

Workshop Report

National Science Foundation's Workshop on Trustworthy Computing

Workshop Topic: NITRD Themes and Science of Cybersecurity

Contributors: Trent Jaeger, *Penn State University*, Adam Smith, *Penn State University*, Marina Blanton, *Notre Dame University*, Kevin Butler, *University of Oregon*, Nicholas Hooper, *University of Minnesota*, John Knight, *University of Virginia*, Tal Rabin, *IBM Research*, Mohamed Shehab, *University of North Carolina, Charlotte*, Patrick Traynor, *Georgia Institute of Technology*

October 27-29, 2010

Arlington, Virginia, United States

Contents

1	Executive Summary	1
2	Keynote Talks	2
2.1	Keynote 1. Trustworthy Computing in the Clouds	2
2.2	Keynote 2. Architectures for Practical Client-Side Security	2
2.3	Keynote 3. Daniel Geer	2
2.4	Keynote 4. Whither Trustworthy Computing? A Short Summary of a Long History and Big Future Themes	3
3	Panel Sessions	4
4	Breakout Sessions	9
4.1	Breakout 1. Tailored Trustworthy Spaces	9
4.2	Breakout 2. Moving Targets	10
4.3	Breakout 3. Cybereconomics	11
4.4	Breakout 4. Science of Cybersecurity	11

1 Executive Summary

A trustworthy cyber infrastructure is vitally important to the future security of the country. It is imperative that researchers focus on the key challenges in building a trustworthy cyber infrastructure. To this end, the National Science Foundation held a Workshop on Trustworthy Computing on October 27-29, 2010 in Arlington, Virginia to bring researchers together to discuss challenges and future directions. The workshop focused on the NITRD Cybersecurity R&D Themes from the Federal Cybersecurity Game-Change Research Agenda (“NITRD themes,” hereafter) and the general challenge of improving the Science of Cybersecurity. The result of this workshop is a set of research challenges identified by the workshop participants proposed as key areas for scientific investigations for a future trustworthy cyber infrastructure.

The workshop was attended by security researchers with experience in research topics related to the NITRD Themes and Science of Cybersecurity, and some of the researchers with the potential to impact these themes in the future. The focus included not only research topics and potential directions, but also information about the government agencies’ interest areas to foster future collaborations. Participants included representatives from academia, industry, and program management.

The main body of the workshop was organized around four keynotes, an NSF TC program overview, four technical panels, and two project management panels. The keynotes focused on the nature of doing computer security research accounting for past, present, and future elements in network, systems, and data security. The NSF Trustworthy Computing program overview presented the history and future of this program. The technical panels targeted the three NITRD themes independently, where researchers presented their experiences in tackling technical challenges in related to these themes and advocated future research and collaborations that could be promising. The program management panels enabled program managers to outline the goals of their programs to motivate research in particular areas and foster collaboration. As the keynote speakers and panelists represented academia, industry, and program management, a variety of perspectives were presented.

In addition to the keynote speakers and panel sessions, a breakout session was held to collect the key research challenges for the NITRD themes and Science of Cybersecurity. The breakouts offered the prospective PIs a chance to ask questions to the panel speakers in a smaller group and brainstorm together to develop sets of key research challenges. See the breakout summaries for the lists.

The workshop agenda, slides, and video are available via links in the electronic version of this document and from the workshop website <http://tc2010.cse.psu.edu/>.

2 Keynote Talks

Four keynote addresses during the workshop served to highlight past challenges in computer security research, our current status on these challenges, and possible future directions.

2.1 Keynote 1. Trustworthy Computing in the Clouds

Speaker: David P. Reed, *SAP Labs*

David P. Reed's keynote examined the future of computing relative to the emergence of utility computing, its future form and uses, and the research challenges it creates. He argued that we have seen many technologies developed for somewhat basic purposes, but these technologies have been applied and subsequently evolved in unpredictable ways. He demonstrated this point through the review of the initial goals and subsequent developments for other technology, such as the Internet itself.

Thus, Reed argued that the utility computing concept will not stop with our current views of cloud computing, but rather the desire to have the right information at the right time will lead to utility computing becoming a platform for context-aware computing. This "3rd cloud" will build on connectivity of the Internet (the "1st cloud") and the "resource liquidity" and "service composability" of cloud computing to provide context-aware computing. The main research challenge in achieving this goal is establishing trust among parties. In context-aware computing, identity will be even more important, so trust becomes even more valuable and difficult to achieve.

