



Cyber Security Research and Experimental Development Program

Issued by the Communications Security Establishment Canada (CSEC)

12 August 2013

Contents

1	Executive Summary	3
2	Background.....	4
2.1	Cyber Security Strategy.....	4
2.2	Research Context of Cyber Security	4
2.3	Structure of this Document	5
3	Cyber Security Research Program Description	5
3.1	Guiding Principles	5
3.2	Programmatic Approach to Address the Problem Space	6
3.3	Cyber security Problem Space	6
4	General Approach.....	11
4.1	An ‘Agile’ Programmatic Response.....	11
4.2	Methodology.....	11
5	Ways of Working	12
5.1	Academic, Private, and Public Engagement.....	12
5.2	Collaboration and Knowledge Sharing.....	17
5.3	Experiments and Data	18
5.4	Programming and Computing Resources	19
6	Considerations	19
6.1	Legal and Ethical Considerations.....	19
6.2	Required Skill Sets.....	21
6.3	Assessing Progress in Science	22
6.4	Technology Transfer.....	24
7	Cyber Security Related Source Material	25
7.1	Summary of Source Material.....	25
7.2	Science of Security Joint Statement of Understanding (2011)	26
7.3	Roadmap for Cybersecurity Research (DHS 2009).....	28
7.4	Science of Cyber-Security (MITRE-JASON, 2010).....	30
7.5	Enabling Distributed Security in Cyberspace (DHS 2011)	32
7.6	Science of Security Hard Problems: A Labet Perspective.....	34
8	Technology Transfer Models.....	34
9	References.....	35

1 Executive Summary

The explosive growth, complexity, adoption and dynamism of cyber space that have enhanced social interaction and expanded our ability to productively utilize our environment have also introduced new adversarial threats and challenges to our society (e.g., [E57]). Cyber-bullying, cyber-crime, cyber-terrorism, adversarial state-sponsored activities, and so on, are all exemplars of malevolent attributes of cyber space. Mitigating these malevolent attributes requires an agile, legal and ethically compliant, interdisciplinary and scientifically-based research and exploratory development program in cyber security.

The cyber security research challenge over-all resides within a particularly complex area, being at the intersection of behavioral sciences, formal sciences and the natural sciences. As stated, cyber space includes a significant adversarial component which has led to a view that the science of cyber security is, in fact, a “Manichean science,” a science in the presence of adversaries, the core components being operations research, cybernetics and game theory. Consistent with this perspective are “nature inspired” approaches that draw upon analogies arising from immunological and biological systems. Other areas that could usefully inform a science of cyber security include cryptography, formal reasoning, machine learning and composition.

In response to Canada’s Cyber Security Strategy (CCSS), this paper offers a research and experimental development program description required to establish a secure, stable and resilient information technology infrastructure. Informed by national and international strategies, roadmaps and problem books (some of which are summarized herein), a research context for investigating the cyber security challenge is presented. In addition, a set of guiding principles are formulated to ensure the cyber security research program addresses the desired improvements, outcomes, and guidance stated in CCSS. Constrained by the context, and satisfying the principles, specific cyber security related capability gaps and operational limitations are described including associated research topics. Given the dynamic attributes of cyber space, this program description is, in fact, an evergreen document and will be periodically updated to reflect cyber space and cyber security evolution.

2 Background

2.1 Cyber Security Strategy

This document serves to establish a Cyber Security Research and Experimental Development Program [E9, E54] regarding the implementation of Canada's Cyber Security Strategy (CCSS) [E25] and associated funding to strengthen the security of federal cyber systems.

Our expectation of a research and experimental development program is to provide the following improvements or outcomes:

- ✓ *ensure that Canada can continue to be in the top tier of countries in terms of leading-edge cyber expertise (RD1); and*
- ✓ *anticipate the techniques and targets of known, new and emerging sophisticated threats and enabling the prevention of a greater number of threats from reaching Government systems (RD2).*

These statements have been tagged Research Driver One (RD1) and Research Driver Two (RD2) respectively and will be used subsequently as justification for the approach taken, and the description given, of the cyber security research and experimental development program.

2.2 Research Context of Cyber Security

This section references work that we believe sets a useful context for establishing an appropriate and relevant research program that addresses challenges that are (i) specific to cyber security or (ii) shared problems that are also particularly relevant to the cyber security domain. The referenced material tends to avoid the issue of defining the term 'cyber' or 'cyber security' and instead together provides a well-considered and useful description of the domain.

Recent work now attempts to establish a doctrine for the domain based upon the view cyber security should be considered as a public good [E1]. Using public health as an exemplar, the doctrine of 'public cyber security' is articulated. This is important contextually because new policy and new institutions are implied. Exploring the shift from public health to public cyber security, [E1] also provides illustrative examples that are useful for evaluating the nature of the cyber security domain as enlightened from this new viewpoint.

From a scientific perspective, the material is also well-founded with respect to emerging research focussed on the grand challenge of establishing a Science of (Cyber) Security [E5, E14-16]. [E19-20] have been cited as early papers that are of sufficient scientific merit to warrant consideration as the Science of cyber security is evolved. Through discussion of classes of attacks, policies, and defences, [E2] references the importance of building upon existing knowledge, particularly formal methods, fault-tolerance, and experimental computer science but [E2] also acknowledges the importance of cryptography, information theory, and game theory. Interestingly, based on safety ("no bad thing") and liveness ("some 'good thing' happens") each being connected to a proof method, [E2, E14] suggests hyper-properties [E4] are a promising candidate for use in a science of cyber security. In [E4] safety and liveness are generalized to hyper-properties which can describe security policies – "indeed we have not been able to find requirements on system behavior that cannot be specified as a hyper-property".

With a focus on technical measures for blocking cyber-attacks, [E3] adopts the human immune system as a metaphor to motivate the need for automated collective action amongst distributed systems to defend individual

computers and networks. [E3] identifies automation, interoperability, and authentication as the building blocks that underpin a five-level focus and convergence maturity model for networked environments. [E3] also describes the attributes and desired end state of a healthy cyber ecosystem (including participants within the ecosystem).

There is also clearly a strong connection between cyber security research and ongoing investigations concerning security analytics and measurements [E6, E21]. According to the founding editor in chief of IEEE Security and Privacy [E6], “Accordingly, we won’t find the appropriate science for understanding the evolving cyber-security landscape in the logic of formal systems or new software engineering techniques; it’s an emerging subarea of game theory that investigates dynamics in adversarial situations and the biases of competing human agents that drive those dynamics.” Based upon game theory, partially observable Markov decision processes and other techniques, [E7] describes a computational approach to the quantitative cyber-security risk assessment of intellectual property in complex systems – we believe this methodology could be augmented/generalized to also address critical infrastructure protection.

Finally, from the perspective of “Reducing Systemic Cybersecurity Risk”, [E13] suggests research responses should adopt a cross-disciplinary approach that combines “hard computer science” with the need to understand social science dimensions since “information system security are achieved only by a fusion of technology and the ways in which people and organisations actually try to deploy them”. Further, [E26] asserts “current approaches to cyber security are ill-suited to detecting or anticipating threats, which increasingly rely on hybrid socio-technical vectors”. [E26] suggests “By identifying and understanding the threat agents as threats themselves, instead of only the technology as threats, we can understand and neutralize other threats before they are created”.

2.3 Structure of this Document

This description of the Cyber security research and experimental development program has been structured to address RD1 and RD2 informed by the research context of cyber security just given.

The next section provides guiding principles, describes a programmatic approach to address the challenge overall and then defines specific cyber security related problems. Part 4 describes the general approach that should be taken when investigating the problems defined in Part 3. Part 5 reviews ways of working concerning academic, private and public engagement; collaboration and knowledge sharing; experiments and data; and programming and computing resources. Part 6 raises other important considerations such as legal and ethical considerations, required skill sets, assessing progress in science and technology transfer. Part 7 highlights important cyber security related roadmaps and problem books that informed Part 3. Part 8 discusses briefly some technology transfer models. Part 9 provides a list of references that compose the body of knowledge which underpins this program description.

3 Cyber Security Research Program Description

3.1 Guiding Principles

A set of guiding principles are formulated to ensure the cyber security program addresses the desired improvements, outcomes, and guidance stated in CCSS.

The following are guiding principles specific to the cyber security domain:

- Coordinate research activities to systematically progress towards achieving the attributes and desired end state of a healthy cyber ecosystem (including participants within the system) [E3].

- Engage social science research labs to understand social science dimensions of cyber security, augmenting “hard computer science” research [E1, E13];
- Focus research on promising scientific approaches which comprehensively and rigorously underpin required security policy [E2, E14];
- Focus research on promising scientific approaches which comprehensively and rigorously underpin the quantitative cyber security risk assessment of complex systems (especially critical infrastructure) [E6, E7];
- Focus research on promising scientific approaches to automate collective action amongst distributed systems to defend individual computers and networks [E3];
- Focus research that recognizes the presence of adversaries in cyberspace, with potential emphasis on the Manichean sciences [E14];
- Engage research labs to investigate cyber security related research gaps and to de-risk scientific approaches and emerging technological solutions [E2, E14-16]; and
- Leverage and influence cyber security related maturity models and standards when investigating hard problems [E3].

The following are guiding principles concerning shared problems and considerations that are also particularly relevant to the cyber security domain:

- Build upon existing knowledge that is relevant to cyber security research [E14-16];
- Leverage research that addresses big data challenges that also addresses cyber challenges [E29, E33];
- Leverage research that addresses the question: What do Data Scientists look like? [E30]; and
- Leverage existing knowledge regarding ways of working (Part 5) and carefully address the myriad of considerations (such as those pertaining to ethics) that influence and are influenced by cyber security (Part 6).

3.2 Programmatic Approach to Address the Problem Space

The following is the recommended programmatic approach to address the cyber security research challenge:

- Adopt an “agile” programmatic response (See Section 4.1);
- Be explicit about the particular methodological approach(es) that will be applied when pursuing a particular investigation but also use a general methodology to ensure a unified response to the cyber security challenge overall (See Section 4.2);
- Adopt appropriate ways of working concerning academic, private, and public engagement; collaboration and knowledge sharing; experiments and data; and programming and computing resources (See Part 5); and
- Understand constraints and dependencies concerning legal and ethical considerations, required skill sets and technology transfer (See Part 6).

