CrossTalk

Sep/Oct 2012 The Journal of Defense Software Engineering Vol. 25 No. 5

Resilient Cyber Ecosystem

RESILIENT CYBER ECOSYSTEMS

Departments

3 From the Sponsor

35 Forum Article

37 Upcoming Events

39 BackTalk

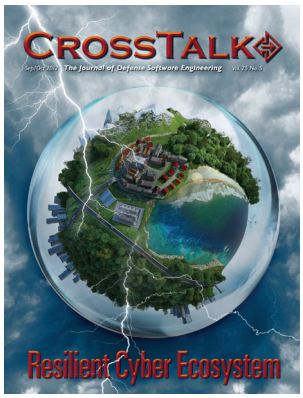

Cover Design by Kent Bingham

Resilient Cyber Ecosystems

4 **The End of the Machine Age**
We will fail to comprehend the essential nature of the cyber domain if we continue to think of it as a technical or engineering system.
by Colin Williams and Tim Watson

8 **Challenges To A Trustworthy Cyber Ecosystem**
It is vital that intrinsic challenges to cyberspace—and software—are recognised and treated such that a trustworthy cyber ecosystem can be formed.
by Ian Bryant and Jasvinder Mahrra

11 **Systems Thinking for a Secure Digital World**
We must learn to think systemically to seek advantage, or at least maintain parity over adversarial threats, as our infrastructure becomes more complexly integrated.
by William D. Miller

15 **Identifying Cyber Ecosystem Security Capabilities**
Strengthening the security and resilience of the cyber ecosystem requires reducing the number of vulnerabilities and the ability to automatically mitigate attack methodologies.
by Peter M. Fonash, Ph.D.

20 **Defining Proactive Software Assurance Practices for Healthier Cyber Ecosystems**
Distributed security in cyberspace can be performed using many real-time components, from intrusion detection systems to incident management systems. However, these components are reactive rather than proactive.
by Brian Badillo and Marc Abrams

25 **Recovery-based Resilient Cyber Ecosystem**
Intrusion tolerance is an approach that treats intrusions as inevitable and shifts the focus from detection and prevention to containing losses and rapid recovery.
by Ajay Nagarajan and Arun Sood

29 **Cyber Mission Resilience: Mission Assurance in the Cyber Ecosystem**
The focus on mission resilience extends the scope of past security practices while simultaneously honing in on mission-critical systems, networks, and processes.
by Chris Peake, Al Underbrink, and Dr. Andrew Potter

CrossTalk

NAVAIR Jeff Schwalb
DHS Joe Jarzombek
309 SMXG Karl Rogers

Publisher Justin T. Hill
Advisor Kasey Thompson
Article Coordinator Lynne Wade
Managing Director Tracy Stauder
Managing Editor Brandon Ellis
Associate Editor Colin Kelly
Art Director Kevin Kiernan

Phone 801-775-5555
E-mail stsc.customerservice@hill.af.mil
CrossTalk Online www.crosstalkonline.org

CrossTalk, The Journal of Defense Software Engineering is co-sponsored by the U.S. Navy (USN); U.S. Air Force (USAF); and the U.S. Department of Homeland Defense (DHS). USN co-sponsor: Naval Air Systems Command. USAF co-sponsor: Ogden-ALC 309 SMXG. DHS co-sponsor: National Cyber Security Division in the National Protection and Program Directorate.

The USAF Software Technology Support Center (STSC) is the publisher of CrossTalk providing both editorial oversight and technical review of the journal. CrossTalk's mission is to encourage the engineering development of software to improve the reliability, sustainability, and responsiveness of our warfighting capability.

Subscriptions: Visit <www.crosstalkonline.org/subscribe> to receive an e-mail notification when each new issue is published online or to subscribe to an RSS notification feed.

Article Submissions: We welcome articles of interest to the defense software community. Articles must be approved by the CrossTalk editorial board prior to publication. Please follow the Author Guidelines, available at <www.crosstalkonline.org/submission-guidelines>. CrossTalk does not pay for submissions. Published articles remain the property of the authors and may be submitted to other publications. Security agency releases, clearances, and public affairs office approvals are the sole responsibility of the authors and their organizations.

Reprints: Permission to reprint or post articles must be requested from the author or the copyright holder and coordinated with CrossTalk.

Trademarks and Endorsements: CrossTalk is an authorized publication for members of the DoD. Contents of CrossTalk are not necessarily the official views of, or endorsed by, the U.S. government, the DoD, the co-sponsors, or the STSC. All product names referenced in this issue are trademarks of their companies.

CrossTalk Online Services:
For questions or concerns about crosstalkonline.org web content or functionality contact the CrossTalk webmaster at 801-417-3000 or webmaster@luminpublishing.com.

Back Issues Available: Please phone or e-mail us to see if back issues are available free of charge.

CrossTalk is published six times a year by the U.S. Air Force STSC in concert with Lumin Publishing <luminpublishing.com>.
ISSN 2160-1577 (print); ISSN 2160-1593 (online)

CrossTalk would like to thank DHS for sponsoring this issue.

Security and Resilience in the Cyber Ecosystem

Similar to a natural ecosystem, the *cyber ecosystem* consists of a community of entities that interact within an environment. The "ecosystem" metaphor, although not perfect, aptly describes important characteristics of cyberspace. Cyberspace is dynamic, and its diverse participants relate to each other in countless complex–and not always healthy–ways.

Adversaries exploit the rich interconnectivity provided by cyberspace, breaching systems with weaker defenses in order to penetrate systems with stronger defenses. The challenge is to use that same rich interconnectivity to collaboratively safeguard users, networks, and devices. In the DHS vision of a healthy and resilient cyber ecosystem, people and devices work together in real time to anticipate and prevent cyber attacks, limit the spread and consequences of attacks, and recover to trusted states. Security capabilities will be built into cyber devices so that preventive and defensive actions can be coordinated within and among decentralized but cooperating communities. By exchanging trusted information, learning and adapting, and coordinating responses in real time, the cyber ecosystem can be open, robust, and healthy.

No one nation or organization owns or controls the cyber ecosystem and as such, cybersecurity is a shared responsibility. DHS recognizes this reality and the *Blueprint for a Secure Cyber Future* < http://www.dhs.gov/files/publications/blueprint-for-a-secure-cyber-future.shtm> is the first strategy focused on the cybersecurity role of the homeland security enterprise, which includes government, nongovernmental, and private sector entities, as well as individuals, families, and communities. Achieving this vision will require determination and cooperation by all of these stakeholders to address the many technical and policy challenges posed by an undertaking of this magnitude. Ultimately, success will depend on the ability to empower all participants in this ecosystem to collectively detect and react faster than adversaries can act.

How will you make your part of the cyber ecosystem more secure? DHS welcomes your thoughts at <cyberfeedback@dhs.gov>, and looks forward to working together with you to build a safe, secure, and resilient cyber ecosystem.

Mark Weatherford
Deputy Under Secretary for Cybersecurity
National Protection and Programs Directorate
Department of Homeland Security

The End of the Machine Age

Colin Williams, De Montfort University
Tim Watson, De Montfort University

Abstract. The totality of human society has become existentially dependent upon the safe and secure operation of the cyber domain. It is impossible to envision a successful military operation without dependable access to the cyber domain. We will fail to comprehend the essential nature of the cyber domain if we continue to think of it as a technical or engineering system. The consequences of such a failure will be catastrophic. Machine Age thinking has become obsolete. There is now an urgent need for the development of a new, fundamentally interdisciplinary and human-centred approach to our understanding of the fifth domain.

Society has become functionally, if not existentially, dependent upon a ubiquitous and pervasive global system of interconnected computer systems. It is no more possible to countenance the human social condition without these computer systems than it is to conceive of life without written and printed language, electricity or the internal combustion engine. This dependence is irreversible and the transformations that it will both enable and command are profound.

Within a few decades the Internet of Things will have become, quite literally, interlaced indivisibly with the material fabric of everyday life. Objects in the home, on the street, in the office and on the battlefield will communicate automatically with other objects, and consequent activity will manifest itself in the corporeal domain in complex and non-linear patterns of cause and effect. The speed of these interactions will increasingly obviate the efficacy of human agency. The observe, orient, decide, and act loop will cycle amplifying data sets at accelerating velocities and human intervention will become an impediment to good outcomes. The complexity of these interactions will make orthodox command and control disciplines dangerously redundant.

At the same time, the fabric and form of computers will transform. Within the lifetime of children now in primary school, humanoid robots will appear in homes and offices. Already, three-dimensional printers have reached the outer fringes of the mass consumer market. By the second half of the current century, most homes and enterprises will have the capacity to transform strings of binary subsisting in the cyber domain into corporeal form and so replicate physical objects as easily and cheaply as they can now print documents. The economy will transform in ways we can only begin to speculate about. As we currently understand them, the boundaries between the real and the virtual will become meaningless.

Relocate the Internet of Things from a civilian to a military context. War fighters and weapons systems will be fully IP addressed. Real time telemetry will be in play. Discharging a round from a personal infantry weapon will, via real time telemetry, trigger actions in the supply chain, including the activation of three-dimensional printers to replicate elements of the depleted stock. Humanoid robots will appear in the battle space. Human agency will be transformed through exoskeletons. Coalition operations will depend upon good metadata as much as on the subtle arts of human liaison. Every link in the sensor to shooter chain, every section of the logistics tail, every item of kit will be interconnected and will be part of a vast amorphous and volatile meta system.

The cyber domain is about far more than a supercharged Command, Control, Communications, Computers, Intelligence, Surveillance, and Reconnaissance (C4ISR) capability or digitising the fog of war. It will effect the most profound transformation in military operations since the invention of gunpowder. It is almost inevitable that this transformation will change completely the relationships and balance between soft and hard power.

The inverse power geometry of asymmetry in kinetic conflicts long recognised in military circles is now apparent in the non-linear matrices of easily anonymised interactions between the cyber and the corporeal domains on a societal scale and in the civil realm. The campaigns against the Stop Online Piracy Act and the Protect IP Act brought us closer than we have been be-

fore to the prospect of orchestrated, massive, anonymised, cross border, civil disobedience. It is only a matter of, ever decreasing, time before one cause or another seeks to transpose the tactics of non-violent mass civil disobedience into the cyber domain. If, or more likely, when, this happens it will happen with lightning speed and with utter disregard to international borders. Conventional law enforcement and engineering solutions alone will be of scant use when it comes to maintaining a resilient cyber ecosystem in the face of this kind of action. The cyber domain is supra national and the tangible expressions of the power of the asymmetric cyber enabled "other" are evidential to the belief that even the might of the most powerful of nation states is confronting a challenge which it is, currently, ill-equipped to face. This is a profound disruption to a narrative forged when computers were an integral element of the Cold War arsenal and the nation state was the epitome of insuperable force.

One of the tragedies of asymmetry in the Information Age is precisely that our own security related business practices and lack of agility continue to inhibit the deployment of IT capability and as such, our own best practices have a predisposition to the perverse outcome of conferring advantage on the irregular opponent. For example, Bring Your Own Device (BYOD) is an opportunity exploited by small, organic firms but, in contrast, its advantages are currently underemployed by more security-conscious, controlling companies, who risk losing competitive advantage and the real benefits of BYOD as a result. Returning to the military context, in a three-block war, the irregular opponent may well be conducting extremely granular and co-ordinated C4ISR through the medium of mobile telephones whilst regular forces are denied an equivalent tactical capability even though the technology exists to grant it.

The Internet will not stop at enabling communication to facilitate the existing forms of the democratic process; it will transform the nature of democracy itself. The current forms of expression of the social contract are rooted in fundamental principles born during the European Enlightenment. Our world would be unrecognisable to Locke, Hobbes, and Rousseau. Yet, we have not yet even begun to discuss matters such as how the legitimate right to protest essential to the democratic process might translate into the cyber domain. Neither have we made sufficient progress toward establishing, let alone codifying, the normative moral and ethical precepts of good behaviour in the cyber domain. Society is now on the brink of having to contend with the formulation of legal definitions of artificial or machine consciousness and intelligence in order to allow law to operate when a computer system, or robot, is cited as the controlling mind. The relationships between the state and the citizen, and perhaps even the shape and nature of these two principal parties to the social contract are set to transform beyond recognition.

The cyber domain is already at the heart of economic prosperity in the sense that without dependable, safe, and trustworthy access to it, even the most conventional of enterprises will struggle to exist, let alone compete. The prospect of running a successful business without computer-based financial accounting, without e-mail, without access to the Internet (which is commonly used for outsourcing payroll operations and the most basic banking services), and without access to the World Wide Web is now as absurd as the prospect of attempting to do so without a telephone or without paper records. This is apart from the use of the Internet in the operation of every aspect of the critical national infrastructure and the dependence upon the Internet of both high street banking and just in time retail logistics. Imagine a turn of events where the cash machines stopped working and bread stopped appearing on supermarket shelves for longer than 24 hours.

We have known that this was coming for some time. In April 1965, barely two decades after Colossus first went into operation, Time carried a lead article that observed that to "process without computers the flood of checks that will be circulating in the U.S. by 1970, banks would have to hire all the American women between 21 and 45." The same article reflected, "Just out of its teens, the computer is beginning to affect the very fabric of society, kindling both wonder and widespread apprehension" and predicted that "swept forward by a great wave of technology … human society is surely headed for some deep-reaching changes." [1]

Our context is now that of the Information Age and although we are a product of all that has gone before, the world we inhabit has been transformed. Over time, the original foundations we used to build the intellectual and cultural constructs, which we still deploy to try to make sense of computers, have dissolved. The overhang of these now derelict constructs is starting to crumble dangerously. We need new and fresh ways of thinking about computers and about the human interactions with them. Our thinking must start from the basis of an examination of the way that computing actually operates in the twenty first century, rather than the way in which the precepts of old tell us that it should. Above all, we are in urgent need of a critical and an interdisciplinary approach to the phenomena of the cyber domain. The story of the cyber domain is principally the story of humans, not that of machines, and humanity is gloriously organic.

Simultaneously, we embrace and celebrate the power and potential of the transformations of the Information Age, whilst fearing both our dependence and the actions of those who would use this vast capability against us and against our way of life. The cyber domain has the potential to be the greatest ally of democracy and its greatest enemy. Which of these it becomes is our responsibility.

As the scale of our dependence becomes ever more apparent and as the awareness of the transformative potential matures, so too does the sense that our current ways of thinking and doing are irrelevant and ineffectual in the cyber domain. We have become terrified by our own creation. There is a palpable and mimetic sense of a cyber-crisis that we express through popular culture, through mainstream journalism, and through increasingly hyperbolic language. Terms such as Cyber Crime, Cyber Terrorism, Cyber War, Cyber Pearl Harbour, and Cybergeddon are commonplace. The paralysis induced by this fear is more apparent amongst the cohort of security experts than amongst the general population.

The successful economies and societies of the Information Age will be built on the assumption that the world is spanned by a safe, secure, and reliable matrix of interconnected computer based information and communications systems operating at speeds and complexities beyond human perception. New economic forms and new types of entrepreneurial behaviour will emerge, not least, as mass access to cyber domain be-

comes even more geographically distributed than it is today. It is unrealistic to assume that economic models spawned by the Western European and Atlantic experiences will endure even the remainder of the current century unchanged.

A real paradox at the heart of all of this is that traditional approaches to security are incapable of generating the trust that must live at the heart of human existence in the cyber domain. Traditional approaches to computing perpetuate fear: fear of the attackers, fear of the insider threat, and fear of the bad effects of doing things with technology. Despite an ostensible move toward risk management, much real world practice displays all the hallmarks of risk avoidance. Security products and services have been sold on the basis of this fear, uncertainty, and doubt. Customers have been cast in a subservient role to the security experts and too much of the sales and marketing activity seems to place the customer under duress to buy. Users, citizens, and business leaders have been taught fear. Fear has eroded trust and encouraged an inertia bordering on paralysis. This absence of trust is a fundamental obstacle to the release of the vast potential of the cyber domain. Worse, this absence of trust plays directly in the hands of our adversaries.

Our current models of computer security, procurement, system design, system implementation, and system management are rooted in a computing model designed originally around the mainframe. As are the core foundations of the economic and business structures of the IT market. Our normative constructs of what a computer system is and how it should behave are rooted in a world when computers filled entire buildings and when the human was the passive subject; business and social interactions were computerised. Attainment, and then rigorous preservation, of a stable state was essential because although already powerful, the early computers were not far removed from their experimental phase.

The time has come for a radical reformulation of the intellectual and conceptual mechanics through which we seek to understand, represent and manage the cyber domain. We are now compelled to question at a fundamental level all of our established norms and precepts. For instance, in a world where even basic defence against the most commonplace malware requires the application of patches and updates which by definition change the system's state, why do we continue to rely upon accreditation, evaluation, and certification methodologies which assume that maintenance of a stable state is a good security goal?

Why do we continue to harbour the view that taking shelter behind digital Maginot Lines is any more effective for us than it was in 1940? We will fail in the task of constructing a resilient cyber ecosystem if we continue to attempt to build it using the frames of reference we have inherited from the Machine Age and the Cold War. In the domains where kinetic power has long been the norm, we have embraced obfuscation, camouflage, misdirection, and freedom of manoeuvre in pursuit of military objectives and the defence of democracy. Perhaps the time has come to adapt and embrace these precepts in the defence of the cyber domain. Our task is to enable the cyber domain to function to support democracy, the rule of law and an economic system based on the ownership of property, including intellectual property. Our responsibility is to learn and adapt in order to do this.

As we move further towards the end of the beginning of the Information Age, it becomes ever more apparent that the narratives through which we are attempting to represent, manage, and make safe computing are being challenged. We continue to attempt to conceptualise computers and computing systems using structures and assumptions predicated on first principals formulated when computers were mainframes and the Cold War was the dominant constituent of the global economic and political context. Our grasp of a rounded and contextualised narrative of the history of computing as a societal, rather than a technical, construct is less well developed than current circumstances require. If we fail to understand our own past, we are doomed to be at the mercy of those who would claim to understand it for us.

We must now transform the way we think and behave about computing. Whilst the technological dimensions of computing are, of course, central to an understanding of the phenomenon; they are subordinate in this regard to the human and social dimensions. Computing for the purpose of comprehending the cyber domain should now be framed as a sociological and anthropological system more than as a technical one. The systems and solutions architects of the future must be as much social as computer scientists.

The relationships between humanity, human society, and information are profound to the point of being definitional, if not existential. Human evolutionary success is predicated on the union of our ability to use tools and our capacity to organise in increasingly sophisticated societies. Our ability to process, store, accumulate, and communicate information is at the heart of this union; it is one of the foundations upon which our tool-using ability and our social capacity themselves depend. Powerful, pervasive, and interconnected computer systems are the most sophisticated tools yet created by humankind and their essential function is to process, store, accumulate, and communicate information. Information is at the centre of our humanity. The cyber domain is the key to the future development of the human condition.

Ours is the Information Age. A period in which computers have transcended the clinical isolation of the mainframe and become equally ubiquitous and interconnected; a period in which computing has become a social, economic, and cultural construct rather than principally a technical one; a period in which the ever deepening and broadening human dependence on pervasive and powerful computing is daily becoming increasingly apparent.

As lawmakers, public policy actors, theologians, business leaders and military strategists grapple with the challenges of the cyber domain; we must now devote focused and sustained effort to the development of a truly interdisciplinary approach to the understanding of the cyber domain and to the challenges of making human activity across it safe and secure. This is an exercise where governments and industry must follow and where academia should lead. Asking academia to be more responsive to the requirements set by government and industry is only credible if these requirements are understood; the evidence is to the contrary. From now on, the ranks of those defending the cyber domain must include sociologists, historians, economists, and psychologists alongside mathematicians, software engineers, and computer scientists. Human history is entering a new epoch and we must now recognise that we are central to the process of setting the course of its development.

ABOUT THE AUTHORS

Colin Williams joined SBL in 1994 and has been an executive director of the company since 1999. In 2011 he was appointed to a visiting lectureship at De Montfort University.
Areas of focus include the history of computing, the development of an interdisciplinary approach to IA, and the development of new forms of collaboration between government, industry and academia.
He holds a BA and an MA in history from the University of York.

Faculty of Technology
Gateway House Room 4.64
De Montfort University
The Gateway
Leicester
LE1 9BH
UK
Telephone: +44 (0) 7714 765 203
E-mail: c.williams@dmu.ac.uk

Dr. Tim Watson is Director of the Cyber Security Centre at De Montfort University. With more than 20 years experience in industry and in academia, he has been involved with a wide range of computer systems on several high-profile projects and has acted as a consultant for some of the largest telecoms, power and oil companies. He is also an advisor to various government departments. Tim is a regular media commentator on cyber security.

Faculty of Technology
Gateway House Room 5.58
De Montfort University
The Gateway
Leicester
LE1 9BH
UK
Telephone: +44 (0) 116 257 7476
E-mail: tw@dmu.ac.uk

REFERENCES

1. "Technology: The Cybernated Generation." Time (Friday, April 2nd 1965).