See David P. Reed's slides at http://tc2010.cse.psu.edu/slides/reed_tc_2010.pdf to see the challenges in establishing trust. See video at <http://tc2010.cse.psu.edu/vid-wed-keynote.html>.

2.2 Keynote 2. Architectures for Practical Client-Side Security

Speaker: Virgil Gligor, *Carnegie-Mellon University*

Virgil Gligor's keynote examined the challenges in deploying trustworthy client systems and some insights to possible directions for addressing such challenges effectively. He started with the axiomatic facts that we have learned about deploying client systems in terms of axioms for insecurity, axioms for usable security, and the chasm between them. He then argued that we have a lot of experience in isolating systems, but problems occur when deploying communicating processes with different security requirements. He cites Lampson's red-green system as an architecture that can be achieved, even on a single platform via physical resource and execution partitioning.

Regarding trustworthy communication, Gligor cited several outstanding challenges. Inputs cannot be sanitized in general and verification could be arbitrarily complex, even when possible. A key methodology that he advocated is "optimistic trust," where input is received optimistically but recourse is possible. In particular, he claimed that optimistic trust requires accountability and recovery, so protocols for trustworthy systems must be designed with those features.

See Virgil Gligor's slides at <http://tc2010.cse.psu.edu/slides/gligor.pdf> to see examples of protocols for optimistic trust. Video to be uploaded.

2.3 Keynote 3. Daniel Geer

Speaker: Daniel Geer, *In-Q-Tel*

Dan Geer's keynote investigated issues in quantifying security and risk. He defines security as the absence of unmitigated surprise, so the challenge is to identify where surprises may occur (risks) and do something about it. Risk is difficult to quantify because more things can happen than actually will happen. So, finding which things may lead to risks involves identifying root causes and subsequent dependencies. These are not well-documented, so we have a dearth of good numbers to work from.

Another area of discussion was concern that adversaries are becoming more professional. Daily new malware is being developed and modified to avoid our current protections. Keeping up with patches may also not be sufficient. Thus, just adhering to compliance procedures is becoming an obsolete approach to computer security.

So, the last part of the talk focused on possible design options. The traditional idea of building security into systems requires too much time delay, so by the time that the system is deployed it is obsolete. Thus, we need to explore other options. For example, Geer cited the notion that fighter jets are inherently unstable, which enables maneuverability under control. Should we develop software that way, enabling the design to change quickly?

See Dan Geer's talk at <http://tc2010.cse.psu.edu/vid-thu-keynote.html>.

2.4 Keynote 4. Whither Trustworthy Computing? A Short Summary of a Long History and Big Future Themes

Speaker: Patrick Lincoln, *SRI International*

Patrick Lincoln's keynote provided a summary to the workshop. He identified a key theme of workshop being that cybersecurity is not evolving fast enough, as attacks are getting more complex and being introduced into new spaces. However, he saw some advantages emerging, such as increasing computing power. A highlight was his play on Landwehr's "swim with sharks" and "seaworthy vessel" to indicate a need to understand security precisely enough to avoid false solutions (e.g., a security Titanic).

The rest of the talk focused on possible directions that PIs may explore. He then listed a number of sources for security studies that may be useful to prospective PIs. An area of discussion was attribution, where Lincoln asked whether recovery and retaliation may be viable defenses. He referenced Dan Geer's question of whether we can build accountable systems from unaccountable components by discussing some instances of secure computing on untrusted infrastructure. Survivability was an important theme here.

See Patrick Lincoln's slides at <http://tc2010.cse.psu.edu/slides/lincoln.pdf> to see his advice for PIs. Video to be uploaded.

3 Panel Sessions

Panel 1. Trustworthy Computing: History and Prospects, Carl Landwehr, *National Science Foundation* Carl Landwehr’s talk kicked off the meeting with an overview of security and trustworthy computing history. His talk began with a number of highlights of computer security work up to the mid-1990s. Defense concerns dominated this part of the history, with MLS systems, security kernels, and evaluation methodologies dominating this time. While a number of significant technical advancements were made, current commercial computing environments have not adopted many of these approaches. However, it is going to be necessary to learn about what was tried, what worked, and what didn’t in order to move forward.

The second part of Landwehr’s talk focused on the Trustworthy Computing program. Key to this is the emphasis on the three NITRD themes: (1) tailored trustworthy spaces; (2) moving targets; and (3) cybereconomics. Also, background on funding and projects was provided.