3.3 Cyber security Problem Space

In the following table we identify nine capability gaps or operational limitations pertaining to cyber security. Briefly, we describe each of these gaps and limitations as follows:

- *Improve Signature Management and Signature Quality*: A signature is a distillation (usually a hash encoding) of a malicious pattern. Signatures are widely used, for example, to tersely identify cyber threats and, most widely, for the identification of viruses. The challenges identified here aim to improve the quality, effectiveness and timeliness of signature-based techniques;
- *Increase effort on anomaly detection and support discovery*: Anomaly detection refers to behaviour that does not conform to expected behaviour or usage patterns. From a cyber security perspective, for example, anomalous traffic patterns in a network could suggest that a system has been penetrated and sensitive data is

being exfiltrated. The challenges identified here target areas where anomaly detection and discovery could be materially improved;

- *Streaming and event driven analytics to reduce time to action*: Streaming analytics refers to the inline analysis of data (e.g., I.P. packets, stock trades, currency trading, health monitoring) so as to rapidly and intelligently respond to evolving situations, potentially in near real time. (There is a spectrum of algorithms, ranging from near real-time algorithms supporting almost instant response to adversarial situations; through to algorithms that take a longer-term, almost forensics-like, perspective. Identifying this algorithmic taxonomy is a research challenge in its own right.);
- *Dynamic defence at the network edge and beyond*: A network edge is the location where the processing and enforcement of organizational policies commences. This hard problem focuses on developing dynamic defence techniques that can rapidly interdict network attacks, using both network and host-based capabilities;
- *Cloud (Virtualization)*: Cloud computing is the delivery of computing resources over a network. Cloud computing brings challenges pertaining to scale, security and privacy. Challenges arise from the evaluation, architecture and design of such systems. Furthermore, there are specific concerns about contagion of malware infections across virtual instances and into the underlying base image. Virtualization is a key technology underpinning Cloud computing;
- *COTS*: Commercial off the shelf (COTS) products are those products that are commercially available, leased, licensed or sold and do not require specific maintenance/modification. COTS products tend to vary in quality, yet also evolve quicker and more usefully in response to broader market forces. The challenges once again pertain to evaluation, architecture and design of such systems as there is a need to scale evaluation capability and the potential to architect systems to mitigate threats arising from specific products. The supply chain is of particular concern with COTS products;
- *Enterprise-level Metrics [E16]*: Such metrics allow us to answer questions that are fundamental to investment and deployment decisions. They allow us to answer such questions as “how secure is my organization?” and “how has my security posture improved through the last set of updates?” To properly manage our systems, scientifically-based metrics and measures are required. Any underpinning “science of cyber security” will require a family of justified measures and metrics. Currently, there are no universally agreed upon methodologies to address the fundamental questions of how to quantify system security;
- *Mobility (including wireless)*: Mobile devices are tending toward ubiquity and there is a strong desire to use capabilities available at home within the work place. Mobility raises unique questions from a TRA perspective and adds potential attack vectors due to the use of wireless and other over-the-air communication mechanisms. Challenges pertaining to evaluation, architecture and design once again arise though within a different context; and
- *Science of Cyber Security*: Here, science is viewed as knowledge that results in correct predictions or reliable outcomes. Successful progress on this capability gap will provide significant science-based foundations for our cyber security techniques.

Table 1 Responding to Capability Gaps or Operational Limitations

CAPABILITY GAPS or OPERATIONAL LIMITATIONS	CHALLENGES	RESEARCH TOPICS
Improve Signature Management and Signature Quality	Metadata-driven Computer Network Operations (CNO) ¹ event prioritization	Machine learning techniques Better correlation of host and network generated events
	Prioritization and arbitration of generated CNO events	Visualization (data structuring; clustering) to support analyst workflow, data management and manipulation Visualization to build an 'analyst workbench' or framework to support anomaly detection at the application, host and network layers
	False positive reduction and suppression techniques	Improve fidelity or build a capability to express more 'stateful' event-driven network-based signatures identifying injection attacks at scale on high-speed links
	Automated signature generation based on data corpus	The ability to refine and/or generate signatures (network or host) with low false positive rates using predetermined data sources
Increase effort on anomaly detection and support discovery	Specification based intrusion detection techniques	Malicious activity detection based on application and network protocol analysis
	Data mining techniques to support anomaly-based detection hypotheses	Big Data - machine learning, statistical analysis, and so on (enhanced by visualization techniques); Large-scale predictive vulnerability analysis Visualization (data structuring; clustering) Visualization
	Mimicry attack detection (network/host protocols)	Network and host activity baselining and predictive analysis
	Host-based anomaly heuristics	Cognitive Radio & Game Theory
	Edge network versus insider threat anomaly detection	Internal protocol analysis to develop anomaly or behaviour-based malicious activity detection at the host and network layers
Streaming and event driven analytics to reduce time to action	Increased context and enrichment of CNO events at ingestion	Big Data - streaming analytics within the cyber defence context and aiming for NRT responses
	Near Real Time (NRT) machine-driven signature generation and data collection	Big Data - streaming analytics within the cyber defence context and aiming for NRT responses
	Machine-driven automated distributed signature generation	Big Data - streaming analytics within the cyber defence context and aiming for NRT responses

¹ A broad definition of Computer Network Operations (CNO) is those actions taken to leverage and optimize digital networks so as to improve human endeavor or enterprise [E49].

CAPABILITY GAPS or OPERATIONAL LIMITATIONS	CHALLENGES	RESEARCH TOPICS
	NRT correlation of events generated at both the host and network layers to support both anomaly detection and dynamic defence	Big Data - streaming analytics within the cyber defence context and aiming for NRT responses
Dynamic defence at the network edge and beyond	Automation	Multi-modal sensor (at the host and network layers) operations – approaches to augment passive observation coupled with in-line interdiction
	Next generation Dynamic Defence at both the host and network layers	Investigate human immune system and biological systems metaphor (nature inspired); Game theoretic strategic planning to predict outcomes of dynamic defence actions
Cloud (Virtualization) ²	Evaluation	Scalable, incremental, composable analysis tools and techniques enabling assurance and traceability amongst system artifacts throughout the evaluation lifecycle; Cost-effective high assurance evaluations; Models and techniques for “on-the-fly” evidence creation; Crypt (COTS); Enterprise Security Architecture; Cross Domain Solutions
	Architecture/Design	Improved operating systems and networking; Provide exemplars of scalable trustworthy cloud systems; Develop building blocks for composing trustworthy cloud systems; Enterprise Security Architecture; Isolation of legacy systems through virtualization
	Security Requirements & Specifications	Containment of infections; TRAs - data centres
COTS ²	Evaluation	Scalable, incremental, composable testing and analysis tools and techniques enabling assurance and traceability amongst system artifacts throughout the evaluation lifecycle; Cost-effective high assurance evaluations; Crypt (COTS) Cross Domain Solutions
	Architecture/Design	Provide exemplars of scalable trustworthy CFC systems; Develop building blocks for composing trustworthy CFC systems; Integration of COTS and open source components
	Security Requirements & Specifications	Creating techniques to mitigate supply chain threats; TRAs - networks, email Characterization of malware techniques
Enterprise Level Metrics ²	Metrics for Information Security	Science of Security (Measurable security)
	Metrics for Trustworthiness	Science of Security (Measurable security)

² Some of the challenges and research topics are drawn from source materials such as [E16].

CAPABILITY GAPS or OPERATIONAL LIMITATIONS	CHALLENGES	RESEARCH TOPICS
Mobility (including wireless)²	Evaluation	Scalable, incremental, composable analysis tools and techniques enabling assurance and traceability amongst system artifacts throughout the evaluation lifecycle; Cost-effective high assurance evaluations; Crypt (COTS)
	Architecture/Design	Improved operating systems and networking; Provide exemplars of scalable trustworthy mobile systems; Develop building blocks for composing trustworthy mobile systems
	Security Requirements & Specifications	TRAs - networks, mobility
Science of Cyber Security³	Common language	Hyper-properties as semantic foundation; Develop new visualizations for risk assessments; Determine whether a language can be developed for expressing core principles (such as trust reallocation); Develop a modelling language for the expression of security aspects of an enterprise architecture
	Core principles	Hyper-properties; Mathematical sound techniques for composition; Role of biological metaphors in cyber security (nature inspired); Fault tolerance, resilience and other safety-critical techniques; Game theory and partially observable Markov decision processes;
	Attack analysis	Partially observable Markov decision processes
	Measurable security ⁴	Lifting low level metrics to a quantitative assessment of the enterprise; Create methods to perform sensitivity analyses to uncertain input values; Create methods to validate metric prediction; Create overall security argument relating business and technical security metrics; Game theory and partially observable Markov decision processes; Benchmarking
	Risk	Game theory and partially observable Markov decision processes; Investigate cross-disciplinary approaches (hard/soft sciences)
	Agility	Process by which tools are inserted into a data-driven quantitative analysis with instant feedback; Models and techniques for “on-the-fly” evidence creation; Composability
	Human factors	Psychology and human factors

³ The research challenges for the Science of Cyber Security are described in the Joint Statement of Understanding (Section 7.2).

⁴ With the exception of game theory and benchmarking, the research topics are drawn from recommendations arising from the Science of Security Labels.

4 General Approach

4.1 An 'Agile' Programmatic Response

Long experience (principally through trial and error) researching, demonstrating, and then deploying complex systems has suggested an iterative and incremental approach is the most effective and efficient method to establishing an agile mission that is responsive to the evolving operational environment [E8]. We advocate a similar approach when formulating, acting upon, and then revising the cyber security research and experimental development program description.

The description given is intended to represent the latest thinking about what the cyber security research program should be at the time it was written. At all times, we strive to be comprehensive, consistent, and concise when articulating a programmatic response to this complex research challenge. We strive to only reference those specific aspects which the authors and their advisors believe to be the most salient concerns relevant to investigating cyber security in a systematic way.

That being said, it is fully expected that the programmatic description will evolve as researchers and experimenters make progress and we better understand the theory, methods, and techniques required to meet the challenge. This evolution will manifest itself as a regular (via annual review) or event-driven (in the case of a game-changing breakthrough) feedback cycle linking ongoing research activities and the latest (revised) programmatic description intended to guide subsequent cyber security related research and experimental development.

4.2 Methodology

4.2.1 Use of Specific Methods to address Specific Problems

Based upon our academic work and our ongoing investigations in the workplace, we understand that it is important to not only have a clear, and well-scoped, understanding of the specific problem under investigation but also to be explicit about the particular methodological approach(es) that will be applied when pursuing a particular investigation. The methods applied should be as 'strong as possible' in the sense some methods may be more applicable than other methods, depending upon the problem.

We advocate always specifying the particular methodological approach(es) that will be adopted coupled with the description of the specific problem of concern. As an investigation proceeds, the methodology should be evaluated along with reporting any research results with respect to the problem itself.

For example, in Table 1, we included research topics pertaining to scalability. To evaluate purported solutions to addressing these topics, scalable laboratory facilities for carrying out rigorous experiments would be beneficial. Two such examples are the DeterLab [E27] (hosted by ISI, Marina del Rey, California) and the botnet laboratory [E28] at École Polytechnique de Montréal.

4.2.2 Use of a General Methodology for a Unified Response to the Challenge

We also recognize the need for applying a general methodology to facilitate a unified response to the cyber security challenge overall. In our view, this methodology needs to be 'strong enough'. As stated, specific methods may vary widely depending upon the problem under investigation. A unifying method must balance rigor with flexibility. The general method must be rigorous enough to provide traceability to top-down and bottom-up objectives, including

the quantification of research and experimental development performance metrics. It must also be flexible enough to accommodate the potentially highly divergent approaches that could be trialed on a problem by problem basis.

We therefore view the general method to be a systematic approach that is strong enough in its formulation to provide explicit traceability, and quantified research performance metrics, amongst cyber security related goals and objectives and specific cyber security related capability gaps or operational limitations, on the one side, and related research initiatives (planned, in progress, or completed), on the other side. Because this will mean linking research and experimental development activities with new and emerging sophisticated threats, this would be a programmatic mechanism to report progress with respect to RD2.

The general methodology should be well-informed by the definitions and methodologies pertaining to research and experimental development as espoused by the Organization for Economic Co-operation and Development's (OECD) Frascati Manual [E9] for measuring scientific and technological activities.

4.2.3 Assessing Progress in Science

Assessing progress and vitality of a science is a key consideration for those involved in scientific policy making, management or investment. However, as noted in [E53], no theory exists that can reliably predict which research activities are most likely to lead to scientific advances or to societal benefits. Nevertheless, [E53] describes various dimensions for assessing both scientific progress internally (measured by intellectual criteria) and externally (measured by contributions to society). In Section 6.3, we further elaborate on the following:

- Scientific progress internally defined (and associated indicators);
- Contributions of science to society (and associated indicators);
- Interdisciplinarity and scientific progress (as cyber security requires interdisciplinary approaches); and
- Implications for decision makers.