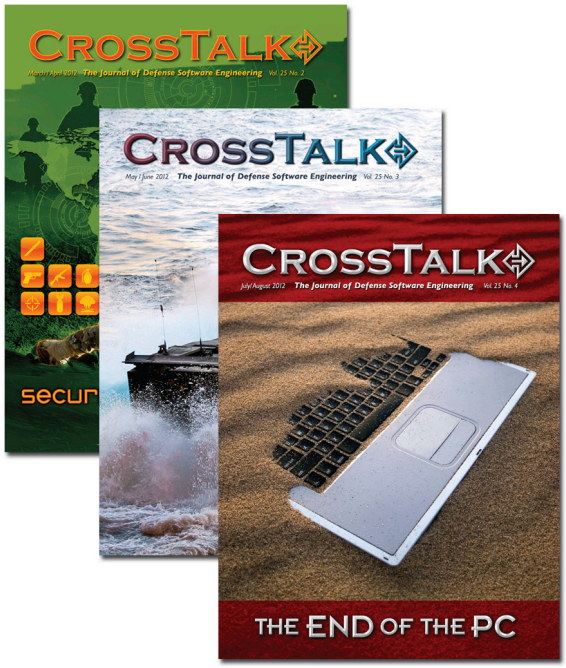

CALL FOR ARTICLES

If your experience or research has produced information that could be useful to others, CROSSTALK can get the word out. We are specifically looking for articles on software-related topics to supplement upcoming theme issues. Below is the submittal schedule for three areas of emphasis we are looking for:

Supply Chain Risk Management
Mar/Apr 2013 Issue
Submission Deadline: Oct 10, 2012

Large Scale Agile
May/Jun 2013 Issue
Submission Deadline: Dec 10, 2012

Legacy Systems Software Sustainment
Jul/Aug 2013 Issue
Submission Deadline: Feb 10, 2013

Please follow the Author Guidelines for CROSSTALK, available on the Internet at <www.crosstalkonline.org/submission-guidelines>. We accept article submissions on software-related topics at any time, along with Letters to the Editor and BackTalk. To see a list of themes for upcoming issues or to learn more about the types of articles we're looking for visit <www.crosstalkonline.org/theme-calendar>.

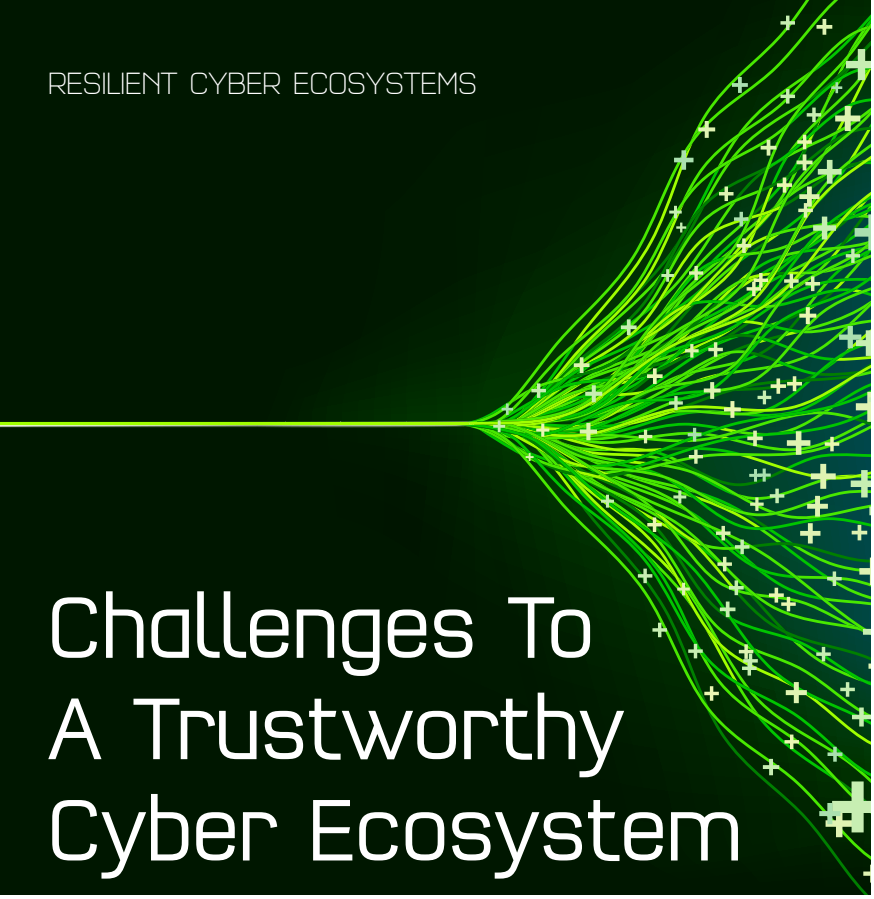

Challenges To A Trustworthy Cyber Ecosystem

Ian Bryant, De Montfort University, Leicester, UK
Jasvinder Mahrra, Institute for Security and Resilience Studies, UCL, UK

Abstract. Cyberspace is recognised as the first man-made environment. Like other natural environments it cannot be controlled. Cyberspace, of which software forms an intrinsic and indivisible element, is ever evolving and an ever growing dependency for defence, yet is contingent upon a variety of diverse participants—private firms, non-profit organisations, governments, individuals, processes, and cyber devices. It is therefore vital that intrinsic challenges to cyberspace—and software—are recognised and treated such that a trustworthy cyber ecosystem can be formed.

The Cyber Ecosystem

Cyberspace is now acknowledged to be the first man-made environment on par with air, land, maritime, and space. Indeed, it weaves all these environments together as never before. Yet, much like these other natural environments, it cannot realistically be controlled.

In doing so, cyberspace does not erase spatial boundaries—rather the transnational dimension opened up by cyberspace allows for anonymity. In contrast to the eons of time the sea has affected life on earth, cyberspace has infiltrated the whole ecosphere in decades. This constantly evolving environment is an emerging national security challenge to all nations. Indeed, the U.S. International Strategy for Cyberspace (May 2011) [1] reported, "Unauthorised network intrusions threaten the integrity of economies and undermine national security." It saw the need for collaboration between the public and private sector as crucial to protect the innovation and secure critical infrastructures such as energy, transportation, finance, and the defence industrial base, central, and local government. The problem of security is inherently complex involving not just national security concerns but commercial interests and privacy.

These characteristics challenge the defining assumptions that underpin conceptions about competent authority, jurisdictions, conflict, criminality, cash, and the use of force. The physical movement of troops through a neutral state's territory would violate neutrality. However, the same is not true for any cyber violation in which communications can pass through another state's infrastructure. How to handle cyber issues is becoming of strategic importance for governments worldwide as they strive for trustworthy and reliable networks.

Protecting the infrastructure becomes all the more essential against the impacts of disruptions and cyber attacks because the forces at work in cyberspace may more readily be asymmetric, that is, unconventional and disproportionate. So far, the new environment has demanded immediate responses, based on inherited tools or technological innovations as we progress. However, these may be necessary but are not sufficient by themselves as they offer only short terms, partial remedies.

Trustworthy cyberspace is vital to the prospects of enhancing a government's reputation for trusted and reliable hubs and networks, but the evolution of cyberspace is uncertain. Conventional approaches to this new ecosystem will not be sufficient and require a new ethos and culture of thinking. Whilst cyberspace can promote freer markets, the proliferation of some knowledge will need greater care. Cybersecurity experts themselves are calling for a radical change of ethos [2].

Whilst there has been a convergence of telecommunications, computer processing and interactive multi-media content, technological convergence is far from complete. Developments of cyber, bio, and nanotechnology are morphing into one another, and the boundaries between users and developers is blurring. But the future lies in cyberspace, and this needs to be trustworthy.

Cyberspace and Software

It is difficult to conceive of any major sector of the economy in the developed world that is not dependent (often critically so) on Information and Communications Technology (ICT) and software. This dependence extends into our private lives; with figures for the UK in October 2011 showing that more than 50% of the population now has a smartphone.

This need for trusted, correct, and reliable operation requires that software be trustable, both in terms of its resistance both to accidental or collateral faults (as exemplified by, but not restricted to, the niche "safety critical" approaches), and to malicious acts (as exemplified by the "security" approaches). This applies both to software and systems developed for specialist markets where trustworthiness is an explicit Functional Requirement (FR), and to all other software and systems, for which trustworthiness is an inherent but often forgotten implicit Non Functional Requirement (NFR).

The difference between these two views of trustworthiness is typically a matter of degree, with those for where these properties are a FR normally having Pareto or comprehensive assurance needs, whereas in the NFR space this is more likely to be a need for due diligence.

Emerging Challenges

The 2010 UK National Security Strategy [3], as approved by the Ministerial National Security Council, identified 15 priority risks across the spectrum of national security risks to the UK. Of the four Tier One risks identified as being of particular concern, one is enumerated [4] as hostile attacks upon UK cyberspace, potential shortcomings in the UK's cyber infrastructure, and the actions of cyber terrorists and criminals: to which end a National Cyber Security Programme [5] has been created.

To address this risk requires a holistic view of the adversities that need to be addressed, as this needs to address both threats (deterministic, deliberate impacts from attacks by hostile actors) and hazards (stochastic, undirected impacts from either natural events and/or collateral damage from other hostile activities). An adversity-driven approach means that not only does an organisation need to understand the threat actors it faces (be they nation states, empowered small agents or cyber-criminals), but also to have an actuarial view of the likelihood of occurrence of other events, such as the chance of climatic or geologic problems causing loss of facilities or communications, or of loss of service from a distributed denial-of-service attack on a completely unrelated organisation with whom bandwidth is shared.

The diverse nature of adversities faced by the cyber ecosystem is in direct conflict with the way in which organisations and nations are normally structured, which historically and continues to be in isolated, and sometimes mutually competitive silos. Taking a nation-state approach as an example, the issues of foreign national-state attacks will typically be handled by the defence/security/intelligence community, the issues from cyber-criminality by the law enforcement/criminal justice community, and the issues from natural hazards by the civil contingency community. Organisations suffer from similar silo effects, with differential degrees of sharing of vital information with governments and their peer community.

The scale of challenge presented by software failures cannot be underestimated, with numerous studies [6][7] identifying problems with software as a major source of project failures, with high costs to the economy, enumerated by NIST as being about $60 billion per year to the U.S. alone, with no definitive figure currently being available for the UK or worldwide.

This dependence of ICT and software can be expected to broaden and deepen in the coming years, with a number of trends already being identifiable to catalyse this dependence and complicate the problem space, including:

- The move to distributed application platforms and services (a.k.a the cloud), where the boundaries of organisation and/or national jurisdiction are increasingly blurred, and the options for either proactive controls and/or reactive measures are similarly constrained.
- Increasing reliance on mobile devices, such as smartphones and tablets, which typically rely on lightweight operating systems with less inherent controls than operating systems of previous generation desktop devices.
- A move in business to consumerization and Bring Your Own Device, where the boundary of ownership is blurred between the organisation and the individuals who work for the organisation.
- Commoditisation in previously closed architectures, such as industrial control systems where, for instance, a step change is being encountered of previously bespoke sensor devices with wireline connections to proprietary control systems are being replaced by configurable, off-the-shelf sensors using wireless connections to generic ICT systems that have onward connections to the global internet.
- The pressure for ICT consolidation for energy efficiency for green reasons (the low carbon imperative) leading to extensive use of software virtualisation to separate previously physically distinct services.

Furthermore, the way in which systems are developed and deployed is changing, with the historic assumption of ICT being engineering artifacts under single organisational control being subverted by factors such as:

- The adoption of open source models for sourcing software, fundamentally disrupting views of single organisational control.
- The growth of multicore processor technologies, which can subvert the risk modelling approaches used in previous generations of hardware.
- Growing questions as to whether hardware platforms used for software can be trusted to execute as expected, with evidence of counterfeit hardware being found in multiple market segments.
- A blurring of the boundary between software and hardware boundary, for instance with the use of software style design languages to implement application-specific integrated circuits and field-programmable gate arrays.
- The increasing use of generic, self-documenting structured data (e.g. XML) to control systems' behaviours rather than rely on pre-defined execution paths.

In terms of software development itself, the classic waterfall model of software development is evolving in a number of ways:

- The adoption of other approaches such as agile and rapid application development by the software industry.
- The growth in small-scale software development, typically carried out by micro-business who will not invest in formal development approaches, as exemplified by the apps movement for smartphones and tablets.
- A plethora of activity which produces artifacts that have the properties of software, as exemplified by the mass of websites which use, to a greater or lesser extent, mobile or active code (such as Java, Javascript and ActiveX). In these cases many of the users will have little, if any, awareness that they are implicitly creating software functionality by their often point-and-click activities.

Creating Trust

These uncertain developments require, "security and resilience for cyberspace to be seen as not just a service but are the services underpinning trust and confidence in an environment that touches all others" [8].

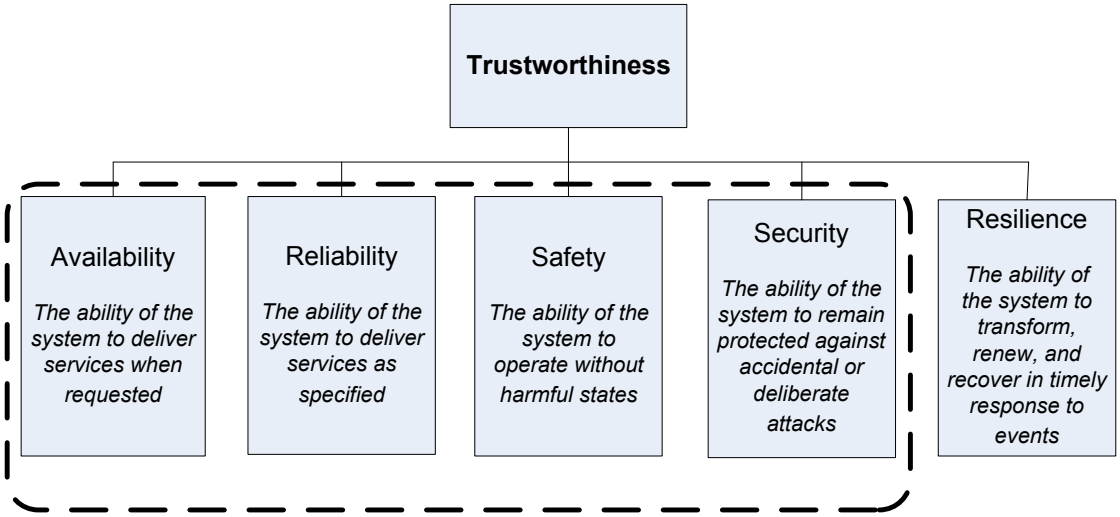

Figure 1

Security and resilience are defined [9] as being complimentary practices required to manage relevant aspects of an organisations operational risk, and have a number of competing definitions, of which the most useful are probably [10]:

- **Security:** the preservation of confidentiality, integrity, and availability of entities.
- **Resilience:** the property of an entity to transform, renew, and recover from the impact of interactions or events.

Investment in cyberspace protection must be increased if we are to move from seeing security as an organisationally focused afterthought and moving towards a more inclusive concept of resilience that is fit for our times, which needs to include consideration of all external and infrastructure dependencies, and the sets of both proactive and reactive controls needed to mitigate risks from such dependencies. It is about transformation first and not about cleaning up after the fact; not bouncing back but bouncing forward and learning to thrive on uncertainty.

But neither security nor resilience gives us holistic trustworthiness, and thus a more expansive model is needed.

Figure 1, adapted from previous work by Professor Ian Sommerville at St. Andrews University [11] attempts to link together the set of existing stovepipes of activity that need to be considered.

Thus in order to get the best from cyberspace and minimize the inherent dangers we need a holistic, ever vigilant, and innovative, approach to trustworthiness:

"A sustainable and trustworthy cyberspace will derive from open sources and standards, driving an internationally coordinated approach to research and development [7]."

Delivering Trust

Whether the focus of concern is the organisation or the nation state, a successful protective regime should regard all adversities holistically so that the most pragmatic, appropriate, and cost-effective treatments can be applied and trustworthy solutions delivered—the option sets available against denial-of-service whether it be from an attack or a natural disaster are likely to be very similar.

Software represents a microcosm of the overall cyberspace, and therefore software engineering must attempt to escape a threat-driven mindset, addressing all adversities to deliver trust.

ABOUT THE AUTHORS

Ian Bryant is the Technical Director for Software Security, Dependability and Resilience at the Cyber Security Centre, De Montfort University, Leicester, UK, where he is on academic attachment from the UK Ministry of Defence. He has more than 20 years of experience in information systems security, dependability and resilience across a number of public sector bodies, and is an active contributor to a variety of standards development organisations in both information security and systems/software engineering.

E-mail: ib@dmu.ac.uk

Jasvinder Mahrra is the Senior Research Fellow at the Institute for Security and Resilience Studies, University College, London, UK. She has more than 10 years of experience in resilience planning and exercising. Before joining ISRS she had her own consultancy and prior to this was part of the UK's Civil Contingencies Secretariat (Cabinet Office) coordinating a programme of events to enhance resilience and preparedness.

E-mail: j.mahrra@ucl.ac.uk

REFERENCES

1. The White House; International Strategy for Cyberspace: Prosperity, Security and Openess in Networked World (May 2011)
2. Evans, K & Reeder, F; A Human Capital Crisis in Cybersecurity: Technical Proficiency Matters, A report of the CSIS Commission on Cybersecurity for the 44t h Presidency, CSIS: Washington DC (2010)
3. Cabinet Offices; A Strong Britain in an Age of Uncertainty: The National Security Strategy; Cmd7953, Cabinet Office (October 2010)
4. Cabinet Office; Securing Britain in an Age of Uncertainty: The Strategic Defence and Security Review; Cmd 7948; Cabinet Office (October 2010)
5. Downing, E; Cyber Security – A new national programme; HC/SN/SC/5832; House of Commons (23 June 2011)
6. Flyvbjerg , B and Budzier, A; Double Whammy – How ICT Projects are Fooled by Randomness and Screwed by Political Intent Alexander; University of Oxford Saïd Business School / McKinsey (2011)
7. Standish "Chaos" Reports (2004 onwards)
8. MacIntosh, JP, Reid J & Tyler, L; Cyber Doctrine: Towards A Coherent Evolutionary Framework for Learning Resilience; Institute for Security & Resilience Studies, UCL (2011)
9. HM Treasury; Management of Risk - Principles and Concepts; HM Treasury (Oct. 2004)
10. ISO/IEC 27000; Information technology -- Security techniques -- Information security management systems -- Overview and vocabulary; ISO/IEC (Working Draft March 2012)
11. Sommerville, I; Software Engineering (9th Ed.); Pearson (2009)

Systems Thinking for a Secure Digital World

William D. Miller, Innovative Decisions, Inc.

Abstract. The practice of cyber security appears to be predominantly a game of Whac-A-Mole, and the moles are winning! Systems are designed and deployed with security such as it is, grafted on, and the standard response to adversarial attacks is to continually patch the IT and burden humans with process and passwords. We must learn to think systemically to seek advantage, or at least maintain parity over adversarial threats, as our infrastructure becomes more complexly integrated.

Introduction

The stage is set by thriving communities of adversaries who seek all possible means to harm cyber systems and potentially to the infrastructure with which they are integrated. The functions currently performed by cyber security should thwart these adversaries but are too often add-ons rather than inherently designed into the cyber systems. A short history of systems thinking and its relevance to thwarting the threat is established. Then specific actions are identified to achieve the vision of a secure digital world, after the fact.

Cops and Robbers in the Digital Age

Adversarial attacks to compromise cyber systems can be broadly categorized as hacks, social engineering, insider jobs and stupid stuff:

- Hacks include malware, e.g. viruses and worms.
- Social engineering includes phishing schemes to trick individuals into divulging private information that can then be exploited. Social engineering can also be used to gain personal knowledge of individuals and to guess passwords.
- Insider jobs occur when adversaries have trusted access to at least some parts of cyber systems and violate the trust they have been granted.
- Stupid stuff occurs when individuals do not take proper care of information systems and/or personal information. Adversaries are the proactive ones that seek to compromise cyber systems and have the advantage to discover and exploit vulnerabilities on Internet time.

Security Engineering, Such as It [1]

Cyber security tends to have a technology focus and provides a defensive, static, security environment versus the dynamic behavior of its adversaries. The behavior of defensive oriented cyber security is asymmetrical, which gives the adversaries the "first move" advantage that must then be detected, identified, and protected against. The result is patch, upon patch, upon patch in response to adversarial attacks. Cyber security behaves as an evolutionary system, not a purpose-designed system.

A Short History of Systems Thinking

"Systems thinking is a discipline for seeing wholes. It is a framework for seeing interrelationships rather than things, for seeing patterns of change rather than static 'snapshots.'"
– Peter Senge, 1990 [2].

