TRUST
Team for Research in Ubiquitous Secure Technology

Annual Report
(2008 – 2009)

May 18, 2009

TRUST is funded by the National Science Foundation
(award number CCF-0424422)

Berkeley UNIVERSITY OF CALIFORNIA
Carnegie Mellon
Cornell University
MILLS COLLEGE
San José State UNIVERSITY
SMITH COLLEGE
STANFORD UNIVERSITY
VANDERBILT UNIVERSITY

TABLE OF CONTENTS

1 GENERAL INFORMATION 4

1.1 SUMMARY 4
1.2 NEW CENTER FACULTY 5
1.3 REPORT POINT OF CONTACT 5
1.4 CONTEXT STATEMENT 6

2 RESEARCH 8

2.1 GOALS AND OBJECTIVES 8
2.2 PERFORMANCE AND MANAGEMENT INDICATORS 8
2.3 CURRENT AND ANTICIPATED PROBLEMS 8
2.4 RESEARCH THRUST AREAS 8
2.4.1 Financial Infrastructures 10
2.4.2 Health Infrastructures 20
2.4.3 Physical Infrastructures 23
2.5 RESEARCH METRICS/INDICATORS 29
2.6 NEXT REPORTING PERIOD RESEARCH PLANS 29
2.6.1 Financial Infrastructures 29
2.6.2 Health Infrastructures 39
2.6.3 Physical Infrastructures 40

3 EDUCATION 44

3.1 GOALS AND OBJECTIVES 44
3.2 PERFORMANCE AND MANAGEMENT INDICATORS 44
3.3 CURRENT AND ANTICIPATED PROBLEMS 45
3.4 INTERNAL EDUCATION ACTIVITIES 45
3.5 PROFESSIONAL DEVELOPMENT ACTIVITIES 57
3.6 EXTERNAL EDUCATION ACTIVITIES 60
3.7 ACTIVITIES TO INTEGRATE RESEARCH AND EDUCATION 62
3.8 EDUCATION METRICS/INDICATORS 63
3.9 NEXT REPORTING PERIOD EDUCATION PLANS 64

4 KNOWLEDGE TRANSFER 66

4.1 GOALS AND OBJECTIVES 66
4.2 PERFORMANCE AND MANAGEMENT INDICATORS 67
4.3 CURRENT AND ANTICIPATED PROBLEMS 67
4.4 KNOWLEDGE TRANSFER ACTIVITIES 67
4.5 OTHER KNOWLEDGE TRANSFER OUTCOMES 73
4.6 KNOWLEDGE TRANSFER METRICS/INDICATORS 73
4.7 NEXT REPORTING PERIOD KNOWLEDGE TRANSFER PLANS 74

5 EXTERNAL PARTNERSHIPS 75

5.1 GOALS AND OBJECTIVES 75
5.2 PERFORMANCE AND MANAGEMENT INDICATORS 75
5.3 CURRENT AND ANTICIPATED PROBLEMS 75
5.4 EXTERNAL PARTNERSHIP ACTIVITIES 75
5.5 OTHER EXTERNAL PARTNERSHIP OUTCOMES 77
5.6 EXTERNAL PARTNERSHIP METRICS/INDICATORS 78
5.7 NEXT REPORTING PERIOD EXTERNAL PARTNERSHIP PLANS 78

6 DIVERSITY 79

6.1 GOALS AND OBJECTIVES 79
6.2 PERFORMANCE AND MANAGEMENT INDICATORS 79

6.3 CURRENT AND ANTICIPATED PROBLEMS 80
6.4 DIVERSITY ACTIVITIES 80
6.5 DIVERSITY ACTIVITY IMPACT 80
6.6 DIVERSITY METRICS/INDICATORS 81
6.7 NEXT REPORTING PERIOD DIVERSITY PLANS 83

7 MANAGEMENT 84

7.1 ORGANIZATIONAL STRATEGY 84
7.2 PERFORMANCE AND MANAGEMENT INDICATORS 84
7.3 MANAGEMENT METRICS/INDICATORS 85
7.4 CURRENT AND ANTICIPATED PROBLEMS 85
7.5 MANAGEMENT AND COMMUNICATIONS SYSTEM 85
7.6 CENTER ADVISORY PERSONNEL 86
7.7 CENTER STRATEGIC PLAN CHANGES 87

8 CENTER-WIDE OUTPUTS AND ISSUES 88

8.1 CENTER PUBLICATIONS 88
8.1.1 Peer Reviewed Publication 88
8.1.2 Journal Articles 92
8.1.3 Books and Book Chapters 93
8.1.4 Non-peer Reviewed Publications 93
8.2 CONFERENCE PRESENTATIONS 93
8.3 OTHER DISSEMINATION ACTIVITIES 96
8.4 AWARDS AND HONORS 97
8.5 GRADUATES 98
8.6 GENERAL KNOWLEDGE TRANSFER OUTPUTS 98
8.7 INSTITUTIONAL PARTNERS 98

9 INDIRECT/OTHER IMPACTS 100

9.1 INTERNATIONAL ACTIVITIES 100
9.2 OTHER OUTPUTS, IMPACTS, AND INFLUENCES 100

10 ATTACHMENTS 101

1 GENERAL INFORMATION

1.1 Summary

Date Submitted	May 18, 2009
Reporting Period	June 1, 2008 – May 31, 2009
Name of the Center	Team for Research in Ubiquitous Secure Technology
Name of the Center Director	S. Shankar Sastry
Lead University	University of California, Berkeley
Contact Information	
Address	320 McLaughlin Hall
Phone Number	510-642-5771
Fax Number	510-642-9178
Email Address of Center Director	sastry@coe.berkeley.edu
Center URL	http://www.truststc.org/

Below are the names of participating Center institutions, their roles, and (for each institution) the name of the contact person and their contact information at that institution.

Institution Name	Carnegie Mellon University, Adrian Perrig
Address	2110 Collaborative Innovation Center Pittsburgh, PA 15213
Phone Number	412-268-2242
Fax Number	412-268-6779
Email Address of Center Director	adrian@ece.cmu.edu
Role of Institution at Center	Carnegie Mellon is a lead research, education, and outreach partner.

Institution Name	Cornell University, Stephen Wicker
Address	386 Rhodes Hall Ithaca, NY 14850
Phone Number	607-255-8817
Fax Number	607-255-9072
Email Address of Center Director	wicker@ece.cornell.edu
Role of Institution at Center	Cornell University is a lead research, education, and outreach partner.

Institution Name	Mills College, Almudena Konrad
Address	CPM 204 Oakland, CA 94613
Phone Number	510-430-2201
Fax Number	510-430-3314
Email Address of Center Director	akonrad@mills.edu
Role of Institution at Center	Mills is an outreach partner to encourage greater female participation in engineering.

Institution Name	San Jose State University, Sigurd Meldal
Address	ENGR 284 San Jose, CA 95192
Phone Number	408-924-4151
Fax Number	408-924-4153
Email Address of Center Director	smeldal@email.sjsu.edu
Role of Institution at Center	SJSU is a lead education partner to spread curriculum and encourage greater underrepresented minority participation in engineering.

Institution Name	Smith College, Judith Cardell
Address	Clark Science Center, EGR 105b, Northampton, MA 01063
Phone Number	413-585-4222
Fax Number	413-585-3827
Email Address of Center Director	jcardell@smith.edu
Role of Institution at Center	Smith is a research and outreach partner to encourage greater female participation in engineering.

Institution Name	Stanford University, John Mitchell
Address	Gates Building 4B-476 Stanford, CA 94305-9045
Phone Number	650-723-8634
Fax Number	650-725-7411
Email Address of Center Director	mitchell@cs.stanford.edu
Role of Institution at Center	Stanford is a lead research, education, and outreach partner.

Institution Name	Vanderbilt University, Janos Sztipanovits
Address	2015 Terrace Place VU Station B 356306 Nashville, TN 37235-6306
Phone Number	615-343-7572
Fax Number	615-343-6702
Email Address of Center Director	janos.sztipanovits@vanderbilt.edu
Role of Institution at Center	Vanderbilt is a lead research, education, and outreach partner.

1.2 New Center Faculty

Please see Appendix A for biographical information on each new faculty member added to the Center during this reporting period.

1.3 Report Point of Contact

Below is the name and contact information for the primary person to contact with any questions regarding this report.

Name of the Individual	Larry Rohrbough
Center Role	Executive Director
Address	337D Cory Hall Berkeley, CA 94720-1774
Phone Number	510-643-3032
Fax Number	510-642-2718
Email Address	larryr@eecs.berkeley.edu

1.4 Context Statement

The Team for Research in Ubiquitous Security Technology (TRUST) was created in response to a growing sense of urgency in dealing with all aspects of cyber security as it affects society. The role and penetration of computing systems and networks in our societal infrastructure continues to grow and their importance to societal safety and the security has never been greater. Beyond mere connection to the internet and access to global resources, information systems form the backbone of our nation's financial services and electronic commerce, are used for controlling critical infrastructures such as power, water, and telecommunications, and enable the rapid evolution in healthcare toward enhanced services increasingly supported by the digital storage of and instant access to patient health and medical data

That said, many such computing and control systems remain untrustworthy. Waves of viruses and worms sweep the Internet and exhibit increasing virulence and rate of speed that is also directly proportional to their growing ease of deployment. Issues affecting privacy are poorly understood and, when they are understood, are often not sufficiently addressed in system design and development. Security is generally inadequate, and some speak of a "market failure" in the domain. Broader issues of software usability, reliability and correctness remain challenging. Industry stakeholders are unable to recruit new employees adequately trained in these technologies. Society is placing computers into critical roles, although they do not meet the requirements of trust.

TRUST is composed of several universities that have joined forces to organize a multifaceted response to these issues. TRUST represents the strongest and most diverse engagement in the area of trusted systems ever assembled. TRUST recognized the breadth of the problems and has combined fundamental science with a broader multidisciplinary focus on economic, social, and legal considerations as well as a substantial education mission. TRUST is enabling dialog with stakeholders whose needs simply cannot be approached in a narrower and purely technical manner or by any single research group. As such, TRUST acts as an intermediary between policy makers and society at large on the one hand, and researchers, academics, and industrial providers of services and technology on the other.

TRUST seeks to achieve its mission through research as well as a global policy for engaging in education of society as a whole. This annual report of TRUST details the experience of the Center along many dimensions—research, education, diversity, and knowledge transfer.

In research, TRUST has achieved success along several fronts and is addressing fundamental scientific and technological problems and advancing the state-of-the-art in a number of areas: security and privacy issues associated with the rapidly increasing use of electronic media for the archival and access of patient medical records; web authentication, end-user privacy, next-generation browser security, malware detection, and improved system forensic techniques to combat online attacks; application defenses for network-level intrusions and attacks including compromised and malfunctioning legacy applications, viruses, worms, and spyware; incentives for research, investment, policies, and procedures for technology that enhance system security, privacy, and trustworthiness; secure embedded sensor networks for large-scale applications critical to the nation's economy, energy, security, and health; and techniques that ensure trustworthy computing by securing hardware, improving software robustness, and increasing the survivability of critical systems.

In education, TRUST is leveraging an existing learning technology infrastructure to quickly enable TRUST courseware and material to be assembled, deposited in a repository, and adapted for wide web-based content dissemination. In addition to developing special courses for undergraduate and graduate curricula, and regular seminars and webcasts, TRUST has hosted a series of workshops on sensor networks, privacy, identity theft, and electronic medical records. A major thrust in the fourth year was

increasing content in the TRUST Academy Online (TAO) and continuing the redesign of the TAO portal. Again, all these are reported below in the section on education.

In diversity, TRUST has an ambitious goal of reaching a diversity goal across the Center of 30% women and 10% from underrepresented minorities. The Center has been very proactive in this regard and expanded several programs for enhancing diversity and broadening the participation of women and underrepresented minorities.

In knowledge transfer, TRUST has continued a robust program of technology transition with industry (from reporting security vulnerabilities to software vendors to various consulting activities) and active engagement with governmental agencies such as the Department of Homeland Security (DHS), the Air Force Office of Scientific Research (AFOSR), the Department of Defense (DoD), and the Department of Energy (DoE) which are all concerned with issues of cyber security and trustworthiness. TRUST also has an active set of industrial partners with whom we are engaging in research and development collaborations of mutual interest. More details are provided in the section on knowledge transfer.

Overall, we are happy to report that the Center is making excellent progress towards its goals, its participants are actively engaged, and the outlook is positive.

2 RESEARCH

2.1 Goals and Objectives

The TRUST vision is to provide a unique opportunity for a wide range of cyber security issues to be addressed from many points of view—technological, scientific, social, policy, and legal. Of paramount importance to TRUST is the creation of a science that will simultaneously address the imperatives of all these points of view and allow scientists and technology developers, policy makers, and social scientists to make informed and rigorous decisions with the full understanding of tradeoffs involved. We think that this new science, though exciting and far-reaching, will come about from an evolution of more traditional areas that impinge on this "science of TRUST" as theory and praxis of these areas co-evolve. In particular, the primary areas of new science creation include cryptographic protocols and supporting systems, high confidence software science, security functionality, policy and management guidance, and complex interconnected networked systems. Furthermore, TRUST will have strong, well proven ties with Information Technology (IT) vendors and commercial infrastructure providers which will serve to both ground TRUST research in real-world problems and enable avenues for knowledge and technology transfer. TRUST will have a significant impact on a national scale as its research results will lead to new concepts and doctrine for (1) public policy issues around privacy, access control, and security; (2) technology for protecting and preventing information security breaches; and (3) increased protection of the nation's critical infrastructures, most notably in the areas of electric power, telecommunication, healthcare, financial services, and military networks.

2.2 Performance and Management Indicators

TRUST projects are both continuously and periodically monitored for meeting the center's overall research objectives and the project's individual research objectives. Periodic monitoring consists of bi-annual meetings of all TRUST personnel where research results are presented and progress in each research thrust is formally reviewed. Continuous monitoring consists of evaluation by both the research thrust area leaders as well as by the TRUST Executive Board. The evaluation metrics are outlined in the table below.

Objective	Metric	Frequency
Scientific Impact	Publications, Presentations, Recognition	Annual
Technological Impact	Transitions, Industry Interest	Annual
Timeliness	Milestone Completion	Semi-Annual
Social Impact	Policy Papers, Legal Policy	Annual

2.3 Current and Anticipated Problems

No significant problems were encountered during the reporting period. No significant problems are anticipated in the next reporting period.

2.4 Research Thrust Areas

TRUST is addressing technical, operational, privacy, and policy challenges via interdisciplinary projects that combine fundamental science and applied research to deliver breakthrough advances in trustworthy

systems in three "grand challenge" areas. During the first three years of the center, TRUST research projects were focused on anywhere from 5 to 11 challenge areas. Evolution of the research areas has occurred due in part to consolidation of similar research interests and a collective agreement among TRUST management and campus principal investigators to focus TRUST researcher efforts in certain areas. Each research thrust was selected to encourage projects that are integrative in nature and provide opportunities for TRUST researchers to work on topics that cross disciplines and allow collaboration across campuses. An overview of the three research thrust areas is below.

- **Financial Infrastructures.** TRUST researchers aim to develop central science and engineering principles to ensure the long-term security, reliability, and ubiquitous usage of the nation's financial infrastructure. This comprises financial service enterprises, online retail businesses, and customers linked together in a trustworthy environment that supports commercial transactions. TRUST is addressing needs and challenges of a trusted financial infrastructure and its key components:
  - *Service Providers.* Financial service providers and online retailers interact with customers through e-mail, operate web servers, carry out back-office operations subject to rigorous security and performance requirements, and have complex partnering agreements and rely on their brand image and reputation for competitive advantage.
  - *Customers.* Individuals interact with financial service providers through e-mail and the web. These individuals are usually not technology experts yet they need to be assured of reliable interaction.
  - *Interconnection.* Financial infrastructure customers rely on open networking standards, browser architecture, and web application development practices. Providers may also communicate through private networks, leverage federated identity management solutions, and outsource functions to other providers through complex networking practices.
  - *Policy.* Financial services and online enterprises are subject to complex and overlapping regulations and evolving levels of customer awareness and sensitivities. Both policy and technology are necessary to drive security in an increasingly decentralized environment in which consumers with limited technical expertise and desire to manage security/privacy play a central role.

- **Health Infrastructures.** Healthcare has been characterized as a "trillion dollar cottage industry" dependent upon paper records and fragmented, error-prone approaches to service delivery. Recently, however, the healthcare industry is changing, including: the dramatic increase in the amount of information required for making health decisions, the rapidly growing use of Internet worldwide, genome research that opens up opportunity to provide personalized healthcare, and medical errors caused by failures in information management.

  Information technology enables the creation of disruptive technologies that can change health care, for example the transition from paper to digital Personal Health Records (PHRs), the growing deployment and use of real-time medical decision support systems and online patient portals, and the emphasis on robust Health Information Systems (HISs). These technologies offer unique opportunities for both improving the delivery of care in medical facilities and shifting healthcare from traditional clinical settings to patient/home-centered settings. That said, adoption of these new, transformational technologies is predicated on the availability of technical solutions and design methodologies to solve problems such as the implementation of privacy requirements and the guarantee of safe operation of HISs. To address this, TRUST researchers are tackling fundamental issues affecting the design of trusted HISs that are composable from component

technologies. A primary concern in HIS design is that privacy and security requirements are frequently expressed in vague, complex and often contradictory laws and regulations. Engineering software systems that are functionally complete, able to adapt to the changing healthcare environment, and can comply with security and privacy laws and regulations is hard, if not impossible, using conventional software and systems design technology. As such, TRUST researchers are using model-based methods to offer a revolutionary way to formally and explicitly integrate privacy and security goals into HIS architectures. While this had led to progress in problem understanding and developing new foundations, TRUST researchers also place strong emphasis on experimental work. Taking advantage of the Center's partnership with the Vanderbilt University Medical Center, researchers have developed a testbed for Model Integrated Clinical Information Systems (MICIS) and home-based health monitoring that integrates TRUST research results in a platform used by the medical community for testing and validation.

- **Physical Infrastructures.** This area addresses next generation Supervisory Control and Data Acquisition (SCADA) and other networked embedded systems that control critical physical infrastructures (e.g., power grid, natural gas distribution, automated railroad control, water, transportation) and futuristic infrastructures such as "smart" buildings and structures" (e.g., active-bridges whose structural integrity depends on dynamic control or actuators).

  In physical infrastructures using new secure SCADA systems and built on top of the emerging new technology of wireless networked embedded systems, substantive issues of ownership and control of the physical infrastructure (whether it is individuals inside their homes or the grid utility provider). Security requirements are traditionally enumerated in terms of confidentiality, availability, and integrity. In this area, confidentiality is not a primary drive. Moreover, availability is often too weak—real-time constraints must be satisfied which changes the approach for defending against denial of service attacks. Ensuring integrity, however, is important as reliable operation of critical infrastructures needs to be ensured even in cases where an adversary controls a subset of the devices (which requires addressing threats such as the physical compromise of unattended nodes deployed in the field). Additionally, privacy issues arise in this area, such as understanding what can be inferred from the use and analysis of infrastructure information (e.g., increased power draw implies somebody is at home). Moreover, when distributed networks of sensors are widely deployed, opportunities for privacy abuse arise through abuse of information that is being collected for other reasons. Future infrastructures such as smart buildings and structures portend immense data collection in places routinely occupied by individuals. TRUST researchers are addressing such privacy concerns by considering them early on in the design and development of technical solutions and in advancing policy and consumer protection awareness and understanding that will support this future.

Specific research activities in each thrust area are described in more detail in the following sections. For each area, overall objectives and a scope of work are provided as well more detailed information about specific research projects conducted.

2.4.1 Financial Infrastructures

Project Leaders: *John Mitchell (Stanford), Doug Tygar (Berkeley)*

In the TRUST Center approach to this area, we view the financial infrastructure as the combination of financial service providers, online retail businesses, and their customers, all linked together in a trustworthy environment supporting commercial transactions. While the World-Wide Web supports a range of financial transactions, we view the financial infrastructure as including Web browser,

applications, and interfaces, and also extending beyond the largely customer-oriented Web infrastructure to include companies that use the Web and back-end systems for financial purposes, their internal and interconnected back-end computer systems, and the cultural and regulatory environments in which they operate.

The complexity of the scientific, engineering, cultural, psychological, and legal challenges facing the financial infrastructure stems from several characteristics of the current environment. Foremost among them is that *attacks against or within the financial infrastructure are prevalent and lucrative*. The FBI estimates that computer crime costs industry $400B/yr, with estimates of $50B for ID theft. Another important characteristic that distinguishes this area from other TRUST grand challenge areas is that *financial systems are not under control of one organization*: web browsers that execute critical parts of current web applications are separately administered by non-experts. In addition, the intra-enterprise financial infrastructure is highly networked. In contrast to traditional computer systems, financial systems *critically involve computers and people*. While authentication of computer systems to each other has been widely studied, websites want to authenticate a person, not a machine. In addition, the importance of the human in the loop leads to significant legal, social, policy, and human factors issues. Finally, the financial infrastructure operates in the face of *rapid technological evolution*. Web technologies are rapidly changing, server development frameworks are similarly rapidly evolving, and the rise of ubiquitous handheld platforms provides a means for development and deployment of new technologies that will replace old ones rapidly.

Based on interviews and discussions with industry leaders and others, TRUST has identified a range of pressing current problems, including:

- *Authentication*. Financial infrastructure enterprises face challenges in reliably authenticating clients (customers) to site and sites (enterprises) to clients, for both email and web.
- *Malware*. Enterprises face sophisticated direct malware-based attacks to their information systems, and indirect attacks through malware on their customer's computers.
- *Internal Operations*. Enterprises face policy, compliance, and risk management challenges as well as continuing exposure to insider threats.

In response to the early 2008 articulation of a TRUST vision for financial infrastructures, TRUST investigators submitted close to 20 separate proposals, many of them initially collaborative and all of them eventually combined with others to form collaborative cross-campus proposals (or dropped). The funded projects were organized into the following areas, listed with the primary investigators active and TRUST-funded in the area:

1. Systems support for financial enterprises (Aiken, Birman, Reiter, Schmidt, Wagner)
2. Network Security (Datta, Francis, Gligor, Mitchell, Perrig, Suh)
3. Web Security (Boneh, Mitchell, Perrig, Song, Tygar)
4. Code analysis, system monitoring; and malware detection (Aiken, Engler, Dill, Mitchell, Seshia, Song, Wagner)
5. Public policy, Decision-making, and Risk Management (Bamberger, Bambos, Burstein, Chuang, Datta, Hoofnagle, King, Mitchell)

Representative accomplishments in each of the areas of described below; corresponding future plans are described in Section 2.6.1.

<u>Systems support for financial enterprises</u>

(i) *Building trustworthy systems using the Live Objects platform.* This effort is focused on creating a secure, trustworthy and scalable technology for constructing a new generation of collaboration tools and applications that can be applied in health and financial settings, as well as in applications such as military search and rescue. This work is cross-cutting and falls into multiple project areas.

There is a growing opportunity to use Service-Oriented Collaboration Applications in ways that can reduce costs, improve productivity, and make possible a world of professional dialog and collaboration without travel. Collaboration applications will need to combine two types of content: traditional web service hosted content, such as data from databases, image repositories, patient records, and weather prediction systems, with a variety of collaboration features, such as chat windows, white boards, peer-to-peer video and other media streams, and replication/coordination mechanisms. Existing web service technologies focus on applications in which all data travels through a data center. Implementing collaboration features using these technologies is problematic because collaborative applications can generate high, bursty update rates and yet often require low latencies and tight synchronization between collaborating users. One can often achieve better performance using direct client-to-client (also called peer-to-peer, or P2P) communication, but in today's SOA platforms, "side-band" communication is hard to integrate with hosted content.

The TRUST Live Distributed Objects platform (Live Objects for short) allows even a non-programmer to construct content-rich solutions that blend traditional web services and peer-to-peer technologies, and to share them with others. This is like creating a slide show: drag-and-drop, after which the solution can be shared in a file or via email and opened on other machines. The users are immersed in the resulting collaborative application: they can interact with the application and peers see the results instantly. Updates are applied to all replicas in a consistent manner. Moreover, in contrast to today's web service platforms, P2P communication can coexist with more standard solutions that reach back to the hosted content and trigger updates at the associated data centers. Thus, when an application needs high data rates, low latency, or special security, it can use protocols that bypass the data center to achieve the full performance of the network.

The Live Objects platform is working and in use by more and more third party efforts that have downloaded and are experimenting with our open-source solutions. These include teams at companies such as Intel, hospitals such as Cornell-Weill, and academic teams using our technology in their own research, such as the work of Jong Hoon-Ahn in Debbie Estrin's sensor networks team at UCLA. We have also written a number of papers, several accepted at highly selective conferences and journals, and have more in the pipeline.