The third part of Landwehr’s talk discussed “What Lies Ahead?” In this section, two strategies were presented: (1) learn to swim with the sharks and (2) building a seaworthy vessel. Also, Landwehr discussed the improvement of scientific method in security. Finally, guidance was provided for those seeking to submit a TC proposal.

See Carl Landwehr’s slides at <http://tc2010.cse.psu.edu/slides/landwehr.pdf> to see the challenges in establishing trust. Video to be uploaded.

Panel 2. NSF Programs Related to Trustworthy Computing Panelists:

Darlene Fisher, *NSF representing NeTS, ICES, NetSE*

Sean Wang, *NSF representing IIS*

Lenore Zuck, *NSF representing CCF*

Sam Weber, *NSF representing SBE*

Helen Gill, *NSF representing CPS*

Kevin Thompson, *NSF representing OCI*

Victor Piotrowski, *NSF representing Directorate for Undergraduate Education*

In this panel, various NSF representatives presented information about other NSF programs that may be of interest to members of the Trustworthy Computing community. Programs discussed included Network Technology and Systems (NeTS), Network Science and Engineering (NetSE), Interface between Computer Science and Economics & and Social Science (ICES), Intelligent Information Systems (IIS), Computing and Communication Foundations (CCF), Cyber-physical Systems (CPS), Social, Behavioral, and Economic Sciences (SBE), Directorate for Undergraduate Education, and Office of Cyberinfrastructure (OCI). News included that the NetSE program is scheduled to be terminated in 2011.

The audience asked about industrial partners for NSF grants. This is possible, although not specifically encouraged. The criteria for funding is somewhat higher in this case.

Panel 3. Tailored Trustworthy Spaces, Chair: Joshua Guttman, *Worcester Polytechnic Institute* Panelists:

William Arbaugh, *University of Maryland*

Carl Gunter, *University of Illinois, Urbana-Champaign*

Ruby Lee, *Princeton University*

Carl Gunter began the panel by discussing two application areas that do monitoring in security-critical environments: (1) healthcare, such as an assisted living service provider and (2) energy, such as a meter data management service. For the former problem, formal verification that the protocol satisfies the desired properties is an important task for this tailored space. For the latter problem, attestation is used to validate the integrity of meters to ensure that they are behaving as expected. Slides: <http://tc2010.cse.psu.edu/slides/gunter.pdf>.

Ruby Lee discussed her work on hardware-enhanced spaces. The goal here is to setup isolated compartments for trusted code in hardware. The challenge is to prevent an attacker from inserting malware below the trusted computing base of the system. However, some configuration needs to be done in software, so a question is can an application create a policy that can be deployed reliably on hardware? Lee described her groups Bastion architecture that provides a minimal trust chain, protection of secrets, and hardware anchors for data. She closed with a number of research questions for the development of hardware for security. Slides: <http://tc2010.cse.psu.edu/slides/lee.pdf>.

William Arbaugh also discussed providing a trusted foundation for computing. In this case, he focused on software integrity. The goal is to enforce invariants that define the integrity of the kernel or VMM. The problem is that attacks can be launched through changes in dynamic data. The key issue is that attacks must have certain behaviors that may be recognized as “tells.” One approach he explored was semantic integrity of the kernel. Using invariant predicates, they developed a tool that could detect bugs known to AV vendors reliably. Slides: <http://tc2010.cse.psu.edu/slides/arbaugh.pdf>.

Panel 4. Moving Targets, Chair: Sal Stolfo, Columbia University Panelists:

Anup Ghosh, *George Mason University*

John Knight, *University of Virginia*

Tal Rabin, *IBM Research*

John Knight began this panel by discussing his biologically-inspired work on security through diversity, called the Metamorphic Shield. This approach mutates the system regularly (i.e., faster than an attack rate) to change its attack surface. By understanding attacks on ISR, Knight predicts how often to mutate. Interestingly, he found that mutation can be done every 100ms with little macroscopic impact. Slides: <http://tc2010.cse.psu.edu/slides/knight.pdf>.

Anup Ghosh described the dimensions of a moving target solution. A designer must think about what to make uncertain, how often, and how this may be used to close attack surfaces. A key finding is that the attack model matters in determine how to proceed. He stated that we need to move from mechanisms to an understanding of how to use such mechanisms. Slides: <http://tc2010.cse.psu.edu/slides/ghosh.pdf>.

