopt a bespoke or open architectural approach in a mission critical system should largely depend on the projected lifecycle for the system. The limited capacity for bespoke systems to evolve over time means that they become increasingly vulnerable to cyber threats following initial design but can present a more holistically hardened system at the outset. On the other hand, open architecture systems are far more reliant on well documented interfaces that provide a possible attack vector but are also able to adapt in response to a cyber threat in a manner that far exceeds that of a bespoke system. In this way an open architecture can rapidly evolve to address new threats as they emerge. In deciding between the two, it is also important for the system owner to consider their capacity to provide a detailed and robust interface standard and therefore accept a higher degree of liability themselves or, delegate it to a vendor to define and control the system's operation. This will have significant impacts on the ongoing system support plan and associated resourcing.

To determine between the two architectural approaches, or indeed to appropriately situate an acquisition and sustainment strategy on the spectrum of these approaches, an early cybersecurity assessment is required into the supply chain through-life; that is the degree of cyber-provenance required. That assessment also should highlight the likely greater in-house test and management overheads of cyber-securing bespoke architectures. Therefore, to support a system's cyber resilience in the most efficient manner the following recommendations are made:

- For major projects with significantly large scope and expected period in-service: an open system architecture is favoured to support system evolution throughout its lifecycle.
- For minor projects, anticipated to provide a short duration or minor capability: a bespoke architecture is favoured to minimise the onus on the customer to own and manage standards or interfaces whilst also outsourcing the ongoing system analysis function.
- Consideration be given in the future (particularly for bespoke and legacy systems) to use cyber-sidecars involving deep-learning artificially intelligent systems, noting however that these must also be to be tightly controlled so as not to act as adversarial accelerants.

Acknowledgments

The initial findings of this research were presented at the Systems Engineering, Test and Evaluation (SETE) conference

in Canberra in April 2019 and delegates contributed numerous questions and lines of further reflection as well as the encouragement to publish the work for wider use.

Disclosure statement

No potential conflict of interest was reported by the authors.

Notes on contributors

Geoffrey Brennan is currently studying at the University of New South Wales Canberra and has a background in deployable information and communications technology.

Keith Joiner, CSC, joined the Air Force in 1985 and became an aeronautical engineer, project manager and teacher over a 30-year career before joining the University of New South Wales in 2015 as a senior lecturer in test and evaluation. From 2010 to 2014 he was the Director-General of Test and Evaluation for the Australian Defence Force, where he conducted numerous joint test and evaluation activities and where he was awarded a Conspicuous Service Cross. He is a Certified Practising Engineer and a Certified Practising Project Director with over 60 major academic presentations and publications.

Elena Sitnikova is an academic and researcher, and the Program Coordinator for the Master in Cyber Security program at the University of New South Wales (UNSW) Canberra. She holds a Bachelor degree in electrical engineering (Honours) and PhD in Communication Control systems. Her current research interests are in critical infrastructure protection area, carrying out research projects in the area of intrusion detection (IDS) for control systems cyber security and Industrial IoT (IIoT). Elena is a Senior Fellow of the Higher Educational Academy (SFHEA), an award winning academic, holding a national Australian Office for Learning and Teaching (OLT) Team Citation award for Outstanding Contributions to Student Learning.

ORCID

Geoffrey Brennan http://orcid.org/0000-0003-0885-808X
Keith Joiner http://orcid.org/0000-0001-6081-3239
Elena Sitnikova http://orcid.org/0000-0001-7392-0383

Future Research

Current doctoral research work by the authors in Australia focus in four main areas. First, the technical regulatory framework (engineering policies and procedures) for cyber-worthiness of developing and developed systems is under first-principles examination and survey among affected engineers. Second, the efficacy of an industry-evolved framework for cyber-security assessments during design acceptance is being evaluated, including this framework's residual utility for on-going through-life cybersecurity risk re-assessments (i.e. support to cyber table-topping). Third, seminal managers and practitioners in cybersecurity security operation centres are being surveyed for the necessary competencies to determine independently if there are foundational differences for a middle-power vice a leader like the United States. Finally, new research is proposed on using behavioural propagation (aka infection) for intrusion detection along epidemiological lines.

References

Alberts, C., J. Haller, C. Wallen, and C. Woody. 2017. "Assessing DoD System Acquisition Supply Chain Risk Management." *CrossTalk* 30 (3): 4–8.