The most popular definition of systems thinking is arguably defined by Peter Senge, who traces its roots to the feedback concepts of cybernetics and servo-mechanisms. Senge gives substantial credit to Jay Forrester's early work beginning in the mid-1950s in system dynamics. Forrester's stock-flow-feedback structure modeling of General Electric appliance manufacturing plants revealed that the observed three-year employment cycle of hiring and layoffs was attributable to the internal structure of the firm and not to the external forces of the business cycle [3].

A key lesson is that answers and solutions to observed phenomena may be non-intuitive without analysis. Stocks define the states of the system, and the variables defining the changes in states are the flows. The stock-flow-feedback metaphor models n^{th} order difference/differential equations that describe the behavior of a system [4]. Nouns represent stocks whereas verbs represent flows. Stocks send out signals representing information about the state of the system to the rest of the system. Stocks have the following characteristics: memory, ability to change the time shape of flows, decouple flows, and create delays.

Forrester's work bloomed into the System Dynamics Society and The System Dynamics in Education Project at MIT, now The Creative Learning Exchange, championing system dynamics and systems thinking in K-12 education. System dynamics has been applied to business management, sustainability studies, policy analysis and design. The Club of Rome embraced system dynamics in its 1972 report, The Limits to Growth. The methodology also supports agent-based modeling. This author applied systems dynamics in the late 1970s to understanding the cost impact of reported but unfound troubles in the telephone network. This provided the basis to justify a cost-effective system to improve the detection and repair of such troubles.

Figure 1. A simple stock and flow model of inventory.

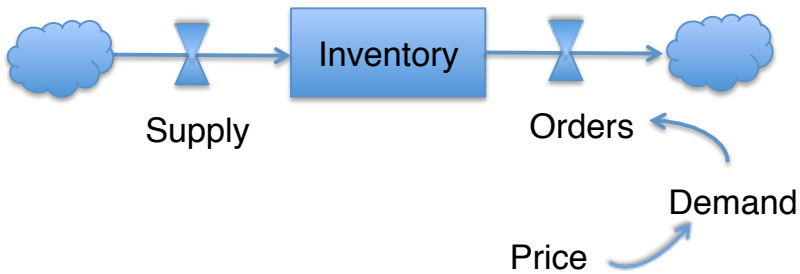

Figure 2: Business inventory causal loop.

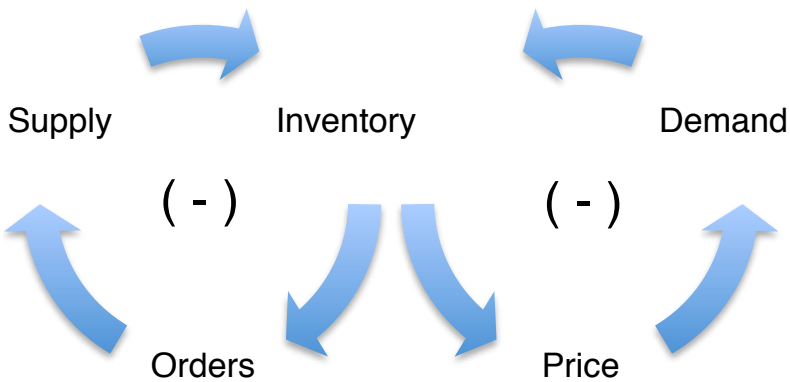

Within DoD, CDR Brett Pierson developed a system dynamics model of the FM 3-24 COIN Manual [5]. There are several popular system dynamics software programs available. A simple inventory stock and flow model is shown in Figure 1.

Published in 1980, a classic document of systems is, "Systems 1: An Introduction to Systems Thinking" by Draper Kauffman and precedes Senge's book by a decade [6]. Kauffman's intent was to translate the ideas of systems and systems thinking, which is full of technical jargon and mathematics. He wanted non-expert educators to be able to teach the concepts to K-12 students.

Kauffman defines systems, the concept of feedback and introduces causal loop diagrams to model their behavior. Figure 2 is an example of a causal loop diagram describing the relationship of a business' inventory to price, demand, orders and supplies.

These causal loops are the precursor to modeling the stocks and flows. Kauffman provides a simple taxonomy of systems and their properties, as well as complex system characteristics and problems as shown in Table 1.

Several of the effects that Kauffman identifies are highly relevant to cyber security:

- Systems cope with problems by reacting to warnings.
- The obvious solution often makes things worse.
- Solving one problem almost always creates others.

The soft systems methodology in, "Systems Thinking, Systems Practice" by Peter Checkland was first published in 1981 and has been republished several times [7]. Checkland acknowledges systems engineers' contributions to the mature understanding of hard systems and then identifies the problems extending those paradigms to the unstructured problems of soft systems. Checkland lays out an action research program that led to the holistic methodology for soft systems, especially human activity systems, such as the British Rail System. He uses causal diagrams that are more free form than the formal causal loops introduced earlier.

Derek Hitchins, a contemporary of Checkland, integrates systems engineering and systems thinking in, "Systems Engineering: A 21st Century Systems Methodology" in 2007, with extensive use of causal loops and system dynamics applied to complex systems [8]. Hitchins focuses on defense capabilities, illustrating concepts in the case study of the World War II Battle of Britain Command and Control System.

Peter Senge popularized systems thinking in, "The Fifth Discipline: The Art & Practice of The Learning Organization" in 1990. Subsequent to its publication, Senge co-authored a series of field books applied to a variety of domains. The Fifth Discipline is systems thinking and completes the four disciplines of personal mastery mental models, shared vision and team learning. Senge's laws of the Fifth Discipline and causal loop system archetypes are shown in Table 2. The archetypes are naturally recurring patterns in systems and are represented by formal causal loop diagrams.

John Boardman and Brian Sauser integrated the concepts of causal loop diagrams, soft systems methodology and social network theory with the introduction of the system diagram, or systemigram, conceptual model [9] The systemigram provides a systemic visualization of system complexity and enables the elucidation of the key attributes of emergence, hierarchy and boundary of complex systems. The application of systems thinking is illustrated by the relevant systemigram example in Figure 3 from the Systems Security Engineering roadmap report published by the Systems Engineering Research Center (SERC), a University-Affiliated Research Center of the DoD [10].

Application of Systems Thinking for a Secure Digital World

The International Council on Systems Engineering's "INSIGHT" publication devoted its July 2011 issue to a special feature on "Systems of Systems and Self Organizing Security." The feature specified that:

"Resilient system strategies may be a more manageable way to counter the asymmetry of attack and defense. In recognition that systems will have vulnerabilities that adversaries will attack, and that system design needs mechanisms to weather successful attack and remain viable, engineers are now placing a new strategic priority on system resiliency. Survivability through resilient design is not a new concept, but still remains largely a research activity."

Table 1. Kauffman's system types and properties, as well as complex system characteristics and problems.

System Types	System Properties	Complex System Characteristics	Complex System Problems
Mechanical Human/ Mechanical Biological Ecological Social	Stability Limitations Loose Systems Reaction Times Anticipation Hidden Systems	Self-Stabilizing Goal-Seeking Program-Following Self-Reprogramming Anticipation Environment Modifying Self-Replicating Self-Maintaining/Repairing Self-Reorganizing Self-Programming	Tragedy of the Commons Cost of Information Distortion of Feedback Loss of Predictability

Table 2. Senge's laws of the Fifth Discipline and system archetypes.

Laws of the Fifth Discipline	System Archetypes
1. Today's problems come from yesterday's solutions	Balancing Process with Delay
2. The harder you push, the harder the system pushes back	Limits to Growth
3. Behavior grows better before it grows worse	Shifting the Burden
4. The easy way usually leads back in	
5. The cure can be worse than the disease	Eroding Goals
6. Faster is slower	
7. Cause and effect are not closely related in time and space	Escalation
8. Small changes can produce big results – but the areas of highest leverage are often the least obvious	Success to the Successful
	Tragedy of the Commons
9. You can have your cake and it it too – but not at once	Fixes that Fail
10. Dividing an elephant in half does not produce two elephants	Growth and Underinvestment
11. There is no blame	

Figure 3. SERC systems security systemigram (used by permission).

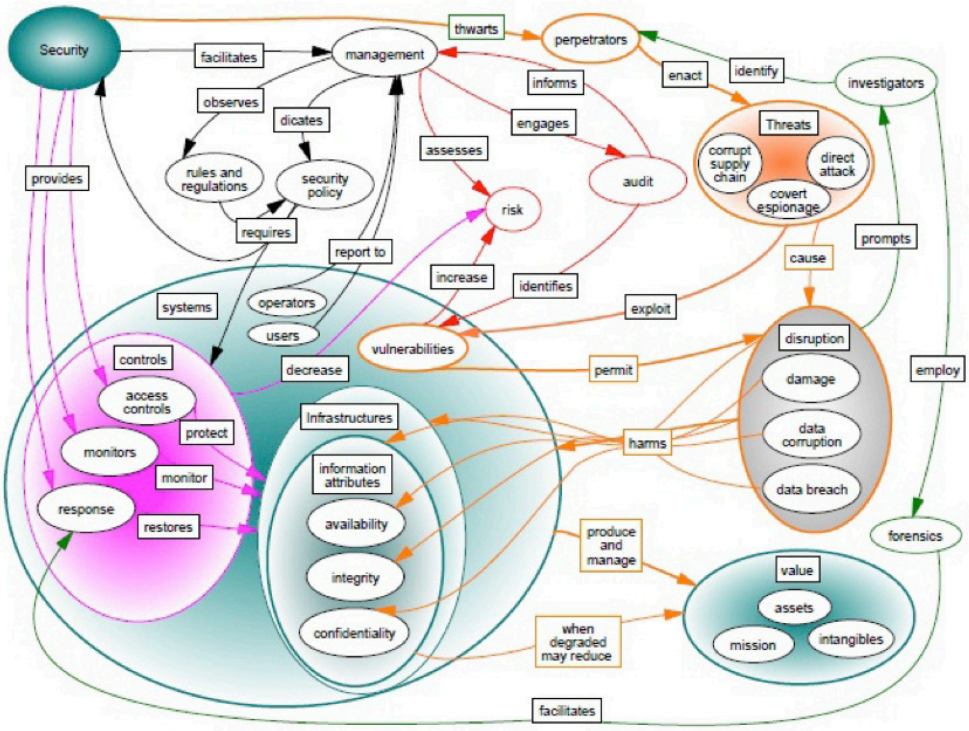

From the systems thinking perspective, the imperative is that cyber security learning loops must be fastest where the stakes are highest, as when systems become high-value targets under attack by determined, intelligent adversaries. Another systems thinking imperative is that people are part of the system, and therefore the human condition, with all its attributes including social systems and social engineering, must be part of the design formulation for cyber security.

Within DoD, the INCOSE "INSIGHT" article goes on to explain:

"Security has focused on keeping critical technology and information from getting out. However, as DoD systems have come to depend on commercial technology and components that are increasingly sourced through complex global supply chains, a new security emphasis is emerging: keeping malicious or compromised system elements or components from getting in."

The SERC Systems Security Engineering Final Technical Report establishes a research roadmap for DoD, with its executive summary summarizing insights from systems thinking:

"The U.S. needs dramatic improvements in systems security. Current defensive strategies, based principally on strengthening system peripheries, inspections, and similar bolt-on techniques add tremendously to cost and do not respond effectively to the growing sophistication of attacks. Systems cannot be assumed to have static boundaries, static user communities, or even a static set of services."

The report goes on to emphasize the application "of scientific and engineering principles to identify security vulnerabilities and minimize or contain the risks associated with these vulnerabilities." The SERC report is available at <http://www.sercuarc.org>.

Two additional works that address cyber security from a systems thinking perspective are "Enterprise Security for the Executive: Setting the Tone from the Top" by Jennifer L. Bayuk [11] and "Cyber Attacks: Protecting National Infrastructure" by Edward G. Amoroso [12]. Bayuk addresses security leadership and Amoroso proposes a comprehensive national infrastructure protection methodology. The reader is encouraged to become involved in INCOSE working groups and the cyber security professional society organizations.

Homeland Security

The Department of Homeland Security, Office of Cybersecurity and Communications, is seeking dynamic individuals to fill several positions in the areas of software assurance, information technology, network engineering, telecommunications, electrical engineering, program management and analysis, budget and finance, research and development, and public affairs.

To learn more about the DHS Office of Cybersecurity and Communications, go to www.dhs.gov/cybercareers. To find out how to apply for a vacant position, please go to USAJOBS at www.usajobs.gov or visit us at www.DHS.gov; follow the link Find Career Opportunities, and then select Cybersecurity under Featured Mission Areas.

Summary

This paper lays out the context of adversarial threats to cyber systems and taking a systems thinking approach to cyber security in the digital world. Past and current practices of patching vulnerabilities as they are discovered leave the initiative to the adversaries and do not solve the underlying structural problems that exist. Systems thinking addresses the wholeness and interrelated, dynamic behavior of this domain. To quote President Abraham Lincoln, "We must think anew, and act anew." Significant research remains to be accomplished, both theoretical and applied.

Acknowledgements

The author gratefully acknowledges the work of the INCOSE Model-Based Systems Engineering Initiative and the following INCOSE Working Groups for their contributions applying systems thinking to cyber security and infrastructure security: 1) Security Engineering, 2) Systems Science, 3) Resilient Systems, 4) Complex Systems, 5) Autonomous Systems Test, 6) Anti-Terrorism International, and 7) Human Systems Integration. The author also gratefully acknowledges the contributions of the Systems Engineering Research Center investigators who established a systems security engineering roadmap for the DoD. In particular, Dr. Jennifer Bayuk, a colleague at the Stevens Institute of Technology, School of Systems and Enterprises, established the Systems Security Engineering graduate program at the school, was a major contributor in both INCOSE and SERC initiatives, and has been an exceptional mentor in relating security engineering to systems engineering for the author.

ABOUT THE AUTHOR

William D. Miller is executive principal analyst with Innovative Decisions, Inc. and adjunct faculty at the School of Systems and Enterprises, Stevens Institute of Technology. Miller is the deputy technical director of the International Council on Systems Engineering (INCOSE), a nonprofit membership organization that promotes international collaboration in systems engineering practice, education and research. He specializes in systems engineering of government and commercial communications systems and services, working at companies including Bell Labs and AT&T.

Innovative Decisions, Inc.
1945 Old Gallows Road
Suite 207
Vienna, VA 22182
Phone: 908-759-7110
Fax: 703-854-1132
E-mail: wmiller@innovativedecisions.com

REFERENCES

1. Dove, Rick and Bayuk, Jennifer, editors. "Special Feature: Systems of Systems and Self-Organizing Security, 14.2 INCOSE INSIGHT (July 2011).
2. Senge, Peter M. The Fifth Discipline: The Art & Practice of The Learning Organization. New York: Currency Doubleday, 1990.
3. Radzicki, Michael J. and Taylor, Robert A. "Origin of System Dynamics: Jay W. Forrester and the History of System Dynamics." 2008.
4. Forrester, Jay W. Industrial Dynamics. Cambridge MA: MIT Press, 1961.
5. Brett Pierson, Brett. "A System Dynamics model of the FM 3-24 COIN Manual." Warfighting Analysis Division J8/WAD, accessed at <http://www.mors.org/UserFiles/file/meetings/07ic/Pierson.pdf on 4/5/2012>.
6. Kauffman, Jr., Draper L. Systems One: An Introduction to Systems Thinking. Future Systems, Inc., 1980. (Originally The Human Environment: An Introduction to Environmental Systems, developed under a grant to the Office of Environmental Education, Office of Education, Department of Health, Education, and Welfare.)
7. Checkland, Peter. Systems Thinking, Systems Practice. Chichester, England: Wiley, 1993.
8. Hitchins, Derek K. Systems Engineering: A 21st Century Systems Methodology. Chichester, England: Wiley, 2007.
9. Boardman, John and Sauser, Brian. Systems Thinking: Coping with 21st Century Problems. New York: CRC Press, 2008.
10. Bayuk, Jennifer, et al. "Systems Security Engineering Final Technical Report." SERC-2010-TR-005, Systems Engineering Research Center, August 22, 2010.
11. Bayuk, Jennifer L. Enterprise Security for the Executive: Setting the Tone from the Top. Santa Barbara, California: Praeger, 2010.
12. Amoroo, Edward G. Cyber Attacks: Protecting National Infrastructure. New York: Elsevier, 2011.

Identifying Cyber Ecosystem Security Capabilities

Peter M. Fonash, Ph.D., DHS

Abstract. Strengthening the security and resilience of the cyber ecosystem requires reducing the number of vulnerabilities and the ability to automatically mitigate attack methodologies. This article draws from various research reports to categorize the underlying attack methodologies and summarizes current perspectives on the capabilities needed within the cyber ecosystem to strengthen its security and resilience, while protecting the privacy of the authorized users of the ecosystem.

Introduction

A general consensus has been forming in the cybersecurity community that cybersecurity defenses must become more automated, less reactive, distributed, and better informed. There have been a number of proposals and ongoing activities to enable automated collective action to strengthen the resilience and security of the cyber ecosystem[1] in the face of the advanced cyber threat. These proposals and activities support a range of automated collective actions, including the sharing of indicators and information, the selection of courses of action, and the coordination of responses. This article uses a three-step process to identify capabilities needed in the future cyber ecosystem to make these automated collective actions possible.

The first step was to understand the types of cyber attacks being faced by today's computer systems. Drawing from reports that help categorize today's attacks, an attack categorization is proposed. The second step was to review recent papers on cyber ecosystem security, including industry and academic comments on a cyber ecosystem paper [1] published by DHS. From these sources, a set of cyber ecosystem security capabilities was proposed. The third step was to analyze the collective cyber ecosystem capabilities and their ability to counter the proposed attack categories. This analysis resulted in a mapping of the cyber ecosystem capabilities against the attack categories.

Categories of Cyber Attacks

Using data from NIST, "Computer Security Incident Handling Guide" [2] and the "2012 Data breach Investigations Report" [3], a list of cyber attack categories was created. The attack categories are attrition, malware, hacking, social tactics, improper use (insider threat), loss or theft of equipment, physical action, and attacks that consist of multiple components. Table 1 provides a description for each cyber attack category, and includes the category "other" for completeness.

To cover current and future attacks, the attack categories have been made very general. For example, hacking is a very broad category of attack, but seems to be sufficient for the purposes of this article. Although other categories of attack can be created, this list is useful for helping to identify capabilities needed within the future cyber ecosystem to improve resilience and security.

The following section briefly discusses recent articles and papers that have proposed automated collective action in the future cyber ecosystem. These proposals will form the basis for the desired capabilities that are identified in a subsequent section.

Proposals for Collective Action in the Future Cyber Ecosystem

DHS has been working with industry, other government agencies, and the research and development community to develop a consensus on desirable future cyber ecosystem capabilities. The DHS National Protection and Programs Directorate (NPPD) published a paper "Enabling Distributed Security in Cyberspace: Building a Healthy and Resilient

Table 1. Categories of Cyber Attack

Attack Category	Description of Attack
Attrition [2]	Use of brute force methods to compromise, degrade, or destroy systems, networks, or services. Includes distributed denial of service attacks intended to impair or deny access to a service or application and resource depletion attacks [4].
Malware [2,3]	Any malicious software, script, or code developed or used for the purpose of compromising or harming information assets without the owner's informed consent, regardless of delivery method. Includes Web and email attacks and attacks executed from removable media or a peripheral device.
Hacking [3,4]	An attempt to intentionally access or harm information assets without authorization or in excess of authorization, usually conducted remotely. Includes data leakage attacks, injection attacks and abuse of functionality, spoofing, time and state attacks, buffer and data structure attacks, resource manipulation, use of stolen credentials, backdoors, brute force and dictionary attacks on passwords, and exploitation of authentication.
Social Tactics [3]	Use of social tactics such as deception, manipulation, and intimidation to obtain access to data, systems or controls. Includes pretexting (fake surveys), solicitation phishing, and elicitation of information through conversation.
Improper Usage (Insider Threat) [2]	Inappropriate use of privileges or inappropriate logical or physical access to data, systems, or controls by a person or persons associated with an organization. Any incident that would violate an organization's acceptable usage policies by an authorized user. Includes installation of unauthorized software and removal of sensitive data.
Physical Action [3]/Loss or Theft of Equipment [2]	Human Driven attacks that employ physical actions and/or require physical proximity. Examples are: stolen identity tokens and credit cards, tampering with or replacing card readers and point of sale terminals, and tampering with sensors. The loss or theft of a computing device or media used by the organization, such as a laptop or smart phone.
Multiple Component [3]	A single attack that encompasses the use of multiple techniques. Advanced attacks would often fall into this category, with various attack components occurring at different steps in the cyber kill chain [5,6].
Other [2]	An attack that does not fit into any of the other categories, such as supply chain attacks and network reconnaissance [4].

Cyber Ecosystem with Automated Collective Action" [1] to encourage a discussion of the cyber ecosystem capabilities. Additionally, the DHS cybersecurity strategy is outlined in the "Blueprint for a Secure Cyber Future" [7].

