

SoS Lablets



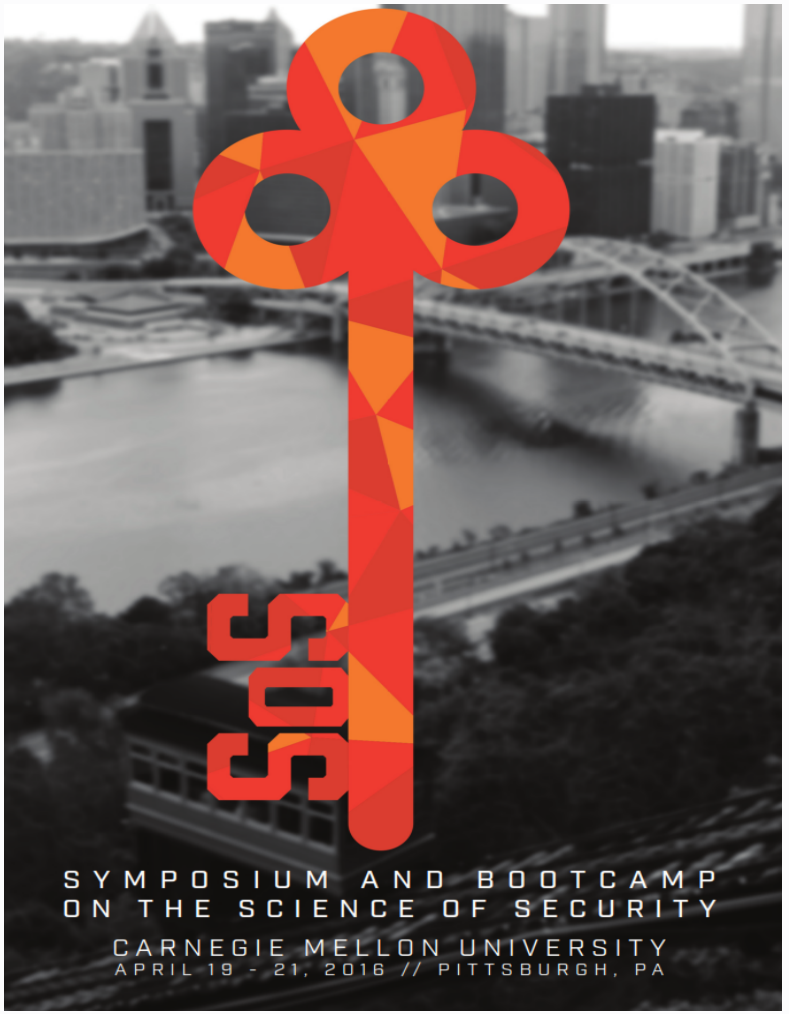
The Carnegie Mellon
Science of Security Lablet

Bill Scherlis (PI), Jonathan Aldrich (co-PI), Travis Breaux (co-PI)

Quarterly Meeting
NSA
2-3 November 2016

The CMU Science of Security (SoS) Lablet

- The team
- Mission
- Significance of SoS
- Hard Problems
- Synergetic benefits
- CMU project portfolio



The team

- The CMU research team
  - David Garlan
  - Anupam Datta
  - Andre Platzer
  - Alessandro Acquisti
  - Christian Kästner
  - Travis Breaux – co-PI
  - Lorrie Cranor
  - Limin Jia
  - Rahul Telang
  - Bill Scherlis - PI
  - Jürgen Pfeffer
  - Jonathan Aldrich - co-PI
  - Bradley Schmerl
  - Nicolas Christin
  - *More than seven academic departments, three colleges*
  - *Diverse disciplines: Mathematical logics and models, software architectures/frameworks, graph-theoretic network analytics, human subjects studies, CPS, policy modeling, software devt and evaluation*
- Partner universities
  - Cornell (Dexter Kozen)
  - GMU (Sam Malek)
  - UC Berkeley (Serge Egelman)
  - U Pittsburgh (Scott Beach)
  - Wayne State (Marwan Abi-Antoun),
  - U Nebraska (Matt Dwyer)
  - UTSA (Jianwei Niu)

Other SoS Lablets
- UIUC
- NCSU
- U Md

Mission concept

(1) Advance specific areas of **cybersecurity research**

- **Scalability and composability**
- **Policy-governed secure collaboration**
- Predictive security metrics
- **Resilient architectures**
- **Human behavior**

