Report on the I Workshop US–Brazil on Cyber Security and Privacy

Brasilia, Brazil[1]

The goal of this series of workshops is to bring together, for the first time, American and Brazilian cyber security researchers to discuss the main challenges of the field, and to promote cross-collaboration and the building of ties among the attendees. This workshop will promote innovation in solving cyber security and privacy problems by having the top researchers in each country from several sub-fields sharing their research and brainstorming solutions for open problems in the area. It will also help researchers to share experiences in education, student recruitment, career development, and entrepreneurship, and to get involved, together with their students, in international projects.

This workshop also had the goal of helping to identify areas where American and Brazilian researchers are working in related areas that might make future collaborations possible, with the potential for future funding opportunities.

The workshop was organized in two meetings, one in Brazil and one in the US, and involved American and Brazilian cyber security researchers. The first meeting was organized by the Brazilian team lead by Prof. Priscila Solis (University of Brasilia) and happened on December 3–4, 2015, at the University of Brasilia in Brasilia, Brazil.

1 Participants

The participants were composed of security researchers at academic institutions (majority), industry researchers, program managers and staff at the National Science Foundation, Department of Homeland Security researchers, representatives of MCTI (Brazilian Ministry of Science, Technology and Innovation), including the Minister of the MCTI, the Secretary of the US Embassy in Brasilia and the Secretary Brazilian Federal Police.

1.1 American Participants

1. Kevin Butler (University of Florida)
2. Anna Squicciarini (Pennsylvania State University)
3. Guofei Gu (Texas A&M)
4. Michelle Mazurek (University of Maryland)
5. Patrick Traynor (University of Florida)

[1] Supported through NSF grant 1552059

6. Mark Tehranipoor (University of Florida)
7. Micah Sherr (Georgetown University)
8. Matthew Bishop (University of California Davis)
9. William Robertson (Northeastern University)
10. Manuel Egele (Boston University)
11. Nadia Heninger (University of Pensylvania)
12. Fabian Monrose (University of North Carolina Chapel Hill)
13. Stephanie Forrest (University of New Mexico)
14. Terry Benzel (University of Southern California)
15. Marco Carvalho (Florida Institute of Technology)
16. Scott Condie (Brigham Young University)
17. Dave Dittrich (University of Washington)
18. Bradley Huffaker (CAIDA, University of San Diego)
19. Daniela Oliveira (University of Florida)
20. Daniel J. Ragsdale (Texas A&M)
21. David Ott (Intel)
22. Parisa Tabriz (Google)
23. Fabio Tagnin (Intel)
24. Jeremy Epstein (National Science Foundation)
25. Wenjing Lou (National Science Foundation)
26. Ann Cox (Department of Homeland Security)
27. Edna N. Rodriguez Torres (Department of Homeland Security)
28. Lesley Blancas (Department of Homeland Security)

1.2 Brazilian Participants

1. Jeroen van de Graaf (Federal University of Parana)
2. Mario Alvim (Federal University of Minas Gerais)
3. Daniel Batista (University of Sao Paulo)

4. Marinho Barcellos (Federal University of Rio Grande do Sul)
5. Wagner Meira (Federal University of Minas Gerais)
6. Ricardo Custodio (Federal University of Santa Catarina)
7. Edmundo Souza e Silva (Federal University of Rio de Janeiro)
8. Paulo Licio de Geus (University of Campinas)
9. Diego Aranha (University of Campinas)
10. Joao Gondim (University of Brasilia)
11. Luis Rust (Federal University of Rio de Janeiro)
12. Jean Martina (Federal University of Santa Catarina)
13. Andre Gregio (University of Campinas)
14. Jorge H. C. Fernandes (University of Brasilia)
15. Adriano Cansian (Sao Paulo State University)
16. Raimundo Macedo (Federal University of Bahia)
17. Marcos Simplicio (University of Sao Paulo)
18. Altair Santin (Catholic University of Parana)
19. Daniel Figueiredo (Federal University of Rio de Janeiro)
20. Priscila Solis (University of Brasilia)

2 Program

Thursday, December 3, 2015

Registration and Badge Pickup

6:00 am–8:30 am Breakfast buffet (Hotel)

8:30 am Bus departs from hotel

9:00 am–10:00 am

Welcome

Celso Pansera – Ministry of State, Ministry of Science, Technology and Innovation, Brazil

Douglas P. Climan – Economic Counselor of the US Embassy in Brazil Eiiti Sato – Advisor for International Relations, University of Brasília

Manoel Augusto Cardoso da Fonseca – SEPIN, Brazilian Ministry of Science, Technology and Innovation

Jeremy Epstein – NSF Directorate for Computer & Information Sciences & Engineering and Secure and Trustworthy Cyberspace Program Director (SaTC)

Ambassador Benedicto Fonseca Filho – Director of the Division of Technological and Scientific Themes, Ministry of Foreign Relations, Brazil

10:00 am–10:30 am

Brazilian Computer Society: Big Challenges in Security

Raimundo Macedo (SBC) and Marinho Barcellos (UFRGS)

10:30 am–11:00 am

Keynote: Global Problems, Global Solutions: Prospecting Opportunities for Joint Research on Computer Security

Marinho Barcellos (UFRGS)

11:00 am–11:30 am

Break with Refreshments

11:30 am–12:15 pm

Panel: Theory × Practice Driven Cryptography: Future Challenges and Opportunities

Panelists: Nadia Heninger (University of Pennsylvania) and Jeroen van de Graaf (UFMG)

Moderator: Marcos Simplicio (University of Sao Paulo)