Two recent Microsoft security documents discuss collective options for improving the security "health" of computer systems. In the first, Scott Charney, Corporate Vice President for Trustworthy Computing, presents [8] a spectrum of computer defense. The computer defense spectrum includes collective defense. Charney recommends that "society needs to explore ways to implement collective defenses to help protect consumers who may be unaware that their computers have been compromised, and to reduce the risk that these compromised devices present to the ecosystem as a whole." In a subsequent Microsoft document, Kevin Sullivan, Senior Security Strategist for Trustworthy Computing, discusses [9] collaboration to secure consumer computers. Sullivan's strategy recognizes that, "As no single entity can defeat global cybercrime by itself, members of the internet ecosystem must take collective action."

Two IBM articles likewise present a case for cybersecurity improvements as a result of information exchange and collaboration. An early IBM Systems Journal article [10] recommends autonomic computing to provide security. The article asserts that computing systems, "like the biological systems that keep our hearts beating and our body chemistry balanced, can take care of routine and even exceptional functions without human intervention." A more recent IBM report [11] makes a similar recommendation, based on a public health and safety model for cybersecurity. "Effective response requires continuous research, open information exchange, and transparency among a wide range of actors. This allows responses to be better individualized to confront the particular nature of the threat and its risk of spreading more widely."

The previously mentioned DHS cyber ecosystem paper [1] discusses automated collaboration to help strengthen the resilience and security of the cyber ecosystem. Drawing a parallel from the practice of continuous monitoring, the DHS NPPD paper proposes to automate collaborative identification, analysis, and responses to strengthen protections against the advanced cyber threat. The DHS cyber ecosystem paper describes a future cyber ecosystem in which computing systems, "work together in near—real time to anticipate and prevent cyber attacks, limit the spread of attacks across participating devices, minimize the consequences of attacks, and recover to a trusted state. In this future cyber ecosystem, security capabilities are built into cyber devices in a way that allows preventive and defensive courses of action to be coordinated within and among communities of devices. Power is distributed among participants, and near—real time coordination is enabled by combining the innate and interoperable capabilities of individual devices with trusted information exchanges and shared, configurable policies." The paper envisions a future in which authentication, automation, and interoperability are the building blocks that enable cyber components to work together.

Based on this understanding of the future cyber ecosystem, the next section identifies capabilities desired in the future cyber ecosystem. The goal is a cyber ecosystem that helps mitigate all categories of cyber attack rather than defending against only known attacks.

Desired Cyber Ecosystem Capabilities

All nine attack categories can benefit from three common capabilities, called cyber ecosystem "building blocks" in the DHS NPPD ecosystem paper [1]. These capabilities are:

- **Automation** – allows the speed of response to approach the speed of attack.
- **Interoperability** – permits dynamic and seamless collaboration by removing technical constraints and barriers.
- **Authentication** – enables trusted online decisions between resources and actors at a distance, preferably in a way that enhances privacy.

The attack categories have additional commonalities, including the need for attack detection and situational awareness [7] and the ability to take advantage of shared information. For the cyber ecosystem to respond to an attack, the attack must be detected. As attacks become more sophisticated, identification of attacks, whether attempted or successful, will become more difficult. Furthermore, to minimize the consequences of an attack, detection should anticipate an attack as early as possible in the cyber attack lifecycle, commonly called the cyber kill chain [5,6]. Once an attempted or successful attack has been detected, the participants in the cyber ecosystem must be able to share and make use of that information. A key value of collective action is the ability to inform other systems of an attack before those systems come under attack. Additionally, a security management system can correlate inputs from various sensors to refine what is known about the attack.

A secure and resilient cyber ecosystem needs to do more than just share information about attacks. Security management systems can use the shared information to develop, evaluate, and implement alternative courses of action, as well as assess the effectiveness of the actions as the actions occur. Risk-based data management [12] will help support these capabilities. The effectiveness assessment can provide inputs for a range of subsequent actions, such as sensor reconfiguration, tightening security configurations, alerts and warnings, and the development of new courses of action. NIST Special Publication 800-61 recommends [2] the capability to document the attack, response and recovery. This is more than just an audit trail. It includes forensics-quality images and records that can subsequently be used to analyze the attack, identify undiscovered attack techniques, and support criminal investigation.

Not all attacks are alike, so the cyber ecosystem must include capabilities that are able to respond to the individual attack categories as well. This includes the capability to:

- Identify and respond to attrition attacks that did not necessarily gain access to an information system. Responses could require action by external participants.
- Identify malware that has no known signature, heuristics, or actions.
- Identify when the performance of systems or components is degraded, preferably before the systems or components fail.
- Perform near-real time risk-based management, so that automated responses are feasible.
- Filter out authorized activity so as to identify unauthorized hacking or insider activity, based on behavior monitoring that incorporates business rules [12].
- Employ actions that will not tip off an adversary, such as (but not limited to) monitoring the attack or using tailored trustworthy spaces [12], moving target [12], or containment (quarantine or honey-pot) to limit the scope of an attack.

The cyber ecosystem will always include well-known existing cybersecurity capabilities. These include user education to increase awareness of the sophisticated attacks, including social and physical attacks; cybersecurity education and training for the IT staff; and the need for secondary capabilities such as reserve power and cooling, backup communications, spare systems, and alternate sites.

The cyber ecosystem must include capabilities that will protect privacy and civil liberties.

Charney wrote, "Privacy concerns must be carefully considered in any effort to promote Internet security by focusing on device health. In that regard, examining health is not the same as examining content; communicating health is not the same as communicating identity; and consumers can be protected in privacy-centric ways that do not adversely impact freedom of expression and freedom of association." [8] The DHS cybersecurity strategy envisions that, "collaboration principles will foster the transfer of specific, actionable cybersecurity information using approved methods to those who need it, while protecting the privacy and civil liberties of the public." [7] Conversely, information systems within the cyber ecosystem will store, but not inappropriately share, data needed by authorized law enforcement officials to perform their duties.

Continued operations and recovery are key resiliency capabilities for the cyber ecosystem. A MITRE report [6] presents the following cyber resiliency goals:

- **Withstand an attack** – continue essential mission/business functions despite successful execution of an attack.
- **Recover from an attack** – restore mission/business functions to the maximum extent possible subsequent to successful execution of an attack.
- **Evolve** – minimize adverse impacts by changing missions/business functions, as well as perhaps changing the supporting cyber capabilities.

In 2011, DHS published the "Blueprint for a Secure Cyber Future: The Cybersecurity Strategy for the Homeland Security Enterprise" [7]. The Blueprint lists a number of objectives to strengthen the cyber ecosystem and enable success against current and future threats:

- **Develop the Cyber Workforce in the Public and Private Sectors:** Maintain a strong cadre of cybersecurity professionals to design, operate, and research cyber technologies.
- **Build a Base for Distributed Security:** Provide individuals with tools, tips, education, training, awareness, and other resources appropriate to their positions that enable them to implement existing cybersecurity features and configurations in protocols, products, and services.
- **Reduce Vulnerabilities:** Design, build and operate information and communication technology to specifically reduce the occurrence of exploitable weaknesses. Enable technology to sense, react to, and communicate changes in its security or its surroundings in a way that preserves or enhances its security posture.

Assess effectiveness
Authentication
Interoperability
Automated Defense Identification, Selection, and Assessment
Build Security In
Business Rules-Based Behavior Monitoring
General Awareness and Education
Moving Target
Privacy
Risk-Based Data Management
Situational Awareness
Tailored Trustworthy Spaces

Table 2. Desired Cyber Ecosystem Capabilities

- **Improve Usability:** Design trusted technology that is easy to use, easy to administer, rapidly customizable, and performs as expected.
- **Appropriately Validate Identities in Cyberspace:** Use risk-based decision making for authentication, raising the level of trust associated with the identities of individuals, organizations, networks, services, and devices involved in online transactions and communication.
- **Increase Technical and Policy Interoperability Across Devices:** On a device-to-device level, strengthen collaboration, create new intelligence, hasten learning, and improve situational awareness.
- **Automate Security Processes:** Employ automated mechanisms for acting collectively in near real-time to anticipate and prevent incidents, limit the spread of incidents across participating devices, and minimize consequences.

The various capabilities discussed above can be combined into a list that takes into consideration similarities and differences. For example, a number of capabilities are related to automation, information sharing, collaboration, and assessment of results. Table 2 presents an alphabetical list of the major capabilities discussed above that are desirable in the future cyber ecosystem.

Mapping Desired Cyber Ecosystem Capabilities Against Attack Categories

The following table (Table 3) maps the desired cyber ecosystem capabilities against the attack categories. It reflects a combination of the recommendations in the literature, the recommendations of the research community [13] and industry review [12] of the DHS cyber Ecosystem paper [1]. It is noted that almost all

Table 3. Compare Attack Categories against Desired Cyber Ecosystem Capabilities

	Categories of Cyber Attack							
Desired Cyber Ecosystem Capabilities	Attrition	Malware	Hacking	Social Tactics	Improper Usage (Insider)	Physical Action; Loss or Theft	Multiple Component	Other
Automation	x	x	x	x	x	x	x	x
Authentication	x	x	x	x		x	x	x
Interoperability	x	x	x	x			x	
Automated Defense Identification, Selection, and Assessment	x	x	x	x	x	x	x	x
Build Security In	x	x	x	x		x	x	x
Business Rules-Based Behavior Monitoring	x	x	x	x	x	x	x	x
General Awareness and Education	x	x	x	x	x	x	x	x
Moving Target	x	x	x	x			x	x
Privacy	x	x	x	x	x	x	x	x
Risk-Based Data Management	x	x	x	x	x	x	x	x
Situational Awareness	x	x	x	x	x	x	x	x
Tailored Trustworthy Spaces	x	x	x	x			x	x

the boxes are filled in. This reflects the thought that the capabilities work together as a system and the probability that a particular capability will help in some way to either help detect or mitigate an attack.

DHS has a number of ongoing efforts that help achieve some of the desired future capabilities. Examples of some of these activities include:

- Early detection of attacks, preferably before an attacker has begun to exploit the attack.
  - Trusted Automated Exchange of Indicator Information (TAXII)
  - National Cyber Protection System
  - Continuous Monitoring activities
- Interoperability that permits maximum collaboration and information sharing by removing technical constraints and barriers.
  - Various Security Content Automation Protocol (SCAP) activities
  - Continuous Monitoring Activities
  - TAXII
  - The National Cybersecurity and Communications Integration Center (NCCIC)
- Authentication that enables trusted collective actions to occur automatically.
  - Support to the National Strategy for Trusted Identities in Cyberspace
- Automation to rapidly share indicators and warnings, possible courses of action, configuration settings and policy updates, and other useful information.
  - TAXII, SCAP, Common Vulnerabilities and Exposures (CVE), Open Vulnerability Assessment Language (OVAL), Malware Attribute Enumeration and Characterization
  - Federal Information Security Management Act
  - Continuous Monitoring
- Develop collaborative courses of action, given available information, policies, tools, procedures, and capabilities.
  - National Cyber Incident Response Plan
  - NCCIC and US-CERT
- Build security into products and components, so that they are able to participate properly and effectively in the future cyber ecosystem.
  - Software Assurance Program
  - Education and Training
  - CVE, OVAL
- Utilize shared information via systems and components that have the ability to produce and consume near-real-time indications and collaborative response information.
  - Dynamic Defense, and Defense-in Depth [12]
  - SCAP
  - TAXII
- Increase awareness of people by providing alerts, tools, tips, guidelines, and resources that are appropriate to a given situation; and of unauthorized activity by business- and operations-based behavioral analysis tools.
  - Education and Outreach Programs
  - NCCIC
- Transparency and Privacy that protects the rights of citizens and system users by sharing data that focuses on the event.
  - DHS Privacy Advocate

Summary and Recommendations

This article presents a categorization of cyber attacks and proposes a set of future cyber ecosystem capabilities to mitigate those attacks. These cybersecurity capabilities, when built into the future cyber ecosystem components and systems, will help strengthen the security and resilience of the cyber ecosystem.

The list of desired capabilities is not expected to change as a result of changes in threats, attack methods, technologies, and processes. This is because our approach is based on broad attack categories, not the specific technical details of those cyber attacks that will change as technology evolves. Although the paper's list of capabilities is not guaranteed to be complete, it does not include characteristics that will become unnecessary in the future. The cyber ecosystem itself is continuously evolving. Recent major evolutional trends are toward mobility and cloud computing. The cyber ecosystem capabilities must be able to adapt to support new environments, such as cloud and mobile. Federal and industry research and development (R&D) are key to the development of many of the desired capabilities.

The federal government's R&D community has developed a plan [13] for the research required to support the development of future cyber ecosystem security capabilities. Among the areas of emphasis in the plan is develop improved metrics for accessing cybersecurity risk and developing cyber security economic investment incentives; tailored trustworthy spaces and moving target [13].

Charney reminds us that collective solutions require collective development and integration. "To build on the current national and industry efforts, we can identify what is working and what is not, and document both to enable more individual action and community building. We can also begin to work through international bodies to standardize what types of information on machine health should be shared and how to exchange it with appropriate security and privacy protections." [8]

Acknowledgement:

The author gratefully acknowledges the assistance of Robin A. Simmons of The MITRE Corporation in preparing this article.

ABOUT THE AUTHOR

Dr. Fonash is the Chief Technology Officer for DHS' Cybersecurity and Communications organization. Dr. Fonash has held several senior positions at the National Communications System (NCS). He was Deputy Manager and Director of the NCS. Prior to that Dr. Fonash was Chief, NCS Technology and Programs Division. He managed special Presidential and priority communications services technology development, nationwide network modeling and analysis, specialized telecommunications research and development, and the deployment of (NS/EP) priority communications services nationwide on all major commercial networks.

Before arriving at the NCS, Dr. Fonash served as the Chief of the Defense Information System's Agency Joint Combat Support Applications Division, providing technical software integration services to the functional communities and guiding functional applications' compliance with the standard common operational environment. He also worked for the Office of the Assistant Secretary of Defense for Command, Control, Communications and Intelligence, and was responsible for Defense communications infrastructure policy and program oversight. He was also Chairman of the Office of the Secretary of Defense Information Technology Architecture Council.

Dr. Fonash has a Bachelor of Science in Electrical Engineering and a Master of Science from the University of Pennsylvania, a Master of Business Administration from the University of Pennsylvania's Wharton School, and a Doctor of Philosophy in Information Technology and Engineering from George Mason University.

REFERENCES

1. "Enabling Distributed Security in Cyberspace: Building a Healthy and Resilient Cyber Ecosystem with Automated Collective Action", DHS National Protection and Programs Directorate, 23 March 2011.
2. "Computer Security Incident Handling Guide" (draft), National Institute of Standards and Technology Special Publication 800-61 Revision 2, March 2012.
3. "2012 Data Breach Investigations Report", Verizon Corporation, March 2012.
4. Common Attack Pattern Enumeration and Classification, http://capec.mitre.org
5. Eric M. Hutchins, Michael J. Cloppert, and Rohan M. Amin, "Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains", Lockheed Martin Corporation, November 2010.
6. Deborah J. Bodeau and Richard Graubart, "Cyber Resiliency Engineering Framework", MITRE Technical Report MTR11023, September 2011.
7. "Blueprint for a Secure Cyber Future: The Cybersecurity Strategy for the Homeland Security Enterprise", Department of Homeland Security, December 2011.
8. Scott Charney, "Collective Defense: Applying Public Health Models to the Internet", Microsoft Corporation, October 2010.
9. Kevin Sullivan, "Collaborating to Secure Consumer Devices: Promoting Device Health for a Safer, More Trusted Internet", Microsoft Corporation, May 2011.
10. David M. Chess, Charles C. Palmer, and Steve R. White, "Security in an Autonomic Computing Environment", IBM Systems Journal, Volume 42, Number 1, 2003.
11. "Meeting the Cybersecurity Challenge: Empowering Stakeholders and Ensuring Coordination", IBM U.S. Federal White Paper, February 2010.
12. Recommendations from the Cyber Ecosystem Working Group, a working group formed by the Cross-Sector Cyber Security Working Group, January 2012.
13. Trustworthy Cyberspace: Strategic Plan for the Federal Cybersecurity Research and Development Program, Executive Office of the President, National Science and Technology Council, December 2011.

NOTES

1. The cyber ecosystem is global, evolving and includes government and private sector information infrastructure; the interacting persons, processes, data, information and communications technologies; and the environment and conditions that influence their cybersecurity.

Defining Proactive Software Assurance Practices for Healthier Cyber Ecosystems

Brian Badillo, Harmonia Holdings Group, LLC
Marc Abrams, Harmonia Holdings Group, LLC

Abstract. Distributed security in cyberspace can be performed using many real-time components, from intrusion detection systems to incident management systems. However, these components are reactive rather than proactive. This article will define the space within resilient cyber ecosystems that represents proactive software assurance during the software development lifecycle, from requirements to design to implementation to testing and beyond. Section 1 defines software assurance and provides some examples of current tools and practices in the government that are used for this proactive security practice. Then Section 2 defines the term cyber ecosystem as presented by DHS so that Section 3 can use this concept to explore the space of a software assurance ecosystem. Then in Section 4 we share our experiences in developing a software assurance infrastructure that implements the principals of a software assurance ecosystem and also bridges the gap between proactive and reactive systems. A healthy software assurance ecosystem is critical. The government needs the capabilities to quickly and efficiently certify and accredit systems to minimize vulnerabilities so they can be connected to networks such as the DoD Global Information Grid.

Section 1: How Does the Government Currently Do Software Assurance?

According to the Committee on National Security Systems, software assurance is the, "Level of confidence that software is free from vulnerabilities, either intentionally designed into the software or accidentally inserted at any time during its lifecycle, and that the software functions in the intended manner" [1]. In this section we mention current software assurance practices and tools used by the DoD, however, many of the principals apply to other agencies and organizations. It is critical that DoD information systems assure software as a proactive security measure before using them in operations. The DoD Information Assurance Certification and Accreditation Process (DIACAP) ensures that risk management practices are put in place during DoD software development. Essentially, DIACAP is a formal framework in which software assurance can take place so that the software assurance process is well documented. DIACAP is a formal, well-defined set of activities, tasks, and management workflow for certifying and accrediting software for the DoD.

There are some software systems in the government that help to enforce workflow in DIACAP including DoD's Enterprise Mission Assurance Support Service and the Air Force's Enterprise Information Technology Data Repository. These systems assist parties undergoing DIACAP by providing management services for workflow among the various roles in the DIACAP process, report generation capabilities for DIACAP requirements, and repositories for generating required reports, and repository capabilities for data pertinent to DIACAP.

In support of secure systems, the Defense Information Systems Agency has introduced Security Technical Implementation Guides that help define specific ways to ensure security in a system. The Application Security and Development STIG is particularly pertinent to software assurance because it defines specific guidelines to be followed by application designers to ensure security (e.g., "The Designer will ensure the application does not display account passwords as clear text" [2].).

The next section summarizes how DHS defines a cyber ecosystem so that we can define a software assurance ecosystem in similar terms in Section 3. Then in Section 4 we share our experiences building a software assurance infrastructure to illustrate the principals.

Section 2: What Is a Cyber Ecosystem?

On March 23, 2011, DHS posted a blog entry with a white paper titled, "Enabling Distributed Security in Cyberspace" [3]. The white paper provides an overview of distributed cyber security approaches and represents the collective vision of 13 federal agencies towards a healthy cyber ecosystem. A cyber ecosystem is defined as the set of diverse participants (which include cyber devices such as computers, software, and communications technologies) that interoperate. However, the ecosystem does not stop at cyber devices, but also includes other participants such as private firms, non-profits, governments, individuals, processes, etc.

The white paper presents three building blocks of cyber ecosystems: Automation, Interoperability, and Authentication. The ecosystem described in the white paper is the operational side of the distributed cyber ecosystem (i.e., running servers, network devices, production software, etc.). The operational side necessitates reactive security measures such as intrusion detection and real-time courses of action. The next section uses the three building blocks to describe a software assurance ecosystem, which is the software development side of cyber ecosystems, and includes requirements, design, implementation, and test stages. In contrast to the operational side, the security in the software development side is proactive in nature and includes security activities such as software assurance.

Section 3: How Does Software Assurance Fit in Cyber Ecosystems?

Software assurance is an important part of any software development project to meet quality, safety, and security requirements. However, in today's software development world, enterprise software assurance capabilities must match the security needs of the growing and diverse cyber ecosystem. For example, many software vendors utilize open source software or acquire COTS software to include in their solution. In this case, each software component is now part of the software assurance ecosystem. Furthermore, the intercommunication between software necessitates standardization of software assurance capabilities to provide a common interface.

The same three building blocks used to describe the secure operations portion of healthy

cyber ecosystems can be used to describe software assurance ecosystems.

Automation

Software assurance tools can aid software developers in making important security, quality, and safety decisions during the requirements, design, implementation, and testing stages. These tools automate parts of the software assurance process by performing much of the brute force work for identifying software weaknesses, which then allows developers to sift through the suspected weaknesses identified by automated tools and decide which weaknesses need further action (later sections describe tool automation as a way to collect evidence for software assurance cases). Software assurance tools can be classified into several analysis approaches and techniques, each of which have specific advantages and identify a specific subset of software weaknesses. A classification of tools is given by NIST Software Assurance Metrics And Tool Evaluation (SAMATE) project <http://samate.nist.gov/index.php/Tool_Survey.html> and includes such tool classes as static analysis, dynamic analysis, pedigree analysis, binary code scanners, disassembler analysis, binary fault injection, fuzzing, etc.