5 Ways of Working

The Government of Canada (GoC) must respond to complex interdisciplinary cyber security research challenges, but with limited internal capacity. We must leverage external research capacity to address these challenges. Hence, through partnerships we leverage the breadth and depth of external expertise and leverage millions of dollars of investment into the billions of dollars being spent externally. Partnering is often a pre-requisite for success. Many large organizations (such as, for example, Apple and Microsoft) have partners numbered in the thousands and provide capabilities that would otherwise be impossible to develop and support internally.

5.1 Academic, Private, and Public Engagement

There are many models for successful partnering of which we describe a few here. The varying forms of partnership respond to varying objectives. For example, some partnerships may purely focus on progressing solutions towards a hard problem, whereas other partnerships may focus on actually productizing and marketing solutions. While by no means exhaustive, and somewhat overlapping, the following forms of partnership are discussed:

- Centres of excellence;
- Meeting grounds;
- Virtual organizations;
- Consortia;
- Strategic networks;
- Incubators;

- For profit; and
- Not for profit.

5.1.1 Centres of Excellence

A Centre of Excellence (CoE) refers to a team, a shared facility or an entity that provides leadership, evangelization, best practices, research, support and/or training for a focus area [E45].

Within the cyber security realm, a number of recent examples of CoEs have appeared. For example, in the U.K. there was a recent announcement [E22] regarding the awarding of academic center status of excellence in cyber security research. The announcement noted that: "Eight UK universities conducting world-class research in the field of cyber security have been awarded 'academic center of excellence in cyber security research' status by UK Government Communications Headquarters (GCHQ) in partnership with the Research Councils' Global Uncertainties Programme and the Department for Business Innovation and Skills (BIS)." These universities included Oxford, Imperial College and University College London.

The announcement went on to note that it was expected that these centres would increase the resiliency of U.K. systems to cyber-attack through:

- Enhancing the UK's cyber knowledge base through original research;
- Providing top quality graduates in the field of cyber security;
- Supporting GCHQ's cyber defence mission; and
- Driving up the level of innovation.

The U.K. is also planning to develop centres of excellence in Cybersecurity education, certification of Cybersecurity training courses and increased sponsorship of PhD research [E56]

In the U.S., the National Security Agency and the Department of Homeland Security jointly sponsor National Centers of Academic Excellence in Information Assurance. The goal of these centers is to reduce the vulnerabilities in the U.S. national information infrastructure by promoting higher education and research in Information Assurance and to produce professionals with relevant expertise. These CoEs are distributed throughout the U.S.

5.1.2 Meeting Grounds

We use the "meeting ground" metaphor to capture the idea of facilities that facilitate the gathering of individuals and capabilities to focus on particular research objectives. We differentiate "meeting grounds" from "incubators" (described below) in that incubators include commercial/investment perspectives.

An example of meeting ground collaboration is the recent announcement [E24] from Sandia National Laboratories on its newly formed cyber research facility at its California site. The facility was opened in June 2012 and offers an "open yet controlled area for Cybersecurity professionals from the Bay Area and across the country to meet and discuss critical cyber research issues." The announcement goes on to state how the new cyber security Technologies Research Laboratory (CTRL) will promote stronger relationships between industry, academia and national laboratories in the research and experimental development of cyber security solutions through technology, practices and policy, with specific objects to:

- Develop the science and computing foundation necessary for robust cyber security research and development;
- Develop critical relationships to help understand the full range of technical threat concerns facing industry, government (non-classified) and academia;

- Develop, test and help implement cyber security approaches in real-world situations;
- Promote the various technical domains that support the advancement of cyber security, essential to the security and stability of the U.S. and the world;
- Develop political and social awareness of the real, imminent threat and the consequences posed by cyber exploits and attacks; and
- Provide a window to the external world on open cyber security and related work throughout Sandia, along with acting as a Bay Area resource for open work performed at Sandia's New Mexico location.

In addition, it was noted that “with CTRL, we can run experiments and talk more freely about a wide range of cyber research activities, and we can do so with a variety of U.S. and international collaborators but without some of the unrelated restrictions that are often associated with a national laboratory.”

5.1.3 Virtual Organizations

For our purposes, we define a Virtual Organization as a “flexible network of independent entities linked by information technology to share skills, knowledge and access to others’ expertise in nontraditional ways” [E43].

The Cyber Physical Systems Virtual Organization (CPS-VO) (<http://cps-vo.org/>) fosters collaboration among Cyber Physical Systems (CPS) professionals in academia, government and industry. CPS-VO fosters such collaboration through its website by supporting the publication of events and announcements, the sharing of documents, supporting collaboration, the formation and management of communities of interest (such as the Science of Security), and the organization of online meetings. CPS-VO highlights important events, provides pointers to research projects and even references research solicitations (such as the National Science Foundation’s 2013 Cyber Physical Systems solicitation). So, for example, a community of interest is the Science of Security community. Within this community there is a chat capability; membership information; recent news; pointers to various labels; current research; and funding opportunities.

5.1.4 Consortia

By definition, a consortium is a group of two or more entities (e.g. individuals, companies, universities, governments) that work together towards achieving a chosen objective. Each member of the consortium retains its independence, but is obligated to the consortium through legal or other obligations. In cyber security, the Team for Research in Ubiquitous Security Technology (TRUST) [E5] is an example of a consortium.

TRUST brings together the University of California at Berkeley, Carnegie Mellon, Cornell, San Jose State, Stanford, Vanderbilt and the U.S. National Science Foundation. Through TRUST, the various stakeholders have joined forces to respond to challenges arising from cyber security and the increasing reliance of society on critical infrastructures such as financial services and economic commerce, power grids, water systems, and so on. As noted in the 2011/2012 Annual Report [E5], security is inadequate; software usability, reliability and correctness are challenging; recruitment of adequately trained employees difficult; and society is relying upon systems that do not meet trust requirements. Hence, TRUST aims to enable dialogue with stakeholders whose needs require multifaceted responses – not only in a purely technical perspective, but also through considering policy and societal issues.

Hence, TRUST defines its mission thusly:

The Team for Research in Ubiquitous Secure Technology (TRUST) is focused on the development of cyber security science and technology that will radically transform the ability of organizations to design, build, and operate

trustworthy information systems for the nation's critical infrastructure. Established as a National Science Foundation Science and Technology Center, TRUST is addressing technical, operational, legal, policy and economic issues affecting security, privacy, and data protection as well as the challenges of developing, deploying and using trustworthy systems.

TRUST activities are advancing a leading-edge research agenda to improve the state-of-the-art in cyber security; developing a robust education plan to teach the next generation of computer scientists, engineers, and social scientists; and pursuing knowledge transfer opportunities to transition TRUST results to end users within industry and the government.

Our second example is an announcement from the U.K. [E23] of the first academic research institute to investigate the "Science of Cyber Security." The institute is funded by a £3.8 million grant and is a part of a U.K. cross-government commitment towards increasing the U.K.'s national academic capability in cyber security. Though identified as an institute, it actually consists of a number of partner Universities, with the institute hosted by University College London. The partner universities are:

1. University College London working with the University of Aberdeen;
2. Imperial College, working with Queen Mary College and Royal Holloway, University of London;
3. Royal Holloway, University of London; and
4. Newcastle University, working with Northumbria University.

The announcement went on to note that addressing these practical challenges requires a blended approach from researchers, drawing from both technological and behavioral disciplines.

The U.K. has set up a second academic research institute [E56], which will focus on new ways of analyzing software automatically to combat cyber threats. The selected projects focus on key issues pertaining to vulnerability discovery, malware analysis and classification of code and improved defenses and mitigations. This institute is funded by a £4.5 million grant and consists of teams from six universities:

1. Queen Mary, University of London, working with University of Kent and University College London;
2. University of Edinburgh;
3. Imperial College London;
4. University College London;
5. University of Kent working with University College London; and
6. The University of Manchester.

Another example, purely governmental in scope, is the U.S. NITRD (the Networking and Information Technology Research and Development) program that provides a framework in which numerous U.S. federal agencies collaborate in their networking and information technology R&ED efforts [E46]. This program outlines three key objectives from the departmental collaboration:

- Provide R&ED foundations for assuring continued U.S. technological leadership in advanced networking, computing systems, software, and associated information technologies;
- Provide R&ED foundations for meeting the needs of the Federal government for advanced networking, computing systems, software, and associated information technologies; and
- Accelerate development and deployment of these technologies in order to maintain world leadership in science and engineering; enhance national defense and national and homeland security; improve U.S. productivity and competitiveness and promote long-term economic growth; improve the health of the U.S. citizenry; protect the environment; improve education, training, and lifelong learning; and improve the quality of life.

Member agencies include DOE, NASA, NIST, NIH, NSF, NSA, DARPA and the Office of the Secretary of Defense (OSD).

Finally, to provide a commercial example, is the October 2012 announcement [E52] of the non-profit Cybersecurity Research Alliance (CSRA) by AMD, Honeywell, Intel, Lockheed Martin and RSA/EMC to form a research consortium to focus on “grand challenges” for cyber security and next generation technologies. The key value propositions are to bridge the gap between government-funded R&ED and commercially available cyber security solutions and to facilitate solutions addressing grand challenges that are bigger than any one company, consortium, sector or nation.

5.1.5 Strategic Networks

For the purposes of this report, a strategic network is a means for increasing research and training in specific areas that could strongly enhance Canada’s economy, society and/or environment within the next 10 years [E44]. The strategic networks are funded for five years and are managed by a university and lead professor along with an advisory board. Typical grants range from \$500,000 to \$1 Million per year. Required partners include Canadian-based companies or government departments/agencies that could potentially apply the results.

Of pertinence to cyber security are two such strategic networks:

Surfnet (<http://www.nsercsurfnet.org>): The goal of Surfnet is to improve the development, performance, and usability of software applications for surface computing environments: nontraditional digital display surfaces including multi-touch screens, tabletops, and wall-sized displays. It is expected that surfaces will naturally support group work and collaboration. Of particular note, it is observed that digital surfaces provide space for working with large data sets. Canadian universities participating in the network are the University of Calgary, University of Waterloo, Carleton University, University of Saskatchewan, Queen’s University, McGill, the University of British Columbia, the Ontario Institute of Technology, University of Manitoba and the University of Ottawa. Sponsoring partners are Smart Technologies and TRILabs.

ISSNet (<https://www.issnet.ca/about-issnet>): The focus of this strategic network is computer and network security emphasizing computer systems research with an experimental or observational approach. The network brings together researchers in computer and network security and related areas. ISSNet’s objectives include addressing the critical shortage of skilled workers in security, the growing of a pool of talented and knowledgeable personnel with relevant expertise, to develop security technologies, and to direct research projects directly related to the understanding and defending against security and stability threats to the Canadian internet infrastructure and related critical infrastructures. Universities involved with ISSNet are the University of British Columbia, University of Calgary, Carleton University, École Polytechnique de Montreal, University of Toronto, Royal Military College and Dalhousie University. Sponsoring partners include RIM, TREND Micro and CA Labs.

5.1.6 Incubators

For our purposes, we focus on “business incubators,” which are “programs designed to support the successful development of entrepreneurial companies through an array of business support resources and services, developed and orchestrated by incubator management and offered both in the incubator and through its network of contacts [E47].

A Canadian example, writ large, of incubation, is Toronto-based MaRS [E48]. MaRS fosters and promotes Canadian innovation through the provision of resources such as people, programs, physical facilities, funding and networks.