12:15 am–1:00 pm

Panel: Cyber Security Technology Transfer and the Education of the Next Generation of Security Professionals

Panelists: Jorge Henrique Cabral Fernandes (UnB), and Patrick Traynor (University of Florida and founder of Pindrop Security)

Moderator: Wagner Meira, UFMG – Federal University of Minas Gerais

1:00 pm–2:00 pm

Lunch

2:15 pm–3:00 pm

Panel: Research Challenges for Secure Electronic Voting

Panelists: Matt Bishop (University of California Davis) and Diego Aranha (UNICAMP – University of Campinas)

Moderator: Daniel Figueiredo, UFRJ – Federal University of Rio de Janeiro

3:00 pm–4:15 pm Breakout Sessions

Participants will choose a theme and can suggest other themes. The idea is to have group discussions on Thursday and Friday in the afternoon about challenges and visions for solutions and discoveries related to the subject. The group will produce a report. Each group will do a 10-minute presentation on Friday with the summary of the findings.

If a theme does not have enough representation, it will be removed. We are looking for session leaders and scribes!

- Cryptography and Cryptocurrency
- Software and Hardware Verification for Security
- Human Factors in Cyber Security and Privacy
- Web security across ALL your devices (PC, smart phones, IoT) – Leader: Parisa Tabriz (Google)
- Integration and Command and Control (C2) of Cyber Defenses – Leader: Marco Carvalho (Florida Institute of Technology)
- Security for Lightweight IoT Devices – Leader: David Ott (Intel)
- Cyber Crime, Malware and Intrusion Detection
- Science of Cyber Security

4:15 pm–4:30 pm

Break with refreshments

4:30 pm–5:30 pm

Rapid-Fire Cross-Collaborations

One-on-one scheduled meetings with other attendees. Pairs will be selected randomly between an American and a Brazilian researcher to provoke inter-field discussions and opportunities.

5:45 pm

Bus departs from University of Brasilia to the Hotel

7:30 pm–10:00 pm

Reception and Social Event – Rubaiyat Restaurant

Friday, December 4, 2015

Registration and Badge Pickup

6:00 am–8:30 am Breakfast buffet (Hotel)

8:30 am Bus departs from Hotel to the University of Brasilia

9:00 am–9:30 am

Opening Session

José Jair Wermann, Technical Scientific Division, Department of Homeland Security, Brazil

Jim Kurose, Assistant Director, Directorate for Computer & Information Sciences & Engineering

Douglas Maughan, Division Director, Department of Homeland Security, USA

9:30 am–10:30 am

Funding Opportunities for US–Brazil Collaborations

Wanderson Paim, CTIC, Center for Research and Development in Digital Technologies for Information and Communication

Wenjing Lou, Program Director, NSF Secure and Trustworthy Cyberspace

Program Moderator: Edmundo Souza e Silva, UFRJ – Federal University of Rio de Janeiro

10:30 am–11:00 am

Break with Refreshments

11:00 am–11:45 am

Panel: Are Privacy and Security Mutually Exclusive?

Michelle Mazurek (University of Maryland) and Mario Alvim (UFMG – Federal University of Minas Gerais)

Moderator: Jean Martina, UFSC – Universidade Federal de Santa Catarina

11:45 am–12:30 pm

Malware, Intrusion Detection and Cyber Crime

Andre Gregio, CTI-RENATO ARCHER Manuel Egele, Boston University

Moderator: Ricardo Custodio, UFSC – Federal University of Santa Catarina

12:30 pm–1:30 pm

Lunch

1:30 pm–2:45 pm

Breakout sessions continue

2:45 pm–3:15 pm

Break with Refreshments

3:15 pm–4:30 pm

Breakout session presents summary of reports

4:30 pm Closing Session

3 Breakout Sessions

Participants organized themselves into seven breakout groups discussing the following topics: (i) web security across all devices; (ii) human factors in cyber security and privacy; (iii) science of security; (iv) security for lightweight IoT devices; (v) cybercrime, malware, and intrusion detection; (vi) hardware and software verification for security; and (vii) integration and command and control (C2) of cyber defenses. All groups, except (viii), produced a report of the discussion.

3.1 Web Security Across All Devices

This group consisted of Parisa Tabriz (Google), Altair Santin (PUC/PR), Patrick Traynor (U. of Florida), Ricardo Custódio (UFSC), Noemi Rodriguez (DHS), Micah Sherr (Georgetown U.), Wenjing Lou (NSF), Bradley Huffaker (CAIDA), and Scott Condie (BYU).

For the purposes of this discussion, the group defined the web broadly as the servers, clients, and content that are connected through network communication over http. Major themes of the discussion included the threats facing web users, the tradeoff between choice and paternalism by web clients, expectations of users for privacy, ease of use when interacting with sites on the web, and how these expectations affect the safety of these individual's behavior with regards to the web.

Several risks to security, both extant and emerging, were identified. While part of the discussion centered on known risks like SQL injection and cross-site scripting, another substantial portion of the discussion focused on the interactions between default web client security behavior, user security expectations, and users' desired client functionality. It was noted that security breaches frequently occur when users' desired web client functionality and security expectations are in conflict with default client security behavior. These circumstances often lead to users “clicking through” client security warnings or switching to other, less secure web clients. Furthermore, since users' security expectations differ across mobile and desktop, browser and native, the merging of UIs across these contexts makes it increasingly difficult for users to understand the risks to which they are exposed and the risks from which they are protected. This difficulty is made even larger by the increased capabilities being provided to desktop browsers. As an example, as desktop browsers increase their capability (e.g., accessing a user's location), the surface available for adversaries to attack may increase, especially if the easiest route for users to accomplish their goal is also the least secure, a case which appears to be quite common. Understanding users' security expectations across different user contexts (desktop/mobile, browser/native) and across countries (e.g., U.S./Brazil) is important since web browsers, native apps, OSes, and especially security risks are cross-national. Solutions to these security challenges should consider the global scope of the risks involved. This and future workshops between Brazilian and U.S. researchers as well as the collaborations that could arise will be important in developing such solutions.