(ii) *Trustworthy and Dependable Platforms for Critical Ultra-Large-Scale Systems:* it has become evident that a diversity of emerging software-intensive systems critical to the US financial infrastructure will need to connect huge numbers of platforms, networked together in ways that may or may not be managed centrally by a higher level point of administrative control. These mission-critical ultra-large-scale (ULS) systems present quality of service (QoS) challenges that go well beyond anything seen in today's systems and systems of systems. QoS properties required by ULS systems include (1) the low latency and jitter expected in conventional real-time systems and (2) high throughput, scalability, and reliability as expected in conventional enterprise distributed systems. Achieving these QoS properties in isolation is hard; it is even harder to achieve them simultaneously in ULS systems composed of heterogeneous and (often) undependable components.

For example, a growing number of stock exchanges, aggregators, trading, and investment firms must solve data distribution problems associated with disseminating ever-increasing message volumes, while

simultaneously reducing end-to-end distribution and trading latency. The need for ultra-low latency and high throughput messaging technologies is widely recognized and is receiving increased attention within the financial industry. Although messaging technologies that promise low latency and high throughput have generated significant hype in the marketplace, the reality is more complex. In particular, a complete solution not only requires low latency and high throughput, but also the proper level of control, stability, usability, and expressiveness to address the need for increased performance, QoS, and reduced time-to-market in ULS financial systems. None of these capabilities exists today. Moreover, the off-the-shelf technology base is completely inadequate for the task summarized above. In the context of this form of broader interaction, our teams at Vanderbilt and Cornell are developing and validating trustworthy and dependable platforms for critical ULS systems. These platforms are based on Service-Oriented Architecture (SOA) technologies and associated educational material that can provide an assured software platform for critical ULS systems whose QoS support enables users and applications to process the right data in the right place at the right time over a much broader range of computers and networks than is possible using conventional SOA technologies.

We use the term QoS here, but in fact our teams currently share a strong interest in time-critical applications with guarantees either that deadlines will be met or, in a new and exciting twist on the usual real-time story, systems that are simply blazingly fast and can preserve their rapid response times even as we scale them up or subject them to stress up to outright attacks. Traditional real-time systems have overlooked this category of "things that need to be very fast," and we are finding this to be fertile ground for research. At Vanderbilt, Schmidt's team led in developing a widely used QoS-enabled middleware platforms, while Cornell's group, under Birman, had led in explorations of fault-tolerance (notably virtual synchrony for process groups), and is now developing Ricochet and Tempest, two time-critical cluster computing solutions that focus on ultra fast, highly reliable communication.

Personnel at Vanderbilt and Cornell have collaborated to develop the Adaptive Network Transport (ANT) framework. ANT presents a flexible architecture whereby transport protocols can be modified or enabled and disabled dynamically. ANT includes support for the Ricochet transport protocol among others. Ricochet is a scalable reliable multicast protocol developed at Cornell University that combines high data rates with strong probabilistic delivery guarantees. We have performed systematic evaluations of real-time event notification middleware integrated with various transport protocols. In particular, we have integrated ANT with the OpenDDS implementation of the Data Distribution Service (DDS). ANT includes support for the Ricochet transport protocol, a scalable reliable multicast protocol developed at Cornell University that combines high data rates with strong probabilistic delivery guarantees. We have characterized the performance of various transport protocols, including Ricochet, TCP, UDP, and IP multicast. Our work has focused on the metrics of latency and reliability – two key areas for QoS-enabled pub/sub middleware. We have currently varied the environment configuration based on the number of receivers, the percentage loss in the network, and the sending rate of the data. Our results show that for any one environment configuration a particular transport protocol provides the best balance of reliability and low latency while for another environment configuration a different transport protocol delivers the best performance. We are using the data collected to evaluate various machine learning techniques to determine which is most appropriate for maintaining specified QoS in dynamic environments.

Network Security

One direction involves trusted computing technologies such as TPMs (Trusted Platform Modules) or secure processors. A second involves network security protocols, which are the core foundation of secure networked systems used in the financial infrastructure.

(i) *Trusted Computing Platforms and Secure Network Enforcement* This project direction aims to develop a trusted computing platform that enables trustworthy enforcement of network operations at each end-host along with network technologies to utilize that platform. Hardware-based mechanisms can serve as a good basis for trust because they cannot be tampered with by any software or even by an owner without substantial effort. The project will investigate four main components for the trusted platform and three components for the network architecture and Operating System network stack.

1) Hardware Authentication: One must be able to distinguish trusted hardware from software simulators or untrusted hardware without proper security mechanisms. The project will investigate ways to exploit timing and physical properties of a processor to authenticate trusted hardware without relying on secret keys.
2) Attestation: Once the hardware is trusted, the microprocessor can attest to a software configuration/state enabling a remote verifier to determine whether the initial state can be trusted.
3) Isolation: During the execution, the system must ensure that the execution state cannot be tampered with by malicious software.
4) Enforcement: As a new trusted platform feature, a trusted module must control network resources so that security policies can be enforced even when the system is compromised.

In the past year, we have developed a technique to distinguish authentic hardware from simulators and emulators based on the performance limits in simulations and emulations (to be published in DAC'09) instead of purely relying on public key infrastructures [1]. We plan to utilize the technique in our network enforcement architecture.

We have also developed a network architecture to defend against Distributed Denial-of-Service (DDoS) attacks, which is currently under submission. DDoS attacks typically originate from exploited endhosts controlled by a remote attacker. Current network-based DDoS defenses can only filter out malicious traffic based on the traffic's inherent properties; they cannot filter based on properties of the endhost that generated the traffic. We observe that the identity of the code that has generated a packet offers powerful predicates for filtering, and we develop a secure, general architecture, for in-network filtering based on endhost properties.

Our proposed architecture leverages hardware-based attestation mechanisms to enable legitimate endhosts to embed secure proofs of code identity in packets. Receivers can specify traffic policies, which are enforced by on-path prioritizers. We design our system to achieve scalability, efficiency, and incremental deployability.

(ii) *Security Analysis of Network Protocols*. There have been many efforts to develop and use methods for finding security vulnerabilities and proving security properties of network protocols. In recent years, most efforts have used the so-called symbolic model, also referred to as the Dolev-Yao model. In the symbolic model, protocol execution and the possible actions of an attacker are characterized using a symbolic model of computation that allows nondeterminism but does not incorporate probability or computational complexity bounds. In addition to many model checking and bug-finding efforts, there have been some significant correctness proofs carried using the symbolic model, including mechanically checked formal proofs, unformalized but mathematical proofs about a multiset rewriting model, and work using compositional formal logic approaches. Several groups of researchers have taken steps to connect the symbolic model to the probabilistic polynomial-time computational model used in cryptographic studies. TRUST researchers at Stanford, Berkeley, and CMU have been at the forefront of these efforts.

Protocol Composition Logic, developed by TRUST faculty at Stanford and CMU and their students and collaborators, is defined and developed or proving security properties of the most widely used network protocols. In recent years, it has been used to prove security properties of versions of Kerberos in the symbolic model and in the computational model, with errors discovered and repaired in the Diffie-Hellman variant of Kerberos. Connections between symbolic trace properties and computational soundness properties have also been developed.

One sample case study in 2008 looked at the GPSK protocol. In spite of significant efforts by the research community to develop foundations for correctness of network security protocols, protocols such as EAP-GPSK are still being designed with flaws and weaknesses that are identical to those found and fixed in previous protocols. A major contribution of TRUST work is to integrate the methods of protocol analysis into the standardization process before the protocol is deployed in a variety of implementations.

In the EAP-GPSK case study, we used finite-state model checking to find errors, and Protocol Composition Logic to prove correctness after errors have been found and repaired. The model checker we use is called Mur*phi*. Mur*phi* has been successfully used in the past on a variety of protocols including Kerberos, SSL, and the 802.11i 4-Way Handshake (for authenticating wireless devices to access points). As a model checker, Mur*phi* is well suited for finding flaws but is insufficient to prove the correctness of a protocol. So to compliment Mur*phi* we use Protocol Composition Logic (PCL) as a proof tool. Mur*phi* was useful in detecting some of the problems with the protocol specification as we first encountered it, while PCL was useful for proving that the fixes we suggested, and which were subsequently adopted, are correct.

Our analysis uncovered three weaknesses with GPSK. The first is a repairable Denial-of-Service attack against the peer, in which the attacker forces the peer to exhaust its memory thereby blocking the protocol. This attack is virtually identical to one which was found on the 802.11i 4-Way Handshake (by TRUST researchers). The second weakness is due to a non-standard use of the key derivation function which is used to create session keys. Although the non-standard use does not provide an obvious attack we cannot exclude the existence of an attack. In addition, we indicate the difficulties which such non-standard usage creates when trying to prove the protocol's correctness. Finally, we identify a ciphersuite downgrading attack in which the attacker can force the peer to choose a weak hash function or encryption algorithm. If the ciphersuite is susceptible to a key-recovery attack then the attacker can learn the session keys and then eavesdrop on all subsequent communications.

In addition to discovering these weaknesses we also suggest ways to fix them and prove that our proposed fixes make the protocol secure. We discussed the weaknesses and our suggested fixes with the authors of the GPSK specification. In turn, the EAP-GPSK authors presented the issues to the EMU working group for open discussion. They recognized the problems and they have since incorporated our suggested fixes to the most recent protocol specification. Our interaction with the EAP-GPSK authors came at a time in which GPSK was mature enough to undergo a thorough analysis, and yet early enough in the standardization process that it was not widely implemented.

Web Security

As the World-Wide Web has exploded, web programming has become more sophisticated and the success of web commerce has led to companies that employ a large fraction of graduating students at all levels. From a security standpoint, a major development of the last decade is that the most prevalent web-based attacks are not aimed at creating mischief, but at stealing money. The web has been a launching pad for phishing attacks, and also a rising number of web-based attacks that exploit vulnerabilities other than gullibility of inexperienced users. The prevalence of advertising now allows a new and very powerful

form of attack -- malicious ads distributed to hundreds of thousands of unsuspecting users' browsers, as the result of relatively inexpensive placements on advertising networks. Third-party advertisements placed in web publishers' content is a simple form of the more general trend, associated with terms such as “mashups” and “Web 2.0," that place executable content from untrusted sources on web browsers. The security of such applications requires isolation mechanisms analogous in effectiveness to those found in modern operating systems, but suited to the more diverse and heterogeneous web environment.

(i) *Browser frame isolation and interframe communication:* TRUST researchers from Stanford and Berkeley have made substantial progress on the core problem of isolation and communication within the browser. In broad terms, we aim to improve two forms of isolation, language-based isolation between JavaScript code from different sources (or with different trust levels), and browser-based isolation associated with the frame mechanism. We have recently studied inter-frame communication, which is subject to browser control. We evaluated existing frame navigation policies and advocated a stricter policy, which we implemented and have seen adopted in the open-source browsers. In addition to preventing undesirable interactions, the browser's strict isolation policy also hinders communication between cooperating frames. We analyzed two techniques for inter-frame communication. The first method, fragment identifier messaging, originally provided confidentiality without authentication, which we repaired using concepts from a well-known network protocol. The second method, postMessage, provides authentication, but we discovered an attack that breaches confidentiality. We modified the postMessage API to provide confidentiality and saw our modifications standardized and adopted in browser implementations. The research paper describing this work was selected as a *research highlight* in Communications of the ACM.

(ii) *Isolation models based on capabilities (a traditional approach from operating systems):* Recently, there has been a growing trend to treat JavaScript pointers as object-capabilities within web browsers in order to build safer mashups and more robust implementations of the browser’s same-origin security policy. In this project, we will evaluate the security of this approach, suggest improvements, and, where appropriate, propose alternative techniques. The main difficulty in reasoning about systems that use object-capabilities is that one capability can lead to another. For example, if a function is given a pointer to one object, then that function is also implicitly given a pointer to all the objects pointed to by that object. We plan to model these transitive grants of capabilities using a capability graph. The capability graph will let us discover that functions are granted more capabilities than expected, leading to attacks, and will let us verify that other functions are not given dangerous capabilities, even transitively.

(iii) *Safer web-based advertising:* On-line commerce is a rapidly growing aspect of our economy, and a lot of that commerce is driven by on-line advertising. Just as other aspects of our financial and commerce infrastructure are vulnerable to phishing attacks, spam, denial-of-service attacks, and so on, on-line advertising is vulnerable to various type of fraud, including click-fraud. Successfully committing ad fraud yields direct monetary gains for attackers at the expense of the victims. Thus, it is natural to consider online ad fraud in an economic context. As part of the TRUST project, and in collaboration with researchers at Google, we have been modeling online advertising and studying various fraud and pricing issues.

During this past year we have continued to investigate online advertising. We use economic analysis because, unlike many online security threats, ad fraud is primarily motivated by financial gain. Successfully committing ad fraud yields direct monetary gains for attackers at the expense of the victims. In one effort, we show how an online advertising network can use filtering, predictive pricing and revenue sharing together to manage the quality of cost-per-click (CPC) traffic. Our results suggest that predictive

pricing alone can and should be used instead of filtering to manage organic traffic quality, whereas either method can be used to deter click inflation.

In considering whether advertisers should fight click fraud, we suppose an ad network decides that a given click-through is invalid (or ``fraudulent"). The implication is that the ad network will not bill the advertiser for that click-through. On the other hand, if the ad network decides that the click-through was valid, they could charge full price. Therefore, arguably, the ad network is ``leaving money on the table" by marking click-throughs invalid. As such, should ad networks even bother fighting click fraud? We analyze a simple economic model of the online advertising market and conclude that the answer is, unequivocally, “yes”.

(iv) *Type-safe programming languages for web applications:* We will study how to improve the security of the web, both on the client side and on the server side, by building upon type-safe programming languages. Securing the web is an important challenge problem for our financial infrastructure. TRUST researchers have designed, developed, improved, and released Joe-E, which is based on a subset of the Java programming language designed to support programming according to object-capability discipline. The Joe-E language guarantees additional security properties by placing restrictions on Java code, but does not modify programs or change their meaning. This allows programmers' existing knowledge of Java to be applied and existing compilers, debuggers, and other tools to be used with Joe-E programs.

Code analysis, system monitoring; and malware detection

(i) *Deep Automatic Error Checking of Critical Software Infrastructure:* TRUST researchers building code analysis tools aim to improve our ability to test large software applications and to statically verify security properties of large-scale, security-critical software infrastructure, such as an entire operating system or web browser. The core challenge is twofold: designing automatic analysis techniques of sufficient precision and scalability to handle real systems with millions of lines of code almost automatically and, for the inevitable small percentage of cases that cannot be fully automated to understand what crucial information the programmer can provide in the form of limited specifications that will render the task tractable.

TRUST researcher Dawson Engler won the 2008 Grace Murray Hopper Award for TRUST-cosponsored work on code analysis and testing tools.

We have been developing a tool, KLEE, that uses a variation on symbolic execution to automatically generate test cases that execute most statements in real programs. We used KLEE to thoroughly check:

(1) all 89 stand-alone programs in the GNU Coreutils utility suite, which form the core user-level environment installed on almost all Unix systems, and, as such, represent some of the most heavily used and tested open-source programs in existence,

(2) 72 application in the BusyBox utilities suite for embedded systems,

(3) the HiStar operating system kernel.

KLEE-generated tests achieved high statement coverage—on average over 90% per tool (median: over 94%)—and in aggregate significantly beat the coverage of the developers' own hand-written test suites. We also used KLEE as a bug finding tool, applying it to 448 applications (over 433K total lines of code), where it found 56 serious bugs, including three in CoreUtils that had been missed for over 15 years. In addition we have demonstrated how to use KLEE to verify path correctness with little manual work in the case where there are two or more implementations of the same interface. With this method, we cross-

checked over 60 purportedly identical utilities finding functional correctness errors and a myriad of inconsistencies.

TRUST researchers extended previous work on user pointer analysis in an effort to eliminate the few remaining sources of unsoundness. These techniques automatically check the entire Linux kernel for user/kernel pointer security holes. We began a renewed effort in statically proving the absence of buffer overflows. While there have been dozens of proposals for eliminating buffer overflows, all current practical approaches either rely on unsound static techniques or purely dynamic techniques. From a static analysis point of view, buffer overruns are all about analyzing array indices, which can be formulated as an integer constraint satisfaction problem. New results we have obtained in integer constraint solving appear to make it possible for the first time to truly prove large portions of code free from overruns automatically. Our integer constraint solver is an orders-of-magnitude more scalable than previous approaches, and our initial results on real programs are very encouraging. All static guarantees of security depend on type safety of the underlying language, but critical infrastructure written in C and C++ is not type safe. We developed a *casting analysis* that is able to prove 75% of the casts in the Linux kernel statically safe. This analysis involves a number of novel techniques and is the first attempt to accomplish anything approaching a full proof of type safety for C programs beyond a few thousand lines of code.

In the past year we also made progress on three additional fronts:

1. We extended our previous work on user pointer analysis in an effort to eliminate the few remaining sources of unsoundness. These techniques automatically check the entire Linux kernel for user/kernel pointer security holes. This part of the work will be concluded in a few weeks (in mid-June).
2. We began a new effort in statically proving the absence of buffer overflows. While there have been dozens of proposals for eliminating buffer overflows. All current practical approaches either rely on unsound static techniques or purely dynamic techniques. From a static analysis point of view, buffer overruns are all about analyzing array indices, which can be formulated as an integer constraint satisfaction problem. New results we have obtained in integer constraint solving appear to make it possible for the first time to truly prove large portions of code free from overruns automatically. Our integer constraint solver is an orders-of-magnitude more scalable than previous approaches, and our initial results on real programs are very encouraging.
3. All static guarantees of security depend on type safety of the underlying language, but critical infrastructure written in C and C++ is not type safe. We developed a *casting analysis* that is able to prove 75% of the casts in the Linux kernel statically safe. This analysis involves a number of novel techniques and is the first attempt to accomplish anything approaching a full proof of type safety for C programs beyond a few thousand lines of code.

(i) *Next Generation Infrastructure for In-depth Malicious Code Analysis and Defense:*

We continue to enhance our state-of-the-art infrastructure for the analysis of malicious and security-sensitive code at the binary level, and to apply it to a variety of challenges in analyzing and defending against malicious code. The flexibility of our infrastructure allows us to apply it to a variety of important security challenges, and the research on those applications provides feedback into what analysis infrastructure features would be most valuable.

Mixed concrete and symbolic execution is an important technique for the analysis vulnerable programs and malicious software, because it can combine the specificity of testing with the ability to automatically reason about alternate executions (for instance, to discover the possibility of a vulnerability while

examining a benign execution). However, existing symbolic execution techniques perform poorly in the presence of loops. We have proposed *loop-extended symbolic execution*, a new approach that is able to reason symbolically about the effects of loops [1]. We anticipate that this will be a valuable enhancement for many applications of symbolic execution; in particular, it allows much more effective analysis of buffer overflow vulnerabilities than previous approaches, which are still a serious problem for legacy infrastructure software. Several recent publications develop malware analysis capabilities such as methods for extracting models of security-sensitive operations using string-enhanced white-box exploration on binaries.

Human Factors, Public policy, Decision-making, and Risk Management

Several subprojects in this topic area were initiated in 2008, with even greater interest and proposed activity appearing in 2009. We summarize some of the 2008 accomplishments here and present a larger range of planned activity in the section on 2009 plans.

(i) *Scaffolding for Human Computer Interfaces in Financial Infrastructures*: It is well known that most computer security failures result from human error, usually attributable to poor user interfaces. We have and will continue our successful work in developing robust user interfaces that are secure from attack. This work includes both study of end-user user interfaces (e.g., electronic banking) as well as institutional user interfaces (both inter-institutional and intra-institutional).

We have made substantial development along the lines of our proposal. We have developed a testbed for both testing usability of financial infrastructure and conducted a unique user test of these principles; and are currently in the final stages of releasing an open-source implementation of our SWOON testbed for wireless security simulation. We have begun developing a new wireless forensics tool. We have also explored the particular case of learning systems in hostile environments and showed that classic techniques for scaffolding in these environments do not work. We have continued our interaction with the Fed Banks.

(ii) *Fraud Detection in Consumer Reports:* This project seeks to determine empirically whether it is possible to detect identity theft by an analysis of a consumer report with no extrinsic information or interaction with the consumer. If such a determination is possible, it could drive policymakers to require consumer reporting agencies (CRAs) to engage in anti-fraud monitoring of reports (CRAs currently have negative incentives to engage in this analysis). With an affirmative fraud monitoring system in place, consumers could learn of identity theft in a positive manner (notice from a CRA) rather than the current situation, where consumers often learn of the problem in a negative way (such as being denied a loan or job, or being pursued by a debt collector, because of a polluted consumer report). Positive notification would mitigate the harms of identity theft and reduce losses to consumers and businesses.

(ii) *Combating Fraud in On-Line Advertising* On-line commerce is a rapidly growing aspect of our economy, and a lot of that commerce is driven by on-line advertising. Just like other aspects of our financial and commerce infrastructure are vulnerable to phishing attacks, spam, denial-of-service attacks, and so on, on-line advertising is vulnerable to various type of fraud, including click-fraud. Successfully committing ad fraud yields direct monetary gains for attackers at the expense of the victims. Thus, it is natural to consider online ad fraud in an economic context.

As part of the TRUST project, and in collaboration with researchers at Google, we have been modeling online advertising and studying various fraud and pricing issues.

During this past year we have continued to investigate online advertising. We use economic analysis because, unlike many online security threats, ad fraud is primarily motivated by financial gain.

Successfully committing ad fraud yields direct monetary gains for attackers at the expense of the victims. Our main results are described in the papers we list in the next section.

(iii) *Security decision-making among interdependent organizations:*
In various settings, such as when customers use the same passwords at several independent web sites, security decisions by one organization may have a significant impact on the security of another. For example, many web users enter the same password at multiple sites. For this reason, compromise of a low security site such as a high school reunion web site may provide valid user names and passwords for sensitive sites such as banks or online merchants. In such situations, it is not immediately clear how a bank should best protect its assets: should it invest more in protecting its own sites from compromise that might reveal its passwords, or donate security services to non-profit organizations that could share user credentials? Perhaps the banking industry would be well served by forming a consortium to provide better authentication and web security to non-profits, who may be measurably free-riding on banking industry consumer education programs already.

We developed and applied a model for security decision-making in such settings, using a variation of linear influence networks. The linear influence model uses a matrix to represent linear dependence between security investment at one organization and resulting security at another and utility functions to measure the overall benefit to each organization. A simple matrix condition implies the existence and uniqueness of Nash equilibria, which can be reached by a natural iterative algorithm. A free-riding index, expressible using quantities computed in this model, measure the degree to which one organization can potentially reduce its security investment and benefit from investments of others. We apply this framework to investigate three examples: web site security with shared passwords, customer education against phishing and identity theft, and anti-spam email filters. While we do not have sufficient quantitative data to draw quantitative conclusions about any of these situations, the model provides qualitative information about each example.

2.4.2 Health Infrastructures

Project Leader: *Janos Sztipanovits (Vanderbilt)*

Over the past decade, many healthcare organizations have started embracing information technology. Since 2002, more than 90% of the approximately 5,000 member institutions of the American Hospital Association have reported having websites, with most having descriptive information about their facilities and services. A relatively small but growing fraction of health care organizations have created “patient portals” that provide secure, personalized customer services via the web. For example, Vanderbilt University‘s patient portal is one of the more advanced healthcare sites, providing a growing set of individualized services to more than 35,000 enrolled patients. In Europe, several national initiatives have been started to provide platforms for shared electronic health data records. For example, health@net is an Austrian initiative to develop concepts and an implementation of distributed cross-institutional health data records. The platform is targeted to support cooperation and information exchange between stakeholders in the healthcare domain like hospitals, family physicians, and pharmacies. These developments and experiences have resulted in the establishment of national goals in health information systems (HIS) that include archiving and accessing personal medical records, evidence-based personalized healthcare, and home-based healthcare delivery.

During this reporting period, four related TRUST projects targeted this area:

- Privacy and Compliance for Healthcare Organizations led by Prof. Mitchell from Stanford,
- Mining Care Provider Behaviors and Anomalies from Electronic Health Record Access Logs led by Prof. Malin of Vanderbilt,

- Experimental Platform for Model-Integrated Clinical Information Systems led by Prof. Sztipanovits from Vanderbilt, and
- Real Time Wireless Monitoring of People for Independent Living and Healthcare led by Prof. Bajcsy of UC Berkeley.

The sections below describe each project in more detail.