Through the provision of these resources MaRS ensures that critical innovation happens. From an R&ED perspective, one critical partnering component of MaRS is called “MaRS Innovation,” a membership-based partnership designed to transform Toronto-based academic research enterprise into a successful commercialization cluster. MaRS Innovation’s mission is to monetize the research assets found within its member institutions, which include Ryerson University, the University of Toronto, York University, St. Michael’s Hospital and the University Health Network.

The Canadian government, through its proposed 2013 budget, recognizes the importance of incubators and “proposes to provide \$60 million over five years to help outstanding and high-potential incubator and accelerator organizations in Canada expand their services to entrepreneurs, and to make available a further \$100 million through the Business Development Bank of Canada to invest in firms graduating from business accelerators.”

5.1.7 For Profit

Of course, there are firms that provide services/products, which is a form of partnering. For example, a recent report [E55] identified and studied the competitive landscape for cyber security companies. These companies included BAE Systems, Boeing, Booz Allen Hamilton, Computer Science Corporation, EADS Group, Finmeccanica Group, General Dynamics, Hewlett Packard, IBM, Intel, Kaspersky, L-3 Communications Holdings Inc., Lockheed Martin, Northrop Grumman, Raytheon, SAIC, Sophos PLC, Symantec, Thales Group and Trend Micro. Some firms are focused on specific cyber security issues. For example, CrowdStrike (<http://www.crowdstrike.com>) takes the perspective that organizations do not have a malware problem; they have an adversary problem and fuse together information from technology and intelligence sources to provide services. CGI recently announced [E58] a Security Centre of Excellence, which consolidated its security expertise and offerings in response to increasing risk security threats pose to Canada.

5.1.8 Not for Profit

In Section 2.2 a perspective was presented in which Cyber security is perceived as a “public good.” Though government institutions are often used to respond to such “public good” challenges (such as defence and health care), “not for profit” (non-governmental) institutions can also play a role in particular areas. A “not for profit” organization is an incorporated entity, which exists, primarily, to meet charitable objectives, and retains income for meeting its expenses and other obligations. There are, for example, no dividend payments to shareholders. Given the societal benefits that often accrue from such organizations; they are often exempted from various tax regimes. An example charitable organization is the MaRS organization described above and in [E48].

5.2 Collaboration and Knowledge Sharing

To start, there are general techniques for knowledge sharing that are as true in disseminating/developing cyber security knowledge as in other domains. These techniques include [E40] communities of practice, after-action reviews and retrospections, action learning sets, challenge sessions, mind maps, online strategies, stakeholder analysis and storytelling.

Additionally in institutions concerned with research and experimental development, knowledge sharing is supported by workshops, conferences, journals, books, curriculum development, courses, seminars, lectures, blogs, wikis, and so on.

Testbeds (see also below) can be used to demonstrate robustness and scalability of technologies and to provide a means for experimenting with and integrating technologies provided by different partners (whether commercial,

academic or government). Testbeds may also provide a means for developing interoperability and also provide a technology transfer role (especially with regards to integration with legacy products) [E42].

Partner relationships are important to knowledge transfer. For example, the Team for Research in Ubiquitous Secure Technology (TRUST) manages relationships between technology developers (innovators), venture capital, industry, government and other stakeholders. TRUST identifies four approaches to facilitating such partnerships [E42]:

- Organizing focused workshops emphasizing technology transfer opportunities;
- Facilitating strategic investment sessions;
- Providing internships for students and faculty; and
- Supporting entrepreneurship by integrating TRUST researchers as entrepreneur-incubators at venture capital partners.

Particularly useful are workshops that bring together investigators from different specialties, different parts of an organization or even different organizations, who are pulled together into teams and then focusing them upon a set of specific problems. The creation of the cross-disciplinary teams enables cross-fertilization of ideas both within the teams themselves and as researchers return to their home bases and provide knowledge transfer. Defcon's [E41] "capture the flag" challenge is perhaps a specific instantiation of this kind of workshop.

Tiger teaming is another means for bringing together expertise to progress research and experimental development initiatives. Such a team is provided with a specific challenge, a relatively short timeline, and then uses their enthusiasm, experience, imagination and technical wherewithal to make progress. The team is often multidisciplinary in its nature.

5.3 Experiments and Data

To perform suitable experiments we need both data and laboratory facilities that representatively scale to the complexity and vastness of cyber-space. Finding representative data at scale and building laboratories that support scalable experiments are ongoing challenges. However, progress is forthcoming. Above, we referenced two such potential facilities: DeterLab [E27] (hosted by ISI, Marina del Rey, California) and the botnet laboratory [E28] at École Polytechnique de Montréal. Regarding representative data, examples include the VAST Challenge 2012 [E31], the Enron email dataset [E32] and the PREDICT (Protected Repository for the Defense of Infrastructure Against Cyber Threats) dataset [E33]. There is also research into the generation of representative synthetic data [E34] since oftentimes data is not available due to proprietary or classification reasons.

Access to data often is impeded by legal, policy and security concerns. Research into fine-grained security that is agile to legal, policy and security concerns has the potential of opening up data to wider dissemination since datasets then could be made available with researchers only accessing that information for which they have been authorized.

Further, larger-scale data analysis for cyber security can be enabled by augmenting general purpose query languages, such as SQL, through utilizing languages that directly support the analysis of the domain in question. For example, if an aspect of the cyber security challenge can be represented as a graph, then a graph-based language could be used to investigate those characteristics of the problem. In more general terms, domain specific languages that support hypothesis formulation and testing of advanced analytics for cyber security (including uncertainty reasoning) could advance analytical capability within the cyber domain.

Obtaining realistic data-at-scale is a hard problem in its own right. For example, making available extensive data gathered from operational activities can lead to privacy concerns; while generating extensive data using algorithmic techniques can lead to criticisms of the fidelity of the data respective to actual operational behaviour. Much data-at-scale is available within closed/secured environments, while much data mining (for example) algorithmic research is performed by academics and industry.

5.4 Programming and Computing Resources

Phrases such as “cyber time,” “cyber space” and “the Internet” conjure up superlatives: rapidly changing; complex technical, social and economic interactions and relationships; and massive amounts of data (Exabytes). These superlatives suggest that to effectively perform research and experimental development will require significant programming and computing resources. Some examples of required programming capability and computing resources include:

- Shared computing resources that are state-of-the art computation systems: from high-end individual workstations through to high performance computing platforms;
- New architectures that respond to new paradigms of programming and analysis, such as those supporting streaming (data) analytics;
- Near-real time programming and computational paradigms that are responsive to the rapidly evolving threat space;
- Advances in embedded, distributed and parallel computing, including the ability to distribute analytics;
- Novel uses of Graphical Processing Units (GPUs) and Special Purpose Devices (SPDs);
- The ability to scale up and to scale down capacity in support of experimentation; and
- Highly capable systems will need high performance I/O and extensive and high performance memory.

6 Considerations

In this section we briefly discuss legal and ethical considerations, required skill sets, assessing progress in science and technology transfer, which are considered to be other important considerations when investigating the cyber security challenge.

6.1 Legal and Ethical Considerations

The “Dark Space Project” (DSP) report [E26] provides interesting guidance on legal and ethical issues. In this project, the methodology used by the researchers was informed by guidance and precedence provided by the Privacy Commissioner of Canada and the Auditor General of Canada for conducting cyber security research from primary sources. Specific mention is made that all research activities were in compliance with applicable Canadian laws and regulations and, additionally, policies and procedures defined by the Bell Canada Code of Business Conduct, Bell Canada Corporate Policies and Ethics, the Bell Competition Law Compliance Handbook, the Bell Code of Fair Information Practices and the Privacy Statement and Ethical Principles regarding cyber security research at the Citizen Lab, Munk School of Global Affairs, University of Toronto. No Personal Identifiable Information (PII) was collected in this project and, in fact, it is noted that intelligence-led proactive defence that interdicts, disrupts, pre-empts and thus prevents emerging threat intent can be achieved without using PIIs.

The DSP commissioned a research paper on the legal and ethical dimensions of fusion techniques based on honeypot, sinkhole and Deep Packet Inspection (DPI) for threat detection. The paper discusses legal challenges to applying network detection/defence practices beyond an organization’s network purview. The key question: Can

these techniques be used by security researchers, academic institutions, corporate entities, or a network defence agency without breaking new ethical and legal ground?

The DSP argues that an effective strategy for dealing with cyber-attacks will require proactive defence that consists of interdiction and disruption of threat activities and notes that fusion techniques (analyses) are particularly useful. However, when one starts interdicting and disrupting, legal/ethical issues arise especially when such techniques are injected outside the bounds of a proprietary network.

DSP takes the approach that for the various techniques applied one should ask four questions:

1. What information is collected?
2. How is information collected?
3. How is the information used?
4. To whom is the information disclosed?

Reference to legislation is made on a sector by sector basis. So, for example, reference is made to the Personal Information Protection and Electronic Documents Act (PIPEDA) with regards to the Private Sector. Reference is also made to the Public Sector; Criminal Law; information sharing and duty to report; and the new anti-spam and anti-spyware law.

The report then applies the four questions to the potential legal challenges that specific fusion techniques could face. The report analyzes sink-holing, honeypots, packet sniffing, and deep packet inspection. The following table summarizes one of the analyses, packet sniffing, so as to give a sense of the approach.

Packet Sniffing	Definition: Looking at data packets as they travel across a network.
What information is collected?	<ul style="list-style-type: none"> • Two forms of data: header data and payload data. • Privacy legislation may apply if the information is not de-identified or aggregated as personal information may be present.
How is information collected?	<ul style="list-style-type: none"> • Packet sniffing is a passive activity; does not modify packets. • If PII present, consent is required unless exemption conditions apply. • In Private Sector, computer-to-computer communication may not be allowed in the absence of consent.
How is the information used?	<ul style="list-style-type: none"> • Varies depending upon purpose. • If PII present, consent is required unless exemption conditions apply.
To whom is the information disclosed?	<ul style="list-style-type: none"> • Varies depending upon purpose. • If PII is present, consent to disclose, for specific purposes, is required unless exemption conditions apply.

Finally, the DSP report recommends other legal areas requiring investigation including Intellectual Property, Torts, Contract Law and the Disclosure of Personal Information.

In September 2011, the U.S. Department of Homeland Security released a paper for comment pertaining to ethical principles guiding information and communication technology research [E50]. The paper is based on the 1979 Belmont Report [E51] for ethical research in the biomedical and behavioral sciences and is applicable to research that has the potential to harm humans. In particular, recognizing that information and communication technology research (ICTR) can harm individuals (e.g., inappropriate access to sensitive databases or malware controlling

compromised machines, embedded medical devices or critical infrastructure systems) the authors of the Menlo Report update the current human subject protection paradigm. The following table, drawn from the Menlo Report, identifies the four principles guiding ethical considerations. The first three principles are drawn from the Belmont Report; the fourth offered by the authors of the Menlo Report.

Principle	Application
Respect for Persons	<ul style="list-style-type: none"> • Participation as a research subject is voluntary, and follows from informed consent. • Treat individuals as autonomous agents and respect their right to determine their own best interests. • Respect individuals who are not targets of research yet are impacted. • Individuals with diminished autonomy, who are incapable of deciding for themselves, are entitled to protection.
Beneficence	<ul style="list-style-type: none"> • Do not harm. • Maximize probable benefits and minimize probable harms. • Systematically assess both risk of harm and benefit.
Justice	<ul style="list-style-type: none"> • Each person deserves equal consideration in how to be treated, and the benefits of research should be fairly distributed according to individual need, effort, societal contribution, and merit. • Selection of subjects should be fair, and burdens should be allocated equitably across impacted subjects.
<i>Respect for Law and Public Interest</i>	<ul style="list-style-type: none"> • <i>Engage in legal due diligence.</i> • <i>Be transparent in methods and results.</i> • <i>Be accountable for actions.</i>

The remainder of the Menlo report characterizes the stakeholders and discusses the application of the various principles. The report ends by noting that: “Proactively and transparently engaging in ethical assessment of ICT research will help move the research community mindset in the direction of embedding ethics into ICTR design as productively and safely as possible, and more practically influence policy and governance at these crossroads.”