Several recent exploits discovered either in Brazil or the U.S. were discussed. Recent attacks on web banking users who had installed the G-Buster plugin, Superfish, and other vulnerabilities highlighted to group members the necessity for security researchers and market participants to understand the entire spectrum of risks occurring across countries. This awareness will help researchers and software providers to develop appropriate mitigation strategies in a more proactive way.

General understanding of the web differs across countries. While for many in America “the web” is Facebook, in Brazil to many “the web” is WhatsApp. While these are narrow (and in the case of WhatsApp,

perhaps unhelpful) understandings of the current web, the difference in understanding of the web across the U.S. and Brazil highlight the varying understanding of the technologies and risks associated with life on the web. The spectrum of risks associated with web participation is vast. Web users in both Brazil and the United States vary in their understanding of these potential risks. Consensus was reached on the desirability of knowledgeable market participants (browser vendors, large sites, etc.), providing at least some information about the risks involved in particular web behaviors when such information can be conveyed helpfully and easily without confusion. In particular, the "green lock" icon was cited as a potential good example of a UI feature that is largely understood to assure a particular level of security (specifically indicating a site's use of relatively up-to-date cryptographic protocols). Further research on effective cross-cultural UI security signaling could add to users' understanding of what constitutes security on the web.

The interaction between end-user and content providers' behavior was further discussed relative to sites' information disclosure policies and security. Particularly, defining customs whereby users understand a site's policies for the use of their personal data and the security risks that accompany these policies was understood to be desirable. While some discussion of both good and bad examples of privacy disclosure and the understanding of security risks were discussed, it was clear that further investigation of optimal site behavior with regard to privacy and security, across both the U.S. and Brazil, should be conducted.

The discussion of the breakout group characterized many of the challenges inhibiting a more secure web while providing avenues to make progress in solving some of these problems. An important feature of these discussions was the global nature of these challenges. Both U.S. and Brazilian researchers came away with increased understanding of the vectors of attack, security challenges, and potential paths of amelioration available in the global security environment.

3.2 Human Factors in Cyber Security and Privacy

This group consisted of Anna Squicciarini (Penn State), Michelle Mazurek (University of Maryland, College Park), Fabio Tagnin (Intel), and Jean Martina (Federal University of Parana – UFPR)

Cyber and information security are recognized as increasingly interdisciplinary topics. In particular, it is now evident that security is no longer a "technical" problem. Rather, humans play a major role in contributing and affecting the security of any system. Accordingly, the notion of sociotechnical security has recently emerged, wherein security issues arising from the interaction of users with systems are considered. The term sociotechnical security encompasses several domains (including, for example, usable security and security economics) affects not only computers but also critical infrastructures, and includes issues that may arise from accidental, careless, or malicious behavior.

3.2.1 Challenges

As technical solutions fail to fully protect users and systems against security incidents, errors, and malicious attacks, we must provide solutions that are not only technically sound but also practical and human-centered. This is a challenging task, which requires bridging the gap between historically distinct disciplines, i.e. sociology and computer science. Below is a non-exhaustive list of sample open challenges

that were highlighted during the workshop. Note that the discussion was primarily based on the perspective of users who are not actively malicious, but whose behavior may lead to security incidents/errors, damaging both themselves and others.

- People who are not security minded are now more than ever creating systems that are affected by security decisions. Yet, one shouldn't have to understand complex information and security concepts in order to operate in a secure fashion. What are appropriate expectations when dealing with users' privacy and security awareness? Users should know of consequences of their choices. Recent studies have shown that there is not a straightforward cause-and-effect/correlation when it comes to security options, making feedback options hard to compile. Is it possible to be "secure" by default without overwhelming users?

    - In particular, how can we reach users who build or manage critical systems (medical, cyberphysical, etc.) to ensure that they account for security properly in their own domains?
    - Simply suggesting more education and training may not be enough, as such users already have many things to manage in their own domains.
    - Security is generally a nonfunctional requirement in any design. As such, it is often a second-class citizen in the design/requirement space of any product. Is it possible to upgrade it to be a main requirement on any system? If not, how can we deal with security issues in an effective way? How can we guarantee that security is accounted for from the beginning and encompasses humans in all stages?

- Privacy and trust are both highly contextual, so it is hard to find useful defaults that can help address users' interpersonal differences and contextual constraints. Further, evaluating user preferences, privacy preferences, and interventions is time consuming, difficult, and inherently limited. What types of studies, designs, and systems can be implemented to help address this inherent complexity?

- Provable or perfect security is impossible in practice. What are the possible trade-offs between needs and opportunities (e.g., a system's functionality and its security capabilities)?

- Security also has some relationships with the notion of trustworthiness, both of users and of systems. If trust can bolster security, the question is how to demonstrate trustworthiness for the lay user.