Interoperability

Given the multitude of tools in the software assurance ecosystem (more than 75 listed on the SAMATE project website), standards for interoperability among these tools is a necessity.

One such standard is the Common Weakness Enumeration (CWE) <http://cwe.mitre.org>. CWE is a dictionary of software weakness types developed by MITRE, intended to facilitate communication about weaknesses in software such as code constructs that are prone to memory leaks, susceptible to injection attacks, etc. From person to person, descriptions of these weaknesses can often be inconsistent; the CWE dictionary gives a standardized reference point as well as levels of specificity for these weaknesses. They are organized in a hierarchy, with general weaknesses (e.g., CWE-710: Coding Standards Violation) at the top level, getting increasingly more specific towards the lower levels (e.g., CWE-259: Use of Hard-coded Password). This hierarchy allows weaknesses to be related to each other with parent/child relationships. Some weaknesses also relate to each other with a precede/follow relationship that suggests that one weakness may be caused by another. Each weakness has a self-explanatory title, accompanied by an index. For example, the weakness described as NULL Pointer Dereference has the index of 476. Many tools available today already reference the CWE indices in their output. CWEs are part of a larger initiative called Making Security Measurable <http://measurablesecurity.mitre.org> to standardize system security.

With tools in the software assurance ecosystem using CWEs to represent their output, developer participants in the ecosystem can use a wider array of tools because using each tool that outputs the familiar CWEs will be easier to learn. Using a combination of tools for software assurance in turn leads to more assurance coverage of software. For example, consider a developer who is already familiar with a static analysis tool of their choice for detecting memory management weaknesses in code. Suppose that their familiar tool maps the weaknesses it identifies to CWEs. Since the developer already has knowledge concerning the weaknesses identified by their tool of choice, they are able to easily use and understand other tools that also produce CWE output.

Another advantage of interoperability is the ability to leverage collective bodies of knowledge concerning common assurance cases. An assurance case is defined as claims, arguments, and evidence that support the contention of particular software requirements [4]. In effect, an assurance case builds confidence in a system given evidence found by automated software assurance tools.

The Software Assurance Evidence Metamodel (SAEM) and the Argument Metamodel (ARM) are standardized models for representing parts of an assurance case, both of which were developed by Object Management Group's (OMG) Systems Assurance Task Force (SATF) <http://sysa.omg.org>. Arguments are logic that combines evidence and other asserted claims in a meaningful way to support or refute another particular claim [5]. The most primitive building blocks arguments are premises and conclusions. The argument asserts that if all the premises are accepted as true, then the conclusion must also be accepted. Arguments can be chained together such that the conclusion of one argument can provide the input to a premise in another argument. A CWE could contribute to evidence in a claim. However, an evidence item in SAEM contains additional useful information to be used in an assurance case, such as the evidence collection method used. Information that SAEM might include is the name and version of the tool used to identify the CWE, the time that the CWE was assessed, or the confidence level given to the evidence item. In addition, SAEM represents whether the evidence strengthens or weakens an assertion made by the evidence (which would in turn support an argument which uses the evidence to make a claim).

CWE, SAEM, and ARM are part of a larger Software Assurance Automation Protocol (SwAAP). SwAAP is a protocol composed of many interrelated standards [6].

Authentication

In the operations realm of the cyber ecosystem, authentication means making sure that the users of cyber devices are who they say they are (which includes both human and machine users). However, in the software assurance ecosystem, authentication means making sure that codebases are in fact the ones that have undergone the extensive assurance processes that they say they have. For instance, suppose that a library for protecting against Cross Site Scripting (XSS) is used in the security of a critical web application. The developers of the web application have decided to use this particular library because it has been vetted by Independent Verification and Validation (IV&V). However, when packaging the web application for production use, the library is not authenticated (i.e., the codebase is not checked to be from the expected supplier) and a malicious look-alike library is used in deployment. Now, the web application contains open vulnerabilities for attack.

Software supply chain integrity is another facet of authentication in the software ecosystem. According to SAFEcode not only must codebases be authenticated to make sure that they are the expected software, but the software must be expected to use secure, safe, and quality assurance processes during development [7]. It is therefore important to consider developer pedigree and policy during sourcing, development, and distribution.

Section 4: Building a Software Assurance Infrastructure

In this section, we discuss our experiences in applying the principles above by combining open standards and open source technologies

Figure 1: Infrastructure for a Software Assurance Ecosystem Overview

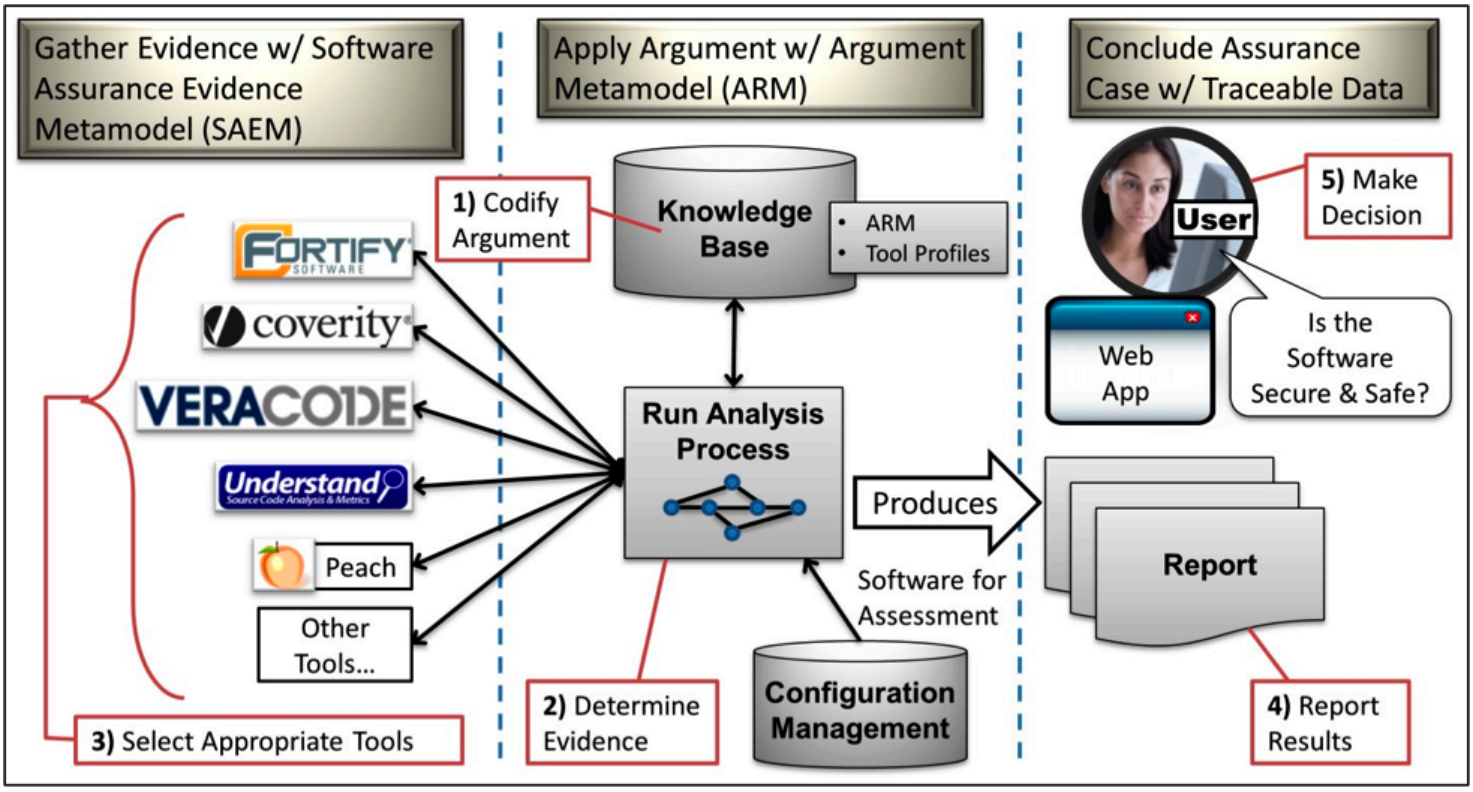

into an infrastructure. Software tooling that supports the three building blocks of software assurance (automation, interoperability, and authentication) is needed to make pre-incident detection practices during the software development lifecycle a reality. For instance, there are many software assurance tools that can be used to identify weaknesses in code, some of which already conform to the SwAAP standards (e.g., produce CWE output). However, traditional tools perform a single class of analysis approach or technique (i.e., static analysis, dynamic analysis, fuzzing, etc.) that provides them certain strengths and shortcomings. In addition, tools are usually focused on finding weaknesses in code developed in a particular language or for a specific platform. Furthermore, any single tool is subject to generating false-positive findings (e.g., a weakness in code that does not lead to vulnerability). In the rest of the section, we discuss our experiences in implementing the principles of a software assurance ecosystem through a software assurance infrastructure called Conforma.

In our experience while building the Conforma software assurance infrastructure we found that the infrastructure can provide the foundation to combine best-of-breed tools from many tool classes that have overlapping CWE coverage to increase confidence and reduce false-positive findings in software assurance. To accomplish this, the Conforma infrastructure contains a Tool Profile for each third-party tool plugged in to the infrastructure. The profile uses Coverage Claims Representation (CCR) [cwe.mitre.org/compatible/ccr.html] from the CWE standard to express which CWEs each tool claims to uncover. This profile enables Conforma to orchestrate the execution of appropriate third-party tools given a set of evidence that must be found to support an assurance case. While some third-party tools already produce CWEs (e.g., Fortify, Veracode, Klocwork), Conforma must map to a CWE each message generated by tools that do not (e.g., Splint, Peach). With the number of tools available, cross checking between tools using the common CWE output provides a base evaluation of the confidence level regarding the results. A software assurance infrastructure, such as Conforma, computes percentages involving the number of tools that found a certain error. For example, some types of tools reliably find particular weaknesses, but if multiple tools report the same weakness then a user's confidence that the weakness is a valid result, and not a false-positive, increases. In this respect, a software assurance infrastructure harnesses an ecosystem of tools to the advantage of the user by increasing confidence and reducing false-positives.

In Figure 1, there are three columns, which represent the three major parts of our software assurance infrastructure design (and the associated OMG standards that they leverage). The left side of the figure shows a list of software assurance tools that are plugged into the infrastructure. Each of these tools performs some sort of analysis, which produces CWEs that the infrastructure wraps into evidence in the form of SAEM, which is sent back to the infrastructure. The middle of the figure depicts a knowledge base containing rules and workflows that

support ARM and that are executed by the infrastructure. In other words, the rules will model claims in the form of premises that must be satisfied in order for certain conclusions to be made. The right side of the figure shows the users of the infrastructure (human participants in the software assurance ecosystem) using a web application UI to make conclusions about the software under assurance assessment and decide whether the evidence, arguments, and claims made in the assurance cases indicate that the software is ready to become an operating member of a healthy cyber ecosystem.

Each of the red boxes in Figure 1, in the order of how each part is used, is described in the list below.

1. **Codify Argument:** Arguments are codified (written in the standard ARM format) and stored in the Knowledge Base.
2. **Determine Evidence:** The codified ARM model is used in the Analysis Process to determine the evidence that is needed for certain claims to be made.
3. **Select Appropriate Tools:** With the evidence identified along with the Tool Profiles stored in the Knowledge Base (which describe tool coverage using CWE Coverage Claims Representation), the appropriate tools are executed by the infrastructure. These tools produce CWEs which should be used to strengthen or weaken evidentiary assertions.
4. **Report Results:** After all tools have been executed and the Analysis Process is complete, the user can initiate the generation of a report. The report shows the resulting claims about the software that can be made using the evidence that has been found using the appropriate tools.
5. **Make Decision:** Finally, the user can answer the question of whether or not the software is secure, safe, and of good quality. They can make a claim that is backed up by the arguments made in conjunction with evidence found by the infrastructure.

An infrastructure such as Conforma can be deployed within an enterprise to support the needs of a single software development house. However, over the course of our work developing a software assurance infrastructure we have learned that the power of the infrastructure is truly realized when deployed on a cloud environment where software assurance community cooperation can be achieved. In a cooperative environment, the infrastructure learns from the software assurance ecosystem participants by continuously expanding its Knowledge Base in real-time, which can then be used across the infrastructure for improved software assurance. This concept is illustrated in Figure 2 where two sets of Assurance Tools (far right and left sides) share through the Collective Intelligence Center some ARM data for assurance cases, Tool Profiles for up-to-date tool data including new CWE mappings and CCR coverage, and COTS Tool Plugins for increased interoperability between tools. The infrastructure deployment depicted in Figure 2 fosters a software assurance ecosystem through knowledge sharing.

Conforma itself is designed to learn how to better assure software when software is in operation in the cyber ecosystem. Conforma reacts to detected vulnerabilities and attacks during operation and learns which parts of the code base were not properly assessed in the assurance process. If there were tools that produced evidence that was originally deemed false-positive, then the tool profile is updated to reflect a different

Figure 2: Community Deployment Overview

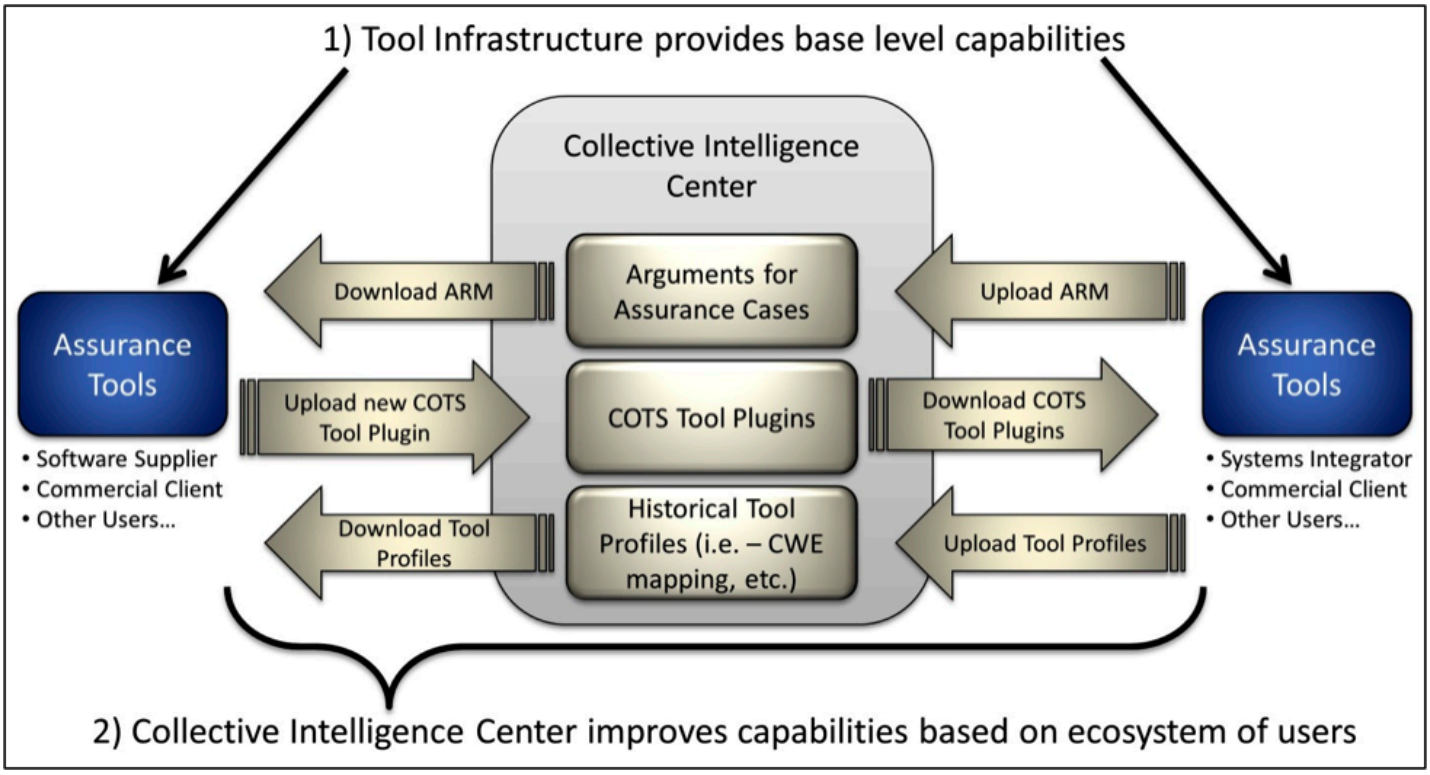

level of confidence in that particular tool. For example, suppose a particular tool (with which Conforma associates a high level of confidence) showed that an input field in a user interface was being properly validated to protect against an XSS attack. If a successful XSS attack on that input field is detected, Conforma would promptly lower the level of confidence associated with that particular tool for assuring input field validation. This information would then be shared across the infrastructure to all users in the software assurance ecosystem as part of the tool's profile.

Securing a cyber ecosystem can be divided into two methods: reactive and proactive. In our experience building a software assurance infrastructure, we found that we could complement reactive security with proactive software assurance, and vice versa. Thus, in addition to fostering a tighter software assurance ecosystem, Conforma bridges the gap between the operations and development lifecycle phases of software in cyber ecosystems by using both reactive and proactive security measures. New vulnerabilities and new attacks continue to be identified every day in the cyber world. It is important that the software assurance community learns how to protect against these vulnerabilities. The Conforma infrastructure is designed to improve its own assurance processes by detecting vulnerabilities and attacks during operation of software that was assured within its infrastructure.

REFERENCES

1. United States. Committee on National Security Systems. National Information Assurance Glossary: CNSS Instruction No. 4009. By Richard C. Schaeffer, Jr. 26 Apr. 2010. Web.
2. United States. DISA for Department of Defense. Application Security and Development: Security Technical Implementation Guide. Version 3, Release 4, 28 Oct. 2011. Web. <http://iase.disa.mil/stigs/app_security/app_sec/u_application_security_dev_stig_v3r4_20111028.zip>.
3. United States. Department of Homeland Security. Enabling Distributed Security in Cyberspace. U.S. Department of Homeland Security, 23 Mar. 2011. Web.
4. Software Assurance Evidence Metamodel (SAEM). Publication no. Ptc/2010-08-37. Object Management Group, Inc. (OMG), Aug. 2010. Web. <www.omg.org/cgi-bin/doc?ptc/10-08-37.pdf>.
5. Argumentation Metamodel (ARM). Publication no. Ptc/2010-08-36. Object Management Group, Inc. (OMG), Aug. 2010. Web. <http://www.omg.org/cgi-bin/doc?ptc/10-08-36.pdf>.
6. Jarzombek, Joe. "Public/Private Collaboration Efforts for Enterprise Security Automation." Speech. Software Assurance Forum: Building Security In. Baltimore Convention Center. 27 Sept. 2010. Security Content Automation Protocol. U.S. Department of Commerce, NIST. Web. <http://scap.nist.gov/events/2010/itsac/presentations/day1/Software_Assurance-PublicPrivate_Collaboration_Efforts_for_Enterprise_Security_Automation.pdf>.
7. Reddy, Dan, Brad Minnis, Chris Fagan, Cheri McGuire, Paul Nicholas, Diego Baldini, Janne Uusilehto, Gunter Bitz, Yuecel Karabulut, and Gary Phillips. The Software Supply Chain Integrity Framework: Defining Risks and Responsibilities for Securing Software in the Global Supply Chain. Tech. Ed. Stacy Simpson. SAFECode, 21 July 2009. Web.

ABOUT THE AUTHORS

Brian Badillo, M.S., (Computer Science, Virginia Polytechnic University) is lead software engineer at Harmonia. He successfully completed seven Phase I Small Business Innovative Research (SBIR) topics, bringing three to Phase II, and has been awarded three Phase I SBIR topics. He has used many community efforts such as DHS' "Building Security In," MITRE's "Making Security Measurable," Open Web Application Security Project, and Microsoft's "Security Lifecycle Development." Using this research, he led Conforma development.

Harmonia Holdings Group, LLC
2020 Kraft Drive, Suite 1000
Blacksburg, VA 24060
Phone: 540-951-5900 Ext. 255
Fax: 540-951-5911
E-mail: bbadillo@harmonia.com

Marc Abrams, Ph.D., (Computer Science, University of Maryland; Post Doctoral Study, Stanford University) is Harmonia's President and CTO, providing technical and business leadership and overseeing all technical activities. He has more than 20 years of professional experience in the design, development, deployment, and maintenance of software and information networks, focusing on user interfaces. He architected Harmonia's LiquidApps® tool suite and led its implementation in the Army's ATIA-M project, US Navy's DDG 1000 destroyer, and Tomahawk weapons control system.