Not comprehensive

(2) Advance the **scientific coherence** of the multidisciplinary body of cybersecurity technical results

- Methods
- Validation
- Productivity

(3) Engage and broaden the cybersecurity **technical community**

- Facilitate community and educational engagement
- Workshops and conference events: HotSoS conference

Significance of **SoS**

- **Premises – operating environment**
  - Growth in urgency and criticality of cybersecurity
    - Common vulnerable tech base
  - Multidisciplinary character of cybersecurity
    - Diverse technical domains: Biometrics, human behavior, crypto math, protocol analysis, language foundations, logics and models, systems architecture, threat analysis, cyber-physical models, networking, hardware, API design, process and measurement, etc.
    - Adversary escapes the abstractions (side channel), so we must continually broaden and make the side channels more expensive
  - Diverse scientific approaches
    - Mathematically based theory
    - Data-driven empirical studies
    - Empirical behavioral studies: observational and interventional
  - Rapid pace of change
    - Software and hardware technical innovation
    - Threat capability and motivation

Significance of **SoS**

- **Premises – engineering environment**
  - Diverse points of potential intervention
    - Requirements, architecture, development, operations, sustainment
    - Evaluation and measurement
  - Complexity and interconnection
    - Rich and diverse supply chains
    - Socio-technical ecosystems
      - Framework-and-app models
      - *Payloads and platforms*
  - Product families – in time and space
    - Configurations (e.g., Linux on Android)
    - Ongoing evolution and need to rapid recertification

Consequence: Three part approach to SoS

1. *Address the most challenging **Hard Problem** areas*

2. *Advance the **process and methods by which science is done***

3. *Engage with the broader research and technical **community** to address these goals*

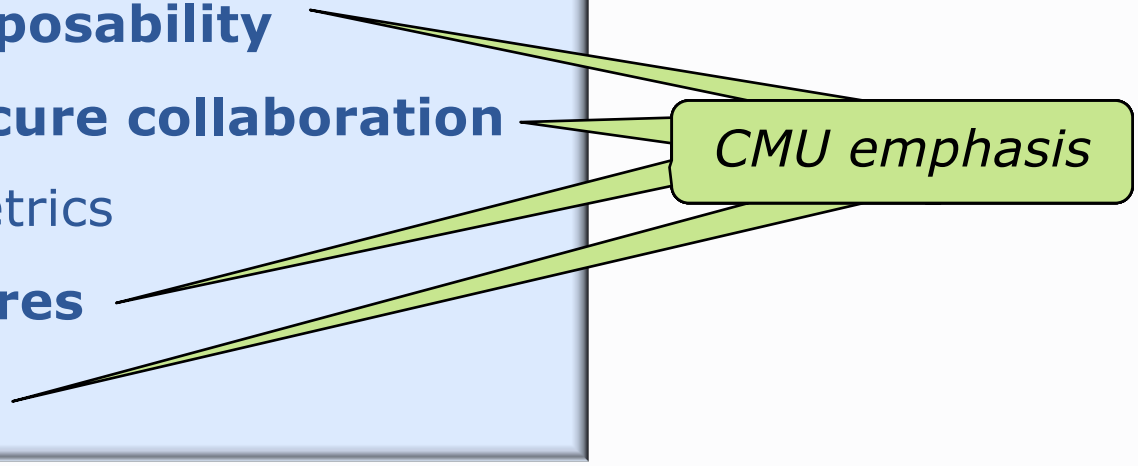
Consequence: Three part approach to SoS

1. Address the most challenging Hard Problem areas

- ***We focus primarily on HP 1 and 5***
- ***Strong activity related to HP 2 and 4.***
- ***(Details in HP Report)***

Five Hard Problems in the Science of Security

1. **Scalability and composability**
2. **Policy-governed secure collaboration**
3. Predictive security metrics
4. **Resilient architectures**
5. **Human behavior**



Not comprehensive:

Focus on the engineering and evaluation of systems

Selection criteria for the problems

- High level of technical challenge
- Significant operational value
- Likelihood of benefiting from emphasis on scientific research methods and improved measurement capabilities