Much of the discussion was devoted to detecting sociocultural differences among Brazilian and US users that can affect users' perceptions of Internet privacy, security, etc. Some differences were potentially identified. In sum, the group acknowledged that users in different cultures vary widely in their opinions of what constitutes the boundary of privacy and in their security-related behavior. Cross-cultural studies could shed light on the key differences between US and Brazil populations with respect to privacy, boundaries management, and overall security behavior. These studies could explore economic, social,

and behavioral factors that affect use and deployment of hosts, networks, and nonconventional technologies.

As an example with respect to privacy, prior privacy research has investigated the degree of privacy perceptions across different cultural dimensions. However, few have examined whether these differences translate directly into greater or less reluctance on the part of users to share personal data or to behave in a security-aware fashion. In Brazil, general data protection principles are contained in the Federal Constitution, providing a solid foundation for privacy regulations. The case in the US is much more fluid, as privacy is a critical concern of both citizens and regulatory parties, subjected to too much controversy. How these fundamental differences translate into users' actual behavior is to be investigated.

Human-centered security is an active topic of research that offers several avenues for interesting international research. Several potential directions that could benefit from international collaborations were mentioned during the workshop. Below are some examples.

- Design user-centered systems, for example, using "security personas" that capture some common behavioral patterns. Personas can reflect prototypical users in their security behavior and their evolution. Systems designers who account for many common personas may be less surprised by aberrant user behavior.

- Develop assistive technologies to help users behave in a privacy-conscious manner. Lay users are asked to make critical decisions in configuring policies, registering with authentication systems, and even setting up networks. Tools and technologies driving users to complete these tasks in a security-minded way, and reducing the cognitive load compared to users making decisions without guidance, would be very useful. These tools could draw information from users' historical data, content, current environment, and context in order to make personalized suggestions.

- Investigate effective ways to improve user feedback and promote security learning. Some basic guidelines are to make clearer connections between actions and their consequences, so that users can learn from mistakes. In addition, making security so transparent that users have no intuition for when something is wrong may backfire. To help, investigate better metaphors for presenting security concepts (e.g., PKI, keys) to users.

- Design "security-minded" systems that limit the risk of security accidents by construction.

3.3 Science of Security

This group consisted of Mário S. Alvim (Federal University of Minas Gerais, Brazil), Terry Benzel (University of Southern California, USA), Anne Cox (Department of Homeland Security, USA), Jorge Fernandes (University of Brasília), Daniel R. Figueiredo (Federal University of Rio de Janeiro, Brazil), Stephanie Forrest

(University of New Mexico, USA), and Daniel Ragsdale (Texas A&M University).

The theory and practice of computer security has too often relied on the expertise of experienced professionals. However competent these experts are, as the field of computer security matures, it becomes natural to question whether the field can be elevated from "art" to "science," and, if so, how. This group discussed these two questions during their breakout session. In particular, they reported on their adopted definition of "science" and of "security," as well as what they would consider to be basic requirements, principles, and means of a "science of security" (SoS). Finally, they enumerated a short list of research questions that they believed fell under the umbrella of SoS.

The group started by identifying key principles of "science," as opposed to "art." A scientific approach to security should (ideally) allow for the following:

- Description and categorization of phenomena;
- Search for ground truth;
- Reproducibility;
- Conceptual frameworks and definitions;
- Principles, theory, and models (qualitative and quantitative);
- Quantitative predictions.

The group recognized that there is no well-accepted definition of "security" and that many definitions confound the problem. They understood, however, that a definition of "security" should account for control of data and systems and should encompass concepts such as availability, integrity, and confidentiality (the classic AIC triad), as well some other concepts that have gained more attention, such as privacy, nonrepudiation, and anonymity.

In the establishment of SoS, a first step is to define, describe, and categorize the key aspects of security. In particular, SoS should allow for the quantification of fundamental concepts such as secrecy, privacy, and threats. SoS should also allow for the creation of models that explain and (if possible) predict behavior. These models should rest on rigorous theory and principles and should be subject to experimentation in the same way as models in other scientific areas. The group understood, however, that SoS may be subject to limits on reproducibility of experiments, as some other sciences are that must rely on historical data (e.g., astronomy) or in which the environment evolves faster than researchers can conduct experiments (e.g., social sciences).

This limitation on reproducibility reinforces the need for rigorous principles, and the principles themselves should be put to experimental test whenever possible. SoS should identify "fundamental laws" that describe the interplay of such principles. The group identified as key principles of SoS:

- Adversaries are a primary;
- Principle of abstraction;
- Emergent processes;
- Control/power/lever points;
- Dynamics/learning/evolution;
- Humans are in the equation (psych, culture, religion, belief, laws, values, norms).

The group understood that SoS would be an intersection of both the so-called "hard sciences" and "soft sciences." SoS should, hence, rely on technology and mathematics, but also on policy, laws, and social engineering to guarantee security properties.

Finally, the group identified the following topics related to research questions for SoS:

- Are there "fundamental laws" about security? If so, what are they?
    - Implications of different definitions, e.g., privacy, nonrepudiation.
    - Quantitative tradeoffs between confidentiality, availability, integrity, etc.

- How well do the mathematical assumptions made along the years fit into the real world?
    - Experimental validation of these assumptions (e.g., distribution of private keys).

- How well do existing interventions work in the real world?
    - Rigorous statistical tests of these interventions.

- How do attacks/defenses scale up with system size?
    - Scenario testing.
    - Largescale models.

- What are the interactions and tradeoffs between policy, technology, and economics?
    - Study of interdisciplinary ways of enforcing security guarantees: when is the best solution technological, and when is it social?
    - Tradeoffs between security and usability: the influence of the human factor.