Harmonia Holdings Group, LLC
2020 Kraft Drive, Suite 1000
Blacksburg, VA 24060
Phone: 540-951-5901
Fax: 540-951-5911
E-mail: mabrams@harmonia.com

Recovery-based Resilient Cyber Ecosystem

Ajay Nagarajan, George Mason University
Arun Sood, George Mason University and SCIT Labs, Inc.

Abstract. Today's approach to security is largely based on perimeter defense and reactive strategies like Intrusion Detection and Intrusion Prevention Systems (IDS/IPS), firewalls and anti-virus products. Past experience has repeatedly shown us that this strategy is not complete and secure. Intrusion tolerance is an approach that treats intrusions as inevitable and shifts the focus from detection and prevention to containing losses and rapid recovery. We suggest that a complete security strategy is one that does defense in depth and involves both traditional security strategies and intrusion tolerance. Security Information and Event Management (SIEM) is a framework that consolidates the plethora of information available from all of the network and security devices into useful information. In this paper, we propose a stand-alone and a collaborative architecture that makes use of information provided by the SIEM framework to perform adaptive intrusion tolerance in unsupervised learning environments. Resilient systems need to be adaptive, and to achieve this goal we show how environmental information can be used to adaptively change system parameters.

1. Introduction

The variety and complexity of cyber attacks are ever increasing. Verizon's 2012 Business Data Breaches Investigation Report [1] shows that customized malware is difficult to detect and data ex-filtration often occurs over a period of days, weeks and months. The current IDS/IPS approaches are reactive in nature and depend on prior information that is inadequate to prevent all attacks. Events such as the VeriSign security breach [2] and the Playstation Network breach [3] reinforce two notions: 1) even the most sophisticated IDS/IPS systems fail to detect/prevent every intrusion and 2) once the system is compromised, the intruder stays in the system doing damage for extended periods of time.

In addition to the shortcomings of IDS/IPS systems, the costs of operating them are high and increasing. To illustrate the issue we take the example of an enterprise with an average of 1 million raw events occurring per day. About 10,000 alerts are generated by perimeter defense systems. Out of these, 100 alerts are correlated on the basis of severity and other considerations. Assuming it takes 1.5 man-hours to handle one alert, a total of 150 man-hours are required per day to handle alerts generated. The cyber security requires 365 days, 24 hours per day support and in general about 30 people are required to carry out this task. How many large companies can afford such an allocation of manpower? In companies we talk to, only two or three people perform this task. What is worse, 50 % of the alerts are false positives—a tremendous waste of resources. With ever increasing bandwidth and millions of new malware items created every day, these numbers are bound to increase.

Despite years of research and investment in developing such reactive security methodologies, our critical systems remain vulnerable to cyber attacks. The reactive perimeter defense approach relies heavily on threat modeling and vulnerability elimination. We suggest that additional attention should be given to the consequences of a successful attack. In our approach, we focus on limiting the consequences, like reducing the losses that are induced. We believe that we must make our cyber systems more proactive and resilient. Such systems will have the property of (1) supporting continuity of operations—working even in the presence of an intruder; (2) losses, if any, must be limited; (3) systems must resume full operations, i.e. system must be restored to a known good state; and (4) the resilient system operations should be independent of the threat.

To design such a system, we assume that intrusions are inevitable. Therefore, we shift our focus from modeling threats/vulnerabilities to developing methods that will minimize the consequences of an intrusion, increase the work effort of the adversary and increase the visibility of the adversary to the defenders. For this, we have developed a moving target defense approach to computer security. We focus on building mission resilient systems that are able to work through an attack. To ensure reliable operations, the system is restored to a pristine state once every short period of time known as the exposure time, thus negating any malicious action performed by the adversary and minimizing consequences. In addition to this, we use redundancy to provide uninterrupted service and increase overall system availability. The more frequent the computer restoration the less likely it is for the intruder to do damage. The restoration frequency can be random to confuse the adversary and increase his work effort. The shortest time between restorations is a trade-off between available system resources and the throughput of the computer. This intrusion tolerant technology is called Self Cleansing Intrusion Tolerance (SCIT) [4]. The recovery driven approach of SCIT is compared to the detection driven and other intrusion tolerance approaches [5].

Consistent with CrossTalk's theme for the September/October 2012 issue, in this paper, we propose a resilient cyber ecosystem in which every member is able to work together and learn from one another in near-real time to predict and prevent cyber attacks, limit propagation of attacks across participating entities, minimize losses occurring from successful attacks and rapidly recover to a pristine state. To build such a system that is resilient to a variety of sustained attacks, we propose a model that integrates tools and mechanisms that provide protection and detection as well as adaptive tolerance. The rest of the paper is organized as follows: Section 2 provides a brief overview of how SCIT works and motivates the rest of the paper by presenting the need for adaptive SCIT, Section 3 introduces SIEM solutions and presents our idea on how information from SIEM solutions can be used to build adaptive intrusion tolerance systems. We will review two scenarios—stand-alone adaptive intrusion tolerance architecture and a peer-to-peer collaborative intrusion tolerance architecture.

2. How SCIT Works

In [4] we presented SCIT, an intrusion tolerant technique that provides enhanced server security. SCIT research has focused

on critical servers that are most prone to malicious attacks. The technique involves multiple virtual instances of servers that are rotated and self-cleansed periodically irrespective of the presence or absence of intrusions. Self-cleansing refers to loading a clean image of the server's OS and application into the Virtual Machine. Rotation here refers to the process of bringing an exposed virtual server off-line, killing it, restarting it and in the meanwhile, bringing another virtual server online to assure availability. By doing so, in the event of an intrusion, the intruder is denied prolonged residence on the server. Once the virtual server's exposure time to the Internet is completed, the virtual server instance is automatically rotated. This virtual instance of the server is what is referred to as virtual server throughout this paper.

Figure 1: SCIT Server rotation

This illustrative example in Figure 1 shows 3 different time periods. At any given time, there are five servers online and three servers being wiped clean. In each case a different set of servers is being cleaned. Eventually every server will be taken offline, cleaned and restored to its pristine state. SCIT technology can be used to build a variety of servers that meet enhanced security requirements. It is best suited to servers that are designed to handle short transactions—the lower the exposure time the shorter the transaction.

2.1 Need For Adaptive SCIT

Resilient systems have to exhibit adaptive and recovery behavior. SCIT is recovery driven, and in this section we show how SCIT can be made more adaptive to the ongoing changes in the environment.

At any point of time, the resilience of a SCIT system is affected by (1) the current attacks; (2) the current workload; (3) the current data integrity level; (4) the current data availability level; and (5) the current behavior of the system [6]. The first four factors together make up the environment of the SCIT system. Two SCIT systems with different behaviors can yield different levels of resilience. This suggests that as the environment and the behavior of the system changes, the effectiveness of SCIT changes as well. To achieve the maximum amount of resilience, the SCIT system must adapt itself to its environment. Through an architecture for adaptive SCIT, we can (1) adapt SCIT to different application semantics; (2) significantly improve the cost-effectiveness of SCIT; (3) prevent dramatic performance degradation due to system environment changes; and (4) maintain trade-off between system security and system performance [6].

In the case of SCIT, the primary metric is exposure time. In [7], we illustrated the relationship between exposure time and security of a system in terms of data compromised. In [8], we discussed the SCIT approach from the perspectives of effectiveness, tunable parameters, performance impact, and integration to application systems. From the derived expression for Mean Time to Security Failures $MTTSF_{SCIT}$, we were able to conjecture mathematically that decreasing the exposure time window will improve the resilience of a SCIT-based system. To adapt SCIT we will need to adapt the exposure time in response to systems parameters. Increasing $MTTSF_{SCIT}$ would require decreasing the exposure window; hence the cycle that a SCIT server has to go through will become shorter. In this space, there is a tradeoff between system security, performance and cost. Adaptive SCIT could help balance this trade-off in real time with the use of a dynamic exposure time window given the current operating environment and system behavior.

3. Use of SIEM Solutions

"The term SIEM, describes the product capabilities of gathering, analyzing and presenting information from network and security devices; identity and access management applications; vulnerability management and policy compliance tools; operating system, database and application logs; and external threat data" [9].

In addition to receiving inputs from IDS/IPS systems, we will use a SIEM solution to collect and correlate data from all the other sources mentioned in Figure 2 to characterize overall network behavior. This behavioral pattern is then compared with a database of normal network behavior patterns to identify irregularities. Based on the findings of this comparison and the severity of the irregularities, the SCIT controller tunes the "exposure time" of the SCIT-ized system to adapt to the current environment. Similar iterative periodic comparisons will help guide the unsupervised learning and automatic adaption of the SCIT-ized system.

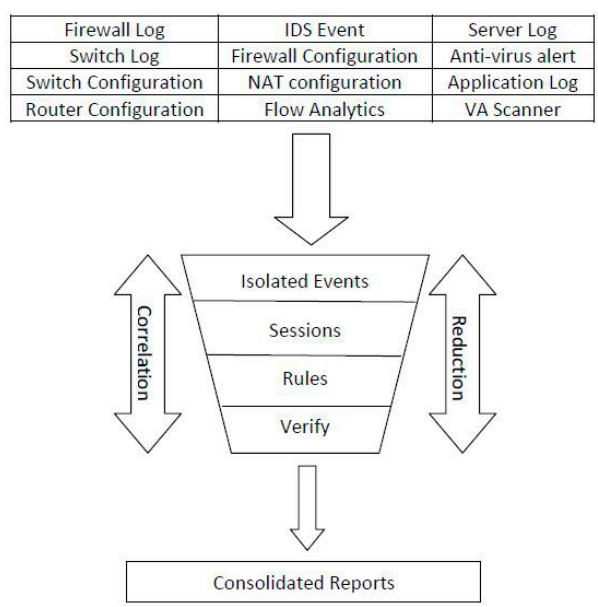

Figure 2: Security Information and Event Management Framework [10]

3.1 Use of Information from SIEM Solutions in Building Adaptive Intrusion Tolerant Systems:

In this section, we expand on the idea of using aggregated information from SIEM solutions to build adaptive intrusion tolerant systems. For the purposes of this paper, SCIT is the intrusion tolerance architecture of choice.

To address the needs outlined in section 2.1, an adaptive SCIT framework must do the following:

1. Employ a dynamic exposure time—the exposure window must keep changing with time as the SCIT environment and the system behavior changes.
2. Constantly receive input from the SIEM framework on the current SCIT environment and state of behavior to make informed alterations to the exposure window.

We present two adaptive SCIT architectures with a common assumption that SCIT is deployed at Enterprise level.

1. Stand-alone adaptive SCIT

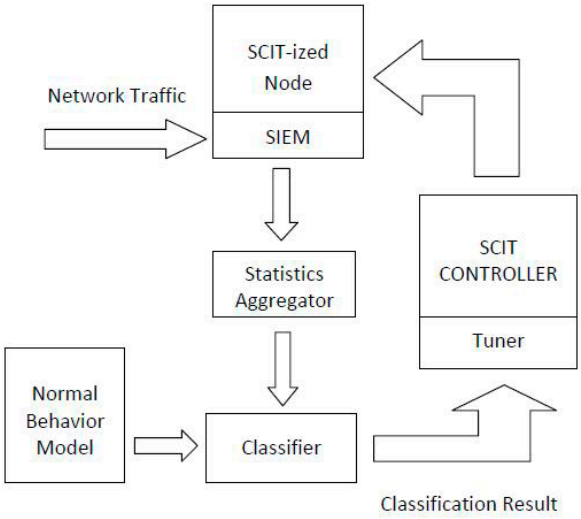

Figure 3: Stand-alone adaptive SCIT

In this architecture, SIEM is constantly monitoring the SCIT-ized node and periodically generates consolidated reports based on the information it has gathered and correlated from varying sources. These reports are fed into the Statistics Aggregator which converts massive information obtained from SIEM into meaningful metrics and their respective values. Further, the classifier compares pre-defined Normal Behavior Model (in terms of metrics and values) with the current values obtained from the Statistics Aggregator. The classifier then feeds the results of the comparison to the Tuner of the SCIT Controller. Based on this, the Tuner makes an informed decision on whether or not to alter the existing "exposure time."

For example, if the results from the classifier identify malicious behavior that points to a Distributed Denial of Service attack, then the SCIT Controller can now reduce the "exposure time" thereby hardening the system against such an attack.

2. Peer-to-peer collaborative SCIT

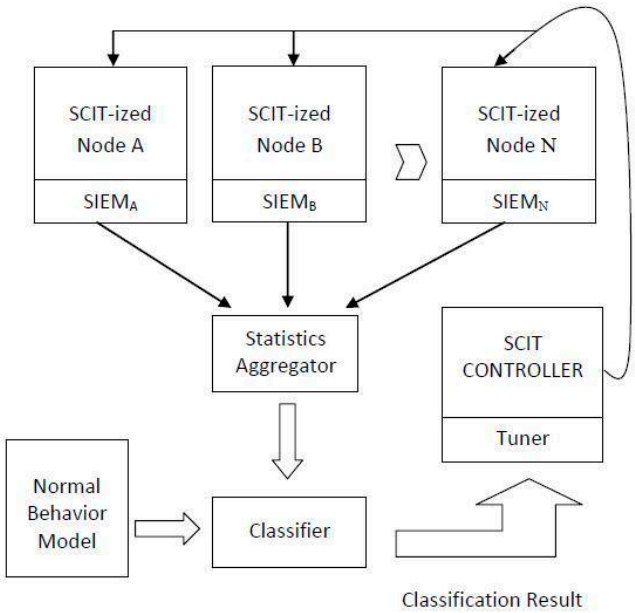

Figure 4: Peer-to-peer collaborative SCIT

This architecture is an extension of the stand-alone architecture. It is meant to mimic a cyber ecosystem with multiple participants in the community that offers recovery-based resilience. In this case, there are 'N' SCIT-ized nodes that are online concurrently. SIEM solutions of each individual node namely $SIEM_A$, $SIEM_B$ so on till $SIEM_N$ generate reports individually and keep forwarding them to the Statistics Aggregator periodically. The advantages of collaborative SCIT are straightforward:

1. There is more information to work with—the statistics aggregator is now fed with useful information from 'n' different SIEM solutions.
2. Acts as a pre-warning system: malicious behavior in any one of the nodes in the community can now be used to warn/harden the rest of the community.
3. Unsupervised Learning—malicious behavior in any one node in the community can help teach an attack pattern to the rest of the community.
4. Fewer chance of false positives since isolated events now carry less weightage.

4. Conclusion

Cyber attacks are becoming more widespread, sophisticated, and consequential with time. However, detecting, handling and identifying the consequences of an intrusion are still persistent problems. This is partly due to the lack of trust between the members of the cyber ecosystem that impedes information sharing and collaboration. If every entity of the cyber ecosystem were to collaborate with one another and took coordinated security decisions, it could lead to unsupervised learning systems that provide hardened proactive defense.

In this paper, we propose two such recovery based cyber resilient adaptive SCIT architectures. One is a stand-alone system and another is a collaborative system that encourages information sharing and promotes cyber health among communities. In addition to the periodic system self-cleansing done proactively, our system constantly partakes in unsupervised learning from other members of the ecosystem to adapt to the current environment and system behavior.

REFERENCES

1. Verizon Business Data Breach Investigation Report 2012
2. "Key Internet Operators VeriSign hit by hackers" Reuters 02/02/2012
3. "Security Experts: Playstation Network breach one of largest ever" USA Today, 04/27/2011
4. Yih Huang, David Arsenault, and Arun Sood, "Incorruptible System Self-Cleansing for Intrusion Tolerance", Proceedings Workshop on Information Assurance (WIA 2006), Phoenix, AZ, 2006
5. Quyen L. Nguyen and Arun Sood, "Comparative Analysis of Intrusion-Tolerant System Architectures", IEEE Security and Privacy, Volume 9 Issue 4, July-Aug 2011
6. Luenam P. and Peng Liu "The design of an adaptive intrusion tolerant database system" Foundations of Intrusion Tolerant Systems, 2003
7. Ajay Nagarajan and Arun Sood, "SCIT and IDS Architectures for Reduced Data Ex-filtration" 4th Workshop on Recent Advances in Intrusion-Tolerant Systems, Chicago, IL, USA, June 28 2010
8. Quyen Nguyen and Arun Sood, "Quantitative Approach to Tuning of a Time-Based Intrusion-Tolerant System Architecture", 3rd Workshop on Recent Advances in Intrusion Tolerant Systems, Portugal, June 29, 2009.
9. Security Information and Event Management – Wikipedia article
10. CISCO Security Monitoring, Analysis and Response System (MARS) Framework

ABOUT THE AUTHORS

Ajay Nagarajan is currently a Ph.D., candidate in Computer Science at George Mason University working under Dr. Arun Sood. He received his M.S. in Computer Science from George Mason University in 2010. He is affiliated with the SCIT Research group at GMU and his main research interests include Intrusion Tolerance, Survivability and Security Evaluation.

Volgenau School of Information Technology & Engineering
George Mason University, MS 4A5
4400 University Drive
Fairfax, Va. 22030
Phone: 540-687-0363
E-mail: anagara1@gmu.edu

Dr. Arun Sood is Professor of Computer Science in the Department of Computer Science, and Co-Director of the International Cyber Center (ICC) at George Mason University, Fairfax, VA. His research interests are in security architectures; image and multimedia computing; performance modeling and evaluation; simulation, modeling, and optimization.

He and his team of faculty and students have developed a new approach to server security, called Self Cleansing Intrusion Tolerance (SCIT). We convert static servers into dynamic servers and reduce the exposure of the servers, while maintaining uninterrupted service. This research has been supported by the U.S. Army, NIST through the Critical Infrastructure Program, SUN, Lockheed Martin, Commonwealth of Virginia CTRF (in partnership with Northrop Grumman).

Recently SCIT technology was the winner of the Global Security Challenge (GSC) sponsored Securities Technologies for Tomorrow Challenge. This technology has been awarded three patents and three additional patents are pending. SCIT Labs, a university spin-off, has been formed to commercialize SCIT technology. Dr. Sood is the founder and CEO of SCIT Labs.

Since 2009 Dr. Sood has directed an annual workshop on Cyber Security and Global Affairs with Office of Naval Research support – Oxford 2009, Zurich 2010 and Budapest 2011.

Dr. Sood has held academic positions at Wayne State University, Detroit, MI, Louisiana State University, Baton Rouge, and IIT, Delhi. His has been supported by the Office of Naval Research, NIMA (now NGA), National Science Foundation, U.S. Army Belvoir RD&E Center, U. S. Army TACOM, U.S. Department of Transportation, and private industry.

He was awarded grants from NATO to organize and direct advance study institutes in relational database machine architecture and active perception and robot vision.

Dr. Sood received the B.Tech degree from the Indian Institute of Technology (IIT), Delhi, in 1966, and the M.S. and Ph.D. degrees in Electrical Engineering from Carnegie Mellon University, Pittsburgh, PA, in 1967 and 1971, respectively.

His research has resulted in more than 160 publications, and his resume including publications list is available at <http://cs.gmu.edu/~asood>.

Volgenau School of Information Technology & Engineering
George Mason University, 4A5
4400 University Drive
Fairfax, Va. 22030
Phone: 703-993-1524
Fax: 703-993-1710
E-mail: asood@gmu.edu

Cyber Mission Resilience

Mission Assurance in the Cyber Ecosystem

Chris Peake, Sentar Inc.
Al Underbrink, Sentar Inc.
Dr. Andrew Potter, Sentar Inc.

Abstract. Cyber Mission Resilience (CMR) is a significant step in the evolution of IT security. Not only does it reduce the complexity and cost of securing today's IT systems, it helps prioritize security-related activities. The focus on mission resilience extends the scope of past security practices while simultaneously honing in on mission-critical systems, networks, and processes. This article explores the concepts and some of the challenges related to CMR and suggests areas for future research and study.

1. Introduction

"Rapid technology advances over the past three decades and the proliferation of computers into weapon systems created a dichotomy of net-centric military superiority and a commensurate reliance on vulnerable technology" [1].

The terms "cyber" and "cyberspace" are used in everyday conversation, as well as in the media, but their meanings are vague. Most definitions describe "cyber" as groups of networks and computers. But that is not all that cyberspace embodies; it is also the "place" where people interact, share, learn, play, work, communicate, explore, buy, sell, and connect. So "cyberspace" is much more than simply a collection of networks and computers; it is also what people do with the networks and computers.

For today's Military, cyberspace is mission-critical; cyber technology is embedded in nearly every part of daily operations. But since cyber technology and information systems are sometimes vulnerable to disruption, the supported missions are also susceptible to disruptions. Current efforts to manage cyber risk focus on preventing attacks on systems and information, but this approach is reactive in nature and cannot keep pace with the threat. Nor does this approach account for the fact that systems are just as susceptible to faults, failures, and accidents that can produce the same effects as cyber attacks. This suggests that new perspectives and approaches to managing operational and cyber risk are necessary.