6.2 Required Skill Sets

The U.S. National Initiative for Cybersecurity Education (NICE) has produced a national cyber security workforce framework [E39]. The purpose of the framework is to help ensure that there is an “enduring capability to prevent and defend against an ever-increasing threat.” The framework provides seven categories of skills and refines the categories so as to provide a comprehensive description of the skill sets required for cyber security operations.

Good research is engendered by an individual’s persistence, openness and curiosity. It requires a good understanding of scientific inquiry and, in particular, of the scientific method. Within cyber security, these generic scientific attributes are then instantiated with knowledge from underpinning disciplines such as computing science, mathematics, engineering, psychology, sociology, and so on. Particular topics will require deep knowledge of specific domains. For example, to perform research in network anomaly detection will require knowledge of numerous networking standards (including the ISO stack); tools for monitoring, provisioning and otherwise managing networks; analytical wherewithal to recognize normal and abnormal behaviours (perhaps using statistical or signature techniques); and data mining. Streaming analytics may play a role so as to respond in near-real time to

potential incidents. Openness to multidisciplinary thinking would be beneficial – there is much to be learned from sociology, psychology and economics.

As described elsewhere in this report, cyber security is viewed as a “Manichean science,” or a science in the presence of adversaries. Such a science requires knowledge of operations research, cybernetics and game theory. Further areas of importance include trust, cryptography, model checking, obfuscation, machine learning and composition.

As noted in Section 5, a number of efforts are focused at increasing the stable of cyber security experts and practitioners. Institutes such as the Tutte Institute for Mathematics and Computing (TIMC) (www.cse-cst.gc.ca/tutte/index-eng.html) provide opportunities for graduate students and post docs to work within a classified milieu and apply their knowledge within a well-informed cyber security sphere thereby enhancing both their capabilities and expanding TIMC’s capacity. Multi-year partnerships between cyber security government agencies with universities also provide a means for skills development in that with increased likelihood of stable funding, the universities are more likely to enhance their curriculum in cyber security and their overall research posture.

6.3 Assessing Progress in Science

The Office of Behavioral and Social Research of the U.S. National Institute for Aging requested a study on how to assess the progress and vitality of behavioral and social research on aging and how to contribute to the likelihood of discoveries in areas of aging research. [E53] was the result of this charter and was developed by a committee of the U.S. National Research Council. It is important to note that while the report came from the National Institute of Aging, the key points are true across scientific disciplines and thereby are of relevance to our cyber security research and experimental development program. The following discussion is drawn largely from [E53].

Science stakeholders need to use multiple dimensions when assessing progress in science and must recognize that scientific progress is nonlinear. The nonlinearity of science means, in part, that the absence of progress on a particular dimension is not necessarily indicative of poor or non-performance. The authors of [E53], based on prior work, distinguish between scientific progress internally defined (which is measured by intellectual criteria) and scientific progress externally defined (which is defined/measured in terms of contributions of science to society).

Internally defined dimensions of progress are:⁵

- **Discovery** – demonstrates the existence of previously unknown phenomena or relationships amongst phenomena, or when it discovers that widely shared understandings of phenomena are wrong or incomplete;
- **Analysis** – develops concepts, typologies, frameworks of understanding, methods, techniques, or data that make it possible to uncover phenomena or test explanations of them. (Improved theory, rigorous and replicable methods, measurement techniques and databases all contribute to analysis.);
- **Explanation** – discovers regularities in the ways phenomena change over time or finds evidence that supports, rules out, or leads to qualifications of possible explanations of these regularities;
- **Integration** – links theories or explanations across different domains or levels of organization (e.g., link understandings emerging from different fields of research or analysis); and

⁵ It is noted that these dimensions overlap and are interdependent.

- **Development** – stimulates additional research, including research critical of past conclusions, and when it stimulates research outside the original field, including interdisciplinary research and research on previously under-researched questions. Also when it attracts new people to work on an important research problem

The following were suggested as sources of early indicators of scientific progress, from the internal perspective, and reflect the vitality of a research field:⁶

- Established scientists begin to work in the new field;
- Students are increasingly attracted to a field, as indicated by enrollments in new courses and programs in the field;
- Highly promising junior scientists choose to pursue new concepts, methods or lines of inquiry;
- The rate of publication in the field increases;
- Citations to publications in the field increase both in number and range across other scientific fields;
- Publications in the new field appear in prominent journals;
- New journals or societies appear;
- Ideas from a field are adopted in other fields; and
- Researchers from different pre-existing fields collaborate to work on a common set of problems.

Major scientific advances are often marked by flurries of research activity and are suggestive of major progress in the scientific area.

Externally defined dimensions of progress (contributions of science to society) are:

- **Identifying issues:** identifying problems;
- **Finding solutions:** developing ways to address issues or solve problems;
- **Informing choices:** providing accurate or compelling information to decision makers; and
- **Educating society:** producing fundamental knowledge and developing frameworks of understanding.

The following were suggested as sources that may indicate that scientific activities may generate results of practical value:

- Research is cited as the basis for patents that lead to licenses;
- Research is used to justify policies or laws or cited in court opinions;
- Research is prominently discussed in trade publications of groups that might apply it;
- Research is used as a basis for practice or training in relevant fields of application;
- Research is cited and discussed in the popular press as having implications for personal decisions or policy; and
- Research attracts investments from other sources, such as philanthropic foundations.

6.3.1 Interdisciplinarity

As noted earlier, cyber security research challenges reside within a particularly complex area, being at the intersection of behavioral sciences, formal sciences and the natural sciences – it requires interdisciplinary research to materially impact our cyber security challenges. [E53] observes that the frontiers of science are generally located at the interstices between and intersections among disciplines and that interdisciplinary thinking has become more integral to many areas of research because of the need to understand “the inherent complexity of nature and society” and “to solve societal problems.”

⁶ In principle, these indicators could be converted into numeric measures of progress.

To support interdisciplinary research it is important that favourable conditions for contact and collaboration amongst researchers be in place. Such favourable conditions can be engendered by:

- Defining challenges in terms of issue-oriented interdisciplinary lines;
- Creating multidisciplinary panels to review proposals in emerging interdisciplinary areas;
- Disciplinary depth and breadth of interests, visions and skills, integrated within the research groups;
- Institutional commitment and research leadership with a clear vision and teambuilding skills;
- Recognizing that effective communication amongst researchers of different backgrounds will take time – need to learn different languages and leadership that encourages contribution and benefits;
- Creating new modes of organization, recruitment and modified reward structures;
- Creating a problem-oriented organization and the ability to reorganize as problems change; and
- Funding organizations may need to change their proposal and review criteria.

6.3.2 Implications for decision making

[E53] captures various implications for decision making as a result of its analysis of progress in science. The following are of particular note:

1. No theory exists that can reliably predict which research activities are most likely to lead to scientific advances or societal benefit;
2. Science produces diverse kinds of benefits: consequently, assessing the potential lines of research is a challenging task;
3. Portfolio diversification strategies that involve investment in multiple fields and multiple kinds of research are appropriate for decision making, considering the inherent uncertainties of scientific progress;
4. Research managers should seek to emphasize investing where their investments are most likely to add value;
5. Types of scientific progress can include related advances in databases and analytic techniques; and
6. For interdisciplinary research – support issue-focused interdisciplinary research.

6.4 Technology Transfer

Technology Transfer pertains to the transfer of knowledge, methods, technology, and so on, amongst institutions so as to ensure that scientific and technological developments are more widely disseminated and can be further exploited and advanced. It is widely recognized that technology transfer is both a key requirement for ongoing institutional success, but also an extremely difficult challenge. The literature is replete with treatises on technology transfer identifying various processes and attributes that may lead to successful adoption. Part 8 provides a brief discussion of and pointers to some technology transfer (diffusion) models.

Instead of writing a lengthy treatise on Technology Transfer, we focus on two technology transfer points that are abstracted from specific experiences. The first point discusses attributes of a Technology Transfer Laboratory. The second point discusses important partnering attributes enabling successful technology transfer.

Above, we introduced the newly formed Sandia National Laboratories cyber research facility and discussed some of its attributes. In general these labs need to support a responsive technology transfer process. We take the perspective that a Technology Transfer Lab (TTL) is a “meeting ground” where organizational critical challenges are addressed, external capacity is leveraged, leading edge technology is brought to the mission, and the mission is brought to leading edge technology. It is a means of responding to the evolving operational environment by transferring leading edge technology. We view a TTL as an open collaborative research environment to demonstrate, evaluate, evolve and document leading edge technology to enhance operational capability.

The key characteristics and outcomes of a TTL are that it is a fully fitted, agile, proximate (to the organization) open lab supporting on-demand, risk-managed integration of external personnel and software, permitting increased availability of organizationally critical staff investigating and rapidly evolving leading edge technologies. This leads to capabilities that are rapidly deployable and more closely aligned with organizational needs and with greater organizational internal architecture fidelity, while simultaneously communicating organizational technology interests, influencing the larger research environment and attracting broader engagement.

Key to the successful evaluation of a potential new product or service is the interaction between the adopting organization and the supplier (for example, a research laboratory, academic institution, or commercial entity). In the best of all worlds, both organizations would benefit from the engagement through dissemination of appropriate information, feature prioritization, and mentoring/assistance with adoptee experiments.

Through [E12], some key observations regarding best practices on software technology transfer are presented and support the comments of the previous paragraph. These practices include:

- Sustainability – e.g., both the adopting organization and the technology provider must see continuing returns from the technology transfer relationship;
- Training and mentoring;
- Scaled project engagement – careful selection of trial projects that are sufficiently bounded in time, are not on the main stream, but exercise the technology. Best talent needs to be assigned to such projects;
- Recruitment – bring in top talent skilled in the new technology;
- Metrics – one must be able to measure the key attributes such as skills improvements, training days per person, performance improvements, new technical problems now solvable, and so on;
- Fail fast; fail often – keep the cost of failure low. Practical suggestions include using short time boxes (a few weeks), light weight agile development process, short steps rather than big leaps, not all prototypes need to be completed; and
- Infrastructure – ideally, similar hardware/software environments internal/external to the organization.

7 Cyber Security Related Source Material

In this section we provide summaries of some of the key cyber-security related source material. The next sub-section provides a high-level summary of the reference material; each subsequent sub-section then focuses on the particular source material.

7.1 Summary of Source Material

While the subsequent subsections provide more detailed information on the source materials; here, we identify the source material that have helped to inform this R&ED cyber security program. The key source materials are:

- **Science of Security Joint Statement of Understanding (2011):** A common perspective reached by agencies of the Canadian, United Kingdom and United States governments in which the lack of a scientific basis for security and some of the deleterious consequences identified. In this context, science is viewed as knowledge that results in correct predictions and reliable outcomes. Seven core interrelated themes are identified: **common language** (express security in a precise and consistent manner); **core principles** (the foundational principles and fundamental definitions of concepts); **attack analysis** (understanding the attacker); **measurable security** (we need valid metrics to inform options); **risk** (improving the quality and consistency of risk assessment); **agility**

(the ability to keep up with the dynamic cyber-environment); and **human factors** (psychological issues, for example).