Tal Rabin discussed the use of cryptography in developing moving targets. Several cryptographic primitives have been developed that are appropriate in this context, including threshold cryptography, proactive key changes, and identity-based encryption. For example, in proactive security the assumption is that all parties will be compromised eventually and a public verification key may be hard to change. Thus, she proposed that new secret keys be generated and keys be distributed into shares to prevent any particular key from getting compromised. Slides: <http://tc2010.cse.psu.edu/slides/rabin.pdf>.

The panel discussed that the pros and cons of moving targets require new methods for mea-

surement and evaluation.

Panel 5. Science of Cybersecurity, Chair: Michael Reiter, *University of North Carolina* Panelists:

Andrew Appel, *Princeton University*

Amit Sahai, *University of California, Los Angeles*

Peter Weinberger, *Google*

Amit Sahai began the panel by defining that security is indeed a cat-and-mouse game, so we need to treat it that way. One insight is that we need to get away from assessing specific systems versus attackers and move towards assessing security models versus attackers. Another goal is to assess security versus a large class of attacks. Such models should motivate the design of powerful tools. The other focus of his talk was that cryptography and systems security could be more integrated. Recent work on homomorphic encryption, functional encryption, and secure computation are areas that may enable trustworthy computation in unverified systems. Slides: <http://tc2010.cse.psu.edu/slides/sahai.pdf>.

Peter Weinberger asked questions about what constitutes a science of cybersecurity. He claimed that this science is different than others because it is a human-made world with a mathematical “essence,” rather than a natural world. Thus, properties of this world may be artificial, yet we still have to understand and defend it from attackers. A particular advantage is the resources that a defender could bring to bear against an attacker. Such resources provide global knowledge, but organization is required. Weinberger recommended setting up multiple cybersecurity base centers to practice science. Slides: <http://tc2010.cse.psu.edu/slides/weinberger.pdf>.

Andrew Appel examined the search for successful abstractions for security. One successful abstraction is type systems. From type systems you can obtain soundness theorems where a program that compiles satisfies a correctness property. Appel argued that this science is well-developed. He discussed an example, where a small trusted computing base could support a proof-carrying code system. He also highlighted problems caused by violations of abstractions that may be available to attackers in practical deployments. Thus, he claimed that it is important to both define and use abstractions, but also understand the vulnerabilities in each. Slides: <http://tc2010.cse.psu.edu/slides/appel.pdf>.

Panel 6. Cybereconomics, Chair: Rebecca Wright, *Rutgers University* Panelists:

Matt Blaze, *University of Pennsylvania*

Jens Grossklags, *Princeton University*

Rafael Pass, *Cornell University*

The panelists discussed three different aspects of the role of economic incentives in trustworthy computing. To over-simplify somewhat, the three speakers addressed the incentive structure we have (or should) set up for technology developers, technology users and attackers, respectively.

Matt Blaze asked why many of the basic insights of security research have not been transferred to broad application (tongue only partly in cheek, he claimed there are two kinds of problems studied in the security literature: problems we don’t yet know how to solve, and problems we know how to solve but to whose solution no one pays attention). He raised the following broad question: where should we put our investment to improve the process of software development?

In particular, is it always true that improving the the general quality of software development improves the security of the resulting system? The question is more subtle than it first appears. For example, a recent paper of Blaze's highlights a dramatic difference between the "life cycles" of exploitable security bugs and other types of software bugs; this suggests that improving security may require a different allocation of resources than reducing the number and severity of other types of flaws. [No slides.]

Jens Grossklags argued for a broad theory of the "economics" of computer security. Individual agents (users, companies developing products, security experts) act to maximize their own benefit; the slow adoption of many basic security technologies and practices can be explained by the conflicting incentives these agents face. Several challenges complicate the application of traditional microeconomic theory to security. For example, security technologies do double-duty as public and private goods (e.g., installing patches to slow the spread of worms helps everyone, while keeping regular backups only directly benefits a particular person or organization). Another challenge is the variability in the expertise levels of different agents; nonexpert agents don't follow "rational" models of behavior because they do not have access to full information about their choices. Slides: <http://tc2010.cse.psu.edu/slides/grossklags.pdf>

Rafael Pass turned to how computational cost is an important component of "incentives" in the security context (especially when cryptography is involved). He gave a number of examples of games where computational considerations clearly affect the players' equilibrium behavior. He argued that any "economic" reasoning about security must take into account the cost of computation, pointing to a number of recent papers in the literature. Pass spent considerable time discussing secret sharing (so far, the best-studied setting): multiple players hold shares of a secret and must find a protocol in which they all have an incentive to reveal their shares rather than keep them secret. This setting highlights the importance of computational cost as well as the difference between game-theoretic, rational modeling and traditional adversarial modeling in the cryptography and security literature. He argued that exploring the interplay between these notions is an important and under-explored area of research. Slides: <http://tc2010.cse.psu.edu/slides/pass.pdf>