3.4 Security for Lightweight IoT Devices

This group consisted of Diego Aranha (UNICAMP), Kevin Butler (University of Florida), Nadia Heninger (University of Pennsylvania), Fabian Monrose (University of North Carolina, Chapel Hill), David Ott (Intel), Edmundo Souza e Silva (Federal University of Rio de Janeiro), Mark M. Tehranipoor (University of Florida), and Jeroen van de Graaf (Federal University of Minas Gerais).

3.4.1 Defining Lightweight IoT Devices

Before discussing issues, the group first took time to define what they meant by lightweight IoT devices. The "lightweightness" of IoT devices is perhaps defined both by:

- Device properties (especially limited resources—see below);
- Usage context (e.g., automotive, medical devices).

Paradigmatic device types include:

- Sensor, actuators;
- Wearables;
- Embedded devices in countless contexts (e.g., refrigerators).

Key properties include (many apply to IoT devices more generally):

- Limited processor capabilities, memory, storage;
- Power (battery) constraints;
- Low-cost;
- Connected (usually wirelessly);
- Designed for use within broader architectures (peer devices, gateway, cloud);
- Often portable, mobile (e.g., wearables);
- Pervasive and deployed at large scale;
- Heterogeneous (platforms, usage);
- Directed functionality (vs. general use).

Almost every one of these characteristics implies a challenging new spin on security and privacy as well as the need for research. For example, how can robust security mechanisms be designed for limited processor capabilities and memory resources? How can robust security and privacy be designed for devices that have only limited power budgets due to battery constraints? And so on.

3.4.2 Lightweight IoT Devices

Research issues may broadly be divided into two categories: those focusing on the devices themselves and those focusing on the broader IoT architecture within which the device will operate. This section summarizes issues associated with the former.

- **Physical exposure.** Devices are physically exposed to attackers by their placement and/or use in public environments (as opposed to being physically secured in a data center). How can tampering and side channel attacks be prevented? How can attacks be detected, and what kind of response schemes are possible? What robust solutions are available given the inherent resource limitations?

- **Trust.** A trusted device is one that provably does exactly and only what it was designed to do. How can roots of trust be provided on lightweight devices? How can trust be established in functionality, data, software, key management, privacy?

• **Cryptography**. Cryptography is a major challenge given lightweight device constraints—low power, limited compute and memory resources, low cost, and so on. What are the key management schemes? Energy-efficient cryptography (e.g., ciphers, key size)?

• **Long-lived devices.** Many lightweight IoT devices will have a long lifetime, like major infrastructure sensors and actuators or automobile appliances and systems. What are the frameworks for reconfiguration and new ownership? Reconfigurable cryptography frameworks?

• **Short-lived devices.** On the other hand, many devices will have short lifetimes as new technology comes out and replaces it—for example, medical devices and many consumer devices. How should end-of-life be managed for data and configuration that is on the device?

• **Mission critical device security.** Another security challenge in this space is designing security for mission critical devices. For example, how do you protect a medical device regulating an individual's heart beat from attack by nearby agents?

• **Common software stacks.** Session participants imagined the emergence of common operating systems, libraries, and software stacks that become widely used across vertical domains and hardware platforms. How can strong designed-in security and privacy be established for these widely used software components and infrastructure elements? How should problems be addressed when they proliferate on a large scale when vulnerabilities emerge?

• **Supply chain security.** How can you prevent and detect cloned and fake devices from entering the supply chain and consumer market? How can counterfeit or malicious integrated circuit components be detected and prevented from appearing and proliferating in the market?

• **Device theft.** Pervasive and portable attributes of lightweight IoT devices implies considerable risk that devices could be stolen. How should the device be enabled to detect when this has happened? Schemes for responding to device theft? For example, what about personal information and data left on the device by the owner?

• **Configuration**. How could you support better user configurability of security and privacy policy on devices? Session participants agreed that current devices are not only not transparent in their operation, they lack user configurability and customization.

3.4.3 Lightweight Devices in the Broader Context of IoT Architectures

Lightweight IoT devices are generally connected to other devices and the Internet in order to be useful. As such, they should be seen as one element within the broader architecture intended by designers. In simple cases, a device may connect to a peer device (e.g., smart phone) to upload data or receive configuration. Many devices will interact with cloud applications that collect information and manage the device in various ways. Many architectures will include collections of devices, a local gateway in physical proximity to the devices, and a cloud service component.

Research issues discussed by session participants were as follows.

• **Attack surface.** How can you manage the large attack surface created by multi-device architectures? How can you contain attacks when they occur?

• **Security architectures.** There is a need for larger security architectures that comprehend differentiated roles of the device, device peers, gateways, cloud. For example, differentiated roles are needed for authentication and key management schemes.

• **Key management.** Session participants agreed that key management is one of the most important and challenging areas of research in this space.

• **Number of attack vectors.** Since devices are networked, how can you avoid the problem of any single device providing an attack vector for the broader architecture? (Weakest link problem.)

• **Device monitoring.** Schemes are needed that monitor device behavior and detect when a device is malfunctioning or misbehaving in a manner indicative of compromise.

• **Secure interoperability.** Many architectures will support heterogeneous devices, or devices that will join and leave dynamically. Security to address interoperability of devices and dynamic reconfiguration are needed. Security is also needed to identify compromised devices and isolate them.

• **Context awareness.** Context awareness is potentially a powerful tool for augmenting security in lightweight IoT device architectures. How can context awareness be incorporated into security solutions for various vertical domains?

• **Real-time response.** In many device architectures, real-time detection and response to attacks will be essential for preserving the integrity of the overall architecture. How should real-time monitoring and response be approached?