Most mission owners/operators realize that merely addressing system-specific vulnerabilities will not assure the mission. And they realize that effective operational risk management must consider a broader range of potentially harmful events that includes protecting systems against cyber-based faults, failures, and attacks. Therefore, achieving mission assurance in the cyber ecosystem means that mission owners/operators have a degree of confidence that their mission-critical systems will be capable of sustaining necessary operational parameters despite cyber degradation. CMR focuses on ensuring that DoD mission owners and operators trust (i.e. have confidence) that the mission-critical systems will perform as required when needed.

The Cyber Ecosystem

Achieving the CMR perspective requires that we first reconsider the cyber ecosystem as a whole. As opposed to hierarchical and stovepipe models, the cyber ecosystem is actually highly interrelated and interdependent. That is, each component both serves and depends on other aspects in the ecosystem. For example, cyber defense without intelligence regarding an adversary's offensive capabilities, and the requisite R&D/engineering capabilities, is ineffectual. Therefore, cyber defense cannot operate independent of cyber offense nor can either operate without trained personnel and governance.

In 2009 an independent study was performed by a group of IT security professionals for the U.S. Army Space and Missile Defense Command/Army Forces Strategic Command in an effort to help depict an understanding of the cyber ecosystem. The study produced a notional view of the cyber ecosystem where each functional area of cyber is highly interconnected with every other area (see Figure 1).

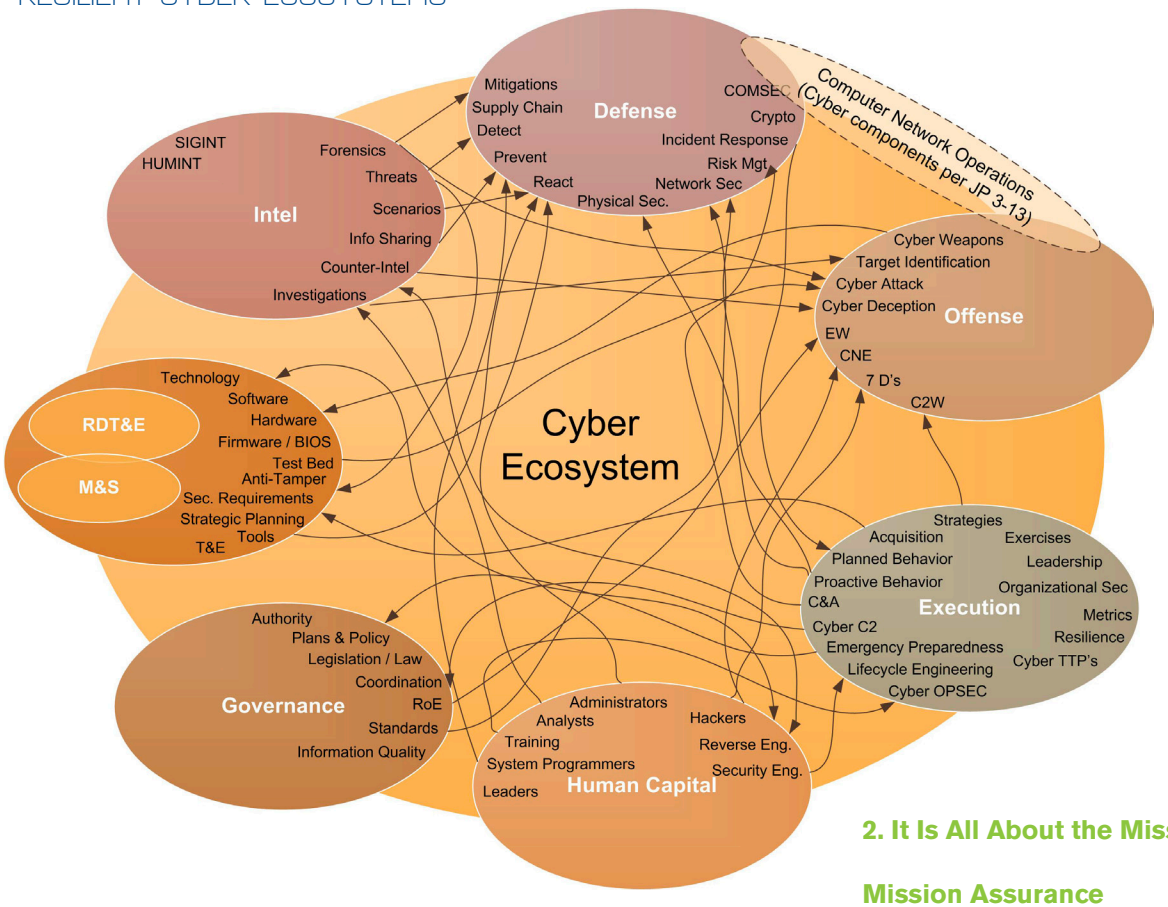

Figure 1: Notional view of the Cyber Ecosystem

Unlike the pillars of Information Operations described in Joint Publication 3-13, this view of the cyber ecosystem attempts to show the relationships among all functional areas in cyberspace. The resulting depiction of the cyber ecosystem is orders of magnitude more complex than what is expressed in current doctrine. Understanding the relationships and dependencies within the cyber ecosystem is a necessary precursor to adopting the CMR perspective.

Viewing Security as a Mission Enabler

The second step to adopting the CMR perspective is to break free from the misconception that security hampers mission functionality, and to start seeing cyber security as a mission enabler.

The mere mention of security gives most program managers and developers heartburn. For years, security has been considered a speed bump in the fast lane to project completion; security controls are thought to minimize capability, complicate architecture, and practically eliminate flexibility in system and software development. But this mindset has to change. Security should be seen as a mechanism to improve threat and fault tolerance in mission-critical system functions. Ideally, security controls should be implemented to ensure the achievement of mission objectives. Although security controls may still complicate the architecture and limit flexibility to some degree, a system developed to be more reliable, available, and dependable will be more efficacious in accomplishing the mission.

2. It Is All About the Mission

Mission Assurance

While the term "Mission Assurance" has only recently been applied to cyber, the concept itself is not new. With the increasing reliance on IT as a medium for carrying out mission objectives, there is a high probability that disruptions to information systems will have serious adverse effects on the overall mission. And despite the speed by which new software is being developed and security updates are made available, new exploits and vulnerabilities are being discovered and used even faster. In short, security professionals are losing the battle to keep our systems secure [2]. The reality is that perfect security is unattainable. Fortunately, mission-critical assets do not have to be perfectly secure; they just have to be secure enough to reliably accomplish their primary goals and objectives (i.e. their mission).

The DoD currently defines mission assurance as, "A process to ensure that assigned tasks or duties can be performed in accordance with the intended purpose or plan… to sustain military operations throughout the continuum of operations" [3]. This definition, while appropriate and applicable at the operational level, does not address the cyber aspects of mission assurance. If military missions depend on cyber technologies, then achieving mission assurance must also account for the mission-critical functions/tasks that are embedded in IT systems.

However, mission assurance does not guarantee mission success. It is a practice to manage operational risks that will increase the probability of achieving mission goals. As such, mission assurance can be expressed as a degree of confidence in mission success as opposed to a certainty of mission success/failure [4]. But identifying, tracking, and addressing risk, as it relates to mission goals and objectives, requires understanding the risk within the operational context (i.e. how the risk relates to achieving the mission).

The point being that mission assurance from an operational perspective cannot be achieved without assuring the cyber technologies upon which the mission depends.

Mission Resilience

The concept of mission resilience is closely related to that of mission assurance. Accenture's paper titled, "Mission Resilience: The New Imperative for High Performance in Public Service" was based on research conducted on 151 corporations and looked at how "routine disruptions" affected the organizations. While the study focused on public service organizations, the concepts presented in the paper are equally applicable to the DoD.

Accenture viewed mission resilience as a, "multi-tiered, life-cycle-focused methodology for understanding, anticipating, mitigating and minimizing the effects of any material disruption." Their model focuses on efficiency of mission both during normal operations and disruptive events. Unlike disaster recovery planning, mission resilience is a proactive approach that systematically prepares for potential disruptions as opposed to waiting for a disruptive event to occur.

To avoid wasting time and effort on trying to predict every possible cause of disruption, Accenture's mission resilience model focuses on protecting the mission from "symptoms" as opposed to specific events. "Building around symptoms (the effects) rather than scenarios (the causes) makes resilience development manageable because it recognizes that many events share characteristics, impacts and (most important) responses." [5] Resiliency is not necessarily about completely eliminating impacts from disruptions (or the disruptions themselves for that matter); it is concerned with minimizing impacts on the mission caused by the disruptions. If an organization or system is predisposed to handle disruptions, it can detect, react, respond, and recover more quickly to minimize the overall impact to the mission. Figure 2 depicts Accenture's resiliency capability concept; resilient organizations are able to minimize the overall impact of potentially harmful events on the mission.

Cyber Resilience

For organizations that rely upon cyber to either support or accomplish the mission, mission resilience becomes dependent

Figure 2: Accenture Resilience Capability

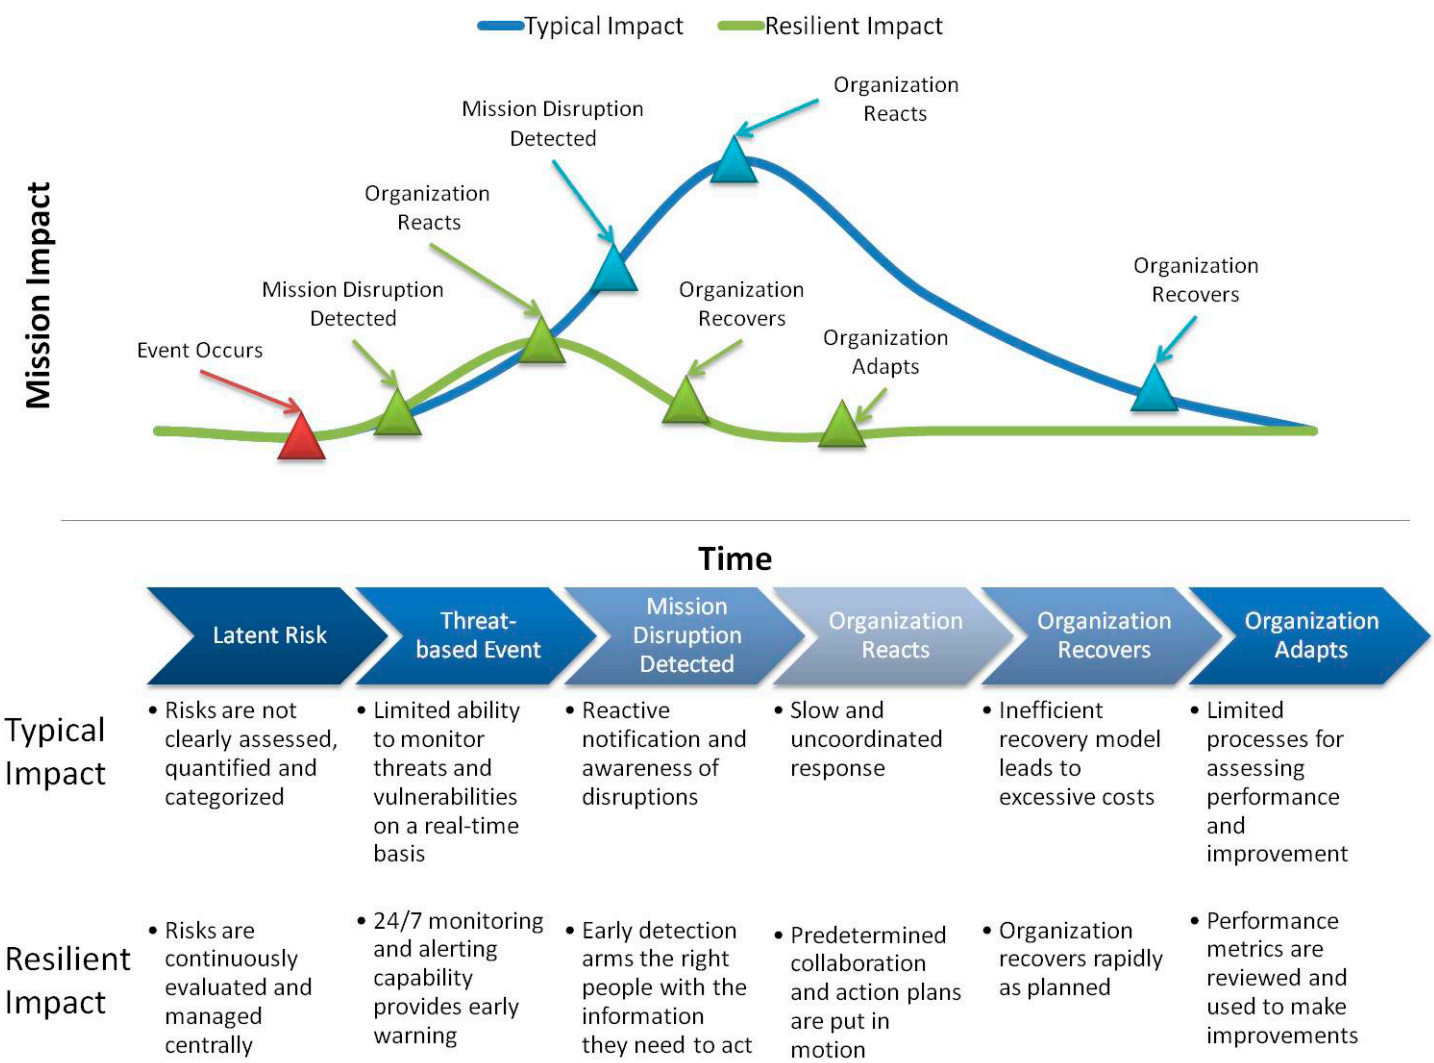

upon cyber resilience. Cyber resilience builds upon the properties of information assurance (i.e. availability, confidentiality, integrity, etc.) by introducing the concepts of maintainability, dependability, safety, reliability, performability[1], and survivability as aspects of system security [6]. The goal of cyber resilience is to sustain mission-critical system capabilities by applying security measures that assure the system can withstand cyber faults, failures, and attacks. Therefore, not only do mission owners need to consider a wider array of threats (i.e. faults, failures, and accidents in addition to cyber attacks), but they also need to assume these threats will affect their mission-critical systems. This assumption switches the focus from preventing cyber threats to minimizing the effects of cyber threats when they occur.

3. Achieving CMR

The concepts of mission assurance, mission resilience, and cyber resilience are, admittedly, confusingly interrelated. And while CMR may seem like just another play on words, in practice it combines the operational aspects of mission assurance and mission resilience with the technical objectives of cyber resilience. But achieving CMR requires a fundamental shift in system security processes, mindsets, controls, and tools. The following subsections discuss three tasks that need to be addressed in order to achieve CMR.

Understanding the Mission

Surprisingly, identification of an organization's mission-critical systems is not immediately self-evident. Complexity, created by the interdependence of systems, can make it difficult to determine which systems and processes are actually critical to the mission. Defining mission criticality requires identifying the impact a particular system has on overall mission success.

One tool used to help define an organization's mission, and mission critical resources, is the Business Impact Analysis[2] (BIA). While a risk assessment considers the threats and vulnerabilities associated with individual systems, the BIA is a comprehensive assessment of system functions that will reveal operational impacts, recovery time objectives, and functional dependencies of the mission critical assets. Once the BIA is complete, each mission-critical system should be assessed for its unique system protection needs.

The system protection needs are based on an examination of the potentially harmful effects (generated from the cyber threat) that can negatively impact the mission [7]. Even though most current risk assessment methodologies were developed according to information assurance-based doctrine, where "threats" referred to malicious adversaries and cyber attacks, the protection needs assessment process is equally applicable to the broader mission assurance definition of threats (e.g. system failures or faults). The key is to understand the impacts and consequences generated by the threat not necessarily the source of the threat [4].

The focus of CMR is on protecting the mission-critical systems against any event or effect that may cause a system disruption that subsequently leads to the failure to achieve mission objectives. To achieve this, it is imperative that mission owners/operators understand the mission dependencies on cyber assets and the operational parameters that are necessary to sustain the mission. In doing so, mission owners and operators are able to improve the confidence in mission success—thereby attaining a degree of mission assurance. But this is only possible if the mission and its dependencies are fully understood.

Resilience Metrics

Assurance practices are fundamentally about establishing confidence and trust, which suggests a need for qualitative and/or quantitative validation of the object being assured (i.e. assurance is not a question of belief). However, as discussed previously, assurance is not a guarantee or certainty either. But confidence and trust can be built through demonstration of reduced variation, improved dependability, consistent performance, and stable reliability. Demonstrating these qualities is a matter of repetition and statistical measurement.

As a result, cyber resilience metrics play a crucial role in achieving cyber-based mission assurance. They can be used to quantify the dependability, maintainability, safety, performability, reliability, and also the overall survivability/resilience of the mission-critical systems/functions as an assessment of mission assurance [8]. However, a comprehensive assessment of resilience requires metrics that address the full spectrum of cyber threats. Therefore, resilience metrics should also include fault measures in addition to security metrics. For example, dependability is a metric based in part on the measures of reliability and maintainability, and can address the performance of mission functions during attack or failure [9]. The aggregated metric actually provides a higher level of assurance understanding, which can be directly applied to mission objectives.

According to the Joint/Coalition Mission Thread Measures Development Standard Operating Procedure [10], the Senior Warfighters Forum (SWarF) prioritized a list of capability attributes that defines metrics in terms of mission-based functions and activities. Figure 3 depicts the SWarF attributes associated with the Net-Centric Joint Capability Area. These attributes were selected specifically because they help define how well mission activities performed. For example, enterprise IT services that are robust, scalable, interoperable, and responsive would be considered effective as a Joint capability. However, enterprise services that are unreliable due to frequent faults, failures, or cyber attacks, would be considered operationally risky. Therefore, in order for enterprise IT services to be mission assured, they must also be resilient to cyber threats.

The needed outcomes of mission assurance quantitative studies are metrics for operational fault tolerance and operational risk tolerance. Although currently not defined, the ideal measure of mission assurance would be a mission survivability or resilience rating, which would combine all other metrics (e.g. robustness, timely, agile, available, secure, etc.) into a single measure that would provide mission owners/operators with a degree of confidence in mission success.

Mission Resilience Engineering

"[Mission assurance] is an engineering process performed over the lifecycle of a program to identify and mitigate design, production, test and field support deficiencies that could affect mission success." [11]

The third task to achieving CMR is based on a lesson learned from information assurance and is summarized in the saying, "cyber security should be built-in not bolted on." Assessing mission impact, and collecting resilience metrics, cannot be accomplished unless resilience attributes are part of the system design. Mission assurance expands the scope of the system development lifecycle by making the mission objectives the driving requirements in the development process (as opposed to security controls). It combines all the necessary components of mission execution and unites them by establishing the mission as the foremost goal in system design, development, and implementation. Security is a secondary objective that is applied to improve the resilience of mission-critical systems and functions. Leveraging the mission objectives to drive system requirements actually serves to reduce overall system complexity by focusing on designing only mission essential components to be resilient [6]. Mission resilience engineering is the overarching discipline that facilitates CMR because it applies an end-to-end lifecycle approach to mission definition, requirements assessment, and metrics.

4. Conclusion

Cyberspace is a mission-critical asset in modern military operations. But the cyber ecosystem has become more complicated due to the interdependent nature of information and systems. And the threat of cyber-related faults, failures, accidents, and attacks not only makes systems unreliable but can also affect the execution of the missions that depend on those systems. Current cyber security models are unable to keep up with the ever-changing threat and as a result, our military commanders lack confidence that the mission-critical systems will be operational when needed.

A new approach is needed to reestablish confidence in the ability to deliver operationally effective and resilient cyber capabilities. CMR seeks to achieve mission assurance through mission resilience by applying engineering discipline and metrics to make cyber-based systems and capabilities resilient to faults, failures, and attacks.

NOTES

1. Performability as used in this paper and by Qian et al. is meant to convey the "ability to ensure performance" as opposed to just performance itself.
2. For additional information on the BIA refer to NIST 800-34 or ISO 27000 toolkit.