Privacy and Compliance for Healthcare Organizations – Privacy is an increasingly important business concern in healthcare, financial services, and other organizations. In this project, we have built on our previous work to develop approaches for modeling systems that handle sensitive information, languages for specifying privacy policies, and algorithms for their enforcement. In particular, the following results and accomplishments were achieved during this reporting period:

- *Differential Privacy for Probabilistic Systems: An Approach to High Assurance Systems for Privacy-Preserving Sharing of Aggregate Information*. Differential privacy is a promising approach to privacy-preserving data analysis. This work is motivated by statistical data sets that contain personal information about a large number of individuals (e.g., census or health data). There is now a well-developed theory of differentially private functions. Despite recent work on implementing database systems that aim to provide differential privacy and distributed systems that use differential privacy as a basis for higher level security properties, there is no formal theory of differential privacy for systems. We have formulated precise definitions of differential privacy within a formal model of probabilistic systems, related these definitions to the original definitions, and developed a proof technique based on an unwinding relation for establishing that a given system achieves this privacy definition. The proof technique has been applied to a representative example motivated by the PINQ system, implemented by researchers at Microsoft Research.

- *HIPAA Compliance Checker*. The Stanford team and collaborators, including S. Sundaram from TCS, have continued to work on formalizing larger segments if the HIPAA privacy legislation in a declarative form, using the Prolog language. The goal of this work is to provide a characterization of this important privacy regulation that can be used by other organizations as the basis of automated operational compliance.

  Substantial progress was made in 2008, including formalization of several main parts of HIPAA, construction of a web interface to the Prolog engine, and construction of a sample messaging system resembling the *MyHealth@Vanderbilt* patient portal built by our medical privacy collaborators at Vanderbilt University Medical Center. The progress established in 2008 provided a basis for integrating a compliance checker into the medical system built and tested in the research group at Vanderbilt.

- *International Outreach*: Mitchell was co-organizer and Datta an invited participate in the February 11-14, 2009 Schloss Dagstuhl seminar on Trustworthy Health Information Systems (called "Model-Based Design of Trustworthy Health Information Systems"). The Dagstuhl Seminar was a way of promoting TRUST work on privacy and soliciting feedback from other researchers and industry participants. Dagstuhl Seminars bring together internationally renowned scientists and practitioners for the purpose of exploring cutting-edge topics. More information about the Seminar is at http://www.dagstuhl.de/de/programm/kalender/semhp/?semnr=2009073.

Mining Care Provider Behaviors and Anomalies from Electronic Health Record Access Logs – This project is developing methods, implemented in working software, that automatically monitor how users (e.g., physicians) access the records of subjects (e.g., patients) and flag potentially privacy-compromising actions (e.g., an unauthorized "peek"). The definition of such behaviors helps inform the research of Sztipanovits and Mitchell, as well as assist healthcare organization officials to understand how their employees function and collaborate, so that policies do not preclude existing complex workflows. So far we focused on building a software toolkit and real medical record access logs repository from the Vanderbilt University Medical Center (VUMC). We also focused on proof-of-concept pattern mining from such a repository. We have obtained five months of access log data from VUMC's longitudinal medical record, and corresponding personnel data about the users. We have also created fully unit-tested Python application for analyzing log files. The application is designed to provide core data structures, plugin chaining and message passing facilities, simple configuration, and utilities for external data sources such as raw log files and modern relational database systems. Furthermore, we have created a set of plugins for the application to specifically analyze the access logs as a relational model. Plugins include the ability to conduct simple association rule mining, network visualization, and abstraction of nodes into higher level groupings. Most importantly, we used the application to analyze the data set to find clinical relationships and workflows. We discovered that the relationships within VUMC are highly volatile suggesting that care is not delivered by a permanent group of people, but rather care is delivered by various people fulfilling certain roles within the organization.

Experimental Platform for Model-Integrated Clinical Information Systems – This project aims to develop the Model-Integrated Clinical Information Systems (MICIS), a software tool suite that assists HCO administrators in the design, verification, implementation and integration of CISs. The MICIS tool suite is capable of graphically representing data, workflows, organizational aspects, and regulatory requirements of the healthcare environment. MICIS translates formal models into the necessary software artifacts and deploys the system on a standard SOA platform. The formal models created in MICIS allow administrators to perform rigorous system analysis and to enforce privacy and security policies within the CIS prior to deployment. MICIS enables a rapid development cycle because CIS prototypes can be quickly generated, evaluated and changed, if necessary, by modifying the system models and regenerating the application. Furthermore, the model-based approach helps over the entire lifecycle of the application by easing maintenance and application evolution, since many aspects of the CIS can be modified via the models without touching the actual code.

The modular architecture of MICIS allows for a flexible and extensible framework. The modularity in MICIS is achieved by adopting a set of SOA standards; both basic and extended functionalities are implemented through communicating web services. The orchestration of these services is managed by a SOA compliant execution engine. MICIS makes use of workflows to capture the business logic of complex applications and to orchestrate the execution of the corresponding services. Our workflows are defined using the Business Process Execution Language for Web Services (WS-BPEL).

Security and privacy enforcement in MICIS is achieved through a reusable application-independent Prolog-based Policy Decision Point and Policy Enforcement Point. In MICIS, selected services can be secured using formally defined policies. This means that if a protected service is invoked, the information flow between the service and its invoking client is carefully monitored. In cases where non-compliance with the defined policies is discovered during the information exchange, the execution of the undesired operation is prevented and an exception is returned to the workflow manager.

The current implementation of MICIS serves as a proof-of-concept that model integrated development and integration of the clinical information systems on top of SOA is a promising way to support health

care organizations. A related effort based on capturing and enforcing evidence-based treatment protocols relying on the same architecture is being prepared for actual deployment at VUMC.

Using evidence-based guidelines to standardize the care of patients with complex medical problems is a difficult challenge. In acute care settings, such as intensive care units, the inherent problems of stabilizing and improving vital patient parameters is complicated by the division of responsibilities among different members of the health care team. Computerized support for implementing such guidelines has tremendous potential. The overall management of complex medical processes requires the formal representation of treatment protocols. This way, the temporal structure and coordination of the tasks are captured explicitly, as opposed to being hidden in the code, enabling protocol validation and verification.

This project is a collaborative effort between Vanderbilt School of Engineering and Vanderbilt Medical Center to apply advanced model-integrated computing techniques to the management of complex clinical processes. The team has completed the beta version of the generic software infrastructure and the sepsis treatment protocol models resulting in the Sepsis Treatment Enhanced through Electronic Protocolization (STEEP) toolset. We are in the process of performing a carefully coordinated, multi-phase experiment to evaluate the presented approach in terms of usability and effectiveness. Phase one of the clinical tests has already started in two ICUs at Vanderbilt to establish the baseline for the comparative study. We anticipate showing that the application will 1) decrease time to detection of patients with developing sepsis, 2) improve physician compliance with evidence-based standards, and 3) result in improved clinical outcomes for patients.

Real Time Wireless Monitoring of People for Independent Living and Healthcare – Rehospitalization is frequent and costly, and associated with gaps in care following hospital discharge. Frequent hospital readmissions are a sign of poor care and an unnecessary expense. Thus, reducing rehospitalization is an important element for health care reform. The current status quo of preventing readmissions by simply discharging patients with instructions to visit their physician within 30 days does little to ensure careful monitoring of vital signs, sodium and fluid intake, and adherence of medications at home, during the time between physician visits. An alternative solution, which is the focus of this research project, is the innovative use of computerized monitoring and management of patients in their home environment.

We are exploring the feasibility of wireless technology for continuous real time monitoring people in indoor and outdoor environments. The critical issues in this paradigm are privacy, reliability, security, and robustness. In the long run we envision an integrated system which will observe/monitor people in their daily activities and interactions with other people and record their location movements with sporadic or continuous feedback from either automated system and/or healthcare supervisor in order to encourage the user's performance goals. The information collected during this process can be stored and correlated with other previously collected information in databases such as medical records, personal data, etc. This kind of integrated system enables not only direct feedback but also the ability to be queried later. While this kind of a system offers many benefits to the users and institutions, it also raises serious questions about privacy, credibility, and integrity of the database, vulnerability from intruders, etc. During this reporting period, we further developed an end-to-end system, called DexterNet and tested it on 10 subjects. DexterNet integrates distributed, body-worn sensor nodes (five of them) with a Nokia smartphone and a PC which connects through Internet to a back-end health information system database. At each point of the system we implemented security measures guaranteeing privacy of the user's data.

2.4.3 Physical Infrastructures

Project Leaders: *Steve Wicker (Cornell), Adrian Perrig (Carnegie Mellon), Shankar Sastry (Berkeley)*

The nation's critical infrastructure—the power grid, telecommunications, water transport, interstate highways, etc.—constitute an immense investment. The financial investment takes the form of sunk costs and ongoing development and maintenance, while human investment is ongoing through development, maintenance, and regulatory organizations at the state and federal level. Infrastructure is clearly critical to the national economy. National modes of production depend on the functionality of critical infrastructures. Furthermore, multiple positive externalities derived from the establishment of critical infrastructure have created secondary and tertiary dependencies (e.g., air traffic control dependence on power and telecom infrastructure).

The TRUST Center recognizes that increasing complexity and 21st century security requirements demand new approaches to control, security, and long-term maintenance. Our work has been based on multi-disciplinary, multi-institutional research projects. The effort also extends across theory, technology, policies, and testbeds. For example, Berkeley, Cornell, and CMU have worked together on threat models, attack detection and attack-resilient models, and control-theoretic approaches to security. Vanderbilt, CMU, and Berkeley have developed an experimental SCADA testbed for use by Center personnel and external researchers. Cornell, Stanford, and Berkeley have worked together on technology and science support for development of privacy policies

The following sections provide more detail on research projects during this reporting period.

Secure and Robust SCADA Systems Testbed – Supervisory Control And Data Acquisition (SCADA) refers to large-scale, distributed measurement and control systems for high-level monitoring and control. Such supervisory systems are layered on top of low-level real-time sensor/control loops that regulate specific process variables. SCADA systems are used to monitor and control entire chemical plants, transport processes, municipal water supply systems, electric power generation, transmission and distribution, gas and oil pipelines, and other distributed processes, often on a societal scale.

Given the critical nature of SCADA systems, ensuring their security is of great importance. Attacks on the SCADA system can have serious consequences, such as endangerment of public health and safety, environmental damage, and significant financial impacts. There is a growing interest in that the current SCADA systems are vulnerable to many cyber attacks. Protection of SCADA systems has traditionally been based on the security by the obscurity concept, where proprietary protocols prevented an attacker from breaking into the system due to insufficient knowledge. Today, when everything is accessible via networks, such protection relies mainly on standards, recommendations, policies, and suggestions for possible countermeasures. In order to better understand how to protect SCADA systems, it is imperative to perform vulnerability assessment on these systems and develop appropriate security mechanisms to protect the SCADA systems against attacks. To do so, we have developed a SCADA system testbed, where realistic experiments can be safely conducted.

The SCADA security testbed provides the TRUST community with the opportunity of designing and analyzing innovative security methodologies and algorithms. The testbed is modular, allowing researchers to seamlessly integrate and test their modules and solutions. The design allows TRUST researchers to access the testbed remotely or to replicate the testbed locally to conduct experiments and to solve accessibility problems. From an educational perspective, the testbed highlights the importance of a secure SCADA system to the larger community by providing examples of damages provoked by attacks that are possible today on the existing SCADA infrastructure.

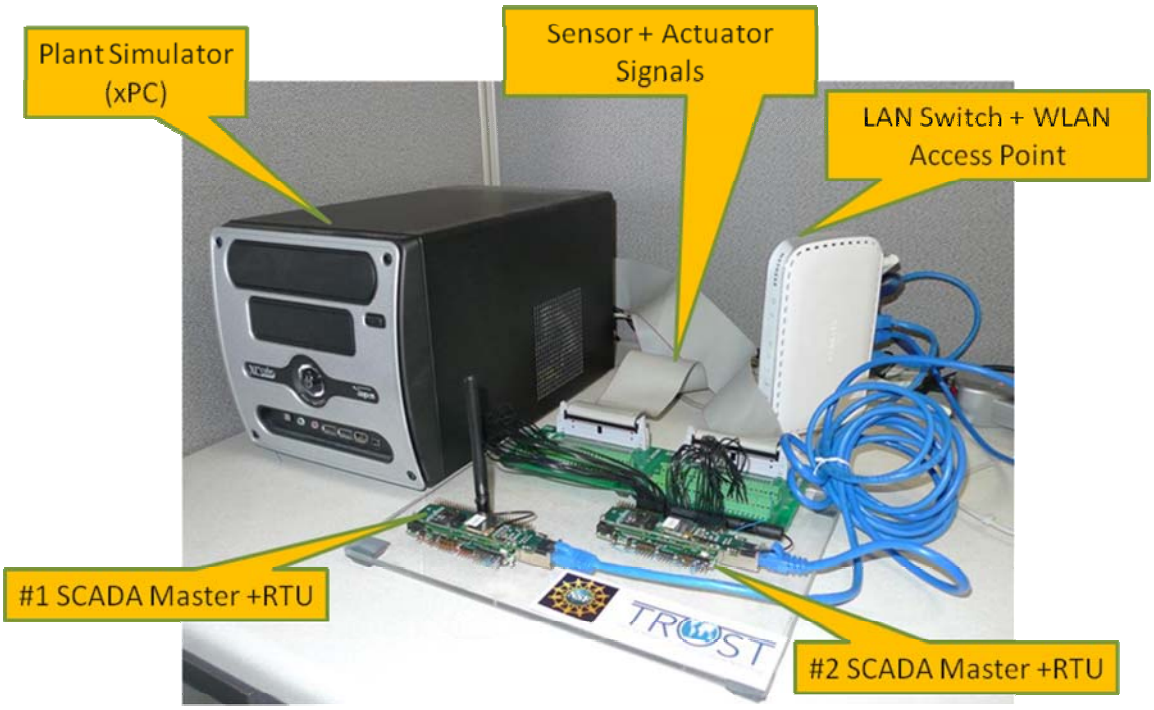

Project wiki site: https://wiki.isis.vanderbilt.edu/SCADA/index.php/Main_Page

Intrusion Detection for SCADA Systems – A hybrid, two-stage intrusion detection system (IDS) for mobile ad hoc networks has been developed to monitor physical infrastructure. This past year we have investigated the deployment of the IDS within SCADA systems. We leveraged the collaborative development of the SCADA testbed within TRUST (described above) to enhance the intrusion detection capability within the remote terminal units (RTUs) of the SCADA architecture. We started with the assumption that the SCADA system has been compromised; our objective was then to identify the deviancy of the compromised nodes and minimize the negative consequences of those nodes on the overall SCADA system.

We developed a hybrid, two-stage intrusion detection system, called HybrIDS, to obtain a flexible method of intrusion detection. Rather than relying on a fixed detection strategy, HybrIDS utilizes a system of two intrusion detection schemes. They operate in a cooperative effort to increase detection efficiency and accuracy beyond the capabilities of each individual method. The first strategy analyzes peaks within the statistical behavior to determine if a deviance exists among networked nodes. The primary advantage of this method is that it quickly establishes a "lock" on the most likely potential deviant node, with zero knowledge of the host system. The second strategy uses cross-correlation to detect multiple intruders in the network. The cross-correlation strategy requires careful selection of a detection threshold to minimize the introduction of false positives. In our method, the first strategy is also used to calibrate the cross-correlation of the second strategy. Together, the two methods provide a capable model of the activities in a distributed network and perform intrusion detection with minimal impact on computational resources.

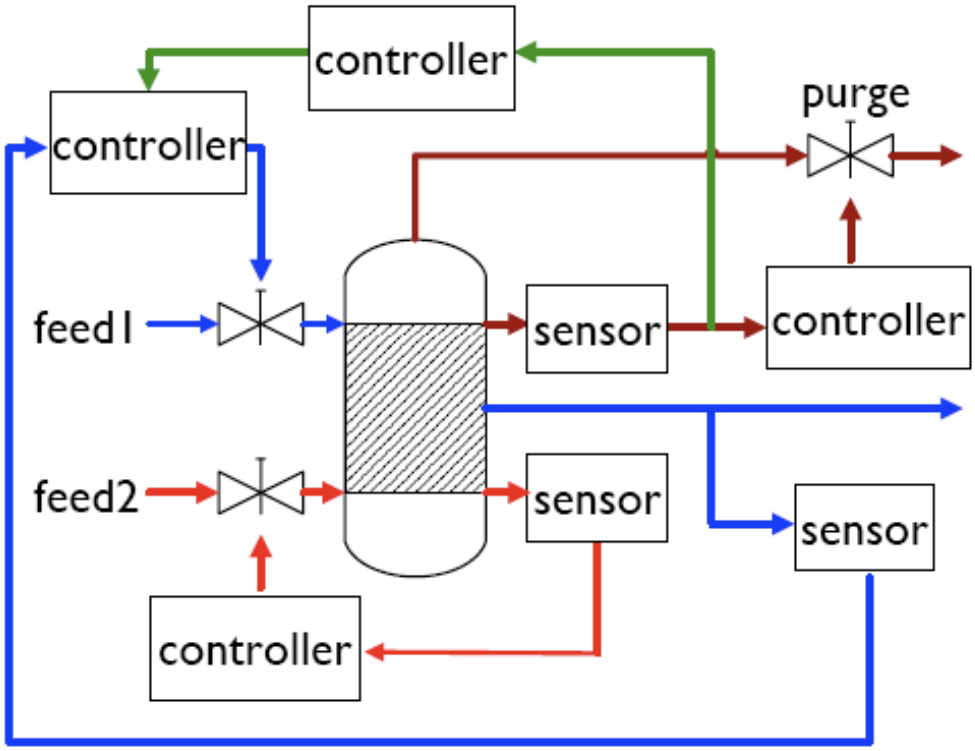

The objective of this work is to use HybrIDS within SCADA systems to enable the additional protection of physical infrastructures. SCADA architectures contain RTUs which could become compromised due to several modes of attack (e.g., mote class, insider). Because these malicious RTUs could affect the entire system, non-compromised RTUs need the capability to identify deviant behavior among their group. HybrIDS was initially developed as a general-purpose distributed IDS for ad-hoc networks but the SCADA architecture represents a specific application which could benefit from intrusion detection. Our initial discussions with TRUST researchers regarding the SCADA testbed have identified a potential collaboration to use HybrIDS to investigate its applicability to SCADA network environments. We will collaborate with TRUST researchers to develop and implement the SCADA testbed. Sensor units, simulated by embedded Gumstick Linux devices, are monitored in a semi-distributed manner by peer nodes as well as at a parent level by the RTUs. Traffic patterns and packet composition are then analyzed by the various nodes running HybrIDS. Different SCADA configurations can then be used to develop a more general solution to intrusion detection. A long term goal of this work is the development of multi-layered intrusion detection for implementation in SCADA and enterprise-level network scenarios.

Privacy Concerns in Upcoming Demand-Response Systems – NSF-funded researchers at Cornell and Berkeley have been engaged in a larger effort by TRUST to explore the confluence of sensor networking, power distribution, privacy, and security issues that will emerge from a substantial increase in power system monitoring at the consumer level.

The Cornell/Berkeley team has successfully completed a study on emerging privacy concerns in upcoming residential and commercial demand-response systems. Their main claim, substantiated by study results, was that in a lax regulatory environment, the detailed household consumption data gathered by advanced metering projects can and will be repurposed by interested parties to reveal personally identifying information such as an individual's activities, preferences, and beliefs.

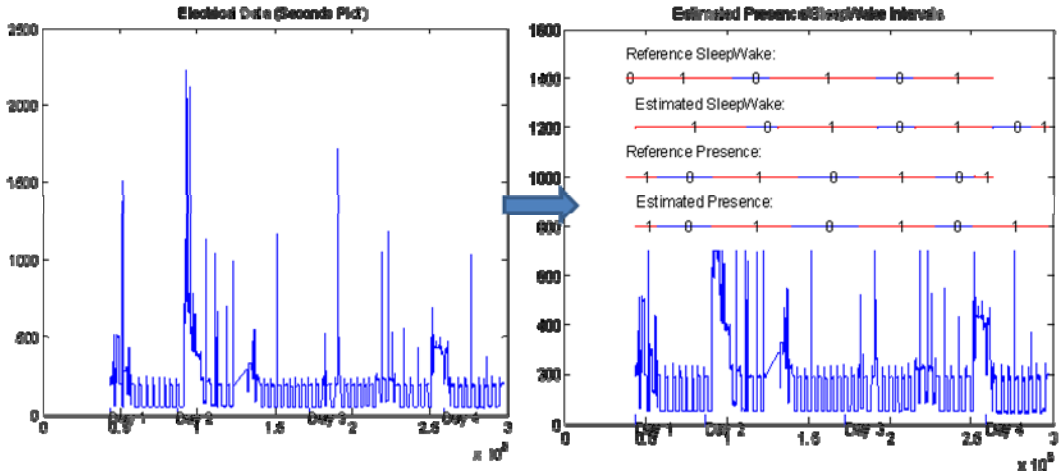

Their study included charting trends in demand-response technologies, establishing relevant legal and historical precedents, and formalizing the notion of privacy using a two-stage privacy metric. Also included was a small-scale monitoring experiment which established that personal information can be estimated with a high degree of accuracy, even with relatively unsophisticated hardware and algorithms.

Since one of the main goals of the study was to raise public awareness of the privacy issue, the Cornell team has also generated an education module, to be made available to the general public through TAO Portal. The module includes an introductory video, a detailed PowerPoint presentation, and an accompanying hands-on demonstration that illustrates behavior extraction algorithms operating on real experimental data.

Planned future work includes further improving behavior extraction algorithms by using Markov Chain and Lempel-Ziv based predictive algorithms originally used and already proven to be effective within the context of home automation. The group also plans to further develop the disclosure metric, which associates data quality (accuracy of readings, time resolution, types of readings, etc.) from a particular source with the information that may potentially be disclosed by the data

Secure Control: Towards Survivable Cyber-Physical Systems – Cyber-Physical Systems (CPS) integrate computing and communication capabilities with monitoring and control of entities in the physical world. These systems are usually composed of a set of networked agents, including: sensors, actuators, control processing units, and communication devices.

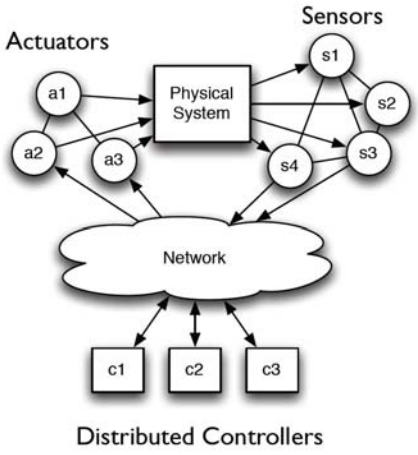

While some forms of CPS are already in use, the widespread growth of wireless embedded sensors and actuators is creating several new applications in areas such as medical devices, autonomous vehicles, and smart structures and increasing the role of existing ones such as SCADA systems.

Many of these applications are safety-critical: their failure can cause irreparable harm to the physical system being controlled and to people who depend on it. SCADA systems, in particular, perform vital functions in national critical infrastructures, such as electric power distribution systems, and transportation systems. The disruption of these control systems could have a significant impact on public health, safety and lead to large economic losses.

While most of the effort in protecting CPS systems, and SCADA in particular, has been done in reliability (e.g., the protection against random failures), there is an urgent growing concern for the protection against malicious cyber attacks. In this project we have studied the problem of secure control by focusing on the following:

1. Security mechanisms in sensor network security have focused on integrity and availability from a communication-network point of view. **They have not considered how integrity or denial-of-service attacks affect the application layer service**; i.e., how successful attacks affect our estimation and control algorithms – and ultimately, how they affect the physical world.

2. Intrusion detection systems are application specific; and they **have not considered the detection of integrity attacks against estimation and control algorithms**. In particular, previous detection of compromised sensor nodes assume a large number of redundant sensors; but **they have not considered the dynamics of the physical system being observed and controlled**, or how this model can be used to detect a compromised node without heuristics. Furthermore, **there has not been any research for detection algorithms to identify integrity attacks launched by compromised nodes**.

3. Most security response mechanisms involve a human in the loop. Because CPS rely on autonomous real-time decision-making algorithms for controlling the physical world, they introduce new challenges for the design and analysis of secure systems: a response by a human may impose time delays that may compromise the safety of the system. Therefore, **we must design autonomous and real-time detection and response algorithms for safety-critical applications**.

4. We defined an adversary model against CPS.

In this project we addressed the problem of "secure control" by addressing each of the issues raised above.