• **Navigating regulation.** How can you develop security solutions and architectures that navigate existing regulatory standards? For example, many utility verticals have long-standing regulations that may not be friendly to new technology solutions.

• **Legal considerations.** Research is needed on the interaction between security architecture design and legal liability. For example, how should liability be handled in modular architectures with multiple providers? How can you incentivize technical advancement despite potential liability?

• **Standards and best practices**. As robust security and privacy frameworks emerge, how may standards and best practices guidelines be used to positively impact the state-of-the-art for widely deployed devices? This includes both domestic and international contexts.

3.4.4 Privacy

As always, privacy issues were of great interest to session participants and generated much discussion.

Transparency. Participants agreed that device functionality should be more transparent to users and stakeholders. Included in this is the need to be more transparent on the manner in which a device collects data and on the scope of data collected.

- What are the right paradigms to foster greater transparency?

Cultural differences. Are the privacy expectations and requirements for Brazil really the same as those in America? Brazilian session participants pointed out that Brazilians often seem more unconditional trusting of technology (e.g., voting machines, fingerprint authentication) and less sensitive to privacy issues (e.g., unrestrained disclosures on Facebook).

- What cultural difference are important? How can the differences be measured?
- How might these differences influence or impact design of privacy for lightweight IoT devices?

3.5 Cybercrime, Malware, and Intrusion Detection

The group consisted of Guofei Gu (Texas A&M), André Grégio (CTI), Adriano Cansian (UNESP), Paulo Lício de Geus (UNICAMP), José Eduardo Brandão (IPEA), Wagner Meira Jr. (UFMG), and Marinho Barcellos (UFRGS).

The ubiquity of Internetconnected systems (including critical infrastructures), the ongoing "cold [cyber] war" among nations, the profit of financial gain, and the easy spread of malware variants and remote exploits motivate attackers not only to compromise systems, but to create more sophisticated cyber weapons. Besides the currently available defense mechanisms not being able to cope with the complexity of some directed attacks/weapons, monitoring systems produce a massive amount of data that need to be analyzed so that researchers and practitioners may better understand how to secure important targets. This session shed some light into new future potential attacks and tried to gather hints on what research directions should be taken to address these current and about to be threats.

The group discussed what threats could arise on computing systems in the near future. They summarized them into two broader topics regarding attack trends and novel defense techniques that should be researched, developed, and/or improved to handle these threats. New attack vectors can rely mainly on the spread of IoT devices (including sensors and vehicular networks), the users' migration from desktops to mobile devices, and the critical infrastructure accessible through the Internet. New defenses should be more proactive and automated as well as make proper use of new technologies (such as Software Defined Networking) that may help minimize the likelihood of certain types of attacks. In addition, researchers will need to have equipment and algorithms to collect and analyze security data so as to generate intelligence for defense purposes.

3.5.1 New challenges, trends, issues related to malware and cyberattacks

- **Attacks on emerging or critical infrastructures**: Malicious codes on smart grids/home/cities, CPS, IoT networks and devices, and vehicular networks are arising and pose challenges related to the diversity of devices that compose these infrastructures, fragmentation of software stacks, proprietary/complex protocols, legacy and vulnerable software still in use that may not be replaceable, and physical-based attacks (e.g., stealing power or cable TV signal).

- **Advanced Persistent Threats (APT)**: Malware is evolving towards sophisticated code whose goal is to attack silently and to keep the target in a permanent compromised state. APTs are stealthier, target-oriented, largely funded, aimed at long-term usage, and hard to detect and analyze. They frequently abuse the weakest link of security: people.

- **Cybercrime and cybercorruption**: We need to consider legal, social, economic, and cultural aspects of each "Internet space." We coined the term "cybercorruption" to address potential attacks that, for instance, make use of digital currency for bribing civil servants. It may also involve money laundering or abuse of public services in a more anonymous way. A big challenge in all these cybercrime and cybercorruption scenarios is how to collect and integrate data from several distinct sources, auditing, tracing, modeling, and measuring the entire procedure. Handling those attacks may require international collaboration and data exchange as well as participation of law enforcement.

- **Threats to emerging mobile/communication devices or channels**: New users, especially teenagers, are migrating their regular computer usage from traditional PCs to smart devices (phones or tablets). Thus, more and more personal, sensitive data will be stored on these devices and more communication and work will be performed using them (e.g., IM instead of email), raising new privacy and security issues.

3.5.2 Potential defenses: ideas/weapons

- **More powerful analytics**: Researchers will need novel techniques for data mining, machine learning, and deep learning at scale. The challenge will be how to integrate huge amounts of data from diverse and noisy sources and then analyze them to make sense of it.

- **Software-Defined Networks and Network Function Virtualization**: As new networking paradigms, SDN and NFV technologies could provide more real-time, intelligent resource allocation, defense function dispatching, and smart response for more effective defenses.

- **Cloud and Crowdsourcing-based security**: Research on these areas should explore how to provide security as a service and crowd-sourced defense mechanisms, e.g., leveraging distributed resources to process lightweight data collected from mobile devices and then trigger security actions back to the user.

- **Intelligent defense**: Adaptable and automated defense systems can lead to automatic intelligence collection and response, as well as adaptation of networks and systems regarding situational awareness. The challenge here is to avoid false-positives while providing enough security, and to bring "personalized" security, addressing the tradeoff between security and privacy.