Consequence: Three part approach to SoS

2. ***Advance the process and methods by which science is done***

   - ***Explicit focus on methods***
   - ***Synergies in the Lablet approach***

Synergies in the Lablet approach

- **Data** *meets* **models**
  - E.g., social network structures, developer usability, end-user usability, API complexity
- **Semantics-based approaches** *meet* **real engineered systems**
  - E.g., hypervisors, Web apps, framework+apps, large components
- **Empirical science** (data, people) *meets* **mathematical reasoning**
  - E.g., language design, API design, model design, tool design

CMU Projects (9 projects)

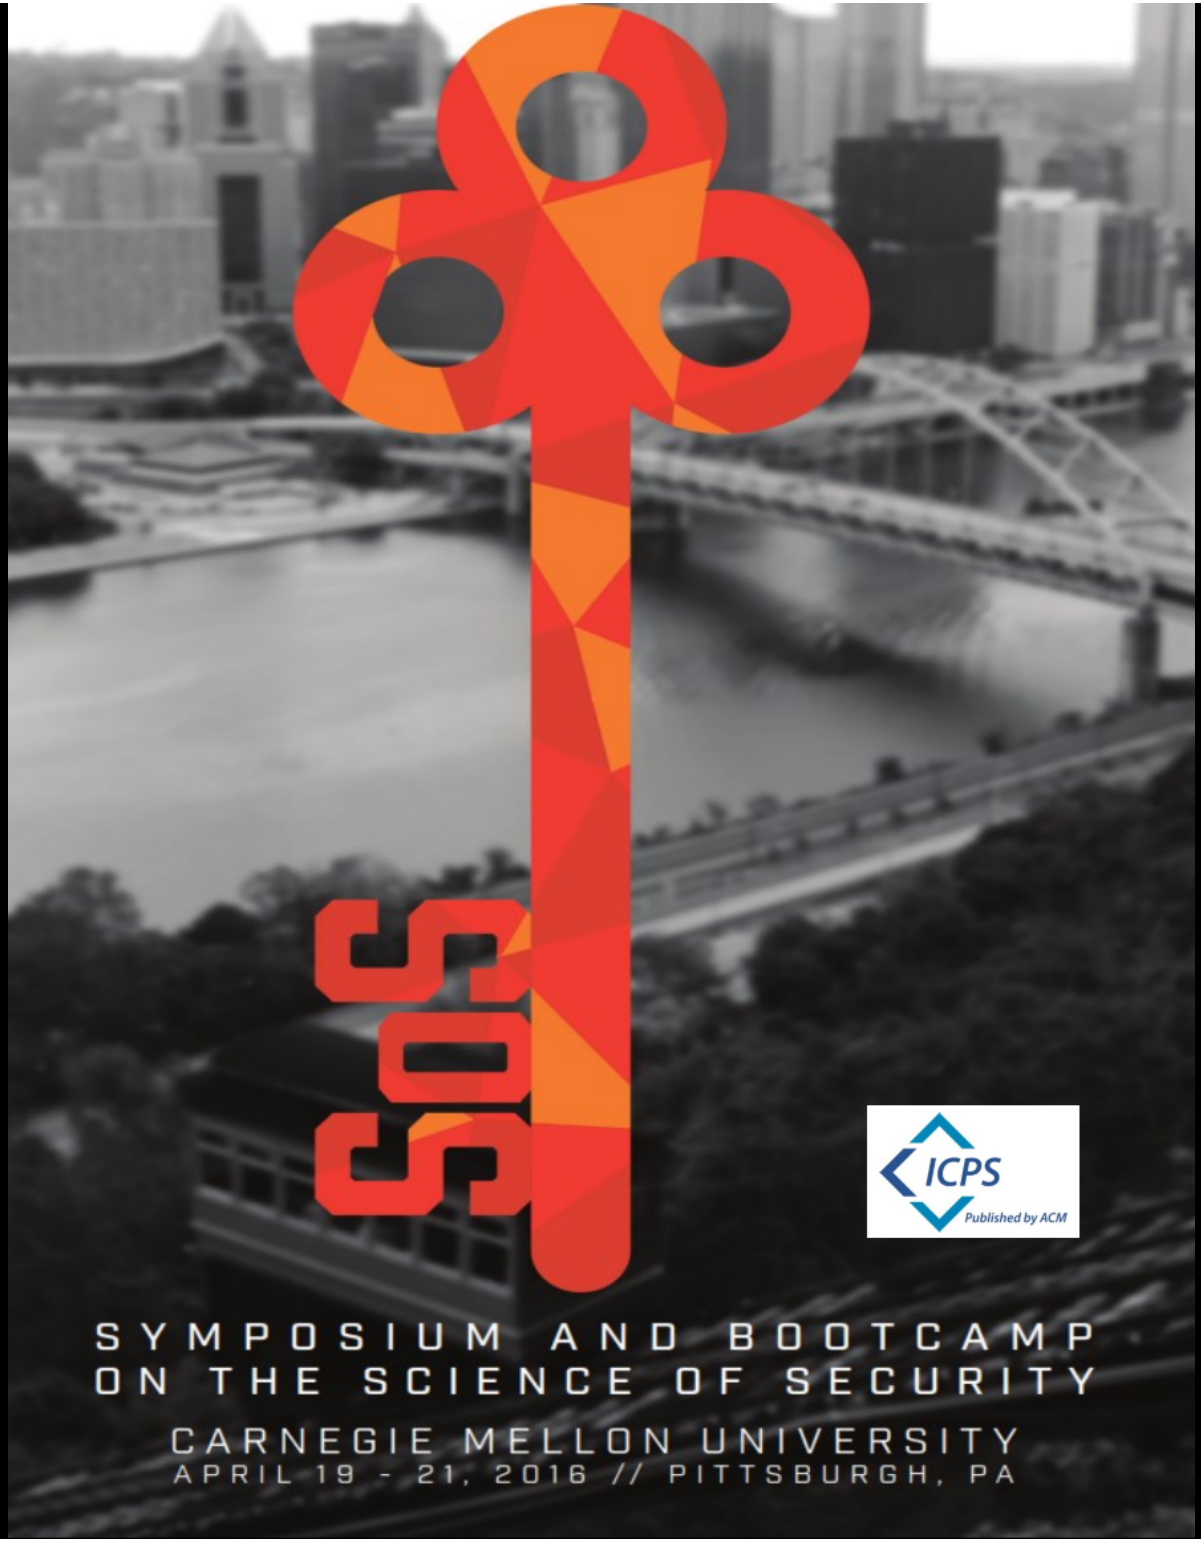
	Com	Pol	Met	Res	Hum		
1	X		X	X		Aldrich, Garlan, Malek (GMU), Abi-Antoun (Wayne State)	Frameworks, APIs, and Composable Security Models
2	X		X			Pfeffer, Kästner	Limiting Recertification in Highly Configurable Systems: Analyzing Interactions and Isolation among Configuration Options
3	X			X		Platzer, Kozen (Cornell)	Security Reasoning for Distributed Systems with Uncertainties
4	X	X				Datta, Jia	Compositional security
6	X				X	Aldrich	A Language and Framework for Development of Secure Mobile Applications
7	X		X	X		Pfeffer, Garlan, Schmerl	Multi-model run-time security analysis
8	X				X	Aldrich, Dwyer (Nebraska)	**Race Vulnerability Study and Hybrid Race Detection**
9	X	X			X	Breaux, Niu (UTSA)	**Usable Formal Methods for the Design and Composition of Security and Privacy Policies**
10			X		X	Cranor, Acquisti, Christin, Telang, Egelman (Berkeley), Beach (U Pitt)	**Understanding user behavior when security is a secondary task**