Australian Senate. 2018a. "Testimony of Vice Admiral Griggs to Senate Estimates for Foreign Affairs Defence and Trade." Accessed 29 May 1900. http://parlview.aph.gov.au/mediaPlayer.php?videoID=399539&operation_mode=parlviews

Australian Senate. 2018b. "Testimony of Mr MacGibbon to Senate Estimates for Foreign Affairs Defence and Trade." 29 May. circa 1800 hours. http://parlview.aph.gov.au/mediaPlayer.php?videoID=399539&operation_mode=parlviews

Bishop, M., L. Futcher, N. Miloslavskaya, and M. Theocharidou 2017. "Information Security Education for a Global Digital Society: 10th IFIP WG 11.8 World Conference." Proceedings, Springer International Publishing, Rome, Italy: WISE 10, May 29-31.

Bodeau, D. J. J., R. D. D. Graubart, and E. R. R. Laderman. 2014. "Cyber Resiliency Engineering Overview of the Architectural Assessment Process." *Procedia Computer Science* 28 (C): 838–847.

Castelle, K. M., A. W. Dean, and C. B. Daniels. 2019. "Benefits and Challenges of Implementing Agile Development in Modular Shipbuilding." *Naval Engineers Journal* 131 (2): 1–11.

Chan, S. 2018. "Prototype Orchestration Framework as A High Exposure Dimension Cyber Defense Accelerant Amidst Ever-Increasing Cycles of Adaptation by Attackers: A Modified Deep Belief Network Accelerated by A Stacked Generative Adversarial Network for Enhanced Event Correlation." CYBER 2018: The Third International Conference on Cyber-Technologies and Cyber-Systems, Athens Greece: IARIA, November 18-22.

Christensen, P. H. 2017. "Cybersecurity Test and Evaluation: A Look Back, Some Lessons Learned, and A Look Forward!" *The ITEA Journal of Test and Evaluation 38 (3);* 221–228.

Dawson, M. 2018. "Australian Generic Vehicle Architecture."MilCIS. https://static1.squarespace.com/static/5274112ae4b02d3f058d4348/t/5c09f64a88251becf5a9e12d/1544156751092/2018-3-6d.pdf

Department of Defense. 2013. "Open Systems Architecture: Contract Guidebook for Program Managers." US Department of Defense Open Systems Architecture Data Rights Team, v.1.1.

Dougall, S. 2018. "A Quick Guide to Essential Types of Software Testing." Computer Business Review, February 1.

Fahey, Kevin. 2015. *Integrating Innovation: Keeping The Leading Edge*. Logistics and Technology: 5-7. https://apps.dtic.mil/docs/citations/AD1015974

Fowler, S., ., C. Sweetman, S. Ravindran, K. F. Joiner, and E. Sitnikova. 2017. "Developing Cyber-security Policies that Penetrate Australian Defence Acquisitions." *The Australian Defence Force Journal* (Issue 102); 17–26.

Hakan, H., and U. Cagal 2016. "A Reputation Based Trust Center Model for Cyber Security." 2016 4th International Symposium on Digital Forensic and Security (ISDFS). https://ieeexplore.ieee.org/abstract/document/7473508

Heininger, C. 2016. *The ABCs of COE*. US Army Acquisition Support Centre. January 11. Accessed 21

October 18. https://asc.army.mil/web/tag/armys-common-operating-environment-coe/

Joiner, K. F. 2018. "Submarine Design Test & Evaluation: Challenging Cyber-Provenance & Cyber Sidecars." Presented as part of Panel 4, 'T&E of Emerging Threats in Complex Environments', 35th ITEA Annual Test and Evaluation Symposium, Oxnard, California, December 10-13.

Joiner, K. F., A. Ghildyal, N. Devine, A. Laing, A. Coull, and E. Sitnikova. 2018a. "Four Testing Types Core to Informed ICT Governance for Cyber-resilient Systems." *International Journal of Advances in Security* 11: 313–327.

Joiner, K. F., and M. G. Tutty. 2018. "A Tale of Two Allied Defence Departments: New Assurance Initiatives for Managing Increasing System Complexity, Interconnectedness, and Vulnerability." *Australian Journal of Multidisciplinary Engineering* 14: 4–25. doi:10.1080/14488388.2018.1426407.

Klemas, T. J., and S. Chan 2018. "Harnessing Machine Learning, Data Analytics, and Computer-Aided Testing for CyberSecurity Applications: Achieving Sustained Cyber Resilience for Typical Attack Surface Configurations and Environments." CYBER 2018: The Third International Conference on Cyber-Technologies and Cyber-Systems, Athens Greece: IARIA, November 18-22.