3.6 Hardware and Software Verification for Security

The group consisted of Manuel Egele (Boston University), Roberto Gallo (KRYPTUS), William Robertson (Northeastern University), Luis Fernando Rust (Federal University of Rio de Janeiro), and Marcos Simplício (University of Sao Paulo).

3.6.1 Motivation

Cyber-security and privacy are major areas of concern in modern computing. The hardware and software that we have designed and deployed is plagued with vulnerabilities, opening the door for adversaries to steal data and disrupt computation. The increasing ubiquity of computing elements with physical sensing and actuating (i.e., cyber-physical systems) only exacerbates the importance of these threats.

Verification and related techniques for analysis and testing hold the promise of significantly improving this situation. Verification can, for instance, (i) prove whether systems adhere to their specifications, (ii) check whether systems behave as expected by users or stakeholders using informal specifications, (iii) prove the absence of known classes of vulnerabilities during the design and development phases, and (iv) provide best-effort identification of more complex vulnerability classes.

Verification is used extensively in hardware design and testing. Unfortunately, despite its promise, it has found limited use in the software domain, where many vulnerabilities continue to manifest themselves. The goal of our breakout session was to identify fundamental obstacles to applying verification techniques to both hardware and software systems where it might be feasible for collaborative efforts to make headway in the near- to medium-term. In the following, we outline several research challenges in this vein that we identified during the course of our discussion. We note that we adopted a broad definition of verification, due to the diverse backgrounds and research interests of the participants.

3.6.2 Research Challenges

Scalability vs. Precision

A classic criticism of verification is the difficulty of scaling to real-world artifacts. The usual method for increasing scalability is to increase the level of abstraction of the analysis domain. For example, in a software static analysis, instead of reasoning about all possible values of integer variables, one could abstract away these values and reason instead about their sign—a much more tractable domain. Or, when performing an inter-procedural analysis, one can reduce the precision of the context that is maintained to disambiguate function invocations, by reducing the length of the call string or by removing it completely to perform a context-insensitive analysis.

While such techniques improve scalability, they also necessarily reduce precision. Therefore, there is an inherent tension between the ability to scale an analysis or verification approach and the precision of the analysis, where the reduction in precision translates to an increase in false positives—i.e., model violations that do not in fact correspond to actual violations of whatever safety property is under consideration.

Balancing scalability with precision is thus a fundamental obstacle to verifying or analyzing a larger range of real-world artifacts. Progress on this front, e.g., by investigation of better modeling abstractions, would be a fundamental enabler for a wide range of analysis and verification approaches and a direction that might be fruitful to pursue in a joint collaboration.

Semantic Gaps

A related research challenge to the scalability vs. precision trade-off is identifying when gaps exist between the abstractions used to model hardware or software system behavior and the actual systems themselves. In particular, of concern is when existing modeling techniques do not capture essential behavior of the underlying artifact, i.e., there is a *semantic* gap. This can arise when new classes of vulnerabilities are discovered at layers of the hardware or software stack that have not yet been modeled. For example, analyzing the low-level behavior of individual software artifacts might obscure or neglect behaviors that are necessary to model and detect vulnerabilities that arise in distributed settings where collections of artifacts compute in parallel. Such mismatches can also have scalability implications, as highlighted above.

Identifying and closing these semantic gaps was therefore identified as another general direction for potential collaborations, drawing upon the unique domain expertise and operational experience of the individual participants.

Abstraction Recovery

One relevant challenge for practical application of verification and analysis techniques concerns legacy artifacts. Though we aim to produce systems that are secure-by-design against known classes of attacks, legacy hardware and software systems that cannot claim this property will continue to be used in many situations. In cases like industrial control systems and other cyber-physical systems where it is difficult to impossible to upgrade hardware and software in a timely manner, deployments can easily stretch to decades.

Therefore, a subset of verification and analysis efforts must be able to operate on legacy artifacts—e.g., binary code in the software world, or existing silicon in the hardware world. A fundamental enabling task in such scenarios is *abstraction recovery*, or the recognition of high-level abstractions in low-level representations. A canonical example is recovering source-level types of data or control abstractions (e.g., `if` conditions, `while` loops) from binary code. Recovering such abstractions can lead to direct improvements in the scalability of analysis or verification efforts, and can broaden the set of techniques that can be brought to bear.

Novel Security Properties and Capabilities

There are security properties for both the software and hardware domains that existing verification and analysis techniques can reason about. For instance, in the hardware domain, much work is focused on preventing the introduction of side channels or the detection of backdoors. As an example on the software side, a plethora of techniques exist for ensuring some degree of memory safety by, e.g., guaranteeing that no stack-based overflows exist in a program.

However, there are certainly other security-relevant properties one might like to check for which no techniques yet exist, and others where it would be beneficial to strengthen our existing techniques. Much progress has been made on these fronts in recent years, expanding the set of safety properties that can be analyzed—e.g., identifying cryptographic misuse, automated program repair, and architecture-level checkers. Developing analyses for novel security properties and improving the existing state-of-the-art are both promising areas for potential collaboration.

Harmonization of Results from Different Sources

During the discussion, it became apparent that a particular issue for Brazilian organizations involved in testing and certification of software and hardware artifacts is the difficulty of combining results produced by different organizations and analysis techniques. For instance, two organizations might analyze the same artifact using different techniques, or one organization might check the hardware while another analyzes the software of what is in principle a single system.

Therefore, we concluded that more research into the *composability* of distributed analyses, perhaps using disparate techniques, is needed to improve the utility and trustworthiness of real-world verification and analysis efforts. As examples, specific research questions that might be asked are whether the results of two given techniques must agree, whether the abstractions used permit certain types of disagreements,

or whether a given inconsistency in the results obtained should be decided by one approach in preference to another.