2.5 Research Metrics/Indicators

A key component of the Center research lifecycle is the monitoring and evaluation of individual projects. TRUST projects are both continuously monitored and periodically reviewed to ensure that they support the Center's overall research goals and make progress against the project's research objectives. The evaluation metrics are described below.

- **Scientific Impact** – How significantly does the project contribute to the knowledge base and general understanding of advances in the research area? This impact is typically measured by the number of published papers, presentations in open research conferences, and awards or other recognition for contributions to the research field.
- **Technological Impact** – How well does the project advance the state-of-the-art or state-of-the-practice in the research area? This impact typically is measured by ways in which research results are transitioned to industry, government, or the end-user community and examples where research results have been leveraged by industry in the creation of commercial or open source technologies.
- **Timeliness** – How effectively does the project meet its planned milestones? This is an evaluation of the actual project progress and advancement against planned activities, milestones, and deliverables.
- **Social Impact** – How well does the project contribute in ways that benefit society as a whole? This impact may be measured in terms of how the project research has influenced the development or refinement of public policies, federal, state, and local legislation, and legal decisions.

The TRUST Executive Committee continuously monitors Center research projects. If it seems unlikely that a particular project will meet its planned goals or objective or is not delivering the desired impact in one or more evaluation areas, that project will be ramped down in a period not to exceed six months from the determination of its lack of viability.

2.6 Next Reporting Period Research Plans

The goal of the TRUST research areas is to set the Center's strategic research agenda and align individual projects in such a way that they support the strategic research objectives. Because trustworthiness is an extremely broad field and TRUST does not have the resources to cover the entire spectrum of challenges, we have annually strived to focus TRUST research in areas where the Center could have the most impact. During the first three years, the research areas enabled TRUST researchers to both pursue specific research directions that the Principal Investigators believed were important and study application areas with an eye towards better understanding the landscape. The sections below provide a description of the planned TRUST research areas for the next reporting period. For each center thrust, the name(s) and institution(s) of the lead TRUST faculty member(s) is included.

2.6.1 Financial Infrastructures

Thrust Leader: *John Mitchell (Stanford University)*

The TRUST projects in this area are organized into the following areas, listed with the primary investigators active and TRUST-funded in the area:

1. Systems support for financial enterprises (Aiken, Birman, Reiter, Schmidt, Wagner)
2. Network Security (Datta, Francis, Gligor, Mitchell, Perrig, Suh)
3. Web Security (Boneh, Mitchell, Perrig, Song, Tygar)
4. Code analysis, system monitoring; and malware detection (Aiken, Engler, Dill, Mitchell, Seshia, Song, Wagner)
5. Public policy, Decision-making, and Risk Management (Bamberger, Bambos, Burstein, Chuang, Datta, Hoofnagle, King, Mitchell)

Representative accomplishments in each of the areas are described in section 2.4.1; representative future plans are summarized below.

Systems support for financial enterprises

(i) *Building trustworthy systems using the Live Objects platform.* This effort is focused on creating a secure, trustworthy and scalable technology for constructing a new generation of collaboration tools and applications that can be applied in health and financial settings, as well as in applications such as military search and rescue. This work is cross-cutting and falls into multiple project areas.

During 2009, w will expand by using real problems derived from our dialogs with the health, financial and military sectors as drivers. We will build simple applications but will extract new challenge questions from them, which can then be tackled through a mixture of theoretical and practical methods and ultimately used to push the envelope on our platform, motivate papers and research talks, and to help other educators get these sorts of ideas and solutions into the hands of their students. We are also hoping to create a wide-area "second life" environment, based on live objects: a potential killer application for our work that could attract a very high level of interest in our effort

(ii) *Trustworthy and Dependable Platforms for Critical Ultra-Large-Scale Systems:* it has become evident that a diversity of emerging software-intensive systems critical to the US financial infrastructure will need to connect huge numbers of platforms, networked together in ways that may or may not be managed centrally by a higher level point of administrative control. These mission-critical ultra-large-scale (ULS) systems present quality of service (QoS) challenges that go well beyond anything seen in today's systems and systems of systems. QoS properties required by ULS systems include (1) the low latency and jitter expected in conventional real-time systems and (2) high throughput, scalability, and reliability as expected in conventional enterprise distributed systems. Achieving these QoS properties in isolation is hard; it is even harder to achieve them simultaneously in ULS systems composed of heterogeneous and (often) undependable components.

Planned future work builds on our current accomplishments by incorporating and enhancing the following technologies and research to support maintaining QoS for pub/sub middleware via autonomic adaptation:

- Supervised machine learning to address timely adaptation to dynamic environments and managing multiple interacting QoS requirements by selecting in a timely manner an appropriate transport protocol and protocol parameters given specified QoS and a particular environment configuration. The machine learning component will include features for several different environment configurations and supervised training to learn the correct protocol and parameters for a given environment configuration. The machine learning will interpolate and extrapolate its learning based on the current environment configuration, which might not have been included in the supervised training.
- Environment monitoring to address timely adaptation to dynamic environments by providing environment configuration information. Relevant environment configuration values will be monitored

as needed such as the number of subscribers, the percentage of network packet loss, and the sending rate of the data. These monitored values will be input to the machine learning component to determine an appropriate network transport and accompanying parameters.

- Autonomic adaptation to address timely adaptation to dynamic environments and managing multiple interacting QoS requirements by (1) querying relevant values from the environment monitoring, (2) activating the machine learning component which will determine an appropriate transport protocol and parameters, (3) retrieving the recommended protocol settings, and (4) transitioning the adaptive network transports to use the recommended settings.
- Incorporation of common general and domain-specific data distribution profiles to determine system behavior and performance, e.g., when certain data types are more popular or in higher demand than others. Cornell personnel (e.g., Ýmir Vigfússon) are conducting research to characterize data distribution profiles, while Vanderbilt personnel are researching behavior of QoS-enabled pub/sub middleware utilizing the profiles.

In the code analysis area, we plan to focus on scaling the new buffer overrun techniques up to both the entire Linux OS and Firefox. Our goal is to automatically check at least 90% of buffer accesses automatically, and to understand what the limits of static analysis are with respect to any remaining, unverified buffer accesses.

Network Security

(i) *Trusted Computing Platforms and Secure Network Enforcement* This project direction aims to develop a trusted computing platform that enables trustworthy enforcement of network operations at each end-host along with network technologies to utilize that platform. Hardware-based mechanisms can serve as a good basis for trust because they cannot be tampered with by any software or even by an owner without substantial effort. The project will investigate four main components for the trusted platform and three components for the network architecture and Operating System network stack. In the next fiscal year, we plan to focus more on the network aspect of the project and further develop the hardware platform based on the needs that are more specific to the network enforcement. We plan to publish a joint paper (Cornell, CMU and Berkeley)

(ii) *Security Analysis of Network Protocols*. There have been many efforts to develop and use methods for finding security vulnerabilities and proving security properties of network protocols. In recent years, most efforts have used the so-called symbolic model, also referred to as the Dolev-Yao model. In the symbolic model, protocol execution and the possible actions of an attacker are characterized using a symbolic model of computation that allows nondeterminism but does not incorporate probability or computational complexity bounds. In addition to many model checking and bug-finding efforts, there have been some significant correctness proofs carried using the symbolic model, including mechanically checked formal proofs, unformalized but mathematical proofs about a multiset rewriting model, and work using compositional formal logic approaches. Several groups of researchers have taken steps to connect the symbolic model to the probabilistic polynomial-time computational model used in cryptographic studies. TRUST researchers at Stanford, Berkeley, and CMU have been at the forefront of these efforts.

Future plans in this area include continued improvements in Protocol Composition Logic (PCL), additional expository and teaching material, such as an invited book chapter describing the system, and additional case studies. We intend to continue toward our goal of producing understandable, rigorous proofs of key security properties of the most widely deployed protocols central to the financial infrastructure and other secure networked systems. On planned case study, already in progress, involves DNS and DNSSEC. After the recent vulnerabilities discovered and popularized by Dan Kaminsky, there

is substantial industrial interest in DNSSEC cost/benefit tradeoffs, and substantial possibility that some additional DNS security measures may be deployed.

Web Security

(i) *Browser isolation and inter-gadget communication:* TRUST researchers from Stanford and Berkeley have made substantial progress on the core problem of isolation and communication within the browser. In broad terms, we aim to improve two forms of isolation, language-based isolation between JavaScript code from different sources (or with different trust levels), and browser-based isolation associated with the frame mechanism. We have recently studied inter-frame communication, which is subject to browser control and evaluated existing frame navigation policies. In future work, we will study language-based isolation applications to social networking sites such as Facebook (which uses FBJS) for user applications), applications for advertising (such as Yahoo! AdSafe) and applications to general mashup gadgets (such as Google Caja). This work will build on the full operational semantics we have recently developed for the ECMA-262 standard (JavaScript) language.

(ii) *Isolation models based on capabilities (a traditional approach from operating systems):* Recently, there has been a growing trend to treat JavaScript pointers as object-capabilities within web browsers in order to build safer mashups and more robust implementations of the browser's same-origin security policy. In this project, we will evaluate the security of this approach, suggest improvements, and, where appropriate, propose alternative techniques. The main difficulty in reasoning about systems that use object-capabilities is that one capability can lead to another. For example, if a function is given a pointer to one object, then that function is also implicitly given a pointer to all the objects pointed to by that object. We plan to model these transitive grants of capabilities using a capability graph. The capability graph will let us discover that functions are granted more capabilities than expected, leading to attacks, and will let us verify that other functions are not given dangerous capabilities, even transitively.

By building these capability graphs, we can address several core questions about object-capability systems. For example, we can compute *reachability*, the set of object-capabilities that can be obtained from a starting set of capabilities. If the capability graph we compute is an under-approximation (as in 1 and 2), reachability lets us find new attacks. If the capability graph is an over-approximation (as in 3), reachability lets us rule out certain classes of attacks.

In addition to reachability, we can *factor the capability graph* into a directed, acyclic graph (DAG) of strongly connected components. In this view, object-capabilities in the same strongly connected component convene equivalent privileges because one capability in the component can be transformed to another capability in the component by following pointers. The DAG part of the factorized graph, then, characterizes the different privilege levels supported by the system.

Finally, we can compute *minimal cuts* in the capability graph. By removing edges in these cuts, the system designer can create new privilege levels in the system. For example, browsers might wish to expose an nsIMimeService object (which maps file extensions to MIME types) without granting access to arbitrary nsIFile objects. However, as currently defined in Firefox, this is impossible to guarantee. If the interface designers removed a single method from nsIMimeService (or created a new superinterface that lacked this method), then the capability graph would guarantee this security property.

In the course of this project, we plan to build a number of tools for extracting capability graphs from web browsers. For example, we plan to instrument the JavaScript heap to record the points-to relation among

JavaScript objects and to implement an XPCOM interface parser and type inference system to deduce the relation between XPCOM types in Firefox. We expect these tools to be useful for future research projects and for automated testing of web browsers.

(iii) *Safer web-based advertising:* On-line commerce is a rapidly growing aspect of our economy, and a lot of that commerce is driven by on-line advertising. Just as other aspects of our financial and commerce infrastructure are vulnerable to phishing attacks, spam, denial-of-service attacks, and so on, on-line advertising is vulnerable to various type of fraud, including click-fraud. Successfully committing ad fraud yields direct monetary gains for attackers at the expense of the victims. Thus, it is natural to consider online ad fraud in an economic context. As part of the TRUST project, and in collaboration with researchers at Google, we have been modeling online advertising and studying various fraud and pricing issues. We plan to continue our research on online advertising and fraud. We will also continue to study a variety of privacy issues related to large scale data management

(iv) *Type-safe programming languages for web applications:* We will study how to improve the security of the web, both on the client side and on the server side, by building upon type-safe programming languages. Securing the web is an important challenge problem for our financial infrastructure. TRUST researchers have designed, developed, improved, and released Joe-E, which is based on a subset of the Java programming language designed to support programming according to object-capability discipline. The Joe-E language guarantees additional security properties by placing restrictions on Java code, but does not modify programs or change their meaning. This allows programmers' existing knowledge of Java to be applied and existing compilers, debuggers, and other tools to be used with Joe-E programs.

We want to provide strong security for the web, including securing both servers and clients. We see type-safe languages as a powerful foundation for this work, and we plan to build upon Java (on the server-side) and Javascript (on clients). Safe languages provide a basis for reasoning about security, and for extending the base language with security functionality. We propose to study both how to retrofit legacy code for security, as well as how to design new systems that are inherently resilient to certain classes of attacks.

On the server side, we propose to develop methods for securing legacy web application code written against a broad variety of data-driven attacks, including cross-site scripting attacks, SQL injection attacks, path manipulation attacks, and others. The key technical tool will be character-level taint tracking for Java. We believe we can track the taint status of each character in every string manipulated by a web application, with little or no overhead. If we are right, this provides a powerful foundation for securing existing code: we can use the character-level taint information to prevent a broad variety of attacks by parsing the output of the web application and comparing the taint information to the structure of the parse tree. For instance, we conjecture that it is possible to detect all cross-site scripting, SQL injection, path manipulation, and HTTP response splitting attacks, with little performance cost and with no changes to the code. There is some data to suggest that fixing these vulnerabilities would fix about 3/4 of the vulnerabilities found in current web sites. Therefore, this work has the potential to have immediate impact.

We would also like to facilitate construction of new web services code, in a way that ensures security against these and other attacks. We note that command injection attacks (such as cross-site scripting and SQL injection) typically arise when web applications manipulate structured data (e.g., HTML or SQL fragments) using string concatenation. Because string concatenation does not respect the structure of these documents, it opens up the possibility of many types of attacks. Moreover, most web templating languages – which are widely used to facilitate construction of web services code – encourage these kinds

of command injection vulnerabilities. We propose to design new web templating languages that are inherently immune from command injection vulnerabilities, so that web applications built using our templating libraries will be free of these kinds of vulnerabilities. We believe that it is possible to design web templating languages in a principled way, to facilitate security; however, there is little prior research on this topic in the literature. We are also interested in building re-factoring tools to convert or migrate existing code over to these new web templating systems, though that will likely need to wait until the second year of this project.

In addition, we propose to demonstrate how new languages and system architectures can provide improved security for server-side web application code. We propose to facilitate the construction of *privilege-separated web applications*, where the web application code is factored into a small trusted security kernel, and the rest of the code receives only a minimal amount of privilege. In particular, we would like to explore architectures where most of the code runs with only the privileges of the user who logged into the web application. As a challenge question, we propose to study how to build a secure webmail system, with a guarantee that one webmail user cannot access other users' email (even in face of potential bugs in the web application system). We are especially interested in *reviewably secure systems*, where it is feasible to convince independent code auditors that the system provides security properties (such as the one mentioned for webmail systems) as part of a code review. This work will build upon our earlier research on Joe-E, and in this respect is a continuation of the project “Capability-based languages for security” initiated last year.

In addition, we propose to study secure on the client side, and in particular, how to provide security for browser extensions. For instance, some people partially attribute the popularity of the Firefox browser to the plethora of extensions available for it. However, there is no security model: a malicious security extension can compromise the entire browser and all of the user's web sessions. To mitigate these issues, Firefox has instituted an extension review process, where Firefox developers manually examine the code of new extensions before they can be listed on the official Firefox extensions web site (similar in spirit to Apple's review process for iPhone apps). However, because this review process is purely manual, it tends to be time-consuming, and there can be a significant delay before new extensions are approved for listing on the official Firefox site. We propose to study how this review process can be made more efficient, by providing tools to partially automate analysis of new extensions. We believe it should be possible to automatically infer an upper bound on what kinds of dangerous privileges a particular extension actually uses, and we conjecture that this can serve as a helpful starting point to help developers prioritize and focus their review of new extensions. In addition, we propose to develop new models for extension construction that better support this kind of review process, using ideas derived from proof-carrying code (particularly, policy-carrying code).

Our work is intended to be practical. We believe it could provide pay-off that cuts across multiple areas, including both financial and health infrastructures.

Code analysis, system monitoring; and malware detection

(i) *Deep Automatic Error Checking of Critical Software Infrastructure:* TRUST researchers building code analysis tools aim to improve our ability to test large software applications and to statically verify security properties of large-scale, security-critical software infrastructure, such as an entire operating system or web browser. The core challenge is twofold: designing automatic analysis techniques of sufficient precision and scalability to handle real systems with millions of lines of code almost automatically and, for the inevitable small percentage of cases that cannot be fully automated to understand what crucial

information the programmer can provide in the form of limited specifications that will render the task tractable.

We have been developing a tool, KLEE, that uses a variation on symbolic execution to automatically generate test cases that execute most statements in real programs. Our long term goal is to be able to take programs of 100K-1M lines and automatically run most statements in them. We are currently in the 10K or less size. Adding another zero or two will mainly require:

- More clever search heuristics. While the set of paths is exponential, the number of "interesting"' paths is not. We have developed (and will further develop) ways to merge equivalent paths (even when they differ superficially) and to reach unexecuted statements.

- More clever constraint solver tricks. While constraint solving in general is NP-hard, people program in particular, not in general. Thus, exploiting the regularities in the constraints generated by code can give exponential speedups.

In the short term (next six months) we are taking 50+ network applications and extending the techniques in KLEE so that we can get 90%+ coverage on them. Given that this code is network exposed, improving its security in a non-trivial way will be a significant practical result. For each bug we find, KLEE is able to generate an attack that will trigger it—i.e., the concrete packet sequence that when sent to an un-instrumented copy of the program will crash it. We will use this ability to focus developer attention on fixing these errors.

In the coming year we also plan to focus on scaling the new buffer overrun techniques up to both the entire Linux OS and Firefox. Our goal is to automatically check at least 90% of buffer accesses automatically, and to understand what the limits of static analysis are with respect to any remaining, unverified buffer accesses. No change in the level of funding is needed for this project.

(i) *Next Generation Infrastructure for In-depth Malicious Code Analysis and Defense:*

We continue to enhance our state-of-the-art infrastructure for the analysis of malicious and security-sensitive code at the binary level, and to apply it to a variety of challenges in analyzing and defending against malicious code. The flexibility of our infrastructure allows us to apply it to a variety of important security challenges, and the research on those applications provides feedback into what analysis infrastructure features would be most valuable. In the coming year, we plan to enhance and apply our analysis infrastructure in several different directions to take on a variety of further security threats.

In the area of automatic directed testing, we will apply our infrastructure for symbolic execution and test input generation to search for vulnerabilities in security-relevant software of both defensive and offensive varieties. For instance, vulnerabilities in defensive software like virus checkers or intrusion detection systems could allow malicious software to escape detection or even allow compromise themselves. On the other hand, vulnerabilities in malicious software such as botnet clients could allow new possibilities for containing their spread or preventing them from doing harm.

Building on our work this year in understanding the protocols that malicious binary programs use for network communication, we want to apply some related techniques to automatically understand the internal communications between functional elements in a malicious code sample. Such information about internal structure could, for instance, allow a decryption function used by a botnet to be extracted

and reused for network monitoring, and understanding the functional decomposition of a binary would allow many other kinds of analysis to be applied in a more scalable way.

Another challenging area for tainting and symbolic-execution based code analysis is control dependencies, in which one part of a program affects the execution of a later part not by directly modifying data that it reads, but by making control decisions that have an indirect effect on later execution. We will investigate moving beyond simple heuristics for control dependencies and developing techniques that apply more generally in both dynamic and static analysis contexts.

We also plan to use our analysis infrastructure to investigate architectures for building more secure systems in the future. For instance, one reason that present systems are difficult to secure is that software at many different levels must access and process sensitive information, and implementation or design flaws at any of these levels might allow compromise (e.g., allowing confidential information to be revealed, or allowing an attacker control of data he should not have). It is widely agreed that these difficulties could be ameliorated by shrinking the amount of code that must be trusted in this way, but it is difficult to do so while preserving all of the complex functionality of modern systems. In order to better understand the design possibilities for future secure systems, we plan to use our analysis infrastructure to study how sensitive information is processed in large systems (e.g., the combination of JavaScript code, a web browser, a windowing system, and an operating system that are involved in using an e-commerce web application). By examining which software accesses sensitive data in current-generation systems, we can look for patterns of unnecessary access and evaluate the prospects for reducing the amount of trusted code in new architectures.

Human Factors, Public policy, Decision-making, and Risk Management

Several subprojects in this topic area were initiated in 2008, with even greater interest and proposed activity appearing in 2009. We summarize some of the 2008 accomplishments here and present a larger range of planned activity in the section on 2009 plans.

(i) *Scaffolding for Human Computer Interfaces in Financial Infrastructures*: It is well known that most computer security failures result from human error, usually attributable to poor user interfaces. We have and will continue our successful work in developing robust user interfaces that are secure from attack. This work includes both study of end-user user interfaces (e.g., electronic banking) as well as institutional user interfaces (both inter-institutional and intra-institutional). We plan to implement prototype systems based on design principles and refine list of forensic techniques. We also plan user studies on various prototype systems (and contrast with existing systems), and user studies on forensic techniques. We will be releasing defense mechanisms for learning systems in hostile environments and we are close to spinning off a commercialized version of our usability testing system and an open source version of our testbeds.

(ii) *Fraud Detection in Consumer Reports:* This project seeks to determine empirically whether it is possible to detect identity theft by an analysis of a consumer report with no extrinsic information or interaction with the consumer. If such a determination is possible, it could drive policymakers to require consumer reporting agencies (CRAs) to engage in anti-fraud monitoring of reports (CRAs currently have negative incentives to engage in this analysis). With an affirmative fraud monitoring system in place, consumers could learn of identity theft in a positive manner (notice from a CRA) rather than the current situation, where consumers often learn of the problem in a negative way (such as being denied a loan or job, or being pursued by a debt collector, because of a polluted consumer report). Positive notification would mitigate the harms of identity theft and reduce losses to consumers and businesses.

We plan to inform the significant policy debate on the role that CRAs should play in preventing identity theft by studying whether CRAs could perform a positive anti-fraud function, and automate the detection of new account identity theft without consumer interaction or extrinsic information. Our objectives include:

- To determine whether it is possible to detect new account identity theft through automated analysis of consumer reports.
- To reduce the harm of identity theft by creating automated detection systems.

We wish to empirically test whether it is possible to detect new account fraud using only the information that a CRA would possess. We propose to test this through our relationship with ID Watchdog. In the FACTA Access study, ID Watchdog provided us access to verified victims of identity theft, a task that proved very difficult through ordinary recruiting. As a result of that collaboration, we learned that ID Watchdog obtains consumer reports from all 3 major CRAs on tens of thousands of its clients. Using ID Watchdog's large dataset of clients, we could examine a large number of consumer reports of both victims and non-victims.

(ii) *Combating Fraud in On-Line Advertising* On-line commerce is a rapidly growing aspect of our economy, and a lot of that commerce is driven by on-line advertising. Just like other aspects of our financial and commerce infrastructure are vulnerable to phishing attacks, spam, denial-of-service attacks, and so on, on-line advertising is vulnerable to various type of fraud, including click-fraud. Successfully committing ad fraud yields direct monetary gains for attackers at the expense of the victims. Thus, it is natural to consider online ad fraud in an economic context.

As part of the TRUST project, and in collaboration with researchers at Google, we have been modeling online advertising and studying various fraud and pricing issues.

During this past year we have continued to investigate online advertising. We use economic analysis because, unlike many online security threats, ad fraud is primarily motivated by financial gain. Successfully committing ad fraud yields direct monetary gains for attackers at the expense of the victims. We plan to continue our research on online advertising and fraud. We will also continue to study a variety of privacy issues related to large scale data management.

(iii) *Behavioral biases in personal information security:* We propose to investigate privacy and security decision making through the theoretical lenses of behavioral economics, and using the tools and methodologies of experimental economics, in a series of human subjects experiments. Our goal is to inform the design of privacy and security technologies through behavioral studies, in order to anticipate and mitigate potential human cognitive and behavioral biases that emerge in the context of privacy and security decisions. In particular, we will focus on an “illusion of control” hypothesis and its impact on privacy and personal information security decision making.

The PI has already obtained IRB approval for two studies and is negotiating approval with CMU IRB for the third study. With sufficient budget to recruit and pay human subjects to participate in the studies, the researchers expect to be able to run the studies within the first 3-5 months from the start of the project, leaving sufficient time for possible follow-up studies (based on the results of the first three studies), as well as the writing and dissemination of the results. The researchers expect to submit the completed results of at least a subset of the studies within the 12 months since the start of the project, and submit a cumulative journal article (combining the different studies) by the end of the project.