- **Roadmap for cyber security Research (DHS 2009) [E16]:** A U.S. roadmap used to define a national R&ED agenda in cyber security, which will respond to the challenges of today and envision those of the future. The roadmap provides detailed research and development agendas relating to 11 hard problem areas in cyber security: scalable trustworthy systems; enterprise-level metrics; system evaluation life-cycle; combatting insider threats; global-scale identity management; survivability of time-critical systems; situational understanding and attack attribution; provenance; privacy-aware security; and usable security. For each of these problem areas critical needs, research gaps and near-, medium- and long-term research agendas are identified.
- **Science of Cyber-Security (2010) [E14]:** A JASON report responding to a U.S. Department of Defense (DoD) request “to examine the theory and practice of cyber-security, and evaluate whether there are underlying fundamental principles that would make it possible to adopt a more scientific approach, identify what is needed in creating a science of cyber-security, and recommend specific ways in which scientific methods can be applied.” JASON refers to an independent group of scientists that advises the U.S. government on issues pertaining to science and technology. The report responded to specific DoD questions, identified computer science sub-disciplines that are relevant (such as cryptography, game theory, model checking and machine learning) and provided recommendations on how to advance a science of cyber-security.
- **Enabling Distributed Security in Cyberspace (2011) [E3]:** A U.S. Department of Homeland Security discussion paper that explores “the idea of a healthy, resilient – and fundamentally more secure – cyber ecosystem of the future, in which cyber participants including cyber devices, are able to work together in near-real time to anticipate and prevent cyber-attacks, limit the spread of attacks across participating devices, minimize the consequence of attacks, and recover to a trusted state.” Key to the DHS discussion is the drawing of analogies and inspiration from the human immune system and public health systems.
- **Science of Security Hard Problems: A Label Perspective (2013) [E59]:** Researchers from Carnegie Mellon, North Carolina State and University of Illinois at Urbana-Champaign identified five “hard problems,” based on “their level of technical challenge, their potential operational significance, and their likelihood of benefiting from emphasis on scientific research methods and improved measurement capabilities.”

7.2 Science of Security Joint Statement of Understanding (2011)

This is a **joint statement of understanding** between agencies of the Canadian, United Kingdom and United States governments on the subject of “Security Science” – known respectively within these nations as “Science of Security”, “IA Science” and “Science of IA.” This statement is based on output from an IA Science Workshop, May 2010. It describes our collective understanding of the problem space; why we believe IA requires a stronger foundation today, and the body of work which is required in order to deliver it.

Security is still an art

With the exception of cryptography, we lack a proper scientific basis for Information Assurance today. Security is now a much broader field, and because of the rising complexity of networks and services, the risk exposure that our operational systems are subject to is increasingly difficult to assess. And yet we need to be able to make this assessment more quickly than ever before. Although we have a great deal of process and methodology, the

underlying appraisal of possible weaknesses is still very subjective - carried out by individuals based on their own level of skill and experience. We often measure security based on process, rather than on objective measures. It is also difficult to carry out trade-offs, to determine how to get the most effective security from limited resources. We cannot currently quantify the effectiveness of different security technologies or policies. We cannot assess the security of systems or organizations.

There are systemic problems in the way security is perceived. This has wide ranging impact, for example on the rate of compliance in the user community; on the motivation of security professionals; on risk decision making by mission leaders.

In addition, we have a mindset which is focused on defending against known attacks. There is a general lack of awareness of the limitations of security technologies, particularly with regard to previously unseen attack. We react to attacks we see, after they have succeeded. There is little appetite to contemplate or debate the extent of attacks which we may not be detecting. We need to move away from the reactive model to be able to apply pre-emptive defence, based on foundational principles.

A foundational science for security

In the context of security, science can be thought of as knowledge that results in correct predictions or reliable outcomes. The "Science of security" resides in a particularly complex area, being at the intersection of behavioral sciences, formal sciences and natural sciences. We identify a set of 7 core themes that together form the foundational basis for our discipline. The themes are strongly inter-related, and mutually inform and benefit each other. They are:

Core theme: *Common language*

This theme seeks ways to express security in a precise and consistent way. This is a vital foundational requirement so that we can express and develop our understanding of security. Languages do not have to be textual; they may be symbolic, graphical or model-based. However, they must have an agreed upon semantics. All other themes will drive requirements for language. Examples include: Can we define a modeling language to express the security aspects of system architecture? Can we develop new visualization techniques to describe the output from risk assessments? Can we find a language to express core principles such as trust relocation?

Core theme: *Core principles*

Security is lacking in foundational principles and fundamental definitions of concepts. There is a body of work to draw upon, and some well-established terms, such as the principle of least privilege or of defence in depth. Definitions, however, tend to vary, and there is little guidance on practical application, where the security principles are often in direct conflict with other design principles. Complicating this matter further is the fact that we often compose principles in an ad hoc fashion without truly understanding the implications. There are other areas where new principles may be developed, dealing with topics such as: trust relocation; and composition of security properties.

Core theme: *Attack analysis*

The deepest understanding of security is obtained when it is informed from an attacker's perspective. As government bodies, we have excellent access to attack information and data. Proper management and analysis of this data could deliver many varied benefits. It could help justify security investment, and in carrying out balanced trade-offs between different security options and the application of core principles. It could help in addressing

problems in the public perception of security – making the deliverable more tangible. Careful analysis may help us to estimate where we are not detecting attacks.

Core theme: *Measurable Security*

A raft of work is required to explore techniques to measure security and develop the economic model. Measurements must be developed which include not just technical measures, but also the influence of security policy and user behaviour. It is important to be able to carry out trade-offs between security, usability, functionality and cost to enable better informed investment decisions. We need to be able to measure and compare the security of: individual products; system architectures; or an entire organization.

Core theme: *Risk*

Work is required to improve the quality and consistency of risk assessment. Much of the work in this field has focused on process and methodology, but assessment is still based on individual expertise. Our desire is for risk assessment to be more consistent and less subjective. Research is required to assess the level of variability in risk decision making and to determine the underlying rationale.

Core theme: *Agility*

IA Services in general need to become more agile, to reflect the more dynamic environment that systems now reside in. We need to be able to respond to an evolving threat landscape and rapidly evolving technology. We need to be able to assess threat and risk much more quickly, and to detect and respond to attacks on our systems in real time.

Core theme: *Human Factors*

This theme tackles factors affecting people's security-relevant behaviour. It tackles issues such as: how to make the intangible benefits of IA visible; how to secure the optimum psychological contract for user compliance; the most effective way to communicate information risk, and how the communication method affects subsequent attitude to risk. It also includes usability issues – particularly in designing security so it incentivizes secure behaviour.

We must deliver short term benefit...

In general, the development of a "science" will be a long term effort, and should be pursued as a collaborative effort among governments, industry and academia. However, all nations are agreed that Security Science must deliver tangible benefit in the short term, without compromising the longer term effort. It should be strongly informed by the needs and insights of practitioners in the field.

7.3 Roadmap for Cybersecurity Research (DHS 2009)

This report provides detailed research and experimental development roadmap for eleven hard problems in cyber security, thereby informing research and experimental development stakeholders (e.g., funding agencies and researchers) of a well-sourced set of critical challenges, research gaps, and research agendas that merit near-, medium- and long-term attention. These hard problems are categorized as being "overarching," "major threats" or "system concepts." A further categorization, "usable security," is separately identified as it is viewed as being cross-cutting through the other three categories. The primary objective of the roadmap is to lay the groundwork for a U.S. R&ED agenda that will ultimately lead to the nation being ahead of its adversaries and protect systems integral to the U.S. It is further recognized that future generations of networking and information technologies must be designed with security built-in and from the ground up.

Overarching: While not particularly defined in the report, “overarching” pertains to core capabilities that must be developed to support progress in almost all of the hard problems; these are fundamental challenges. Absent success in the “overarching” hard problems, material progress elsewhere will be difficult to impossible. The “overarching” hard problems are:

- **Scalable Trustworthy Systems:** Three key components to this hard problem: attaining trustworthiness (a measure of the extent that a system satisfies multiple aspects of requirements), attaining scalability (the ability to continue to satisfy requirements as systems grow in size, complexity and functionality), and achieving composability (the ability to create systems and applications with predictable behavior from components). Three roadmap categories are identified: improving trustworthiness in existing systems, clean-slate approaches and operating successfully despite the presence of untrusted environments.
- **Enterprise-level Metrics:** Such metrics allow us to answer questions that are fundamental to investment and deployment decisions. They allow us to answer such questions as “how secure is my organization?” and “how has my security posture improved through the last set of updates?” To properly manage our systems, scientifically-based metrics and measures are required. Any underpinning “science of cyber security” will require a family of justified measures and metrics. Currently, there are no universally agreed upon methodologies to address the fundamental questions of how to quantify system security.
- **System Evaluation Life Cycle:** Evaluation encompasses testing/evaluation methods/tools that are deployed to evaluate whether a system (or component) satisfies particular security requirements. Current evaluation practices are unable to systematically and cost-effectively evaluate systems in a timely manner. Within this hard problem, it was identified that a major gap arises due to our lack of knowledge of the threat domain and this impairs the development of realistic security requirements.

Major Threats: These hard problems focus on the adversary and the techniques (malware) that are utilized. The “major threats” hard problems are:

- **Combatting Insider Threats:** An “insider threat” is one that is attributable to individuals who abuse granted privileges, either inadvertently or intentionally. The focus here is on the individual; the technical mechanisms are the focus of “combatting malware and botnets.” Current approaches (e.g., background checks, user authentication, application-level profiling and monitoring) are not consistently or stringently applied due to cost, limited effectiveness and low motivation. Research recommendations were presented within the following categorizations: collect, analyze and detect; deter and protect; and predict and react.
- **Combatting Malware and Botnets:** Malware refers to attack software that is loaded onto a machine and compromises that machine to the benefit of the adversary. Bots are a class of malware that allows for unauthorized remote control of a machine for malicious purpose. The technical challenges being addressed are characterized by prevention, protection (from malware extent in a system), detection (as it propagates), analysis (of infection, propagation and destructive mechanisms) and reaction (e.g., remediation and attribution)

System Concepts (“ilities”): These hard problems refer to more specific system attributes and concepts required to implement the above topics. The “system concepts” hard problems are:

- **Global-scale Identity Management:** This hard problem focuses on identifying and authenticating entities (e.g., people, hardware, sensors) when accessing critical information technology systems from anywhere. It does not pertain to universal access or single identity. Identification and Authentication systems are being regularly attacked. Research directions are categorized as being either mechanisms (e.g., authentication, revocation, accountability, evaluation) and policy-related (e.g., privacy, administration, social/cultural mores). A key gap in identity management is the lack of transparent, fine-grained, strongly typed control of identities, roles, attributes and credentials.