Figure 3: Net-Centric Joint Capability Attributes (JCA)

Information Transport	Enterprise Services	Net Management	Information Assurance
Accessible	Accessible	Accessible	Security
Capacity	Interoperable	Dynamic	Available
Accurate	Survivable	Flexible	Timely
Timely	Timely	Agile	Accurate
Throughput	Reliable	Integrated	Visible
Expeditionary	Accurate	Maintainable	Responsive
Latency	Relevant	Complete	Controllable
	Scalable	Reconfigurable	Complete
	Responsive		
	Robust		

REFERENCES

1. Jabbour, K. "CyberVision and Cyber Force Development." Strategic Studies Quarterly, Vol 4, No 1 2010. May 2010.
2. Lindstrom, P. "Metrics: Practical Ways to Measure Security Success." 2005. techtarget.com. 2008.
3. DoDD 3020.40. "DoD Policy and Responsibilities for Critical Infrastructure." 14 JAN 2010. FAS.org. JUN 2010. <http://www.fas.org/irp/doddir/dod/d3020_40.pdf>.
4. Alberts, C. and A Dorofee. "Mission Assurance Analysis Protocol (MAAP): Assessing Risk in Complex Environments." Sept 2005. SEI.CMU.EDU. 25 Apr 2010.
5. Accenture. "Mission Resilience: The New Imperative for High Performance in Public Service." 2008.
6. Qian, Y., D. Tipper and P., Joshi, J. Krishnamurthy. Information Assurance: Dependability and Security in Networked Systems. Morgan Kaufmann Publishing, 2007.
7. National Security Agency. "Information Assurance Technical Framework." 2002. IAD.gov. 30 Apr 2010. <https://www.iad.gov/library/iacf.cfm>.
8. Payne, S. "A Guide to Security Metrics." JUN 2006. SANS Reading Room. JUN 2010. <http://www.sans.org/reading_room/whitepapers/auditing/guide-security-metrics_55>.
9. Jaquith, A. Security Metrics: Replacing Fear, Uncertainty, and Doubt. Boston, MA: Pearson Education, Inc., 2008.
10. DISA. "Joint/Coalition Mission Thread Measures Development Standard Operating Procedure." 2010.
11. Grimm, J. "The Role of CMMI in Mission Assurance." 16 Nov 2004. DTIC.mil. 25 Apr 2010.

ABOUT THE AUTHORS

Chris Peake has spent the last 15 years in the IT field studying and practicing cyber security. During that time, he has supported DoD, Federal, and commercial customers by providing network, system, and security expertise. Currently, he serves as a Cyber Assurance Strategist for Huntsville-based Sentar Inc. where he supports MDA, SMDC, DARPA, SPAWAR, and other DoD/Federal customers with Information System Security Engineering, Mission Assurance, and Cyber R&D.

Sentar Inc.
315 Wynn Dr.
Huntsville, AL 35805
E-mail: chris.peake@sentar.com
Phone: 256-430-0860

Al Underbrink has been a Senior Analyst with Sentar for nine years and has served as a PI and as a technical contributor on many projects involved with research and development of secure software systems. His technical areas of expertise include computer security, information assurance, artificial intelligence, robotics, automated planning, distributed systems, model-based diagnosis and reasoning, and agent-based frameworks and systems.

Sentar Inc.
315 Wynn Dr.
Huntsville, AL 35805
Phone: 256-430-0860
E-mail: al.underbrink@sentar.com

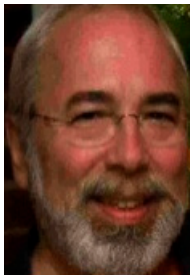

Dr. Andrew Potter is Director of Research and Development with Sentar. Current research interests include the application of a wide range of knowledge-based and quantitative techniques to the problems of cyber security and malware analysis.

Sentar Inc.
315 Wynn Dr.
Huntsville, AL 35805
Phone: 256-430-0860
E-mail: andrew.potter@sentar.com

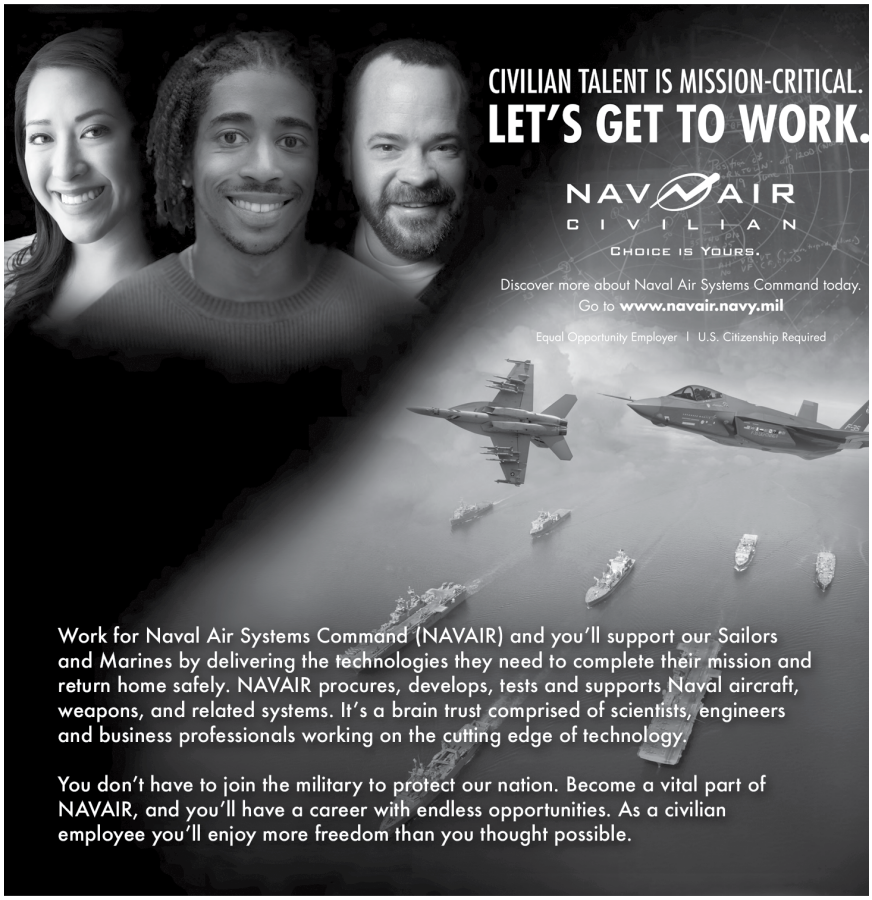

Forum Article

Crawl–Walk–Jog–Run: Evolving Measurement Capabilities

David P. Quinn, MOSAIC Technologies Group

Abstract

Many organizations try to jump into measurement with some wonderful measures that the organization cannot generate. Unsurprisingly, the measurement program fails because they were not ready for some very advance measurements. Organizations need to understand that you need to walk before you can run with measures. You might even need to crawl before you can walk when addressing measures.

I have to confess that I am a big fan of measures like defect density, earned value, and requirements volatility. When measures are introduced and management asks about them, behaviors start to change in the organization. To quote Martha Stewart, "This is a very good thing."

Unfortunately, most organizations starting down the process improvement road are not ready for these advanced measures. While some extraordinary efforts could get these measures into place, usually organizations need more base measures before they can proceed to the more advance, derived measures. In other words, they need to walk before they can run.

Actually, saying an organization must walk before it can run for measurement removes a couple of stages of measurement evolution. Measurement evolves as the organization gains access to data and gets comfortable with the measures. There are some measures that lead to other necessary measures that create a full set of derived measures. This is another way of saying that, for measurement, organizations must crawl before they walk, walk in order to jog, and jog in order to run.

A base measure is defined as an attribute and the method of quantifying it [ISO-15939]. Examples of this are number of functional requirements in a project, number of hours expended on a task, and number of defects in a unit of code. These base measures contribute to derived measures. A derived measure is a measure obtained from combining other measures. For instance, defect density is a derived measure that takes the number of defects and divides it by the number of lines of code or number of function points.

Let us use requirements as an example of measures evolving. For an organization that has no requirements measurement capability, it may have to start with a simple measure like how many projects have written requirements. An organization usually would show this as a percentage of all projects in the organization but that assumes they know how many projects they have. That is a different problem that should be explored some other time.

After getting the data on how many projects have written requirements (crawling), the organization should next look at determining how many requirements each project has. This equates to moving from crawling to walking. By knowing how many requirements the organization has, it is now able to quantify its workload. True, not all requirements are equal and some requirements require more work than other requirements; however, simply having a quantitative understanding of the organization's workload is a major step in learning how to walk.

To move from walking to jogging, the organization determines how requirements change each month. By monitoring whether more requirements come in each month than are closed and moni-

toring the number of requirements the customer modifies each month, the organization gets a sense of whether its workload is increasing or decreasing each month. It also provides an opportunity to justify preserving staffing levels or adjusting them up or down.

Finally, to start running, the organization uses the requirement volatility derived measure to understand the rate of requirements changes on a project. This derived measure has proven to be very effective in communicating to management and customers why projects are running late. It is just very difficult to get to this derived measure without having gone through cycles of determining information needs, measures that satisfy them, and indicators that address the information needs.

While in the various phases of measuring, the organization may need to do some additional analysis and qualification of the measure. For instance, an organization may start out crawling by finding out how many projects have written requirements. The organization may need to know a) how many projects had requirements written once but were never updated, b) how many projects had incomplete requirements, or c) how many projects' requirements were not approved by the customer. It is this type of analysis and qualification of data that leads to the evolution of measures.

This evolution of measures relies on managers being able to ask the right questions about measures. Unfortunately, managers may not respond appropriately to answers given by project teams to measurement questions. Managers should have a "question tree" that guides them on how to follow up responses to measurement questions based on likely responses they will hear. The "question tree" acts as a pseudo-script for the manager during a questions session on measures.

Using requirements measures as an example of a "question tree," a manager may start by asking how many requirements a project has. There are a set of likely responses that a project could give. One response is "We do not know." For this response, the tree would recommend a response of "When will you know?" This becomes an action item that the project and organization tracks to closure. Another possible response is "Do you mean customer requirements, system requirements, functional requirements, or what?" Again, an appropriate response would be part of the "question tree" and lead back to the original question. If the project team answers with the specific number of requirements, the manager moves on to the next prepared question, which could be "How many requirements are open and how many are closed?" This continues until the organization gets all the information of interest or until the project receives an action item to address.

Getting started in measurement can be scary, especially when introduced to some very advanced derived measures. No one should expect organizations to run before they walk in terms of measurement. Organizations need to understand that it is okay to crawl before walking. Measurements will evolve over time. And over time, organizations go from crawling to walking to jogging to running with measures.

About the Author:

David P. Quinn is the Director for Process Services at MOSAIC Technologies Group, Inc. He has more than 25 years software and systems development, maintenance, and management experience. He has more than 15 years experience as a process improvement consultant, helping large financial firms, defense contractors, telecommunications companies, aerospace contractors, and healthcare providers. He is certified by the Software Engineering Institute as a SCAMPI℠ Lead Appraiser and Instructor for CMMI® for Development and CMMI for Services.

MOSAIC Technologies Group, Inc.
8161 Maple Lawn Blvd, Suite 430
Fulton, MD 20759
Phone: 301-725-0925 x724
Fax: 301-725-0895
E-mail: dquinn@mosaicsgroup.com

Upcoming Events

ASIS/(ISC)2 Security Congress
10-13 September 2012
Philadelphia, PA
https://www.isc2.org/congress2012/default.aspx

AUTOTESTCON 2012
10-13 September 2012
Anaheim, CA
http://www.autotestcon.com/general/autotestcon-2012

TSP Symposium 2012
17-20 September 2012
St. Petersburg, FL
http://www.sei.cmu.edu/tspsymposium/2012/

Software Assurance Forum – Fall 2012
18-20 September 2012
McLean, VA
https://buildsecurityin.us-cert.gov/bsi/1364-BSI.html

ICW: International Certification Week
15-19 October 2012
Englewood, CO
http://www.iist.org/seminar/registration.php

PTI Technology Leadership Conference
21-23 October 2012
San Diego, CA
http://www.pti.org/index.php/ptiee1/more/776/

15th Annual Systems Engineering Conference
22-25 October 2012
San Diego, CA
http://www.ndia.org/meetings/3870/Pages/default.aspx

Events (continued)

OWASP AppSec USA 2012
22-26 October 2012
Austin, TX
https://www.owasp.org/index.php/Category:OWASP_AppSec_Conference

International Conference on Software Quality
29-31 October 2012
Indianapolis, IN
http://asq-icsq.org/index.html

12th Annual CMMI Technology Conference
5-8 November 2012
Denver, CO
http://www.ndia.org/meetings/3110/Pages/default.aspx

Software Assurance Working Group Sessions – Winter 2012
27-29 November 2012
McLean, VA
https://buildsecurityin.us-cert.gov/bsi/events/1406-BSI.html

Annual Computer Security Applications Conference
3-7 December 2012
Orlando, FL
http://www.acsac.org

Software Assurance Forum - March 2013
12-14 March 2013
McLean, VA
https://buildsecurityin.us-cert.gov/bsi/events/1417-BSI.html

Stressed Out Systems

Do you manage software? Perhaps you develop software? Then you already know what stress is. As a former developer, program manager, and software engineer, I understand the stress of writing, developing, and managing software systems. There is stress in gathering requirements, identifying users, designing the system, writing the code, and managing changing requirements. There is stress in testing the code, and—once delivered—there is stress in managing the inevitable change upon change upon change.

Stress refers to the pressure, pull, or force exerted upon an object. The resilience of an object, therefore, refers to its ability to recover from stress. Other sources refer to the resilience of an object as its ability to adjust to stress. There are interesting parallels in comparing software to the brain. In fact, the well-studied field of Psychological Resilience is a good starting point.

Paraphrased from Wikipedia:

Psychological Resilience refers to, "The idea of an individual's tendency to cope with stress and adversity. This coping may result in the individual 'bouncing back' to a previous state of normal functioning, or simply not showing negative effects. Another more controversial form of resilience is sometimes referred to as 'post-traumatic growth' or 'steeling effects' wherein the experience of adversity leads to better functioning (much like an inoculation gives one the capacity to cope well with future exposure to disease). Resilience is most commonly understood as a process, and not an individual trait."

Question 1: How do you make resilient software? Answer: you write resilient code by starting out writing non-resilient software, and learning how to keep it running. As part of my software engineering class, my students have to write a bulletproof program (typically, a simple one that prompts for names, hours worked, and hourly rate, and then prints out a simple payroll). I warn them that I will actively try and crash it. Even knowing that I plan on being malicious—I usually manage to crash about 50% of the programs. I run them in front of the class, and ask the class to join in and help me find and exploit flaws. Students initially are somewhat proud of their code, then watch in dismay as I find inputs that will crash their code: invalid inputs, extremely large numbers, zeroes for all inputs, strings for numbers, or very large strings. It is usually their first experience with actively evil input. They learn. They learn to bulletproof their code, to check all inputs, and to test for valid inputs all the time. They learn to trap and handle exceptions. And the viewpoint of writing really resilient good code is learned. You learn to write good resilient code by writing bad resilient code—and improving it over and over (…and over). And then you learn to write code that, when presented with inconsistent or invalid conditions, gracefully recovers, and returns to a consistent and usable state, without destroying data and without invalidating previous work.

Question 2: How can you maintain "normal functionality" in software? Answer: by taking economically reasonable steps to ensure that the user can perform normal operations under almost any type of system stress. In Question No. 1, it was the code that needed to be good. However, in this question, you see that your control over the environment needs to be good, too. Network down? You better have some local cached data to permit emergency functionality. Is the network really slow? Maybe have a good pre-fetch to reduce network latency. Worried about Denial of Service because of overloading or attacks? Use firewalls, redundancy, multiple servers, honeypots, etc. Do you have a single point of failure when contacting remote devices? Maybe you need to have multiple redundant routes to reach them. Mind you; you just can't throw hardware at the problems—you have to analyze the needs of the user, evaluate how the environment will be compromised, and take economically feasible preventative actions to minimize or prevent compromise. Assume your system is constantly under attack—and write not just good but defensive code. In my classes on Enterprise Security, students learn that paranoia is a good trait for network administrators. They are out to get you.

Question 3: How do you get a system to bounce back from failure? Answer: you need to have a process in proactively updating your system in response to constantly changing environments and conditions. Every day there is a new onslaught of viruses, hacks, threats, system vulnerabilities, etc. You cannot just write a program and expect it to be resilient for very long. It takes proactive planning and constant work. It is a continual process, not a single effort. One of the traits of a cyber system is a high degree of interaction between your computer hardware and other physical elements. These physical elements can be networks, remote hardware, and a large collection of physical devices. Cyber systems try to control all of this, and at the same time possibly interact with many other systems. Cyber systems sometimes need extremely high levels of reliability, precision, and coordination among the components—think air traffic control, unmanned vehicle operation, robotic surgery, and healthcare monitoring. Every piece you add gives yet another opportunity for the overall system to exhibit negative behavior (a nice euphemism for fail). There is no sane way to approach this as a single software-writing exercise performed as a solo exercise. You need a high-integrity process to create and update the software. Complex systems require complex processes—processes that are comprehensive, tested, and updated frequently. They need processes that are continually updated as new weaknesses or deficiencies are found.

I never said it was easy. In fact, developers agree—this is hard work. Creating reliable, resilient, robust, high-integrity cyber systems is probably one of the hardest development efforts in the field of software engineering. It is hard to do.

On the other hand, it is a lot easier than living with the potential consequences of not doing it.

David A. Cook, Ph.D.
Stephen F. Austin State University
cookda@sfasu.edu

Collaboratively advancing security automation to enable a resilient cyber ecosystem

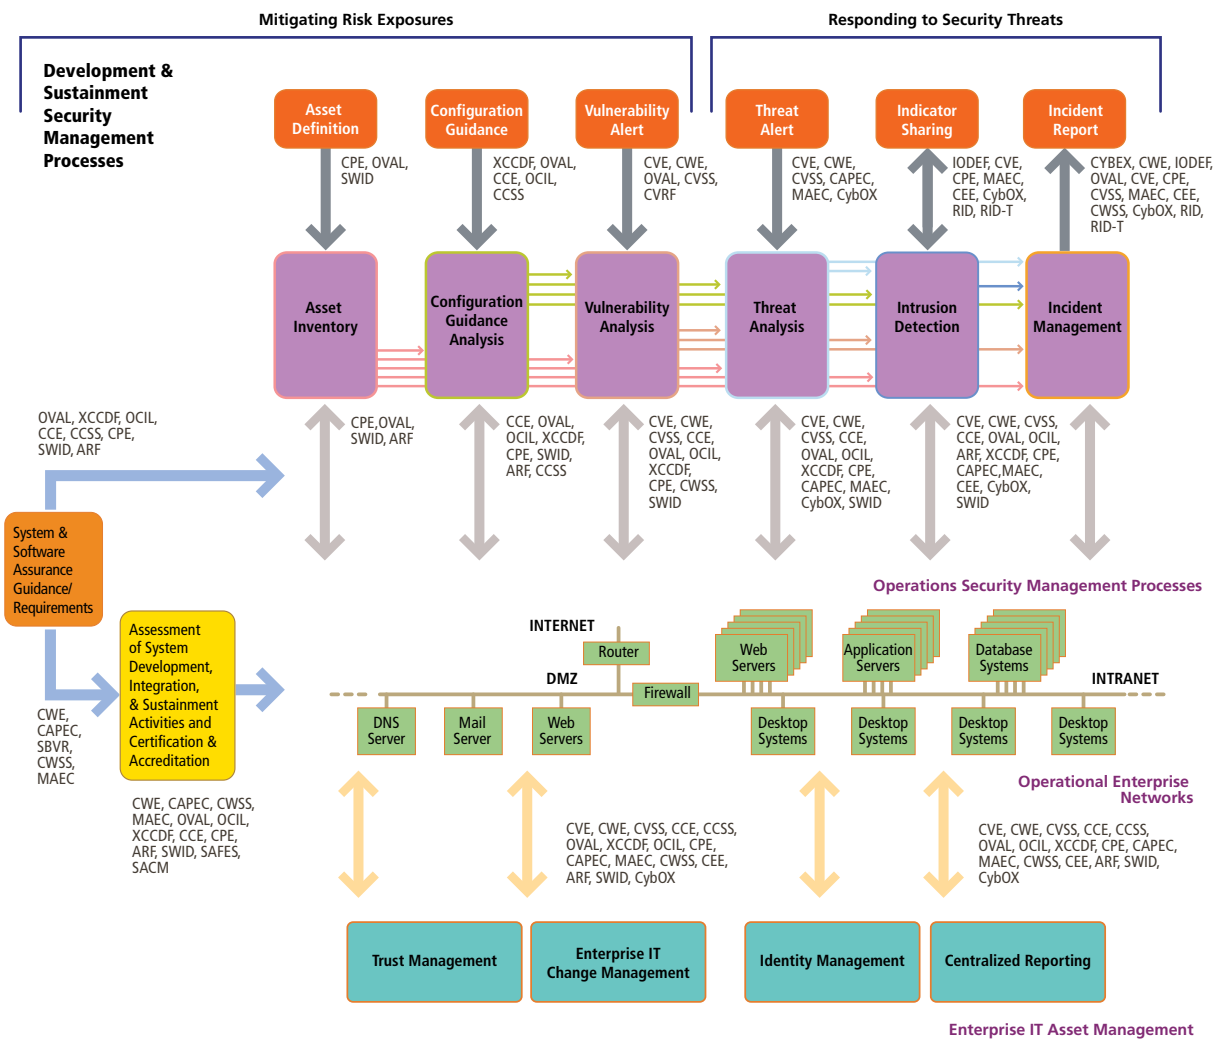

Homeland Security

Gain insights about security automation and measurement at the Software Assurance Community Resources and Information Clearinghouse at https://buildsecurityin.us-cert.gov/swa/measurable.html

CrossTalk thanks the above organizations for providing their support.