(iv) *Characterizing negative externalities and their effect in security decision-making*: Deployment of security technology or practices in a network of non-cooperative agents suffers from strong negative externalities, which attackers can use to gain the upper hand. Indeed, the security investments of each agent impact the network as a whole, but do not necessarily translate in increased security for the agent investing. For instance, an individual who spends significant time and money patching and securing her machine before connecting it to the Internet nevertheless remains at the mercy of attacks that rely on other, unprotected machines, over which she has no control.

This project will seek a formal characterization of the impact of negative externalities on global network security, both from the attacker and the target's perspectives. We will combine formal, game-theoretic, modeling with behavioral experiments (user studies), and with data collection and analysis to make original contributions to the economics foundations of information security, and to demonstrate the practical benefits of this research.

(v) *Path of Identity Theft* This project will deconstruct the "path" of identity theft. It will explore the steps taken by identity thieves in actual situations where they attempt to take control of a victim's identity. The goal is to identify and create a taxonomy of early indicia of fraud, in order to prevent and mitigate the harm of identity theft. Once understood empirically, this knowledge could be used to develop effective early detection systems for fraud, and guide federal regulators in the specification of identity theft "Red Flags," which are now required under the Fair Credit Reporting Act.

(vi) *Economics of Managing the Interdependent Security (IDS) Risks.* When networked parties (individuals and organizations) make decisions about their systems' security, they impact the security of the overall networked infrastructure. Thus, Internet security is interdependent security (IDS). We will investigate the following questions: 1) What are the networked parties' incentives about security choices, given the interdependent nature of network security; 2) How would the introduction of new policies and regulations - to mitigate the divergence of individual and socially optimal incentives - affect the networked parties' incentives. 3) How legal and regulatory channels could be jointly used to improve information structure? The latter question is relevant for establishing the framework that will facilitate enabling prosecution of international crimes (such as bank fraud and identity theft) driven by Internet insecurity; and reduce inefficiencies driven by the separation of rights for information ownership and its control, in application to privacy and data collection issues.

We believe that Internet IDS risks management will remain sporadic unless major improvements of information structure occur. In this project, we will use game-theoretic modeling to evaluate several options (such as disclosure rules, liability regimes, introduction of mandatory user certification, and other public policies and regulations) that could improve information. Our goal is to investigate which policies will be the most effective for improving information structure. We will focus on several information inefficiencies and discuss how to alleviate them.

This project will study how IDS risks affect the networked parties' incentives to invest in network security in the settings where informational asymmetries (and thus, moral hazard and adverse selection) are focal. Our project will combine economic and legal analyses. In the former, we will apply the advances of game and contract theory and microeconomics to industrial organization and regulation of the Internet. We will make original contributions to the foundations of economics of information security, and specifically into theoretical insights about how IDS risks could and should be managed. We will explore the public policies and assess their impact on reducing incentive conflicts via improvements of information structure.

2.6.2 Health Infrastructures

Thrust Leader: *Janos Sztipanovits (Vanderbilt University)*

The health infrastructures research area has the following three objectives:

1. Privacy Modeling and Analysis: Development of logical foundations and practical methods for expressing and enforcing privacy policy and other policies and regulations adopted by hospitals, clinics, and other healthcare organizations.
2. Architectures for Trusted HIS: Development of formal, model-based design technology that integrates architecture modeling with privacy and security modeling for generating trusted HIS.
3. Secure Home-Based Health Monitoring Systems: Development of systems that address consumer privacy protection recommendations (to increase consumer adoption) and integration of these systems in the HIS of healthcare providers.

During the next reporting period, we plan the following activities in each of the three areas:

Privacy Modeling and Analysis – We will investigate two fundamental questions: (1) How does one describe formally and precisely privacy and security policies such that they faithfully capture the meaning of laws and regulations, clinical practices and provide foundation for formal reasoning and enforcement?; (2) How does one extract formal operational rules (interactions, information flows and restrictions) from audit logs of data access in a complex, dynamic clinical environment?

Privacy is an increasingly important business concern in healthcare, financial services, and other organizations. We will build on our previous work to develop *approaches for modeling systems that handle sensitive information*, *languages for specifying privacy policies*, and *algorithms for their enforcement*. We also plan to develop more extensive policy examples, such as a significant portion of HIPAA in a machine-processable format, and begin a new direction *incorporating risk management concepts into security analysis and decision making*.

Architectures for Trusted HIS – The goal is to develop a formalized, model-based design approach to trusted HIS. The proposed approach is based on the theoretical foundations and tools for Model Integrated Computing (MIC) developed at Vanderbilt, Platform Based Design (PBD) advocated at UC Berkeley and the results of the foundational work on privacy modeling and analysis at Stanford, CMU and Vanderbilt.

The combination of MIC, PBD, and SOA techniques can enable the design of complex HIS to ensure reliability, performance, privacy and security beyond what can be achieved by current ad hoc practices. We will answer the following questions: (1) How to provide deep integration for models used in functional design and for privacy and security policies?; (2) How to create a run-time architecture that can be model driven and include components for enforcing dynamic security and privacy models?; (3) How to validate the proposed approach and how to show its effectiveness using relevant metrics?

The Vanderbilt team has developed a theory for the structural semantics of DSMLs. Using this theory, metamodels are translated into constraints so additional constraints can be injected directly into the resulting interpretation of the metamodel. We will extend this work in new directions on the separation of static and dynamic structural and policy constraints, where static constraints can be checked at design time while dynamic constraints are checked at run time.

Answering the second question requires progress in extending standard SOA execution platforms (such as Apache ODE) with a policy evaluation point that is Prolog based. We will create a suite of model transformation tools that can translate dynamic structural and policy constraints into Horn logic.

To answer the third question, we have designed and implemented the Model Integrated Clinical Information System (MICIS) platform that serves as an evolving collaboration platform between TRUST researchers and the Vanderbilt University Medical Center. The first real-life experiment using MICIS has been the Sepsis Treatment Enhanced through Electronic Protocolization (STEEP) which is currently being tested and validated at the VUMC Simulation Center and scheduled for clinical trials this year.

Secure Home-Based Health Monitoring Systems – We address fundamental questions of scalability and extensibility; robust, reliable and QoS-aware service delivery in a dynamic wireless environment and privacy and security requirements for a medical practice. Our preliminary work has designed and built DexterNet, a small-scale wireless sensor networking environment for remote healthcare. Privacy and security are primary considerations in our system design. Our current system design features built-in secure communication components which are adaptively implemented for different networking environments to protect the physiological information and video streams that carry the footage of the human movement. As a next step, we will investigate detection mechanisms to identify anomaly situations caused by malicious attacks or device malfunctions.

We will also extend the system to a new medical domain, the post-operative home-based monitoring of Congestive Heart Failure (CHF) patients. We will enhance the system to collect, analyze, and securely transmit heart failure patient medical data from different home medical devices to the clinical information system. We will develop a decision support system for the treatment management of CHF patients based on the STEEP toolset. We will also support the delivery of treatment recommendations from decision support system to the patients. Research in this area will be supported by Phil Kuryloski, a Cornell Ph.D. student who will join Berkeley as a post doctoral researcher.

2.6.3 Physical Infrastructures

Thrust Leaders: *Steve Wicker (Cornell University)*

During the next reporting period, we plan activities in the following projects:

TRUST SCADA Testbed Infrastructure and Experiments – The objective of this project is to provide an experimental testbed with well documented examples for SCADA research in the TRUST community based on tool-supported experimental machinery used to build prototypical experiments. The SCADA testbed will be equipped with new tools and capabilities to support easy experimentation. Experiments will be conducted, documented, and results published for other researchers to use as examples and templates to build their own experiments. The testbed will have remote experimentation functionality. Key properties will be modularity, composability (i.e., plug-n-play capabilities), and configurability via graphical interfaces.

Analysis of Networked Embedded Control System Attacks and Defenses – Growing concern has recently risen on the vulnerabilities of the country's strategic physical infrastructures to security attack. The extensive use of information and communication technology has made easier to gain access to system components, increasingly connected to the internet. Distributed Control Systems (DCS) and SCADA infrastructures are of particular interest, as they are usually the basis for sensing and control of large critical infrastructures such as power, gas, water, and industrial plants. In this project, we plan to develop system theoretic tools for the design and analysis of attack detection schemes and attack-resilient estimation and control algorithms together with an evaluation of the potential consequences of an attack.

Significant advances are required in modeling attacks, developing model based detection schemes specific to cyber-physical systems and finally in designing attack-resilient estimation and control algorithms.

Empirical Investigations of Privacy – Individuals' privacy concerns emanate from both online and offline sources: information sharing on social networking sites, new location-based services such as Google StreetView, and offline venues such as video surveillance and other systems that network physical places. Often the privacy concerns and objections of individuals fall outside what existing privacy law protects. For example, public and government objections to Google StreetView motivated the company to institute facial and license plate blurring within images of public streets, sidewalks, and street side facades throughout the StreetView product. Similarly, individuals object to police use of CCTV on public streets despite a legal framework that offers little to no protection. While multiple theories of privacy—and its relationship to technology—have been proposed, little empirical effort has been undertaken to document and understand how individuals conceptualize privacy on the ground. What problems do individuals perceive as "privacy" problems? How do they frame and articulate their objections? What animates their concerns? What does this tell us about the sufficiency of existing policy and technical approaches to privacy protection? To that end, this project is an empirical investigation of several datasets that contain information about the privacy objections individuals raise in relation to Internet applications and services. Current US law provides limited protection for privacy and individuals experiencing privacy harms. It often shies away from the added publicity that generally attends litigation, therefore, examining law suits problems limited insight into contested privacy issues on the Internet. Fortunately, through the collection of data and access to existing data sets about interpersonal and group efforts at norm enforcement in social networks, other forms of private ordering, dispute resolution, and other "below the radar" efforts to protect privacy we can assemble a rich understanding of privacy's meaning in an everyday life influenced by these new technologies.

Privacy Concerns in Upcoming Demand-Response Systems – TRUST researchers have been engaged in an effort to explore the confluence of sensor networking, power distribution, privacy, and security issues that will emerge from a substantial increase in power system monitoring at the consumer level. Future work in this area includes further improving behavior extraction algorithms by using Markov Chain and Lempel-Ziv based predictive algorithms originally used, and already proven to be effective within the context of home automation. It is also plan to further develop the disclosure metric, which associates data quality (accuracy of readings, time resolution, types of readings, etc) from a particular source with the information that may potentially be disclosed by the data.

Smart Meters as a Security Threat – This project will leverage the TRUST SCADA testbed to show how distribution stations can serve as an entry point for hackers looking to shut off power to thousands of homes. By incorporating smart meters into the data acquisition phase of an electric power SCADA system, it can be shown that if the hacker is able to gain access to the utilities back office systems via the smart meter, he/she can disrupt the normal operation of substations and distribution of power. As it stands now, utilities across the U.S. envision a smart meter network consisting of a back office system that polls smart meters. Using this as a vector of entry, the hacker could stage false warnings or initiate false alarms stating that energy consumption levels, within a portion of the utilities service territory, have gone beyond the limits capable of being supported by the utilities infrastructure thus causing the utility to perform a rolling blackout (load shedding) for this area. If the utility uses a single portal to control—remotely—the substations as well as the smart meter networks within their service territory, then the hacker not only has the ability to falsely convince the utility and grid operators that load needs to be shed, but also has the ability to shut off or alter the behavior of entire substations. Having demonstrated the

vulnerability, we propose to explore schemes for securing this entry point and demonstrate the potential for our security schemes through the use of the TRUST SCADA testbed.

Data Aggregation Schemes for SCADA – This project will develop a theory of aggregation of SCADA data through in-network processing and combine it with a routing scheme. Given the potentially enormous quantities of data collected by SCADA, it would be useful to utilize an aggregation scheme that considers spatial correlation of information. Variations of direct diffusion and related content-aware routing schemes will be considered. In addition, the theory could consider a network using nodes with IP addresses, given the ample support to such networks and the all-IP features of next-generation cellular systems.

Defense-in-Depth Intrusion Detection and Intrusion Tolerant Control for SCADA Systems – The cyber-physical security of real-time, continuous systems necessitates a comprehensive view and holistic understanding of network security, control theory and the physical system. Ultimately, any viable technical solutions and research directions in securing SCADA systems must lie in the conjunction of computer security, communication network and control engineering. However, the very large installed base of such systems means that in many instances we must for a long time to come rely on retrofitted security mechanisms, rather than having the option to design them in from scratch This leads to a pressing need for deployable, robust, SCADA-specific intrusion detection systems (IDS) and intrusion tolerant control techniques. The goals of this effort are to develop IDS technology and intrusion tolerant control techniques that can (1) efficiently detect and block cyber intrusions into SCADA systems in real operational environments, and in real-time, (2) without interrupting the control performance of the protected system, (3) without creating extra operational burden or operational reservations due to false alarms, (4) in the presence of both malicious and messily benign network traffic. The system must operate in a real-time, robust fashion, with performance adequate to meet the demands of the dynamic cyber-physical interactions inherent to SCADA systems.

Intrusion Detection for SCADA Systems – A hybrid, two-stage intrusion detection system (IDS) for mobile ad-hoc networks has been developed through previous TRUST funding. The framework for this IDS has the flexibility to also monitor physical infrastructures. The goal of this project is to investigate the deployment of the IDS within Supervisory Control and Data Acquisition (SCADA) systems. Our approach will leverage the collaborative development of the SCADA testbed within TRUST (UCB, CMU, Vanderbilt) to enhance the intrusion detection capability within the remote terminal units (RTUs) of the SCADA architecture. Our work will start with the assumption that the SCADA system has been compromised; our objective is to identify the deviancy of the compromised nodes and minimize the negative consequences of those nodes on the overall SCADA system.

Reconciling Cyber Security Policy for Government and Private Networks – This project will undertake a comprehensive examination of the role of public-private partnerships (PPPs) in cyber security policy. PPPs have been a central feature in cyber security policy discussions as a means of coordinating government and private-sector efforts in cyber security. The notion of public-private partnership is attractive because it avoids direct regulation of the private firms that operate IT and communications infrastructure, but examples of effective partnership structures for cyber security are scant. Indeed, basic questions about cyber security PPPs' purposes and functions remain unanswered. This project would offer answers to these basic questions by (1) providing a comprehensive history of cyber security PPP proposals; (2) comparing PPPs in cyber security to those in other domains; and (3) identifying PPP structures that would address technical needs of cyber security while fitting into the broader fabric of cyber security law and policy.

Privacy Preserving Estimation Algorithms in Participatory Sensing Systems – This project will investigate the development of estimation algorithms for information systems based on participatory sensing data. The systems of interest are those in which users are given incentives to share their data from cellular devices. Examples include traffic monitoring systems, location based services, and participatory environmental monitoring systems. The specific problem of interest is the reconstruction of the state of a distributed cyber-physical system from sparse measurements shared by users. The approach considers the use of a fixed amount of perturbation to data available to distributed estimation techniques. We will perturb individual measurements with known noise (hence permanently destroying their individual value and thus protecting the privacy of the users) while preserving aggregate quantities of interest which we are of interest to state estimation algorithms. We will also illustrate the algorithms developed with the Mobile Millennium system, a traffic information system jointly developed by Nokia and UC Berkeley, currently operational in Northern California.

Privacy-Aware Design Rules for Networking Infrastructure – This project will develop a framework for privacy-aware engineering design. The Fair Information Practices proposed in Records, Computers, and the Rights of Citizens (HEW 1973) can be translated into privacy-aware engineering design rules. These rules begin with an absolute imperative to limit information collection to explicit and publicly expressed mission requirements. This simple imperative flows into a mandate for distributed information processing, anonymity-preserving information routing and tracking functions, and strong distinctions between identifying active equipment and identifying operators and owners. The outcome of this project will be a set of clear design rules and several test cases, including 3G cellular, which demonstrate that full functionality can be retained without the massive accumulation of personal information.

A Low Power Hardware Platform for Secure Embedded Systems –The goal of this project is to develop an ultra low power hardware platform for secure embedded sensing. We plan to use a combination of expertise in ultra low power asynchronous processor design and rapid prototyping using an FPGA-based approach to evaluate the necessary trade-offs between hardware and software. The net result of this project will be a security-aware ultra low power asynchronous microprocessor suitable for embedded sensing.

Scalable Code Reuse Detection in Binary Code – This project will develop low-level code analysis techniques that scale analysis by taking advantage of software component reuse. There is tremendous evidence that software components are reused in new programs. Unfortunately, unless code reuse is detected, software security analysis will end up analyzing the same code again and again. However, if we could detect code reuse, we could make use of that information to prevent unnecessary analysis and thus increase scalability. Although automatic code reuse detection can potentially benefit security analysis, there are significant challenges. First, we need to develop techniques that are appropriate for low-level code since physical systems often use a significant amount of low-level programming. In addition, malware authors are unlikely to provide source code in order to aid malware classification and categorization. Copyright infringers are also unlikely to provide source code that demonstrates direct plagiarism. Second, the value of automatic code reuse detection in analysis is likely proportional to the size of the dataset. The more code we analyze, the bigger the database of summaries we can use in subsequent code reuse detection-enabled analysis.

3 EDUCATION

3.1 Goals and Objectives

In education, TRUST is generating learning materials, providing dissemination structures, and establishing broad educator communities. Our education activities have reached undergraduate and graduate students, postdoctoral scholars and junior faculty, and industry professionals to address the technical, policy, and economic issues essential to improving cyber security and trustworthy systems.

Affiliated with TRUST is a multi-disciplinary team of students, post doctoral scholars, research scientists, and faculty from a world class research group of universities providing a unique breadth and depth of research expertise and accomplishment in cyber security and critical infrastructure protection. The Center research team is supported by students and faculty from partner institutions with whom the Center collaborates to provide unique opportunities for female and underrepresented minority students and faculty to engage in cross-institutional activities.

The TRUST education mission is to educate the next generation of computer scientists, engineers, lawyers, policy makers, and social scientists in the field of cyber security and trustworthy systems. Specific TRUST education goals are to:

1. Provide graduate students with research opportunities in cyber security and trustworthy systems topics.
2. Provide academic-year and summer research opportunities to undergraduate students.
3. Increase the number of women and underrepresented students that pursue graduate education in cyber security and trustworthy systems.
4. Provide academic courses and degree programs supporting TRUST research and education mission.
5. Prepare and support HSIs, MSIs and HBCUs faculty in the teaching of TRUST related research topics.
6. Develop technology to assist with the dissemination and outreach efforts of the Center.

Research and education are interwoven into all Center activities. TRUST summer programs, workshops, technical series, seminars, and internships leverage the materials and tools developed in our research projects. These materials and tools also become module content and project profiles distributed on the TRUST Academy Online (TAO). A goal of TRUST is to disseminate education materials for engineering, computer science, law, public policy, economics, and social science students working in cyber security. In the TAO we have developed teaching modules that can be incorporated into diverse curricula, ranging from privacy modules that can be taught to engineers working on SCADA control systems to cryptography modules that introduce digital rights management concepts to law students.

3.2 Performance and Management Indicators

During this reporting period, a compressive review and evaluation of TRUST Academy Online (TAO) was performed. This review was fundamental to the portal's current interface design, information architecture, and Metadata technology. Beginning in April 2008, the TAO's reconfiguration represented the second launch of the repository—emphasizing TRUST research thrusts and course materials. Additional user survey feedback helped refine the portal's technology and user functions. We are in the process of implement data collection strategies that will track the use and dissemination of TRUST education materials from the TRUST Academy Online (TAO). Analysis of the TAO online access statistics indicate that approximately 25% of people accessing the TAO download a resource in the

repository. That said, we will further develop portal survey and user-rating technologies to help us better understand our online community and their usage of the TAO.

To support both quantitative and qualitative analysis of TRUST education programs, we will continue to use participant and mentor surveys, focus groups, in-depth interviews, rubrics, program metrics, and electronic portfolios as methods for data collection. These are intended to capture the effectiveness of our programs and the educational and professional development value added to participants. We will also expand our participant tracking efforts, especially for Center students after graduation, to continue contact with participants, monitor where they are in their careers, and better understand the impact affiliation with TRUST had on their professional development and advancement. Working with organizations like the National Center for Women and Information Technology (NCWIT), the Anita Borg Institute for Women and Technology, and the Assessing Women and Men in Engineering Project (AWE), will support our assessment efforts while disseminating our results to a broader audience as well as our TRUST Academy Online community.

3.3 Current and Anticipated Problems

No significant problems were encountered during the reporting period. No significant problems are anticipated in the next reporting period.

3.4 Internal Education Activities

During this reporting period, the Center education activities were focused in the following areas:

1. Learning Technology Infrastructure
2. Undergraduate Programs
3. Graduate programs
4. TRUST Summer Seminars
5. TRUST Domain Workshops
6. Recurring and Significant Presence at Key Education Conferences.

Significant impact was seen in areas 3-6 and progress was made in areas 1-2 by establishing the infrastructure for the TRUST learning modules repository, establishing a set of pilot course modules within this repository, and bringing together material from the various TRUST partner institutions in an integrative learning material generation exercise.

The items below describe in more detail specific education activities of the TRUST Center during this reporting period.

Activity Name	TRUST Academy Online (TAO) Portal
Led by	Larry Howard (Vanderbilt)
Intended Audience	Students, Faculty and Industry Professionals
Approx Number of Attendees (if appl.)	Unlimited; portal and content is open access via the Internet

The Trust Academy Online (TAO) Portal (http://tao.truststc.org) is a vehicle for online community outreach for the TRUST Center. Its initial emphasis was to provide educators access to sets of learning materials contributed by center investigators, institutions, and partners and is used to disseminate learning materials developed or contributed by educators participating in the TRUST Center.

TAO content is bundled into "profiles" that provide descriptions, metadata, and complementary scaffolding resources such as guides to their use for teaching and learning in the classroom, lab, or online. The profiles include a variety of learning materials such as PowerPoint presentations, lecture notes, case studies, class assignments, related web site links, video clips, and "rich" media content.

Figure 1: The TAO Portal Front Page

During this reporting period, we continued our effort in the development of "visual storytelling" as the vehicle for this communication. In project profiles, lightweight multimedia shorts are used to quickly present the essential details of a project's work in a way that is accessible to a broad audience. We used a small group of TRUST research projects at Vanderbilt to "prototype" and refine this concept. TAO media designers then collaborated with project graduate students to identify story elements and produce multimedia resources. The profiles were then established and fully populated by the project teams. Given that TRUST projects comprise a fairly stable portfolio, we feel this strategy is scalable to incrementally include all TRUST projects, resulting in a rich information flow.

Accompanying this extension in audience, we enhanced the user experience on TAO. A keystone element in our strategy was the introduction of "visual browsers" as an alternative way of presenting and selecting profiles from collections. This navigation vehicle was influenced by innovations such as Apple's "cover flow" browsers and its distinct quality makes a significant contribution to the visual impact of the portal. At the same time, we have retained the tabular, text-based browser of the courseware profiles as a navigation alternative. These changes resulted in an increase in TAO portal usage of 71% over the last year and an increase in the number of learning modules and courseware, currently 176 resource files produced by 46 TRUST researchers.

Figure 2: The Visual Browser for TAO Courseware Profiles

To further our continuing commitment to provide educators and other users online access to materials and resources produced by TRUST researchers, the TAO has been registered as a collection in the National Science Digital Library (NSDL), the Nation's online library for education and research in Science, Technology, Engineering, and Mathematics. Using the OAI-PMH metadata harvesting protocol, the NSDL now will routinely import metadata from the TAO's courseware and project profiles and will support searching this metadata from within the digital library. Registration of the TAO in the NSDL is part of our ongoing efforts to broaden awareness of the Center's research and education missions. Membership in the

NSDL and support for the OAI-PMH metadata resulted in a three-fold increase in the number of search engine visits (e.g., accessing the TAO via a Google search result) and promotion of the portal at TRUST education seminars and workshops has significantly increased the exposure of the TAO and greatly increased usage.

Activity Name	TRUST Courseware Modules and Projects
Led by	Larry Howard (Vanderbilt), Kristen Gates (UC Berkeley)
Intended Audience	TRUST portal users: students, faculty, researchers, and industry professionals
Approx Number of Attendees (if appl.)	Unlimited. Portal and content is open access via the Internet.

The TRUST Academy Online (TAO) is an online repository for TRUST Courseware Modules and Projects. Accessible by the public, the TAO contains learning materials available at no cost and enables educators access to leading-edge research and teaching materials specific to trusted systems technology and policy issues. The purpose of the courseware modules is to provide learning materials that are assessable via the TAO portal that are usable by teaching faculty as course content, lecture material, and supporting information for higher education courses. Modules consist of a variety of learning materials, including PowerPoint presentations, lecture notes, case studies, class assignments, related web site links, and video clips.