The Carley (#11) and Harper (#5) projects are concluded.

Consequence: Three part approach to SoS

3. Engage with the broader research and technical community to address these goals

- ***HotSoS 2016 Conference***
- ***Conference on Safety and Control for AI (with OSTP)***
- ***Workshop on Safety and Control for AI (with Lablet)***


SOS
ICPS
Published by ACM
SYMPOSIUM AND BOOTCAMP
ON THE SCIENCE OF SECURITY
CARNEGIE MELLON UNIVERSITY
APRIL 19 - 21, 2016 // PITTSBURGH, PA

HotSoS focus

- **Science of Security Hard Problems**
  - **Scalability and composability** in the construction of secure systems
  - **Policy-governed collaboration** in handling data across different domains of authority for security and privacy protection
  - **Predictive security metrics** to guide choicemaking in security engineering and response
  - **Resilient architectures** that can deliver service despite compromised components
  - **Human behavior**, modeling users, operators, and adversaries to support improved design and analysis

- **Scientific Methods**
  - Data gathering and analysis
  - Experimental methods
  - Mathematical modeling and reasoning











Program Committee

- Marwan Abi-Antoun
- Alessandro Acquisti
- Jonathan Aldrich
- Michael Bailey
- Lujo Bauer
- Travis Breaux
- David Brumley, *co-chair*
- Kathleen Carley
- Michel Cukier
- Serge Egelman
- David Garlan
- Limin Jia
- Jonathan Katz
- Stuart Krohn
- Michael Maass
- Sam Malek
- David Malcolm Nicol
- Andre Platzer
- William Sanders
- William Scherlis, *co-chair*
- Munindar Singh
- Witawas Srisa-An
- Joshua Sunshine
- Adam Tagert

Additional Reviewers

- Atul Bohara
- Carmen Cheh
- Hongying Du
- Fawazt Ahmed
- Joshua Garcia
- Ozgur Kafali
- João G. Martins
- Uttam Zawadzki
- Erik Thakore

Additional thanks

- Sharon Blazevich
- Monika DeReno
- Tom Pope
- Josh Quicksall
- Bernadette Dayak
- Katie Dey
- *Our research sponsors*

HotSoS 2016

TUESDAY, APRIL 19	
6:00pm - 9:00pm	Evening Reception at Phipps Conservatory
WEDNESDAY, APRIL 20	
8:00am - 9:00am	Breakfast at Cohon University Center
SESSION 1	
9:00am - 10:30am	Welcome and Introduction **Keynote: Lorrie Cranor** *Adventures in Usable Privacy and Security: From Empirical Studies to Public Policy* Paper: Hui Lin, Homa Alemzadeh, Daniel Chen, Zbigniew Kalbararczyk, and Ravishankar Iyer *Safety-critical Cyber-Physical Attacks: Analysis, Detection, and Mitigation*
10:30am - 11:00am	Break
SESSION 2	
11:00am - 12:30pm	Paper: Aron Laszka, Waseem Abbas, S. Shankar Sastry, Yevgeniy Vorobeychik, and Xenofon Koutsoukos *Optimal Thresholds for Intrusion Detection Systems* Paper: Atul Bohara, Uttam Thakore, and William Sanders *Intrusion Detection in Enterprise Systems by Combining and Clustering Diverse Monitor Data* Paper: Jeffrey Carver, Morgan Burcham, Sedef Akinli Kocak, Ayse Bener, Michael Felderer, Matthias Gander, Jason King, Jouni Markkula, Markku Oivo, Clemens Sauerwein, and Laurie Williams *Establishing a Baseline for Mesauring Advancement in the Science of Security – an Analysis of the 2015 IEEE Security & Privacy Proceedings*
12:30pm - 2:00pm	Posters and Lunch

HotSoS 2016

SESSION 3	
2:00pm - 3:00pm	**Keynote: Greg Shannon** *Science Challenges in the Federal Cybersecurity R&D Strategic Plan*
3:00pm - 3:30pm	Posters and Break
SESSION 4	
3:30pm - 5:20pm	Paper: Bradley Potteiger, Goncalo Martins, and Xenofon Koutsoukos *Software and Attack Centric Integrated Threat Modeling for Quantitative Risk Assessment* **Tutorial: Hanan Hibshi** *Systematic Analysis of Qualitative Data in Security*
5:30pm - 6:00pm	Break
6:00pm - 9:00pm	Dinner at Carnegie Museum of Natural History
THURSDAY, APRIL 21	
8:00am - 9:00am	Breakfast at Cohon University Center
SESSION 5	
9:00am - 10:30am	**Keynote: Christos Faloutsos** *Anomaly Detection in Large Graphs* Paper: Kevin Bush, Antonio Roque, and Christopher Degni *Security is About Control: Insights from Cybernetics* Paper: Zhenqi Huang, Yu Wang, Sayan Mitra, and Geir Dullerud *Controller Synthesis for Linear Time-varying Systems with Adversaries*
10:30am - 11:00am	Break

HotSoS 2016

SESSION 6	
11:00am - 12:20pm	**Tutorial: Tao Xie, William Enck** *Text Analytics for Security*
12:20pm - 1:20pm	Lunch
SESSION 7	
1:20pm - 3:10pm	**Keynote: Matt Briggs** *The Threat Environment* Paper: Carl Pearson, Allaire Welk, William Boettcher, Roger Mayer, Sean Streck, Joseph Simons-Rudolph, and Christopher Mayhorn *Difference in Trust Between Human and Automated Decision Aids* Paper: Phuong Cao, Eric Badger, Zbigniew Kalbarczyk, Ravishankar Iyer *A Framework for Generation, Replay and Analysis of Real-World Attack Variants*
3:10pm - 3:30pm	Ice Cream Social