Panel 7. Other Government Funding Agencies Panelists:

Steven E. King, *Office of the Director, Defense Research & Engineering*

Sandy Landsberg, *DoE*

Bill Newhouse, *NIST Information Technology Lab*

Brad Martin, *NSA*

Douglas Maughan, *DHS*

Sandy Landsberg described DoE's interests in the mathematics of cybersecurity. Of particular interest is understanding the network structure and cybersecurity in their complex, networked systems. See Landsberg's slides for projects of interest at <http://tc2010.cse.psu.edu/slides/landsberg.pdf>.

Bill Newhouse discussed projects at NIST, where a key areas of interest is cloud computing, key management, usability, identity management, and cryptography. Newhouse provided an overview of NIST's goals and approach. See Newhouse's slides for addition information and NIST's National Initiative for Cybersecurity Education at <http://tc2010.cse.psu.edu/slides/newhouse.pdf>.

Brad Martin discussed areas of interest to the NSA. These focus on deployment of high con-

confidence software systems, where trust can be established in design and maintained through implementation and deployment. Of particular interest are trusted computing approaches, including hardware designs, attestation methods, and virtualization.

Steven King described goals with the Defense Research & Engineering organization. Their focus is on high assurance systems, and he highlighted the areas of interest to the different DoD organizations. They support a number of programs aiming at research to improve research in those areas. See King's slides to see details on those programs at <http://tc2010.cse.psu.edu/slides/king.pdf>.

Kevin Thompson, formerly of DHS, presented Douglas Maughan's slides on the DHS Cyber Security R&D Program. They are interested in R&D, test, and evaluation of security capabilities, so that they may be deployed in a timely fashion. He described the DHS process and active cyber security program areas, including two new programs. Maughan is particularly interested in the NITRD Theme of Tailored Trustworthy Spaces. See Maughan's slides for active programs at <http://tc2010.cse.psu.edu/slides/maughan.pdf>.

4 Breakout Sessions

4.1 Breakout 1. Tailored Trustworthy Spaces

Scribe: Kevin Butler, *University of Oregon*

This breakout began by discussing what a tailored trustworthy space means. While constructing an isolated system is feasible, no system is truly isolated, the discussion quickly focused on obtaining trust in a distributed system. Two disparate thoughts dominated this discussion: (1) obtaining a trustworthy space from imperfect components and (2) building up trustworthy spaces incremental in secure-boot manner isolated from others. The scientific challenges were seen to be in establishing an isolated environment, controlling communication effectively, providing verifiable assurance for the TTS result, and dealing with trust or lack thereof. Orthogonal to these the development of practical secure computation was seen as a mechanism that may impact the priority of the scientific challenges above. Methods for building spaces were also highlighted as important, as building a finely-crafted spaces may be preferable if it can be done easily.

The group generated a set of detailed research challenges to address in building tailored trustworthy spaces (TTS).

1. How to develop and manage trust for controlling communication among TTS's? What are the degrees of trustworthiness and metrics for finding them?
2. Who gets to articulate policies of TTS and how are multiparty policies composed?
3. How is provenance and tracking and enforcement wrt what kind of security operations performed on it?
4. What are the safety properties for TTS (e.g., type safety and info-flow) and how are they defined and enforced?
5. What are the potentially successful abstractions of TTS, such that they can successfully compose (e.g., peer-to-peer and/or hierarchical)?
6. How does a TTS deal with attacks against it? How can it protect valuable information and recover in a dangerous environment? Can we make a TTS for self-protecting data?
7. Can we make these spaces dynamic and adaptable through configuration or crowd-sourcing?
8. Can we define models of integrity and test against threat models, see what risks you have and figure out mediations o risks to minimize the risk you are facing subsequently? And once you find threat model, how do you resolve while finding least, most manageable risk?
9. How can content be described to see whether information can be assured or not - characterizing the meaning of information as to its relevance? Who decides relevance? Machine or user?
10. What are the range of options to user, where the range of misbehaviors could perhaps be limited?

11. How to integrate dependency/workflow among distributed systems into cyber defense? For example, what if we can't reach certain trustworthiness?
12. How do we protect the provenance of metadata and what infrastructure necessary to support its processing and storage?