TRUST researchers incorporate their findings and methods, whenever possible, into the standard curricula addressing operating systems, programming languages and compilers, analysis of algorithms, networking protocols, and databases. The primary goal for the TAO Portal is to make a body of these curricular materials available to the larger educational community. Courseware development aims at three areas of research: Security Technology, Systems Science, and Social Sciences. It is anticipated that curriculum development based on this courseware will follow different trajectories resulting in materials of different granularities, from individual modules to complete courses and lower division to the advanced graduate level.

Building on our third-year inventory, the portal now host: 49 contributing members, 11 projects, 36 courseware profiles and 235 file and link resources. The table below lists TRUST courseware modules available on the TAO Portal, grouped by category. The table includes both TRUST-developed modules as well modules created by non-TRUST members who have contributed content based on collaboration through Center outreach programs.

Information Security	
Information Security: Principles and Practices	As businesses and consumers become more dependent on complex multinational information systems, the need to understand and devise sound information security systems has never been greater.
Role-Based Access Control	The purpose of this module is to teach the concept of Role-Based Access Control to chemical engineering students in a senior capstone course. This module makes use of anchored learning and the legacy cycle in particular as a pedagogical tool to engage students in learning.
Cryptography	Introduction to basic models and principles of cryptography. Topics covered include symmetric key cryptography, asymmetric key cryptography, key distribution and management, hash function, message authentication code, digital signature and authentication protocols.

Access Control in Distributed Systems	Treatments of access control and authorization typically focus on individual mechanisms (access control lists, capabilities) or types of policies. This module seeks to isolate the basic concepts underlying many types of authorization in distributed systems, and then to show how this perspective can lead to more flexible and sound access-control implementations.
Information Security	This security course covers 4 main topics: Cryptography, Access Control, Protocols and Software. The focus is hands-on practice by tackling real-world examples. While not sidestepping the theory, the emphasis is on developing the skills and knowledge that security and information technology students need to face their challenges.
Information Security for Everyone	Information Security for Everyone is a course that is designed to teach the principles and practices that all computer users need to keep themselves safe, both at work and at home. By presenting best practices along with a small amount of theory, trainees are taught both what to do and why to do it.
Information Security Basics	Information Security Basics is designed to teach entry and mid-level IT workers the technological fundamentals of information security. The goal of this course is to provide trainees some preliminary knowledge of computer security to help in identifying and stopping various cyber threats.
Business Information Continuity	Business Information Continuity is a course that will train business managers to respond to varying threats that might impact their organization's access to information. Business Information Continuity provides requisite background theory and recommended best practices needed by managers to keep their offices running during incidents of different types.
Information Risk Management	This is an intermediate level course covering topics on information assets, identifying risks, and management processes highlighting best principles and practices.

Network Security	
Internet Worms and Botnets	Internet Worms and Botnets: In the past decade, large-scale Internet epidemics have profoundly demonstrated the threat posed by self-propagating programs ("worms"). More recently, attackers have refined these techniques to develop huge sets of compromised Internet hosts collected together into “botnets”.
Network Security	This course provides an introduction to the principles and practice of network security. Topics include: security threats in networks, principles for providing security mechanisms (cryptography, key management, message authentication), practice of securing systems (PGP, IPsec, SSL), and recent research topics in security.
Web Programming and Security	The web uses complex applications that run on heterogeneous browsers that may be built using programming technologies such as Javascript, AJAX, Google Web Toolkit, Apache Struts, Java Server Faces, and Rails. This course covers how core web technologies work; common security vulnerabilities; and how to build secure web applications that avoid them.
Network Security War Stories	Networking War Stories is a historical survey of the Internet and Web security attaches from the 1970 to the more recent Zotob Virus of 2005
Networking: Background and Overview	Networking: Background and Overview is an introduction to a communication network, common network taxonomy and network protocols used by the Internet and the Web.
Web Security and Online Identity Theft	This lecture describes the phishing problem, which involves deceptive email and malicious web sites that steal user passwords, and examines several current defenses against phishing attacks. The defenses developed by TRUST researchers include browser extensions that detect malicious web sites, create customized "hashed" passwords, leverage user-chosen images to identify

	servers, and modify browser…
Computer and Network Security	The course covers principles of computer systems security. We will discuss various attack techniques and how to defend against them. Topics include network attacks and defenses, operating system holes, application security (web, e-mail, databases), viruses, social engineering attacks, privacy, and digital rights management. Course projects will focus on building reliable code.
Network Security Programming	Network security protocols and applications, cryptography algorithms, authentication systems, intrusion detection, network attacks and defenses, system-level security issues, and how to build secure systems. This semester is web programming oriented.

Trustworthy Systems	
Trustworthy Computing Curriculum	Most universities teach elements of trustworthy computing through upper-class specialist courses and through topic-level coverage in systems courses. Unfortunately, the requirements and curriculum flexibility for CS undergraduate majors mean that a significant number of undergraduates miss much of the materials.
Fault-tolerant Distributed Computer Systems	Distributed systems are difficult to build and understand, for all sorts of reasons. Failures are common in large systems, and we can't let them shut the application down; a well-engineered system will tolerate failures and repair any damage they caused.
Pattern Oriented Software Architecture	Software patterns have revolutionized the way developers think about how software is designed, built, and documented, and this unique book offers an in-depth look of what patterns are, what they are not, and how to use them successfully.
Secure and Scalable Event Notification	Large enterprises lack technologies for building scalable trust platforms -- namely, technologies for building policy databases and spreading the policy enforcement actions across a potentially large number of nodes.
Security in Embedded System Design	Today, security in embedded system design is increasingly being recognized as an important aspect of such applications. The material contained here provides an introduction to two relevant security technologies: role-based access control and partitioning. Two additional courseware materials describe case studies based on material presented here.
Access Control in Embedded System Models	This is a case study that follows the introductory courseware on security in embedded system design. Specific examples highlight how the RBAC technique can be incorporated into a model-based development process.
Security Modeling and Analysis	This is a case study that follows the introductory courseware on security in embedded system design. The case study shows a simple domain-specific modeling language that could be used create of embedded systems and their security aspects, and an analysis tool that validates that the system design satisfies the desired properties.
Secure Software and Network Assurance	This course covers secure programming practices necessary to secure applications against attacks and exploits. Topics covered include fundamental concepts of secure software development, defensive programming techniques, secure design and testing, and secure development methodologies.

Social Sciences	
Data Privacy in Biomedicine	The integration of information technology into biomedical environments has enabled unprecedented advances in the collection, storage, analysis, and rapid dissemination of patient-specific data. Many organizations need to share data for various purposes, such as quality assurance, public health, and basic research.

Information Law and Policy	This course is divided into three sections that correspond to information and law and policy issues iSchool students are likely to face in their role as individuals, as designers and developers, and as employees in an information society.
Internet Policy Challenges in a Global Environment	The internet - as a global, "always-on" platform - poses unique challenges to legal and political frameworks premised on territorial jurisdiction. Operating in this global marketplace exposes companies, and sometimes individuals, to conflicting normative, legal and political commitments.
Legal and Ethical Issues in Cybersecurity	The primary legal obstacles to conducting cyber security are not outright prohibitions but rather the difficulty of determining which of a large set of complex statutes might regulate a given research project. Privacy, computer abuse, tort, and contract law are all potentially applicable.
Security and Other Values	This module uses two case studies to examine how the security of computers and networks relates to other policy values. The first case study examines the relationship between security and privacy in the specific context of cyber security research.
UCB Unblinking Symposium	Privacy is a complex and often abstract topic: this symposium addressed "visual privacy," a subset of the much broader topic of data privacy. In a rapidly evolving environment of unblinking eyes, technologically perfected recollections, and permanent visual records, what will it mean to have privacy?
BFOIT	Berkeley Foundation for Opportunities in Information Technology, a project of the International Computer Science Institute, supports historically underrepresented ethnic minorities and women in their desire to become leaders in the fields of computer science, engineering and information technology.
Software And Internet Law	Software And Internet Law, Third Edition, is an excellent choice for courses that cover all aspects of computer law. This careful exploration of computer software, the Internet, and e-commerce focuses on intellectual property, licensing, and antitrust law to give students a solid introduction to the full range of the field.
Privacy's Relationship to Technology	The Fourth Amendment guarantees the right of the people to be secure in their persons, houses, papers, and effects, against unreasonable searches and seizures.
Cyber Ethics	Cyber Ethics is designed to teach students the proper techniques with which to approach the difficult ethical dilemmas that arise from using the modern Internet. In addition to providing students with the skills to assess future ethical dilemmas for themselves, Cyber Ethics also looks at some of the more pressing concerns related to Internet usage today.
Cyber Law and White Collar Crime	This training will highlight the various computer crimes and appropriate response by first defenders (line officers). Participants learn legislation and law enforcement efforts to control such crimes.

The table below lists TRUST projects showcased on the TAO portal. TRUST projects demonstrate research and technology initiatives across the Center.

TRUST Seminar Series 2009	The TRUST seminar series is a weekly series of technical talks given by leading academics and industry experts in the field of cyber security, trustworthy systems and privacy.
Sting: Defense System against Worm Attacks	Worms such as CodeRed and SQL Slammer exploit software vulnerabilities to self-propagate. They can compromise millions of hosts within hours or even minutes and have caused billions of dollars in estimated damage.
Panorama: Analysis	Information Access and Processing behavior is the fundamental trait of

Infrastructure for Malware	numerous malware categories (including keyloggers, password thieves, network sniffers, stealth backdoors, spyware and rootkits), which separates these malicious applications from benign software. Panorama is a system that detects and analyzes malware by capturing this fundamental trait.
CareNet: Environment for Remote Healthcare	Recent advances in wireless sensor networks have made it possible to deploy wearable sensors on the bodies of patients in a residential setting, allowing continuous monitoring of physiological signals (such as ECG, blood oxygen levels) and other health related information (such as physical activity levels).
Foundations for Intrusion-Tolerant Services	This project seeks to develop fundamental theories and frameworks that support intrusion-tolerance in distributed systems, and to use these advances to construct more highly assured and/or more scalable and efficient intrusion-tolerant services than have been achievable to date.
Electronic Medical Records	The objective of this project is to develop cyber security, privacy science and technology needed to support the development of patient portals, a revolutionary new way for the interaction between medical patients and their doctors.
Active Localization for Safety and Security	This project explores how automated technologies using sensor and camera networks, coupled with systems to process this data in real-time, can be used to automatically assist with Safety and Security. As these systems are typically ubiquitous that can observe people, we also explore approaches to provide privacy while accomplishing these goals.
SCADA	The Supervisory Control and Data Acquisition System (SCADA) monitor and control real-time systems. SCADA systems are the backbone of the critical infrastructure, and any compromise in their security can have grave consequences. Therefore, there is a need to have a SCADA testbed for checking vulnerabilities and validating security solutions.
SUPERB	The Team for Research in Ubiquitous Secure Technology (TRUST) sponsors undergraduate students from diverse backgrounds and cultures, to participate in the Summer Undergraduate Program in Engineering Research at Berkeley (SUPERB).
Trustworthy QoS Policy Management	Content includes the Distributed QoS Modeling Language (DQML) in GME's XML format (i.e., .xme) and a paper accepted at the Distributed Event-Based Systems (DEBS) Conference 2007 (along with the presentation given).
BitBlaze: Binary Analysis for Computer Security	The BitBlaze project aims to design and develop a powerful binary analysis platform and employ the platform.

Activity Name	Women's Institute in Summer Enrichment (WISE)
Led by	Kristen Gates (UC Berkeley), Steve Wicker (Cornell)
Intended Audience	Graduate students, post-doctoral fellows, and professors from all disciplines that are interested in Ubiquitous Secure Technology and the social, political, and economical ramifications that are associated with this technology, with a focused recruitment effort toward underrepresented minority groups and women.
Approx Number of Attendees (if appl.)	23 participants with 12 speakers

WISE is an annual one-week residential summer program that brings together graduate students, post-doctoral fellows, and professors from all disciplines that are interested in Ubiquitous Secure Technology and the social, political, and economical ramifications that are associated with this technology. WISE participation is open to U.S. professors and post-doctoral fellows, and Ph.D. candidates studying at U.S. universities. Participation is limited to 30 people and will be selected

from a nationwide pool of applicants, who have demonstrated, outstanding academic talent. The WISE target audience is underrepresented minority groups and women in information technology. Learning and presentation materials were cataloged on the TAO Portal for reference.

The program was held June 8-13, 2008 on the campus of TRUST partner institution Cornell University. For the summer 2008 program, the topic was sensor networks with a healthcare and policy emphasis. Presentations and discussion covered several areas, including:

- Sensor networks within healthcare
- Radio frequency identification
- Electronic medical records
- Privacy enhancing software
- Networks and policy rights
- Responsibilities associated with data, data owners, and data users.

WISE 2008 speakers were:

- Annie Anton: Computer Science, NC State
- Ruzena Bajcsy: Electrical Engineering and Computer Sciences, UC Berkeley
- Judy Cardell: Computer Science, Smith College
- Maryanne Davidson: Oracle
- Deirdre Mulligan, Law, UC Berkeley
- Wendi Heinzelman: Electrical and Computer Engineering, University of Rochester
- Sheila Hemami: Electrical and Computer Engineering, Cornell University
- Susan Landau: Sun Microsystems Laboratories
- Christine Shoemaker: Environmental and Civil Engineering, Cornell University
- Yuan Xue: Electrical Engineering, Vanderbilt University.

Tuition for WISE 2008 was $2,500; however, TRUST fellowships were available to U.S. professors, post-doctoral fellows, and Ph.D. candidates studying at U.S. universities. A total of 24 fellowships with travel stipend were awarded.

Program Evaluation: Each WISE fellow completed a program evaluation. WISE participants are tacked over a several year period to evaluate the program impact on research, teaching, professional development, job placement, and retention. This was the third year for WISE and the first time hosted at Cornell University. WISE 2009 will be hosted by UC Berkeley.

An evaluation of first-year WISE participants was conducted with a follow-up survey scheduled for years one, three, and five. Recommendations from the WISE 2006 survey were put into place for the WISE 2007 program. The WISE 2007 cohort was surveyed at the end of the program and again one year out. The cohort will also be surveyed again at three and five years out. TRUST is also setting up a program group on LinkedIn for the WISE 2009 cohort. The LinkedIn site, along with survey instruments, will facilitate the tracking of the WISE cohorts to help determine if participants leveraged workshop information into their professional and career development goals. For example, they will be asked if they initiated a course or research activity, incorporated research ideas from the workshop, initiated collaboration with WISE speakers, and/or maintained contact with the network of WISE participants.

Activity Name	Summer Undergraduate Program in Engineering Research at Berkeley-

	TRUST (SUPERB-TRUST)
Led by	Kristen Gates (UC Berkeley), David Wagner (UC Berkeley), Ruzena Bajcsy (UC Berkeley)
Intended Audience	Undergraduate students, underrepresented minority groups, and women.
Approx Number of Attendees (if appl.)	10

The Summer Undergraduate Program in Engineering Research at Berkeley–TRUST (SUPERB-TRUST) offers a group of talented undergraduate engineering students the opportunity to gain research experience. The program's objective is to provide research opportunities in engineering to students who have been historically underrepresented in the field for reasons of social, cultural, educational, or economic barriers. The program provides students with the opportunity to gain research experience by participating in research projects with TRUST faculty and graduate students and affirms their motivation for graduate study and strengthens their qualifications. Upon completion of this program, it is expected that SUPERB-TRUST students will be better prepared and motivated to attend graduate school.

Students work with graduate student mentors throughout the summer performing research and supporting activities in the area of information technology and TRUST related topics. Past SUPERB-TRUST research topics have included:

- Design of a Distributed Tracking System for Camera Networks
- Camera Networks and Computer Vision
- Time Synchronization Security in Sensor Networks
- Implementation of an Electronic Medical Record System
- Analysis of Wireless Connectivity in Sensor Network Deployments.

SUPER-TRUST is an annual eight-week program, last year conducted June 9 – August 1, 2008. In 2008, SUPERB-TRUST had eleven students participating in TRUST related research topics. Each student was given a $4,000 stipend for the period, travel allowance, and provided on-campus housing. In addition to the research experience, SUPERB-TRUST students participated in educational activities including lab tours and industry field trips, received graduate school advising, and took part in a subsidized GRE test preparation course.

SUPERB-TRUST 2008 was a novel and innovative approach to summer undergraduate research by creating a thematic cohort where students conducted research using a prototype tool for finding security bugs in desktop applications.

The table below lists the names of the SUPERB-TRUST 2008 participants and their research projects:

Participant, Home Institution	**Project Title**
Marjan Aslani (George Washington)	Software Security Project
Nga Chung (San Jose State)	Comparative Study of Black-Box and White-Box Fuzz Testing
Jason Doherty (San Jose State)	Comparison of Commercial and Experimental Fuzzing Tools to Find Security Vulnerabilities in Software
Nichole Estrella (Polytechnic University of Puerto Rico)	Cyber Security: Finding Bugs and Error Reporting
Matiwos Gebre (Morgan State)	Software Security
Katherine Gilani (University of Texas, Dallas)	Mobile Body Sensor Networks with GPS Tracking

Shabana Khan (CSU Pomona)	Cyber Security: Finding Bugs and Error Reporting
Phoebi Lai (Lehigh)	Securing the Cyber World
William Quach (San Jose State)	Identifying Security Vulnerabilities in Open-Source and Commercial Software Applications Using Smart Fuzz Testers
Nichole Stockman (Mills)	Comparing the Fuzzing Tools "zzuf" and "catchconv" by Using them to Discover and Report Exploitable Bugs in Widely-Used Software

Program Evaluation: SUPERB-TRUST students are evaluated at midterm and at the end of the program. They also report their research progress at the regular weekly meetings. They receive feedback on their work from faculty advisors at the weekly meetings and after the midterm evaluation.

At the end of the program, the SUPERB-TRUST students evaluate the program via a questionnaire, the results of which are distributed to faculty advisors and graduate student mentors as feedback and for program development. SUPERB-TRUST participants are also tracked over time to identify those students considering graduate school and those that have been accepted into graduate school programs.

For the summer 2009, TRUST will have 10 students participating in the SUPERB-TRUST program. SUPERB-TRUST 2009 will continue the thematic cohort and students will work together in small groups to conduct research on specific problems related to computer security and trustworthy system technology.

Activity Name	TRUST Academic Courses
Led by	Various TRUST Faculty
Intended Audience	Undergraduate and graduate students
Approx Number of Attendees (if appl.)	Varies by course

During this reporting period, a number of academic courses were developed or updated by TRUST faculty across the Center partner institutions. Listed below is information on each course, including the title, faculty teacher(s), intended audience, enrollment (per semester), when the course was or will be first offered, and a brief description.

Course Name:	Information Technology in Society "Trustworthy Systems: the societal/ethical impact of information technology applications"
Taught By:	Maryanne McCormick (UC Berkeley) and Ruzena Bajcsy (UC Berkeley)
Audience:	Undergraduate majors in computer science and engineering
Enrollment:	24 per semester
First Offered:	Spring 2008
Description:	This course provides an interdisciplinary introduction and overview of the societal and ethical implications of trustworthy systems in information technology in society. It will cover the positive and negative consequences of IT on individuals, neighborhoods, schools, commerce, and democracy. Prerequisites: None (appropriate for all undergraduate majors, and particularly encouraged for computer science and engineering majors). Course Objectives: The goal of this course is to provide a unified introduction to the ramifications of IT design and deployment on individuals and society. The

	course provides a broad exposure to IT applications and systems, at a level of detail aimed at both the beginning technical student and the social science student. For the technical student, this course will provide a societal context for their studies, placing the objectives and results of their design and deployment decisions in a larger context. At the same time, for the social science student, this course will provide a basic understanding of the technology and provide an opportunity to focus on the intersection of policy and technology. For all students, this course will provide a venue to consider IT issues in an interdisciplinary context, and in so doing, we hope to provide good foundational training for the next generation of cyber-security professionals.

Course Name:	Internet Policy Challenges in a Global Environment (INF290)
Taught By:	Deirdre Mulligan (UC Berkeley)
Audience:	Graduate majors in Information Science, Law & Computer Science
Enrollment:	24 per semester
First Offered:	Spring 2009
Description:	The internet—as a global, "always-on" platform—poses unique challenges to legal and political frameworks premised on territorial jurisdiction. Operating in this global marketplace exposes companies, and sometimes individuals, to conflicting normative, legal and political commitments. Through case studies, this course considers the options in (i) developing technologies and business strategies to address the varied, and sometimes competing, laws of different countries; (ii) amending laws and otherwise engaging in policy development for the global internet; and (iii) explaining these choices and limitations to regulators, business partners and users. It will consider the implications of these various strategies on an issue-by-issue basis in the areas of content regulation, intellectual property, information security, and privacy, and explore the cross-cutting consequences and dependencies between choices in these various issue areas.

Course Name:	The Digital World and Society (CMPE25)
Taught By:	Russ Smith (San Jose State University)
Audience:	Undergraduate majors in computer science and engineering
Enrollment:	20 per semester
First Offered:	Spring 2008
Description:	This course is designed to enable students to understand how digital technology impacts the world in which we live. Emphasis is placed on how the Internet and emerging digital technologies are changing fundamental and traditional elements of society. Course Objectives: The course provides an overview of the technologies covered will be given followed by sections on the impact they have on governments, society and the individuals. Students will become aware of the far reaching impacts that the technologies that we so easily have embraced. Class will consist of both lecture and class activities.

Course Name:	Software Security Technologies (CMPE279)
Taught By:	Mark Stamp (San Jose State University)
Audience:	Upper Division and graduate level majors in computer science and engineering
Enrollment:	50 per semester
First Offered:	Spring 2008
Description:	The course provides the fundamental concepts, methods and tools used to design and implement software security technologies for constructing trustworthy centralized, distributed or enterprise-wide software systems.

Course Name:	Web Programming and Society (CS142)
Taught By:	John Mitchell (Stanford University)
Audience:	Undergraduate majors in computer science and engineering
Enrollment:	100 per semester
First Offered:	Spring 2009
Description:	The web uses complex applications that run on heterogeneous browsers that may be built using programming technologies such as Javascript, AJAX, Google Web Toolkit, Apache Struts, Java Server Faces, and Rails. This course covers how core web technologies work, common security vulnerabilities, and how to build secure web applications that avoid them.

Course Name:	TechLaw with Progressive Minds (CS302)
Taught By:	John Mitchell (Stanford University)
Audience:	Graduate majors in computer science and engineering
Enrollment:	20 per semester
First Offered:	Spring 2008
Description:	How the advent of computing technologies is reflected in the confluence of law, public policy, and technology. Issues relating to civil liberties, consumer protection, e-voting, copyright law, patent law, international patent law, trade secrets, political processes, and litigation are covered.

Course Name:	Mobile Communications (ECE5680)
Taught By:	Steve Wicker (Cornell University)
Audience:	Graduate majors in computer science and engineering
Enrollment:	50 per semester
First Offered:	Spring 2009
Description:	Theory and analysis of mobile communication systems, with an emphasis on understanding the unique characteristics of these systems. Topics include cellular planning, mobile radio propagation and path loss, characterization of multipath and fading channels, modulation and equalization techniques for mobile radio systems, source coding techniques, multiple access alternatives, CDMA system design, and capacity calculations. Course Objectives: This course will provide an overview of wireless communications, with an emphasis on un-tethered transceivers. It covers traditional topics– channel modeling, demodulation in the presence of noise, and error control coding–and then moves on to recent developments in multicarrier modulation, spread spectrum, and space-time modulation and coding. It emphasizes applications to successful wireless telephony and LAN systems and also considers higher-layer system concepts such as mobility management, with an emphasis on third-generation (3G) cellular systems. The course will conclude with a brief overview of communication and privacy law, with a discussion of recent research into privacy-aware network design techniques.

Course Name:	Fault-tolerant Distributed Computer Systems (CS514)
Taught By:	Ken Birman (Cornell University)
Audience:	Upper division majors in computer science and engineering
Enrollment:	80 per semester
First Offered:	Fall 2008
Description:	Distributed systems are difficult to build and understand, for all sorts of reasons. Failures are common in large systems, and we can't let them shut the application

	down; a well-engineered system will tolerate failures and repair any damage they caused. A single event, perceived at multiple locations, may not be totally ordered with respect to other, conflicting, events. Networks have annoying connectivity and bandwidth properties that force the designer to confront challenging engineering tradeoffs. Course Objectives: The focus of CS5410 is on the principles and techniques that one can use to achieve a high-quality, trustworthy, fault-tolerant distributed system. Lectures will present the principles; programming assignments will enable students to put these principles into practice. The course project will expose students to state of the art technology platforms (notably web services) but will also involve using cutting-edge techniques that are not available (yet) in products. The goal is to understand what these kinds of platforms can be expected to do "without help" but also to get used to the idea that one can push beyond their limitations when necessary, and that doing so can open the door to all sorts of creative possibilities.