Knyish, N. 2019. "Why Software Testing Is Worth The Investment."March 28. Business.com

Koontz, R., and D. Johnson. 2015. "Apache Mission Processor Software Architecture: Future Airborne Capability Environment (FACE™) Considerations." *Journal of the American Helicopter Society* 60 (1): 011004. doi:10.4050/JAHS.60.011004.

Littlejohn, K., V. Rajabian-Schwart, N. Kovach, and C. Satterthwaite 2017. "Mission Systems Open Architecture Science and Technology (MOAST) Program." Proc. SPIE 10205, Open Architecture/Open Business Model Net-Centric Systems and Defense Transformation 2017, Anaheim, Accessed 25 April 2017. 1020504

Mack, C. 2015. "The Multiple Lives of Moore's Law." *Spectrum* IEEE 52 (4): 31. doi:10.1109/MSPEC.2015.7065415.

Matteson, S. 2017. "DevSecOps: What It Is and How It Can Help You Innovate in Cybersecurity." ZDNet. April 3. https://www.zdnet.com/article/devsecops-what-it-is-and-how-it-can-help-you-innovate-in-cybersecurity/

Nejib, P., D. Beyer, and E. Yakabovicz 2017. "Systems Security Engineering: What Every System Engineer Needs to Know." 27th Annual INCOSE Int. Symp., Adelaide, July 434–445. doi:10.1002/iis2.2017.27.issue-1.

Pearson, G., R. Smith, H. Tripp, and O. Worthington2015. "A Systems Approach to Achieving the Benefits of Open and Modular Systems." Proc. SPIE 9479, Open Architecture/Open Business Model Net-Centric Systems and Defense Transformation 2015, 94790A, Baltimore, May 21.

Pearson, L. 11 September 2015. *The Four Levels of Software Testing*. Segue Technologies.

Pittet, S. 2018. "The Different Types of Software Testing." Atlassian CI/CD. https://www.atlassian.com/continuous-delivery/software-testing/types-of-software-testing

Rose, L., J. Shaver, Q. Young, and J. Christensen 2014. "Open Architecture Applied to Next-generation Weapons." Proc. SPIE 9096, Open Architecture/Open Business Model Net-Centric Systems and Defense Transformation 2014, Accessed 17 June 2014. 90960K.

Schreider, T. 2017. "Building Effective Cybersecurity Programs: A Security Manager's Handbook." edited by A. Noakes-Fry. Rothstein Publishing, Brookfield.

Schwartz, M., and H. M. Peters January 2018. "Acquisition Reform in the FY2016–FY2018 National Defense Authorization Acts (ndaas)." In Congressional Research Service Report for Members and Committees of Congress.

Serbu, J. 2013. "DoD Brings Culture of Open Architecture to a World of Proprietary Systems." Federal News Network. November 13. Accessed 21 October 18. https://federalnewsnetwork.com/defense/2013/11/dod-brings-culture-of-open-architecture-to-a-world-of-proprietary-systems/

Sledge, C. 2015. *A Discussion on Open-Systems Architecture*. Carnegie Mellon University Software Engineering Institute. November 23. Accessed 21 October 18. https://insights.sei.cmu.edu/sei_blog/2015/11/a-discussion-on-open-systems-architecture.html

Stevens, J. 2008. "The How and Why of Open Architecture." *Undersea Warfare* (Issue 37). Spring 2008, Washington: 6–9.

Toohey, K.2017. "Address to the Defence + Industry Fundamental Inputs to Capabilities Conference by Head Land Capability." Australian Army. Accessed 21 October 18. https://www.army.gov.au/our-work/speeches-and-transcripts/address-to-the-defence-industry-fundamental-inputs-to-capabilities

U.S. DoD. 2015. "Cybersecurity T&E Guidebook." Version 1.0, July 1. available online in numerous locations.

U.S. DoD Defense Science Board (DSB). 2016. "Summer Study on Autonomy." Washington: Office of the Under Secretary of Defence for Acquisitions, Technology and Logistics: 28–30.

U.S. DoD, Chief of Naval Operations. 2018. *Transforming Our End to End Information Environment - Compile to Combat in 24 Hours Implementation Framework*. Washington, District of Columbia. December 26. https://www.public.navy.mil/bupers-npc/reference/messages/Documents/NAVADMINS/NAV2018/nav18315.txt

U.S. NIST. April 2015. "Special Publication 800-161 Supply Chain Risk Management Practices for Federal Information Systems and Organizations." https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-161.pdf

Vassilev, A., and C. Celi. 2014. "Avoiding Cyberspace Catastrophes through Smarter Testing." *Computer* 47 (10): 102–106. doi:10.1109/MC.2014.273.