SĀF
ART
INT
WORKSHOP ON
SAFETY & CONTROL FOR
ARTIFICIAL INTELLIGENCE

The workshops on Safety and Control for AI

- Carnegie Mellon University 27-28 June 2016
  - 27 June – Exploratory Technical Workshop, co-sponsored with the SoS Lablets
  - 28 June – Public Workshop, co-sponsored with the White House Office of Science and Technology Policy (OSTP)

- Goal:
  - How can we advance a practice of ***AI Engineering*** with safety and security considerations integrated from the outset?
  - How to engineer AI-intensive systems with ***safety built in?***

June 27, 2016 – Carnegie Mellon Exploratory Workshop

Context	Ed Felten (White House OSTP) Bill Scherlis (CMU) Eric Horvitz (Microsoft)
Perspective	Andrew Moore (CMU) Richard Mallah (Future of Life Inst) Tom Mitchell (CMU
Focus	Dario Amodei (Google Brain) Claire Le Goues (CMU) Robert Rahmer (IARPA)

June 28, 2016 – White House OSTP and Carnegie Mellon Workshop

Introduction	Ed Felten (White House OSTP) Bill Scherlis (CMU)
AI applications, safety requirements, big ideas	Manuela Veloso (CMU) John Launchbury (DARPA) Jason Matheny (IARPA) Andy Grotto (White House NSC)
AI Algorithms, Intrinsic Safety, Explanations, Evaluations	Claire Tomlin (UC Berkeley) Tom Dietterich (Oregon State) Emma Brunskill (CMU) Michael Littman (Brown)
Mathematical Models and Reasoning	Jeannette Wing (Microsoft) Kathleen Fisher (Tufts) Anupam Datta (CMU) Sarah Loos (Google)
Systems and Safety Engineering	Mike Wagner (CMU) Drew Bagnell (Uber) Reid Simmons (NSF) Brian Murray (ZF) Doug Schmidt (Vanderbilt)

Setting Context

- What are the **mission contexts**
  - What are the systems **doing** in the world?
- What are they ways we **manifest AI** capabilities?
  - Statistical machine learning, but also planning, language processing, search, etc.
- What is the **dimensionality of safety** for systems with embedded AI?
  - Analogous challenge to defining the *dimensionality of security*
  - Scope, Framing, modeling, reasoning
- How might we think about **designing in** safety?
- How might we define **measures** to support evaluation with respect to safety attributes?

Triangulation 1 – AI Algorithms, Intrinsic Safety, Explanations, Evaluations

- Can we usefully **pierce the opacity** of modern systems and create explanations?
- How can develop better -- **objective fns** -- to assess what is learned?
- How can we manage the learning process, including **selection of data**?
- How can we avoid pitfalls due to **frame problems** --
  - That is, how much of the operational reality is embraced by our abstractions and models?

Triangulation 2 – Mathematical models and reasoning

- How can we **model** the many kinds of safety-related properties?
- How do we **reason** within the models?
- What are we **not** modeling?
  - The choice of models is important -- this goes back to that framing issue.

(This area has its roots in early symbolic AI, theorem proving, formal methods, verification, etc.)

Triangulation 3 – Safety engineering and systems engineering

- What are useful architecture and design concepts?
  - What architectural features can bound the consequences of AI accidents?
- How can we mitigate the framing problem without modeling the entire world?
- What about traditional systems ideas such as big red buttons and hot back-up?
- How can we build on established systems engineering and safety engineering practices?
  - Robot and autonomy people know the limits of their practices.
- (Except in limited domains, we are aren't very good at this. How can we improve?)

Consequence: Three part approach to SoS

1. *Address the most challenging **Hard Problem** areas*

2. *Advance the process and **methods by which science is done***

3. *Engage with the broader research and technical **community** to address these goals*

The CMU Science of Security (SoS) Lablet

- The team
- Mission
- Significance of SoS
- Hard Problems
- Synergetic benefits
- CMU project portfolio