3.5 Professional Development Activities

During this reporting period, TRUST students were involved in a number of professional development activities within the domains of computer science, information technology, law and social policy as well as additional activities such as internships, entrepreneurial business course, career preparation workshops, and professional societies. The following sections list the various professional development activities of TRUST students.

The TRUST Center provides a unique opportunity for a wide range of cyber security issues to be addressed from many points of view—technological, scientific, social, policy, and legal. The diverse academic and professional interests of TRUST students are a major contribution to the Center's success. TRUST students have a wide range of academic and professional interests reflected by the conferences attended, workshops supported, personal development courses taken, and social and professional society memberships. These professional development activities increase student cross-domain and multi-domain knowledge, professional growth, academic success, and overall retention–all of which benefit TRUST and the student learning experience and impact provided by the Center.

TRUST students have participated in the following business development courses, training, internship, and fellowship programs:

- Coursework: "Customer and Business Development in Hi-Tech Enterprise", Hass Business School, UC Berkeley, CA
- Coursework: "Marketing Strategies for Entrepreneurs in Silicon Valley" Hass Business School, UC Berkeley, CA
- Internship at Microsoft Research, Redmond, WA
- Internship at IBM Research, Haifa, Israel
- Internship at Telecom Italia WSN Lab Berkeley, CA
- Internship at Yahoo! Research, Santa Clara, CA
- Invited talk at Yahoo!, Santa Clara, CA

TRUST students have membership in the following organizations:

- ACM: Association for Computing Machinery
- California State Bar Association
- District of Columbia Bar Association
- Future of Privacy Foundation's Advisory Council
- Global Environment for Network Innovations (GENI) End-User Opt-In Working Group

- HKN: Eta Kappa Nu National Electrical Engineering honor society
- IEEE: Institute of Electrical and Electronics Engineers
- International Association of Privacy Professionals
- ISAlliance-ANSI Workshop on Developing a Framework to Analyze and Manage Financial Risk for Cyber Security
- Java Community Process
- Network Research Review Committee, Internet2
- SIGMA XI: International Honor Society of Science and Engineering
- Silicon Valley Software Quality Association
- SWE: Society of Women Engineers
- Tau Beta Pi: The Engineering Honor Society
- Transaction Processing Performance Council (TCP)
- USENIX: Advanced Computing Systems Association
- W3C: The World Wide Web Consortium
- WICSE: Women in Computer Science and Electrical Engineering

TRUST students have participated in the following workshops, conferences, and symposiums:

- AAAI Conference on Artificial Intelligence, Chicago, IL
- ACM International Symposium on Software Testing and Analysis (ISSTA 2008), Seattle, WA
- ACM Computer and Communications Security Conference, Alexandria, VA
- ACM Computer and Communications Security Symposium, Tokyo, Japan
- ACM International Conference on Distributed Event-Based Systems (DEBS 2009), Nashville, TN
- ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing (PODC 2008), Toronto, Canada
- ACM Workshop on Hot Topics in Networks (HotNets-VII), Calgary, Alberta, Canada
- ACM/IEEE International Conference on Information Processing in Sensor Networks, San Francisco, CA
- ACM/IEEE International Conference on Model Driven Engineering Languages and Systems, Toulouse, France
- Air Force Office of Scientific Research (AFOSR) Complex Networks Program Review, Arlington, VA
- Annual Network & Distributed System Security Symposium (NDSS), San Diego, CA
- Carnegie Mellon Conference on the Electricity Industry, Ithaca, NY
- Carnegie Mellon University, Detection of Complex and Multistage Computer Network Attacks, Pittsburgh, PA
- Cyber Security and Information Intelligence Challenges and Strategies (CSIIRW09), Oak Ridge National Lab, Oak Ridge, TN
- Electronic Crimes Task Force, DHS, United State Secrete Service, San Jose, CA
- European Conference on Object-Oriented Programming (ECOOP 2008). Paphos, Cyprus
- GE Global Research Center PhD Women's Research Networking Event 2008
- Grace Hopper Celebration of Women in Computing conference, Keystone, CO
- ICO: An Introduction to Secure Control, Pisa, Italy
- Identity Theft Technology Council (ITTC), DHS–SRI International, San Mateo, CA
- IEEE Conference on Automation Science and Engineering (CASE 2009), Washington, DC
- IEEE Consumer Communications & Networking Conference, Las Vegas, NV
- IEEE Energy 2030 Conference, Atlanta, GA

- IEEE International Conference on Information Assurance and Security, Naples, Italy
- IEEE Symposium on Security and Privacy, Oakland, CA
- IEEE/IFIP International Conference on Dependable Systems and Networks (DNS2008), Anchorage, AK
- IFIP Working Group 11.10 International Conference on Critical Infrastructure Protection, Dartmouth College, Hanover, NH
- Infosec Technology Transition Council (ITTC) Meetings, Menlo Park, CA
- International Compulog/ALP Summer School on Logic Programming and Computational Logic, New Mexico State University, Las Cruces, NM
- International Conference on Automated Planning and Scheduling (ICAPS), Sydney, Australia
- International Conference on Computer Communications and Networks (ICCCN'08), Saint Thomas, U.S. Virgin Islands
- International Conference on Distributed Smart Cameras, Palo Alto, CA
- International Conference on Embedded Software (EMSOFT), Atlanta, GA
- International Conference on Intelligent Robots and Systems (IROS), St. Louis, MO
- International Conference on Object Oriented Programming, Systems, Languages and Applications (OOPSLA 2009), Orlando, FL
- International Conference on Pervasive Computing Technologies for Healthcare, London, England
- International Conference on Telecommunications (ICT'08), Saint Petersburg, Russia
- International Conference on Web Services (ICWS 2009), Los Angeles, CA
- International Joint Conference on Automated Reasoning, Sydney, Australia
- International Symposium on Distributed Computing (DISC '08), Arcachon, France
- International Symposium on Distributed Objects, Middleware, and Applications (DOA'08), Monterrey, Mexico
- International Symposium on Recent Advances in Intrusion Detection (RAID 2009), Saint-Malo, Brittany, France
- International Symposium on Wireless and Pervasive Computing (ISWPC'09), Melbourne, Australia
- International Workshop on Cyber-Physical Systems, Beijing, China
- International Workshop on Sensor Networks (SN 2008), U.S. Virgin Islands
- International World Wide Web Conference (WWW 2009), Madrid, Spain
- Kavli Institute Symposium on Computing Challenges. Ithaca, NY
- Large-Scale Distributed Systems and Middleware (LADIS 2008), White Plains, NY
- Measuring Identity Theft at Telecommunications Companies, Consumer Issues Conference, University of Wyoming Consumer Issues Conference, WY
- MIT Lincoln Labs Information Session 2009: MIT Lincoln Laboratory, Worcester Polytechnic Institute, Lexington, MA
- MOTHIS '08 Workshop, Vanderbilt, TN
- Northern California Undergraduate Mathematics Conference, Sonoma State, CA
- O'Reilly Emerging Technologies Conference, San Jose, CA
- Object Management Group (OMG) Workshop on Distributed Object Computing for Real-time and Embedded Systems, Washington, DC
- Privacy Law Scholars Conference, Berkeley, CA
- Privacy Law Scholars Conference, Washington, DC
- Richard Tapia Celebration of Diversity in Computing Conference, Portland, OR
- RSA Security Conference 2009, San Francisco, CA
- Securing Sensor Networks, Stockholm, Sweden

- Annual Security Breach Notification Seven Years Later, BCLT/BTLJ, Berkeley, CA
- SSEAT, Workshop on State-space Exploration for Automated Testing, Chicago, IL
- Sun Microsystems, A Body Sensor Network Application, Menlo Park, CA
- Super Computing (SC) High Performance Computing (HPC), Networking, Storage and Analysis Conference, Austin, TX
- Targeted Malware, Workshop on Interdisciplinary Studies in Security and Privacy, Polytechnic Institute of NYU, Brooklyn, NY
- The IT Security Entrepreneurs' Forum (ITSEF III), Palo Alto, CA
- USENIX Annual Technical Conference, Boston, MA
- USENIX Conference on File and Storage Technologies (FAST '09), San Francisco, CA
- USENIX Security Symposium, San Jose, CA
- USENIX Symposium on Operating Systems Design and Implementation (OSDI 2008), San Diego, CA
- USENIX Workshop on Hot Topics in Security, San Jose, CA
- USENIX Workshop on Large-Scale Exploits and Emerging Threats (LEET), Boston, MA
- USENIX/ACM Networked Systems Design and Implementation Symposium (NSDI), San Francisco, CA
- W3C Workshop on Security for Access to Device APIs from the Web, London, England
- Web 2.0 Security and Privacy (W2SP 2009), Oakland, CA
- Women's Institute in Summer Enrichment (WISE), Cornell, NY
- Workshop on Privacy in the Electronic, Alexandria, VA

3.6 External Education Activities

The items below describe in more detail specific external education activities of the TRUST Center during this reporting period.

Activity Name	Curriculum Development in Security and Information Assurance (CDSIA)
Led by	Sigurd Meldal (San Jose State University)
Intended Audience	California State University System and Hispanic Association of Colleges and Universities member institutions
Approx Number of Attendees (if appl.)	37

On May 1, 2009 TRUST hosted the second annual Workshop on Curriculum Development in Security and Information Assurance (CDSIA 2008) at San Jose State University.

The objectives were to (1) reach out to the many universities of the California State University system and to other universities whose mission is focused on work-force preparation and undergraduate education, (2) to share with faculty members of these institutions material and support structures developed by the TRUST partners, (3) to strengthen the TRUST-related community of educators, and (4) to facilitate the education of members of underrepresented communities in the domain of secure technologies.

CDSIA 2009 had 37 participants from 19 universities. Half of those universities are Hispanic Serving Institutions (HSIs) and the remainders are all Associate members of the Hispanic Association of Colleges and Universities (HACU). Fifteen of the 23 CSU schools were represented. Three TRUST partner institutions (San Jose State (host), Stanford, UC Berkeley) also participated in CDSIA 2009. In 2009, to reduce overlap of activities, the Information Assurance Capacity Symposium Program

(IACSP) was incorporated into the CDSIA workshop. The IACSP was an outreach to HSI and Historically Black College and University (HBCU) faculty members, to work with them to introduce and strengthen the Information Assurance components of their curriculum.

The workshop topics included:

- Security, information assurance, and policy in the general education curriculum
- Tools support for teaching IA and security curriculum components
- Sharing and delivering curricula through the TRUST Academy Online (TAO)
- What preparation does industry require?
- Certification and accreditation - where are we with respect to security?
- What role (if any) should the teaching of "malware" play in the curriculum?

Program materials generated by this program were cataloged on the TAO Portal.

Activity Name	Summer Experience, Colloquium and Research in Information Technology (SECuR-IT)
Led by	Kristen Gates (UC Berkeley), Sigurd Meldal (San Jose State)
Intended Audience	Graduate level (MS & Ph.D.) students in computer science
Approx Number of Attendees (if appl.)	6

SECuR-IT is a ten-week residential program with paid internship co-located at Stanford and San Jose State. The program dates were June 9 – August 15, 2009.

SECuR-IT is a graduate student academic immersion with internship program. In addition to working with an industry sponsor during the program, participants benefit from the following programmatic components:

- Seminars conducted by faculty and industry experts that expose students to a wide range of information technology and computer security research instruction
- Faculty participation from Stanford, UC Berkeley, and San Jose Sate
- Informal social gatherings that provide a relaxed setting for students and faculty to exchange ideas and share experiences
- Residential housing at San Jose State
- Paid, 40-hour-per-week internship.

Graduate student internship opportunities have been available in areas such as Security Architecture, Security Awareness and Security Management, Host and OS Security, Application Security, Network Security, Secure Software Engineering, Risk Management, and Policy and Legal Compliance. TRUST industry partners supporting this program in 2008 were eBay, PayPal, Sun Microsystems, and Salesforce.com. Learning materials generated for this program were cataloged on the TAO Portal.

This is a 40-hour per week obligation to internship, research, and learning activities. Students who participate in SECUR-IT view this program as a full-time summer experience and participate in a residential cohort, attend TRUST-run education courses, and are employed as an intern by a participating industry partner.

Program Evaluation: Each SECuR-IT student completed a pre- and post-program evaluation. SECuR-IT participants will be tacked over a two-year period to evaluate the program's impact on their education, professional development, and job placement. Industry partners and mentors will also be evaluated as to the program's structure, effectiveness, and means for improvement. The number of new hires resulting from this program will also be tracked.

Activity Name	TRUST Seminar Series
Led by	Annarita Giani (UC Berkeley), Alvaro Cardenas (UC Berkeley)
Intended Audience	Graduate level (MS & Ph.D.) students in computer science, faculty and industry professionals
Approx Number of Attendees (if appl.)	1,355 over 28 week series

The TRUST Speakers Series began in the fall of 2007. The program is a weekly event on the UC Berkeley campus that brings in well-known speakers who are experts in the fields of security, privacy, and trustworthy systems. The event is focused toward graduate students in computer science, industry professionals, and campus community at large.

In the fall 2008, TRUST hosted 11 speakers with a total number of 355 people attending and in the spring 2008, TRUST hosted 14 speakers with a total number of 490 people attending. Next year we are investigating both broadcasting the TRUST Seminar talks live via the Web and archiving the talks for offline viewing—both of which will make the talks available to a much wider audience.

3.7 Activities to Integrate Research and Education

Education deliverables were tied to all TRUST research, education and outreach projects. Learning materials and modules were distilled from the TRUST research trust and archived on the TRUST Academy Online TAO portal. Workshops and symposiums such as TIPPI are available via the TAO portal. WISE 2007 archived presentations to the TAO portal. SUPERB-IT students worked on TRUST research topics. The SECuR-IT summer immersion program with internship presented a computer security focused curriculum and SECuR-IT seminars featured TRUST faculty from UC Berkeley, Stanford, and San Jose State presenting topics related to TRUST center research and activities.

Activity Name	DHS-SRI Infosec Technology Transition Council (ITTC)
Led by	John Mitchell (Stanford), TRUST liaison to the ITTC
Intended Audience	Academics and Industry Professionals
Approx Number of Attendees (if appl.)	Total of 280 over the three meetings in 2008-2009

The DHS-SRI Infosec Technology Transition Council (ITTC) is a working forum that brings together experts and leaders from the government, private sector, financial industry, information technology services, venture capital, and academia and science sectors to address the problem of information security and related activity.

Workshops are held quarterly (during this period they were held in June 2008, October 2008, and February 2009) and are used to identify proactive IT security solutions and assist in the acceleration of its development and deployment into the marketplace. Seasoned IT security practitioners, law enforcement professionals, and representatives from academia and science have strategically aligned themselves with subject matter experts and organizations to accomplish this goal. A key component to the success of this public-private partnership is the ability to actively work with leaders in the

community who are principals of change in an effort to better protect our communities and corporations from attacks against their critical infrastructures. The subject matter experts of the ITTC seek to share information that will assist in the discovery, due diligence, development, and deployment of next generation technologies best suited to protect our critical infrastructures and serve our communities.

John Mitchell from Stanford is the TRUST liaison and participant in the DHS-SRI ITTC and meetings are attended by various TRUST researchers.

Activity Name	IT Security Entrepreneurs' Forum (ITSEF)
Led by	John Mitchell (Stanford)
Intended Audience	Academics and Industry Professionals
Approx Number of Attendees (if appl.)	251 for the March 19, 2009 meeting

The Department of Homeland Security and Kauffman Foundation IT Security Entrepreneurs' Forum (ITSEF) is a Public Private Partnership initiative is designed to "bridge the gap" between IT security solution providers and the end users of our nation's IT and Telecommunications critical infrastructures. The ITSEF believes that innovative solutions developed by entrepreneurs' can best be promoted through collaborative efforts between the public and private sectors.

A key component to the success of such relationships is to identify and bring together public and private sector "change agents" who can drive education and awareness programs through forums that will promote lasting and permanent relationships between all levels of government and the full range of emerging and established private sector companies. This year's forum occurred during a critical time as attacks and emerging threats continue to increase in sophistication and frequency against our nation's IT and telecommunication infrastructures. The ITSEF strives to accelerate the search for and implementation of "best of class" solutions to address such threats.

John Mitchell from Stanford is the TRUST liaison and sponsor of the IT Security Entrepreneurs' Forum.

3.8 Education Metrics/Indicators

The items below describe how the Center is doing with respect to the education metrics and indicators and data that have been collected during this reporting period. Information is provided for both Learning Materials and Technology and Professional Workshops and Symposiums.

Learning Materials and Technology

During this reporting period, there was a continued effort to reconfigure the TAO Portal, including metadata technology and information architecture, as well as the further development of TAO courseware modules and projects. This effort has resulted in a usage increase of 71% over the last year. We have also increased the number of learning modules and courseware, currently 176 resource files produced by 46 TRUST contributors.

We will also implement data collection strategies that will track the use and dissemination of TRUST education materials from the TAO. Further analysis of the TAO online access statistics presented at the site visit indicate that approximately 25% of people accessing the TAO download a resource in the repository. That said, we will further develop portal survey and user rating technologies to help us better understand our online community and their usage of the TAO. Additionally, the TAO has been registered

as a collection in the National Science Digital Library (NSDL), the nation's online library for education and research in Science, Technology, Engineering, and Mathematics. Registration of the TAO in the NSDL is part of our ongoing efforts to broaden awareness of the Center's research and education missions.

Professional Development Workshop and Symposiums

TRUST professional development activities are designed for graduate students, post-doctoral scholars, industry researchers, and faculty from various disciplines working and conducting research in cyber security and trustworthy systems. In addition to education and learning opportunities, these programs support professional growth, especially for female and URM faculty, with the goal of ultimately expanding the number female and URM researchers in cyber security and trustworthy systems.

TRUST faculty and staff have participated at education oriented conferences through panels, associated workshops or a series of presentations, including: Engineering Education NSF Awardees Conference, International Conference on Engineering Education, Computer Alliance for Hispanic Serving Institutions, Richard Tapia Celebration in Diversity in Computing, Grace Hopper Celebration of Women in Computing, TechLeaders: Anita Borg Institute, Executive Women's Forum, Information Security, Risk Management and Privacy, Broadening Participation in Computing Community Meeting, Berkeley EECS Annual Research Symposium (BEARS), International Information Integrity Institute, San Francisco Electronic Crime Task Force Meetings, and the Department of Homeland Security SRI Infosec Technology Transition Council.

SUPERB-TRUST, the TRUST Research Experiences for Undergraduates (REU) program, will continue at UC Berkeley will continue for summer 2009. Beginning summer 2010, we plan to expand that program to TRUST partner institutions Carnegie Mellon, Cornell, Stanford, and Vanderbilt to increase the number of undergraduate students exposed to research in general and the TRUST Center in particular. The TRUST REU will continue to support the Center's goal of increasing the number of underrepresented minority groups and women that are conducting research in cyber security and trusted systems.

The assessment process is both qualitative and quantitative and will include pre- and post-evaluation surveys, focus groups, participant assessments, and program evaluations for the education, human resource development, and underrepresented minority student uptake initiatives of TRUST. Evaluation rubrics will be developed for assessment of course materials, electronic portfolios, and research activities.

3.9 Next Reporting Period Education Plans

The education initiatives detailed in this document will continue into the next reporting period. No major changes in the direction are anticipated but the level of activity will increase.

The TAO will continue to develop. Course modules and learning objects will be developed as educational deliverables of each TRUST research area. As the review process continues, refinement will be made to the module design and the portal. The TAO is making an impact by providing TRUST Center learning materials for use by teaching faculty as course content, lecture materials, and program support for the development of their computer science or related higher education courses. TRUST has created 38 learning modules across 21 topical areas, providing educators access to a substantial amount of leading-edge research and teaching material. The Center will continue to place materials generated by our education, outreach, and diversity programs on the TAO to be shared with other teachers and researchers.

TRUST visibility and influence in education community continues to grow as TRUST researchers and staff participation in educational conferences, workshops, panel discussions, and industry workgroups take hold.

The Women's Institute in Summer Enrichment (WISE) is a signature program for TRUST and consistently receives excellent evaluations form participants. WISE is hosted at TRUST partner institutions and the Center will continue to offer this program each summer. To meet the program's increasing demand, TRUST will expand WISE to 30 participants per summer with a greater emphasis on recruiting female URM scholars.

CDSIA has received excellent reviews from faculty participants. CDSIA is creating a community of TRUST scholars and has merged it with the IACBP into one annual event called the Annual Symposium on Curriculum. Going forward, we will continue to leverage CDSIA to engage community colleges and broaden participation.

The SECuR-IT summer program has created a great deal of interest among CSOs of Silicon Valley computer security companies and we expect to expand the SECuR-IT program from six graduate students during the summer of 2008 to 10 graduate students for the summer of 2009. Summer 2010, the SECuR-IT Silicon Valley will grow to include 10 industry partners and 30 participants and the internship model will be expanded to TRUST partner campuses Cornell University (Financial Infrastructures) and Vanderbilt University (Health Infrastructures).

We will continue our relationship with the **Student Transitional Alliance for Research in STEM (STARS)**. STARS is a NSF-sponsored program and the partnership is designed to provide faculty and students from minority serving institutions (MSIs) with increased access to undergraduate and graduate research opportunities at six NSF-sponsored STCs (those funded in FY 2005 and FY 2006). The goals of this program are to: 1) increase the number of students from MSIs completing graduate degrees on STC campuses, 2) increase the number of students and faculty members from under-represented groups to obtain research experience at STC sites, 3) increase the involvement of MSI researchers on STC projects, 4) provide an expanded forum for STCs to share their education and knowledge transfer initiatives, and 5) increase faculty and staff diversity at STCs.

TRUST is actively collaborating with the STARS program and its lead, Dr. William McHenry, the Project Director of the Science and Diversity Center and Executive Director of the Mississippi e-Center at Jackson State University. During this reporting period, STARS funding supported two undergraduate students in the SUBPERB-TRUST REU program and we anticipate receiving STARS funding for additional SUPERB-TRUST program participants in the summer 2009.

One new education initiative that is under development is an **Academic Program Specialization**. TRUST will create a graduate-level academic specialization in Cyber Security and Trustworthy System at all TRUST partner institutions. This specialization will be open to all students and require the completion of a minimum of eight semester units of advanced courses from a predefined list of topics. TRUST partner institutions will not only leverage existing graduate-level courses but will augment each institution's course offerings with the new courses described above. This, in turn, will enable each TRUST partner institution to expose a larger number of students to those topics.

4 KNOWLEDGE TRANSFER

4.1 Goals and Objectives

The Center's knowledge transfer goal is to establish TRUST as a true public private partnership—namely a trusted intermediary between industry, government, infrastructure stakeholders, and the research community.

TRUST knowledge transfer objectives are to: (1) develop strong liaison with the concerns of industry and infrastructure stakeholders; (2) produce legislative and legal policy papers and amicus briefs; (3) leverage testbeds for demonstrating Center research project results; (4) enable student internships and support entrepreneurial clubs; and (5) convene meetings, summits, and workshops to share the results and knowledge gained through Center research activities.

The structure of TRUST lends itself to a comprehensive approach to knowledge transfer. Since TRUST addresses well defined and long term societal needs, the results in computer security, privacy, and critical infrastructure protection can be easily communicated to decision makers, policy makers, and government agencies. With respect to industry, the Center's integrative testbeds represent focal points for interaction and dialog with major stakeholder industries (e.g., power, telecommunication, embedded systems). In fact, several integrative testbeds are being provided by the stakeholders, which offer significant leverage for the Center and support technology transfer from the research community to government and industry partners. Finally, TRUST researchers are leaders in their scientific communities. Their broad cooperation to achieve the TRUST objectives will serve as a catalyst to turn attention of the community toward the emerging science of secure systems.

TRUST comprises multiple institutions, technology vendors, and infrastructure users and providers. Broad participation from leading research universities, undergraduate colleges serving under-represented groups, computer vendors (e.g., Cisco, HP, IBM, Intel, Microsoft, Symantec), and infrastructure providers (BellSouth, Boeing, Qualcomm, Raytheon) will result in wide spread dissemination, adaptation and continued evolution of ubiquitous secure technology. TRUST research will learn and evolve with our results using an iterative investigate-develop-educate-apply cycle. We will develop science, technology, and proof of concept prototypes that will be tested through models that emerge from a series of analytical and case studies, experimentation, and simulations. We plan to use periodic updates of living reports and community workshops throughout the life-cycle of TRUST.