- **Survivability of Time-Critical Systems:** The report defines survivability as the capability of a system to fulfill its mission, in a timely manner, in the presence of attacks, failures or accidents. This is a key component of trustworthiness. A “time-critical system” must respond in a “faster-than-human” timeframe to maintain mission requirements. Currently, survivability is achieved through replication, redundancy, automated recovery, smart load sharing, and so on. Metrics are based on probabilities and do not typically take into account intentional attacks, cascading failures or other correlated causes/effects. Both the current state of the research and proposed future directions are categorized by “understanding the mission and risks,” “survivability architectures, methods and tools,” and “test and evaluation.” Issues include rigorous definition of properties, interdependencies, composable/scalable trustworthiness, composable survivability and evaluable metrics.
- **Situational Understanding and Attack Attribution:** “Situation Understanding” is information pertinent to one’s level and areas of interest; it encompasses role, environment, adversary, mission, resources, and so on. It includes the state of one’s own system from a defensive posture, regardless as to the presence/absence of attack. “Attack attribution” determines the identity/location of an attacker or an attacker’s intermediary. Intrusion detection/prevention systems are a key current technology used for “situational understanding,” with various attempts at the use of visualizations and other analytical tools to improve comprehension in the face of large amounts of data. Attack signatures are failing. Key research areas include the analysis of massive data scales, novel approaches to presentation in large-scale data, cross-boundary sharing and visualization (using both small handheld devices and wall-size screens).
- **Provenance:** Provenance refers to the chain of successive custody and includes pedigree (historical dependencies) and tracking. Provenance assists in the determination of trustworthiness and reliability of data and thereby informs the decision-making process. A key dependency is trustworthy systems so that provenance can be determined accurately. Key research categories are representation (e.g., data models), management (e.g., creation and access) and presentation (e.g., query, alert). For management, consideration of trustworthy distributed embedding with integrated analysis tools is a potential approach.
- **Privacy-aware Security:** The goal of privacy-aware security is the enablement of users/organizations to better express, protect, and control the confidentiality of their private information, even when they are required to share information. This area includes the topics of anonymity, confidentiality, query protection, monitoring and accessibility. There is a complex mix of legal, policy and technical considerations. Currently, there are no widely adopted frameworks for privacy-aware security. Significant gaps exist in selective disclosure and privacy-aware access (e.g., sound frameworks), specification frameworks (for expressing privacy guarantees), and numerous policy issues.

The final hard problem category is Usable Security. Bottom line here is that security must be usable by persons ranging from nontechnical users to experts and systems administrators. Unfamiliar technology, changing security landscape, hidden and/or inflexible controls, lack of education/training, complexity are all significant challenges facing the various stakeholders. Current approaches with passwords, mail authentication, client-side certificates, CAPTCHAs, and so on, are all inadequate in terms of usability. Research recommendations were provided in the context of interface design, science of evaluation for usable security, and tool development.

7.4 Science of Cyber-Security (MITRE-JASON, 2010)

JASON refers to an advisory group of scientists that advises the U.S. government on issues pertaining to science and technology.

The Science of Cyber-Security report discusses the interplay of science with cyber-security. Over the past few years there has been an increasing interest in developing a scientific foundation for cyber-security so as to improve from

the current sets of ad hoc practices. As discussed elsewhere, the U.S., U.K. and, to a lesser extent, Canada, have been building competencies in the science of (cyber-) security.

In the DoD's charge to JASON it was noted that "we do not even have the fundamental concepts, principles, mathematical constructs, or tools to reliably predict or even measure cyber-security." The charge goes on to note that "it is difficult to determine the qualitative impact of changing the cyber-infrastructure (more secure now or less secure?) much less quantify the improvement on some specific scale."

JASON recognizes that cyber-security consists of a peculiar set of features not shared by any other area of study. These features include that cyber-space is a human built entity; it is digital; and it is adversarial. The view is synopsized by: "... the reasoning in cyber-security has to be both about the constructed universe and the actions and reactions of the adversaries." Hence, Peter Galison of Harvard, views the science of cyber-security as a "Manichean science," or a science in the presence of adversaries.

The JASON report outlines various areas that could usefully inform a science of cyber-security. These areas include Trust, Cryptography, Game Theory, Model Checking, Obfuscation, Machine Learning, and Composition of Components.

The importance of metrics as fundamental to scientific progress is identified. The JASON report suggests that there are many statistics available from which measures could be built. They also take the view that metrics must change over time to reflect a changing environment (e.g., changes in defence or attacks).

The report has two chapters that focus on model checking (the role of model checking and formal methods in computer security) and on the analogy of cyber-security to the immune system. These chapters reflect presentations by two leading researchers (Gerald Holzmann and Stephanie Forrest). Reference is also made to a presentation by Fred Schneider regarding hyper-properties, which has been identified as a potentially common language for reasoning not only about properties of program, but executions of systems and enforcement of policies. JASON concludes that "there is a connection between cyber-security and its underlying science, which in turn is connected to fundamental issues in model verification." The report goes on to state that "there is no panacea for analysis of cyber-security but the foundations, and, to some extent, the methodologies for addressing specific questions do exist in the computer science community." What is missing is the direct connection to day-to-day programming!

The analogy with immunology allows for useful general guidance for cyber-security, but is of limited help once the analogy is driven down into fine-grained levels as important technical differences arise and merely mimicking a biological system is unlikely to be a workable approach. However, at the higher level, immunology is suggestive that cyber-security must support adaptive response, consist of a mix of sensing modalities, controlled experiments (though the complexity of modern systems is a prohibitive challenge), time scale differences and response to malware detection.

JASON concludes that there is a science of cyber-security. They also respond specifically to a number of questions posed by DoD (of which we summarize briefly JASON responses):

What elements of scientific theory, experimentation, and/or practice should the cyber-security research community adopt to make significant progress in the field? How will this benefit the community? Are there philosophical underpinnings of science that the cyber-security research community should adopt?

- Define a common language and a set of basic concepts about which the security community can develop a shared understanding.
- Such a language along with agreed upon experimental protocols will facilitate the testing of hypotheses and validation of concepts.

Are there “laws of nature” in cyber space that can form the basis of scientific inquiry in the field of cyber-security? Are there mathematical abstractions or theoretical constructs that should be considered?

- No, since cyber-security is an applied science informed by the mathematical constructs of computer science such as automata theory, complexity, and mathematical logic.

Are there metrics that can be used to measure with repeatable results the cyber-security status of a system, of a network, of a mission? Can measurement theory or practice be expanded to improve our ability to quantify cyber-security?

- JASON believes that there are many metrics that can be productively employed, but that they are empirically based and statistical in nature – they do not apply well to ill-defined scenarios.
- Unobserved new attack vectors will not contribute to metrics and repeatability is an issue.

How should a scientific basis for cyber-security research be organized? Are the traditional domains of experimental and theoretical inquiry valid in cyber-security? Are there analytic and methodological approaches that can help? What are they?

- Experimental and theoretical inquiry does apply to the study of cyber-security.
- The highest priority is the establishment of research protocols that enable repeatability of experiments.

Are there traditional scientific domains and methods such as complexity theory, physics, theory of dynamical systems, network topology, formal methods, mathematics, social sciences, and so on, that can contribute to a science of cyber-security?

- Model checking, cryptography, code obfuscation, type theory, game theory, secrecy were all mentioned.

How can modeling and simulation methods contribute to a science of cyber-security?

- Continually test security assumptions on running systems.
- Use virtual machines, and so on, to provide well-defined test beds.

Repeatable cyber experiments are possible in small closed and controlled conditions but can they be scaled up to produce repeatable results on the entire Internet? To the subset of the Internet that support DoD and the IC?

- Premature to ask this question due to the lack of repeatability of extent experiments.

What steps are recommended to develop and nurture scientific inquiry into forming a science of cyber-security field? What is needed to establish the cyber-security science community?

- Establish interdisciplinary centres that connect academia, industry, national labs, and the DoD.

Is there reason to believe the above goals are, in principle, not achievable and if so, why not?

- Significant progress can be made towards the above goals, though one must first understand the nature of scientific enterprise for cyber-security and characterize the objects under discussion.

7.5 Enabling Distributed Security in Cyberspace (DHS 2011)

This report, through analogies with the human immune system, pictures a future cyber ecosystem that is healthy, resilient and fundamentally more secure. It is envisioned that the participants in such a cyber-ecosystem will jointly work in near real-time to respond to, anticipate or prevent cyber-attacks. In particular, the spread of attacks would be limited, the consequences minimized, and the cyber ecosystem able to return to a trusted state.

Increasing virulent, cyber-attacks tend to follow a systematic escalation path: reconnaissance, gaining entry, establishing persistence, creating external exfiltration paths, and conducting attacks. If cyber defence were to advance from today's generally manual and ad hoc responses by near real-time communication regarding attacks, coordinating security hardening consistent with policy, then enterprise objectives could be more substantively sustained. In other words, a healthy cyber ecosystem would interoperate broadly, collaborate effectively, respond agilely, and recover rapidly.

Three interrelated building cyber ecosystem building blocks are identified and discussed:

Automation: Automated Courses of Action (ACOAs) are strategies that incorporate decisions and actions taken in response to cyber events. Drawing analogies from the human immune system, the report goes on to suggest that a healthy cyber ecosystem could employ an automation strategy of fixed local defenses (analogous to cell mediation) supported by mobile and global defenses at multiple levels (analogous to the humeral system of the body). (This analogy with health is also raised to the public health sector, whereby the report suggests benefits accruing from the creation of a cyber-CDC (Centres for Disease Prevention and Control)).

Interoperability: In the presence of interoperability, cyber communities may be defined by policies and allows for seamless and dynamic collaboration. It enables common operational pictures and shared situation awareness. Three forms of interoperability are defined and described in some detail: semantic interoperability (shared understanding of communicated data), technical interoperability (sharing based on well-defined and widely adopted interface standards, and policy interoperability (pertaining to common business processes related to the transmission, receipt, and acceptance of data among cyber participants). Specific mention was made of security content automation (and the security content automation protocol - SCAP) as enabling all three forms of interoperability. A roadmap relating to strategic considerations of cyber security content automation (waves of increasingly capable security functions) was summarized:

- **first wave:** vulnerability assessment, configuration assessment, compliance management, asset inventory
- **second wave:** e.g., malware analysis incident reporting, event management, software assurance
- **third wave:** e.g., collaborative threat intelligence, forensics and damage assessment, modeling and simulation, supply chain assurance

Authentication: Authentication provides assurance that participants are authentic or genuine and is critical to cyber defence since communication and content attribution are essential components of security decisions.

Authentication is foundational to many cyber capabilities. Significant considerations underpinning authentication techniques are the ease of integration into emerging and deployed devices and software applications and the ease for which authentication can be exchanged or federated across networks and organizations.

The report goes on to discuss an evolving approach to replace command and control with agility, focus and convergence. The report briefly defines agility, focus and convergence thusly: agility is the critical capability to meet complexity and uncertainty challenges; focus provides the context and defines the purposes of an endeavor (agnostic to underlying technology); and convergence pertains to the goal-seeking process that guides actions and effects. In its most fulsome form, the most mature focus and convergence networked environment would be

characterized by a robustly-networked collection of devices having widespread and easy access to information, sharing information extensively, interacting in a rich and continuous fashion, and having the broadest possible distribution of decision rights. The environment will be able to self-synchronize in an agile and adaptable manner.

From an outcomes perspective, a healthy cyber ecosystem should have the following attributes: information is connected across time and space; rapid and essential universal (across all cyber participants) learning; greater attribution; new analytics (thereby producing new intelligence); greater network reach; new defensive tactics (such as moving target defence); and lifecycle feedback. The building blocks of the ecosystem should be inclusive, effective, smart, barrier-free, optimized, understandable, useable and assured.

7.6 Science of Security Hard Problems: A Lablet Perspective

This report [E59] provides a perspective on the Science of Security as expressed by three Science of Security Lablets, which are located at Carnegie Mellon, North Carolina State and University of Illinois at Urbana-Champaign. These lablets perform focused research activities and “share a broad common goal, which is to develop the foundations for security science, with a focus on advancing solutions to a selection of the hardest technical problems. The goal is to develop foundations for the science of security ...”