4.2 Breakout 2. Moving Targets

Scribe: Patrick Traynor, *Georgia Institute of Technology*
with additional notes from John Knight and Tal Rabin

The breakout started by stating that the idea of a “moving target” is to try to improve system security by morphing a system in some way over time so as to change the target to which the adversary is trying to gain access. The expectation is that the changes will defeat some attacks because the adversary will not have essential details needed for a successful attack.

This breakout began with a question from Blaze's presentation. “Complexity of a system is advantageous to attackers.” Does this contradict the basis of moving targets? The group found that moving targets is about moving complexity from the defender to the attacker. For example, this may be done by making the attacker think your system is an active defender (a bot from Google or Microsoft). Thus, nothing may need to be “moved” or “randomized”, that perhaps it is just the work/complexity that is moved. It was suggested that a metric of success may be to determine how much work you've pushed back to them.

Prior research in cryptography was found to be a useful analogy. The nice thing about cryptography is that it forces the complexity onto the attacker. Most cryptographic attacks simply aren't executed; instead adversaries go after the system, whether technically or socially. However, cryptographic protocols are often proven in isolation, and when run with other protocols, their assumptions may be violated. Composition is therefore critical in moving targets as well.

Finally, are we approaching this in an ad hoc fashion or is there a more general framework to understand moving target as a methodology? If we understand our chances of being attacked, can we use that to influence how often we should “move”? This is closely linked to cyber-economics. Measurement will likely also be important.

The group generated a set of research challenges to address in developing systems that are moving targets.

1. Can we move the complexity to the attacker without impacting regular users?
2. Under what conditions does a moving target make sense?
3. Can the principles used in insurance (risk) estimation be used to determine when to randomize?
4. Where in the hardware/software stack can we apply this? Hardware reconfiguration? Software re-writing?
5. What are the general issues associated with risk management/economics?
6. Can we create better models of our adversaries so that we better understand how, when and why to move? Can we quantify any of this?

7. Can hiding, obfuscation and mimicry be effective strategies?
8. What is the role of disinformation?
9. What role might “moving target” play in denial of service prevention?

4.3 Breakout 3. Cybereconomics

Scribes: Marina Blanton, *University of Notre Dame* and Mohamed Shehab, *University of North Carolina, Charlotte*

The discussion touched on three main themes, each of which reflects a major research challenge: (1) understanding the incentives (or utility, in game-theoretic language) of regular user population as well as attackers, (2) once this understanding is gained, utilizing it in technical solutions, and (3) developing metrics for measuring security. As in several of the other discussions, point 3 (the development of good measures of security) was seen as a key obstacle to tackling the first two points.

The analogy to conventional crimefighting recurred several times throughout the discussion. There is already a lot of research on this and related topics in behavioral psychology, sociology, and related disciplines. It is not always obvious, however, how to apply those findings in computer science. In particular, the assumptions built into the design of a particular experiment are crucial. Tying work from these fields (generally, not only on the topic of crime) into mainstream research on computer security was seen as an important step towards addressing the challenges above.

Cryptographic protocols based on rational models of participant behavior also recurred several times in the discussion. They provide concrete examples of technical solutions that were explicitly motivated by incentive-compatibility considerations.

Blanton summed up the discussion, “Overall, tying economic and behavioral components into our work is going to make computer science harder, but likely in a fruitful way.”

4.4 Breakout 4. Science of Cybersecurity

Scribe: Nicholas Hopper, *University of Minnesota*

This breakout began with a discussion of what is the current state of science in cybersecurity research. The group identified that there are several theories, including those found for cryptography, formal methods, model checking, type systems, that are used by security researchers. However, much systems research is largely experimental.

The group then discussed the desire for a science for modeling adversarial behavior. It was noted that this could benefit from a science, such as a “science of war” or “science of crimefighting”, but these have not been developed.

Another topic was the development of principles for security. Questions included whether it will be possible to develop impossibility results, determine meaningful metrics, or quantify trade-offs. It was posited that security needs principles that can compose or enable composition, so there can be advancement. However, others noted that the principles may be very difficult to identify, so can start with abstractions.

In summary, the group found that the two scientific areas that emerged from this breakout were: (1) adversarial models and (2) methods to unify models. Also, it was suggested that the security

community should consider the role that other fields play in providing scientific foundations for security. However, group members argued that a science of security can be broad, encompassing more than traditional theory or leveraging of theory from other fields.