3.7 Federated Command and Control for Cyber Operations

The group consisted of Daniel Batista, Matt Bishop (UC Davis), Marco Carvalho (Florida Institute of Technology), Dave Dittrich (University of Washington), J. H. Fernandes (University of Brasilia), Joao Gondim (University of Brasilia), and Mauricio Leite.

The challenge problem identified and discussed by the group was centered on the study of federated cyber command and control (C2) frameworks. In the context of the discussions, Cyber C2 refers to the infrastructure, services, and logic responsible for the integration and coordination of sensors and defenses in cyber operations. The range of products and services available for coordinated operations is fairly broad, ranging from commercial cyber security orchestrators to experimental self-adaptive infrastructures for resilient command and control. Our challenge problem focuses on the coordination and interactions between different enclaves managed by their own C2 infrastructures.

Motivation

The concept of Federated C2 was motivated by prior experiences of some of the members of the group who have worked on the coordination of cyber defense operations across multiple organizations in large-scale events (such as Rio 20, in 2012), and also by the experience of group members involved in related projects sponsored by the DHS S&T. In all cases discussed, the need for a Federated C2 emerged from the characteristic interdependency between most modern enterprise environments. As illustrated in Figure 1, most critical systems and infrastructures today depend on information and services provided by different enterprises and enclaves that are likely operating under different administrative domains with different operational procedures and often with different goals and priorities.

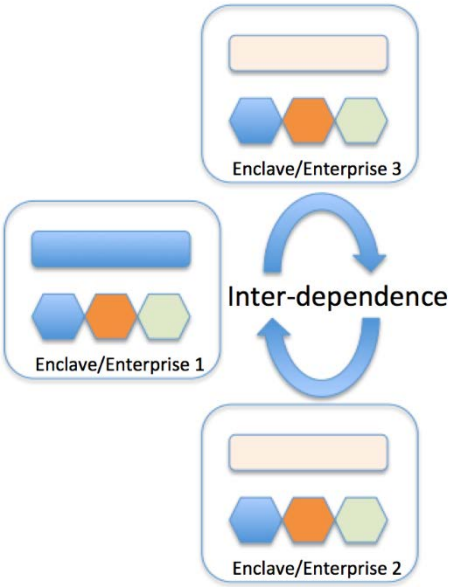

While independent from a monitoring and management perspective, the enclaves illustrated in Figure 1 often have similar vulnerabilities and common adversaries. Furthermore, the very interdependency that defines the business process of the different organizations is often exploited as a new attack vector.

The Federated C2 Challenge Problem

It is intuitive that the coordination and the sharing of information between the enclaves illustrated in Figure 1 could be beneficial to the defense of each individual enclave and to the collective defense of the system. The challenge lies in the technical design and implementation, and the administrative and operational policies, procedures, and constraints of the infrastructure, services, and capabilities required

to enable the practical formation and maintenance of the Federated C2.

A Federated C2 would allow each enclave to determine what information to share under different contexts. Furthermore, it would also allow enclaves to improve their individual defensive postures and coordinate their responses to a specific threat.

As part of our group discussions, we identified some of the following specific challenges that would need to be addressed to enable such a capability.

- The nature of the coordination
  - The coordination of Federated members can be achieved through multiple ways, for example, through the establishments of incentives or regulatory means. There is a need to investigate effective ways to facilitate and incentivize the interactions between federation members.
  - Is there a fundamental limit to the ability to enable trusted command and control beyond a single Federated trust group? That is to say, is sharing of operationally actionable information as far as federation members are willing to go, reserving decision-making about command and control actions to individuals within each federation member organization?
- The semantic problem
  - It is necessary and important to investigate practical and scalable approaches to facilitate on-demand information sharing between environments that use incompatible models to represent specific data and events.
- The trust problem
  - There is a need to investigate effective ways to establish and maintain trust among participating members and to share information between multiple Federated trust groups (i.e., operate more like a Confederated group of groups rather than a single large Federated group of individuals).
  - As the size of a trust group increases, the cohesiveness of trust bonds decreases between group members. When sharing sensitive data form one trust group to another (and possibly again relayed to yet other trust groups) the issue of *transitivity of trust* comes up, again impeding the motivation to share sensitive data of operational value.
  - As a means to overcome certain scenarios where granting access to another party to sensitive data is problematic, novel methods of performing *private set intersections* can overcome the limits of trust.
- The combination, disambiguation, and/or prioritization of local vs. global goals
  - Related to the incentives issue, there is also a need to investigate mechanisms to allow federation members to operate across multiple Federated domains (overlapping federations), possibly under different incentive models or regulatory constraints.
  - Alternatively, is there a mixed Federated/Confederated model that can take advantage of mechanisms such as *attribute-based encryption* or individuals having multiple vetted identities that simultaneously satisfies both local and global objectives? For example, a distributed and hierarchical mixed model of multiple Federated groups at the state/local level, or within sector-specific verticals, operating as a Confederated meta-group, or distributing encryption keys to individuals through a federal level vetted mechanism (e.g., the DHS Cyber

Information Sharing and Collaboration Program [CISCP] portal) that would support sharing of encrypted information usable for command and control through lower-trust channels?

Conclusions

While clearly recognizing the challenge associated with the proposed problem, the team agreed on the value that such a capability could provide to improve the resilience and security of complex infrastructures involving multiple operation domains. There is ongoing research in closely related areas that could help support a possible research agenda towards that goal.