Five “hard problems” were identified, based on “their level of technical challenge, their potential operational significance, and their likelihood of benefiting from emphasis on scientific research methods and improved measurement capabilities.” The five challenge problems are as follows:

- **Scalability and composability:** Develop methods to enable the construction of secure systems with known security properties from components with known security properties, without a requirement to fully re-analyze the constituent components.
- **Policy-Governed Secure Collaboration:** Develop methods to express and enforce normative requirements and policies for handling data with differing usage needs and among users in different authority domains.
- **Security-Metrics-Driven Evaluation, Design, Development and Deployment:** Develop security metrics and models capable of predicting whether or confirming that a given cyber system preserves a given set of security properties (deterministically or probabilistically), in a given context.
- **Resilient Architectures:** Develop means to design and analyze system architectures that deliver required service in the face of compromised components.
- **Understanding and Accounting for Human Behavior:** Develop models of human behavior (of both users and adversaries) that enable the design, modeling, and analysis of systems with specified security properties.

8 Technology Transfer Models

As noted earlier, Technology Transfer pertains to the transfer of knowledge, methods, technology, and so on, amongst institutions so as to ensure that scientific and technological developments are more widely disseminated and can be further exploited and advanced. It is widely recognized that technology transfer is both a key requirement for ongoing institutional success, but also an extremely difficult challenge. Here, we provide a brief discussion of and pointers to some technology transfer (diffusion) models.

For example, Everett Rogers [E35] introduced a technology diffusion model which, slightly adapted here, includes the key criteria of:

- **relative advantage** – an analysis of the technical and business superiority of the innovation over technology it might replace;
- **compatibility** – an analysis of how well the innovation meshes with existing practice;
- **complexity** – an analysis of how easy the innovation is to use and understand;
- **trialability** – an analysis of the type, scope and duration of feasibility experiments and pilot projects;
- **observability** – an analysis of how easily and widely the results and benefits of the innovation are communicated;
- **transferability** – an analysis of the economic, psychological and sociological factors influencing adoption and achieving critical mass. These include:
 - “prior technology drag” – the presence of large and mature installed bases for prior technologies,
 - “irreversible investments” – the cost of investing in the new technology,
 - “sponsorship” – champions, and
 - “expectations.”

Further, it is noted that early adopters are also at risk because of “transient incompatibility” (the adopter is moving ahead of their community/market) and “risks of stranding” (the adopter is stranded because the community/market did not adopt the technology).

Other examples include the work by:

- Geoffrey Moore on Crossing the Chasm [E36] in which he distinguishes between innovators (who pursue new technology aggressively), early adopters (who appreciate the benefits of a new technology and relate the benefits to their concerns), early majority (relate to the technology, but are driven by practicality), late majority (similar concerns to early majority, but dislike technology products) and laggards (do not like technology);
- Clayton Christensen [E37] in which he notes that disruptive products must be simple, convenient and fool proof and provides strategies for advancing the product (such as setting up an autonomous organization, run by a senior manager and positioning the innovation as a threat); and
- Richard Gabriel [E38] who takes the perspective that it is best to start with a minimal creation (product) and grow it as needed. Gabriel argues a “worse is better” metaphor in which simplicity and correctness are key, consistency must be largely attained, but completeness can be sacrificed.

9 References

- [E1] Deirder K. Mulligan and Fred B. Schneider. Doctrine for Cybersecurity. Technical report, University of California (Berkeley) and Cornell University, May 15, 2011.
- [E2] Fred B. Schneider. Blueprint for a Science of Cybersecurity. Technical report, Cornell University, May 24, 2011.
- [E3] Enabling Distributed Security in Cyberspace: Building a Healthy and Resilient Cyber Ecosystem with Automated Collective Action. Department of Homeland Security, March 23, 2011.
- [E4] Michael R. Clarkson and Fred B. Schneider. Hyperproperties. Technical Report, Cornell University, January 25, 2008.
- [E5] Team for Research in Ubiquitous Secure Technology (TRUST). See: <<http://www.truststc.org/>>.
- [E6] George Cybenko and Carle E. Landwehr. Security Analytics and Measurements. In IEEE Security and Privacy, May/June 2012.

- [E7] Lawrence Carin, George Cybenko, and Jeff Hughs. Cybersecurity Strategies: The QuERIES Methodology. In IEEE Computer, August 2008.
- [E8] Dave Thomas. Lean and Agile in the Large: Principles, Practices and Experiences for Large Scale Software Development. Bedarra Research Labs. Presented at the JA00 conference 2008. May/June 2008.
- [E9] Frascati Manual. Proposed Standard Practice for Surveys on Research and Experimental Development. OECD. 2002.
- [E12] Brian Barry and Dave Thomas. Private communication, July 20, 2012.
- [E13] Peter Sommer and Ian Brown. "Reducing Systemic Cybersecurity Risk". OECD/IFP Project on "Future Global Shocks". OECD Multi-Disciplinary Issues International Futures Programme. January 14, 2011.
- [E14] Colin D. McMorrow. Science of Cyber-Security. The MITRE Corporation, JASON Program Office. JSR-10-102. November 2010.
- [E15] NSF, IARPA and NSA sponsored a workshop in November 2008 on the Science of Security. See: <http://sos.cs.virginia.edu/>.
- [E16] A Roadmap for Cybersecurity Research. Department of Homeland Security, November, 2009.
- [E19] Dorothy Denning. A Lattice Model of Secure Information Flow, Communications of the ACM, May 1976.
- [E20] Michael Harrison, Walter Ruzzo, Jeffrey Ullman. Protection in Operating Systems, Communications of the ACM, August 1976.
- [E21] George Yee. The State and Scientific Basis of Cyber Security Metrics, DRDC Ottawa, March 2012.
- [E22] Engineering and Physical Sciences Research Council. Academic Centres of Excellence in Cyber Security Research, <http://www.epsrc.ac.uk/funding/cetnres/Pages/acecybersecurity.aspx>.
- [E23] Engineering and Physical Sciences Research Council. UK's First Academic Research Institute to investigate the "Science of Cyber Security," <http://www.epsrc.ac.uk/newsevents/news/2012/Pages/scienceofcybersecurity.aspx>.
- [E24] Sandia National Laboratories. Cyber research facility opens at Sandia's California site, https://share.sandia.gov/news/resources/news_releases/ctrl/.
- [E25] Canada's Cyber Security Strategy, www.publicsafety.gc.ca/prg/ns/cbr/fl/ccss-scc-eng.pdf.
- [E26] The Dark Space Project, in preparation.
- [E27] DeterLab: Cyber-Security Experimentation and Testing Facility, www.isi.deterlab.net.
- [E28] Joan Calvet, et al. The Case for in-the-lab botnet experimentation: creating and taking down a 3000-node botnet. ACSAC'10, Proceedings of the 26th Annual Computer Security Applications Conference, December 2010.
- [E29] Fact Sheet: Big Data Across the Federal Government (March 29, 2012), www.whitehouse.gov/sites/default/files/microsites/ostp/big_data_fact_sheet_final.pdf.
- [E30] What is a Data Scientist, www-01.ibm.com/software/data/infosphere/data-scientist/.
- [E31] VAST Challenge 2012 (Bankworld), www.vacomunity.org.
- [E32] Enron Email Dataset, www.cs.cmu.edu/~enron/.
- [E33] PREDICT (Protected Repository for the Defense of Infrastructure Against Cyber Threats), www.predict.org.
- [E34] Mark Whiting, et al. Creating Realistic, Scenario-Based Synthetic Data for Test and Evaluation of Information Analytics Software, Proceedings of BELIV'08 (Workshop on Beyond time and errors: novel evaluation methods for Information Visualization), ACM, dl.acm.org/citation.cfm?id=1377977.
- [E35] Everett Rogers, Diffusion of Innovations, Free Press, New York, 1983.
- [E36] Geoffrey A. Moore, Crossing the Chasm, Harper Business, 1995.
- [E37] Clayton M. Christensen and Michael E. Raynor, The Innovator's Solution: Creating and Sustaining Growth, Harvard Business School Press, 2003.
- [E38] Richard Gabriel, Worse is Better, www.dreamsongs.com/WorselsBetter.html.
- [E39] The National Cybersecurity Workforce Framework, csrc.nist.gov/nice/framework.

- [E40] Denise Melvin, One Page Practical Guides on Knowledge Sharing Techniques, www.km4dev.org/profiles/blogs/one-page-practical-guides-on.
- [E41] Defcon, www.defcon.org.
- [E42] Team for Research in Ubiquitous Secure Technologies, Knowledge Transfer, www.truststc.org/kt/index.html.
- [E43] ISACA, Understanding Virtual Organizations, www.isaca.org/Journal/Past-Issues/2001/Volume-6/Pages/Understanding-Virtual-Organizations.aspx.
- [E44] Natural Sciences and Engineering Research Council of Canada, Strategic Network Grants, <http://www.nsercpartnerships.ca/FundingPrograms-ProgrammeDeSubventions/index-eng.asp>.
- [E45] Wikipedia, Article on Centre of Excellence, http://en.wikipedia.org/wiki/Center_of_excellence.
- [E46] NITRD, www.nitrd.gov.
- [E47] Wikipedia, Article on Business Incubator, http://en.wikipedia.org/wiki/Business_incubator.
- [E48] MaRS, Building Canada's Next Generation of Growth Companies, www.marsdd.com.
- [E49] http://en.wikipedia.org/wiki/Computer_network_operations.
- [E50] The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research, U.S. Department of Homeland Security, www.cyber.st.dhs.gov/wp-content/uploads/2011/12/MenloPrinciplesCORE-20110915-r560.pdf
- [E51] The Belmont Report: Ethical Principles and Guidelines for the Protection of Human Subjects of Research, National Commission for the Protection of Human Subjects and Biomedical and Behavioral Research, U.S. Department of Health and Human Services, April 1979. www.hhs.gov/ohrp/humansubjects/guidance/belmont.html.
- [E52] http://www.cybersecurityresearch.org/news_and_events/press_releases/pr_20121023.html.
- [E53] A Strategy for Assessing Science: Behavioral and Social Research on Aging, Irwin Feller and Paul C. Stern (Editors), National Academies Press (U.S.), 2007, ISBN-13: 978-0-309-103976-8, www.ncbi.nlm.nih.gov/books/NBK26380/pdf/TOC.pdf.
- [E54] Innovation Canada: A Call to Action (Review of Federal Support to Research and Development – Expert Panel Report), October 2011, [http://rd-review.ca/eic/site/033.nsf/vwapj/R-D_InnovationCanada_Final-eng.pdf/\\$FILE/R-D_InnovationCanada_Final-eng.pdf](http://rd-review.ca/eic/site/033.nsf/vwapj/R-D_InnovationCanada_Final-eng.pdf/$FILE/R-D_InnovationCanada_Final-eng.pdf).
- [E55] Global 20 Leading Cyber Security Companies 2013: Competitive Landscape Analysis. <http://www.reportlinker.com>.
- [E56] GCHQ sets up £4.5m cyber vulnerability research institute, <http://www.computerweekly.com/news/2240179924/GCHQ-sets-up-45m-cyber-vulnerability-research-institute>.
- [E57] Assessing Cyber Threats to Canadian Infrastructure, Report prepared for the Canadian Security Intelligence Service, Angela Gendron and Martin Rudner, March 2012. http://www.csis-scrcs.gc.ca/pblctns/cdmtrch/20121001_cssnlpprs-eng.asp.
- [E58] Cyber security poses ongoing threat to Canadians: CGI clients leverage local and global security expertise to protect against escalating risk. <http://www.cgi.com/en/Cyber-security-ongoing-threat-CGI-clients-leverage-local-global-security-expertise-protect-risk>.
- [E59] Science of Security Hard Problems: A Lablet Perspective. David Nicol, William Sanders, William Scherlis and Laurie Williams. Draft of November 27, 2012.