The research output of the Center will be disseminated in four ways: (1) publications in the open literature and on the web, (2) Seminars and workshops held at major conferences and infrastructure protection meetings, (3) public lectures and meetings with the general public concerned about security and privacy issues on the internet and critical infrastructure protection, and (4) curriculum development and courses taught at the partner institutions as well as the outreach institutions.

During the reporting period, we believe that TRUST has been solidly on track with respect to its knowledge transfer objectives. Success is measurable in many ways: technologies that are being commercialized, TRUST researchers who are working hand-in-hand with industry and standards groups to help improve trustworthiness of major infrastructure systems, activities aimed at educating the public and exploring non-technical ramifications of TRUST themes, development of significant TRUST spin-offs, and exploratory discussions regarding additional activities such as a center for research on trustworthy electronic health records and a trusted financial systems center under discussion with the U.S. Department of Treasury.

4.2 Performance and Management Indicators

TRUST knowledge transfer activities are periodically monitored for meeting the Center's overall knowledge transfer objectives and the individual activity's knowledge transfer objectives. Periodic monitoring consists of meetings of the TRUST Executive Board where progress of each knowledge transfer activity (or sets of activities) is formally reviewed. The evaluation metrics are outlined in the table below.

Goals	Objectives	Evaluation Criteria	Frequency
Economic, Legal, Social Impact of TRUST	Policy paper, amicus briefs, legislation	Scholarly impact, Societal impact, Legislative impact, Judicial impact	Bi-Annual
Testbeds	Demonstrations to scale of TRUST technology on realistic platforms	Industrial interest, Industrial adoption, Stakeholder interest, Stakeholder adoption	Annual
Financial infrastructures	Identify generic/unique features of TRUST issues, propose solutions, privacy issues	Stakeholder interest, stakeholder support	Annual
Electric power demand side infrastructures	Identify vulnerabilities of SCADA systems, propose secure network embedded systems solutions	Stakeholder interest, Stakeholder support	Annual
Secure Global Information Grid Architectures	Examine and critique proposed architectures, propose security architectures and solutions	Stakeholder interest, Stakeholder support	Annual

4.3 Current and Anticipated Problems

No significant problems were encountered during the reporting period. No significant problems are anticipated in the next reporting period.

4.4 Knowledge Transfer Activities

The TRUST industrial collaboration and technology transfer initiatives support the goals and objectives of the Center's knowledge transfer component. Within TRUST, knowledge transfer is enabled by (1) using partner knowledge and experience to focus research on real-world problems; (2) verifying our science and technology at partner sites to ensure they work in practice; (3) including partners in every stage of the research, science and technology development process; and (4) aggressively licensing TRUST intellectual property to corporate partners for commercialization. (In particular, the Center has developed an interesting open source software IP model to facilitate interactions with industry.)

The items below describe in more detail specific knowledge transfer activities of TRUST researchers. Items are grouped by the lead institution(s).

Technology Transition to the U.S. Air Force	
Led by	Cornell University
Organizations Involved	

	Name	Address
1	Cornell University	Ithaca, NY 14850

At the request of the Chief Information Officer of the U.S. Air Force, Mr. Tilotson (and the AF/XC, Mr. Werner), Birman and Schneider organized a workshop to study risks associated with Air Force deployment of Windows Vista as a single solution on client platforms. Although the workshop did identify some risks, we also identified a number of cutting edge risk management options that seem to address most issues. For example, TRUST research on artificial diversity seems to be a powerful remedy for the potential creation of a viral "target" associated with the very homogeneous deployment model, and indeed Windows Vista itself incorporates stack randomization, which is a very important first step. AF/XC was extremely pleased with the outcome and is acting on our recommendations for next steps, including early deployment suggestions and longer term research proposals. Contact: Dr. Sekar Chandersekaran (cchander@ida.org)

Financial Services Industry Research and Development		
Led by	Cornell University	
Organizations Involved		
	Name	Address
1	Cornell University	Ithaca, NY 14850
2	Stanford University	Stanford, CA 94305
3	University of California, Berkeley	Berkeley, CA 94720

TRUST has started a dialog with the financial sector, coordinated through the Financial Services Technology Consortium (FSTC), a group of about 150 organizations running out of New York City. FSTC has a committee, the Business Continuity Standing Committee (BCSCOM), which prioritized enterprise continuity solutions as one of their top needs. In response, Cornell's research effort studied challenges of doing enterprise backup for entire datacenters over high-speed optical networks and concluded that there are serious technical obstacles to overcome. Our new Maelstrom protocol (NSDI 08) is a first step to a comprehensive solution, and the Smoke and Mirrors File System (submitted to Middleware 08), which runs over Maelstrom, a second step. These systems make possible a completely transparent enterprise backup story, in real-time, even with the backup at geographically remote locations. TRUST researchers John Mitchell, Fred Schneider, and Doug Tygar briefed the FSTC Security and Infrastructure Standing Committee and TRUST researcher Ken Birman spoke at the FSTC annual meeting in Napa, CA in June 2008.

Research Dissemination via Conferences and Workshops		
Led by	Cornell University	
Organizations Involved		
	Name	Address
1	Cornell University	Ithaca, NY 14850
2	University of California, Berkeley	Berkeley, CA 94720
3	Stanford University	Stanford, CA 94305
4	Carnegie Mellon University	Pittsburgh, PA 15213
5	Vanderbilt University	Nashville, TN 37235

The TRUST research team has had prominent roles such as keynote and other invited talks, both at major research conferences, industry-oriented conferences, and at some of the largest platform vendors, such as IBM, Microsoft and Cisco and are infusing these talks with TRUST themes. Such activities are good opportunities for dialog with folks "on the ground". Additionally, multiple TRUST members often support the same government workshops. For example, several TRUST researchers participated in a series of NSF sponsored workshops associated with the national cyber

security research and development strategy, embedded sensors, and other small real-time devices. NSF is now exploring the creation of a new research program in this area. Finally, TRUST researchers have taken the lead to start new workshops and conferences focused around TRUST research themes. Of note during this period was the second annual Model-Based Trustworthy Health Information Systems (MOTHIS) workshop which was established by TRUST research Janos Sztipanovits.

Industry Technology Transition and Product Adoption		
Led by	Cornell University and Stanford University	
Organizations Involved		
	Name	Address
1	Cornell University	Ithaca, NY 14850
2	Stanford University	Stanford, CA 94305

Under the direction of Professor Ken Birman at Cornell, work is underway on helping the Red Hat Linux community develop a new, open-source technology for time-critical event-driven computing. Many applications, such as financial systems or medical systems, are "event driven" in that some form of external data source (a ticker plant, or medical telemetry) must drive a reaction by the system. Today, there are surprisingly few technical options for building such systems: users are forced to purchase message middleware products from vendors and complain that the solutions are complex, expensive, and unstable in scaled-out deployments. Cornell's Ricochet protocol (NSDI 07) addresses these requirements in a simple, lightweight manner that offers extremely good real-time properties and involves minimal infrastructure. We're now working to produce a version matched to the needs of the Red Hat community, with the hope that the IP might enter their public-source distribution early in the 2009 timeframe. Patents on Ricochet would be transferred to OIN and licensed, for free, to any organization wishing to implement a new solution using the same ideas, and the Ricochet platform itself would become an open source component. We're also working on a new research paper reflecting some of the innovations needed to address practical deployment issues posed by the folks at Red Hat. Our main contact is Carl Trieloff (cctrieloff@redhat.com), the Chief Technology Officer of Red Hat.

Researchers from Stanford University collaborated with RSA Security on integration with the RSA SecurID hardware token. SecurID generates a one-time password that is still vulnerable to "attacker-in-the-middle" password stealing attacks. With the server-side software developed as a result of this collaboration, RSA SecurID one-time passwords are protected from phishing attacks.

Open Source Software Dissemination		
Led by	Stanford University	
Organizations Involved		
	Name	Address
1	Stanford University	Stanford, CA 94305
2	University of California, Berkeley	Berkeley, CA 94720

Pwdhash, SafeCache, SafeHistory, and SpyBlock are all available as freely downloadable open-source software. At least tens of thousands of downloads have occurred, and there has been continuing media attention through 2006-07. Additionally, we have made available open source software releases of our Doppelganger code (http://www.umeshshankar.com/doppelganger/).

Privacy Issues in Electronic Medical Records	
Led by	Stanford University

Organizations Involved		
	Name	Address
1	Stanford University	Stanford, CA 94305
2	Vanderbilt University Medical Center	Nashville, TN 37235
3	Vanderbilt University (ISIS)	Nashville, TN 37235

Currently, the Stanford model of the MyHealth system is a simple workflow graph on the roles in the portal (patient, secretary, nurse, doctor, etc). Based on our analysis of this simplified workflow, we have made several design suggestions to the MyHealth team at the Vanderbilt Medical Center. Specifically, we have suggested (1) MyHealth include tags for messages, (2) use these tags to enforce privacy requirements, and (3) use these tags to route messages more accurately. The Vanderbilt team at ISIS is currently creating a hi-fidelity model of the MyHealth system, including its workflow. We will use this model to further evaluate MyHealth.

DexterNet		
Led by	University of California, Berkeley	
Organizations Involved		
	Name	Address
1	Vanderbilt University	Nashville, TN 37235
2	Cornell University	Ithaca, NY 14850

Berkeley, Cornell, and Vanderbilt researchers have jointly developed an integrated wireless sensor networking environment for remote healthcare. DexterNet is a demonstration platform for TRUST technologies for robust, reliable, and privacy-aware remote healthcare service using robust and privacy-aware wireless sensor network mesh network routing, minimum-disruption service recovery in ad hoc networks, and digital right management for sensor information.

Industry Technology Collaboration and Consulting		
Led by	University of California, Berkeley and Stanford University	
Organizations Involved		
	Name	Address
1	University of California, Berkeley	Berkeley, CA 94720
2	Stanford University	Stanford, CA 94305

David Wagner from the University of California, Berkeley has partnered closely with Hewlett Packard Labs researchers on the Joe-E project. HP Labs researchers are serving as the first users of Joe-E, and two internal HP projects have decided to adopt Joe-E. In particular, the Waterken server is implemented using 18K lines of Joe-E code and 3K lines of Java code. HP Labs researchers have helped us ensure that our techniques work in practice and to improve the Joe-E programming language. HP Labs researchers have been closely involved in the development of Joe-E; we have held day-long meetings approximately once each month. In addition, Wagner's research group at UC Berkeley and researchers at HP Labs jointly organized a security review of the Waterken server, to assess our experience with how well Joe-E was able to support the security goals of the Waterken project. Wagner also consults for Fortify Software, a startup producing software security tools, on their security products. Fortify Software is in the process of commercializing research into program analysis from several TRUST participants, including research by Aiken, Dawson, Song, Wagner, and others. Wagner has helped Fortify to transition his own research into their commercial products, as well as to transition research by other software security researchers from TRUST and elsewhere.

Dan Boneh and John Mitchell from Stanford University were advisors to Passmark, which was acquired by RSA. Rachna Dhamija from the University of California, Berkeley started a company based on the Berkeley Dynamic Skins technology.

Model Integrated Clinical Information Systems (MICIS)		
Led by	Vanderbilt University	
Organizations Involved		
	Name	Address
1	Vanderbilt University Medical Center	Nashville, TN 37235

Vanderbilt researchers have developed MICIS, a software toolkit that is based on model-based design techniques and high-level modeling abstractions to represent complex clinical workflows in a service-oriented architecture paradigm. MICIS models are enriched with formal security and privacy policy specifications, which are enforced within the execution environment. One of the application domains of MICIS is the management of sepsis in acute care settings at the Vanderbilt Medical Center. The Sepsis Treatment Enhanced through Electronic Protocolization (STEEP) is a joint effort between TRUST at Vanderbilt and the Health Tech Lab of Vanderbilt Medical Center. MICIS is also being applied is the Emergency Department (ED) of the Vanderbilt Medical Center. The goal in any ED is the rapid turnaround of patients while maintaining a high quality of care and reducing cost by not ordering unnecessary tests. Privacy and security is achieved using the policy languages developed by TRUST.

Sepsis Treatment Enhanced through Electronic Protocolization (STEEP)		
Led by	Vanderbilt University	
Organizations Involved		
	Name	Address
1	Vanderbilt University Medical Center	Nashville, TN 37235

STEEP is a joint effort between TRUST at Vanderbilt and the Health Tech Lab of Vanderbilt University Medical Center. It is an on-line patient management and advisory system using evidence-based guidelines for managing septic patients in Emergency Departments. The use of model-based techniques for specifying and implementing guidelines as coordinated asynchronous processes has proved to be a promising new methodology for providing advanced clinical decision support. STEEP is currently tested at the Simulation Center of Vanderbilt University Medical Center.

Health Education Relational Network Extraction Toolkit (HORNET)		
Led by	Vanderbilt University	
Organizations Involved		
	Name	Address
1	Carnegie Mellon University	Pittsburgh, PA 15213
2	Stanford University	Stanford, CA 94305
3	University of California, Berkeley	Berkeley, CA 94720

Vanderbilt, Carnegie Mellon, and Stanford researchers have developed this open source electronic medical record access surveillance toolkit which can detect suspicious behavior with respect to usage of medical records. HORNET incorporates a suite of algorithms and statistical techniques for building social, or interaction networks in a temporal setting which is platform independent and can be integrated with existing health records infrastructures. It is currently being piloted with real-world access transaction logs from the Vanderbilt University Medical Center.

Architectural Modeling and Policy Languages		
Led by	Vanderbilt University	
Organizations Involved		
	Name	Address
1	Stanford University	Stanford, CA 94305
2	Vanderbilt University	Nashville, TN 37235

Vanderbilt and Stanford has been having regular telecons where they explore the ways how the temporal logic based policy language developed at Stanford can be integrated into the Model Integrated Computing toolsuite of Vanderbilt. The modeling environment, model analysis and model transformation tools support the precise specification of workflows in the system, while the policy language captures the policies that influence the execution of those workflows as well as guarantee the privacy, confidentiality and integrity of the data involved. The ongoing regular meetings have been helping both groups to gain better understanding of each other's technology.

Model-Based Trustworthy Health Information Systems (MOTHIS) Workshop		
Led by	Vanderbilt University	
Organizations Involved		
	Name	Address
1	Stanford University	Stanford, CA 94305
2	Vanderbilt University	Nashville, TN 37235
3	UC Berkeley	Berkeley, CA 94720
4	Cornell University	Ithaca, NY 14850

The objective of the workshop was to discuss model-based methods for the design of Health Information Systems (HIS) offering a revolutionary new way for the interaction between medical patients and Health Care Providers. While other information-intensive industries have developed and deployed standards-based, secure information infrastructures, healthcare has been characterized as a "trillion dollar cottage industry" that is still dependent upon paper records and fragmented, error-prone approaches to service delivery. The primary concern is security and privacy that needs to be organically integrated into HIS architectures. This workshop brought together computer scientists, medical experts, and legal policy experts to discuss research results in the development and application of model-based methods for representing, analyzing, and integrating architectures, privacy and security policies, computer security mechanisms, web authentication, and human factors engineering. A central focus of the discussions was a Design Platform which will provide a suite of modeling languages, modeling tools, model verification tools, and model-based generators for building HIS and integrating HIS with Electronic Medical Record systems and the business processes of providers. The second annual MOTHIS workshop took place September 30, 2008 in Toulouse, France.

Security Co-Design Toolbox		
Led by	Vanderbilt University	
Organizations Involved		
	Name	Address
1	Vanderbilt University	Nashville, TN 37235
2	Cornell University	Ithaca, NY 14850

We have developed security co-design tools that couple security with the initial design stages of sensor networks. The basic idea is that embedded (a.k.a. cyber-physical) systems must be designed with security considerations in mind. At its core, interactions are established between embedded system properties (response-time, bandwidth, data lifetime) and computer security issues. Co-design

then takes the form of interweaving security and para-functional aspects in the design process. Ongoing work is focused on security property verification of design-models and metamodel composition for integrating security modeling into embedded system design languages. The final objective is a toolbox with application-specific extensions that can be used to develop secure sensor networks in a wide variety of application domains.

4.5 Other Knowledge Transfer Outcomes

No additional knowledge transfer outcomes to report.

4.6 Knowledge Transfer Metrics/Indicators

Knowledge transfer provides the means by which research results are transitioned from Center faculty and students to society. TRUST knowledge transfer activities are both continuously monitored and periodically reviewed to ensure that they support the Center's overall knowledge transfer goals and make progress against the activity's knowledge transfer objectives. The evaluation metrics are described below.

- **Economic, Legal, and Social Impact of TRUST** – How does the activity improve the understanding of economic, legal, and social aspects of cyber security and critical infrastructure protection technologies? This impact is measured by the number of policy papers and amicus briefs produced as well as efforts to provide subject matter expertise that helps shape legislation and influences judicial decisions.
- **Testbeds** – How well does the activity leverage testbeds to promote industry and stakeholder interest and adoption? The role of the testbeds is to integrate and evaluate technologies in specific and realistic systems, keep the research on track to answer societal objectives, and demonstrate technologies to stakeholders in real systems.
- **Financial Infrastructures** – How does the activity address the unique security, privacy, and data protection challenges of the financial services industry? While a number of the problems encountered in financial infrastructures are generic to the development of trusted systems, there are several unique problems having to do with strong needs for privacy, selective revelation, and forensics.
- **Electric Power Demand Side Infrastructure** – How does the activity address the unique challenges being faced by electric power service providers, SCADA operators, and government organizations and research laboratories? The problems associated with securing electric power systems, and their associated network of SCADA components, is demanding and complex and requires solutions that solve specific issues in the security of SCADA networks.
- **Secure Global Information Grid Architectures** – How does the activity address challenges within the Department of Defense as it strives to interconnect enterprise networks, information exchange networks, and tactical networks via the Global Information Grid (GIG)? In particular, there are opportunities to provide impact in information assurance, specifically in the areas of multiple levels of security, real time information sharing architectures, and command and control architectures.

Knowledge transfer activities are periodically monitored by the TRUST Executive Board where progress of each activity (or sets of activities) is formally reviewed. Knowledge transfer activities are expected to produce specific deliverables or results such as amicus briefs, position papers, industrial liaison consultations, solution repositories, summits, and case studies.

4.7 Next Reporting Period Knowledge Transfer Plans

For the next reporting period, the Center will increase dialog with major stakeholder industries and specific companies within those industries. In particular, the Center is hoping to leverage its growing relationships with industry via the many research and education activities that have been established in the first four years of the Center.

Additionally, the Center plans on expanding the collaborative research being conducted in support of the International Collaboration for Advancing Security Technology (iCAST) program. For iCAST, TRUST researchers are not only collaborating with international researchers to develop information security technologies, they're also working on ways to increase information security public awareness and foster information security partnership among government organizations, academic institutions, and private sector companies.

The hope is to see similar sets of TRUST researchers form mini-centers in the areas of SCADA computing, electronic health care records, and trusted computing for financial applications. These mini-centers will bring additional resources to TRUST enabling the Center to leverage the government investment being made in core TRUST research and provide concrete application areas on which TRUST researchers can focus their efforts.

5 EXTERNAL PARTNERSHIPS

5.1 Goals and Objectives

One of the goals of the Center is to serve as a trusted intermediary between academics, industry, and policy makers, while simultaneously addressing long term societal needs in its research and education activities, and pursuing knowledge transfer. To integrate these objectives together, TRUST has sought to partner with representatives from the Information Technology (IT) industry and national laboratories. These partnerships not only facilitate the transfer of TRUST research results to industry but they provide an opportunity for TRUST to receive guidance in the Center's overall strategic planning and implementation through senior industry personnel on the TRUST Scientific Advisory Board (SAB).

5.2 Performance and Management Indicators

Several performance indicators are used to track progress in meeting the overall metric of global impact of the Center. As with other areas, TRUST partnerships are periodically monitored for their effectiveness in supporting the Center's partnership goals objectives. The evaluation metrics are outlined in the table below.

Objective	Metric	Frequency
Increased External Partnerships	Number of TRUST partners	Annual
Increased Amount of External Funding	Level of funding from industrial partners	Annual
Growth in Base of Knowledge Transfer Collaborators	Number of Knowledge Transfer collaborators	Annual
Joint Research Impact	Number and magnitude of joint research activities with National Laboratories	Annual
Policy and Legislation Influence	Level of interaction with Policy/Legislative organization	Annual

5.3 Current and Anticipated Problems

No significant problems were encountered during the reporting period. No significant problems are anticipated in the next reporting period.

5.4 External Partnership Activities

Partnership Activity		Industrial Research Partnership	
Led by		Shankar Sastry	
Organizations Involved			
	Name of Organization	Shared Resources (if any)	Use of Resources (if applicable)

1	University of California, Berkeley (Lead Organization)		
2	Carnegie Mellon University		
3	Cornell University		
4	Mills College		
5	San Jose State University		
6	Smith College		
7	Stanford University		
8	Vanderbilt University		

TRUST researchers and staff at all partner institutions are working with a number of industrial companies. The Industrial Research Partnership initiative strives to strengthen ties between TRUST and industry. Through this initiative, a number of industrial partners participate in knowledge transfer, serve on the Center's Scientific Advisory Board, or collaborate actively with TRUST researchers. Current TRUST industrial partners are:

- BT
- Cisco Systems
- DoCoMo USA Labs
- EADS
- ESCHER Research Institute
- Hewlett Packard
- IBM
- Intel
- Microsoft
- Oak Ridge National Laboratory
- Pirelli
- Qualcomm
- Sun
- Symantec
- Tata Consultancy Services
- Telecom Italia
- United Technologies.

The primary means of supporting the Center through the Industrial Research Partnership is for a company to become an official corporate partner at one of the Center's sponsorship levels (Affiliate, Small or Minority-Owned Business, Partner, or Premium Partner) and provide the associated level of funding to the Center. Sponsorship benefits and types of collaboration with Center faculty vary by membership level.

Partnership Activity		International Collaboration for Advancing Security Technology (iCAST)	
Led by		Shankar Sastry	
Organizations Involved			
	Name of Organization	Shared Resources (if any)	Use of Resources (if applicable)

1	University of California, Berkeley		
2	Carnegie Mellon University		

iCAST is a team consisting of members from the Taiwan Information Security Center (TWISC), the Institute for Information Industry (III), the Industrial Technology Research Institute of Taiwan (ITRI), and the Chung Cheng Institute of Technology at the National Defense University (NDU). iCAST collaborates with international institutions in various fields related to information security. In particular, TRUST currently works closely with TWISC to expand information security research and development activities, to increase information security public awareness, and foster information security partnership among government organizations, academic institutions, and private sector companies. TWISC research is in the areas of cryptology, network security, multimedia security, software security, and information security management. For this proposal, we will partner with the TWISC Education & Training Division which is focused on creating material for educational programs on information security, offering training courses and promote information sharing and public awareness of information security, and hosting training workshops in information security for academic and industrial professionals.

Partnership Activity		Air Force Team for Research in Ubiquitous Secure Technology for GIG/NCES (AF-TRUST-GNC)	
Led by		Shankar Sastry	
Organizations Involved			
	Name of Organization	Shared Resources (if any)	Use of Resources (if applicable)
1	University of California, Berkeley		
2	Cornell University		
3	Vanderbilt University		

AF-TRUST-GNC is funded by the U.S. Air Force Office of Scientific Research (AFOSR) and is researching challenges associated with the Global Information Grid (GIG) and Network Centric Enterprise System (NCES). AF-TRUST-GNC focuses on top research priorities identified in a recent study of plans to unify three major Air Force enterprise subsystems and to link the Air Force network with networks operated by other Department of Defense (DoD) services. The objective of AF-TRUST-GNC is to advance the state-of-the-art on cyber-assurance to address key trust- and QoS-related properties simultaneously throughout the lifecycles of large-scale Air Force systems via a novel combination of analytical and experimental techniques. Researchers on AF-TRUST-GNC are exploring innovation in the following areas:

- Guaranteed scalable, real-time, and fault-tolerant quality of service (QoS) for network-centric AF operational and tactical systems
- Techniques for large-scale information assurance and security policy management
- New algorithms and tools for secure scalable, information discovery, information architecture, and mediation.

5.5 Other External Partnership Outcomes

None to report.

5.6 External Partnership Metrics/Indicators

During this reporting period, there was significant progress made in the area of external partnerships. TRUST faculty and staff worked closely with a number of companies through the Center's Industrial Research Partnership program to obtain support for TRUST research projects as well as education and outreach activities. Industrial partners new to TRUST during this reporting period are DoCoMo USA Labs, EADS, and Tata Consultancy Services. Several technology companies in the Silicon Valley area allocated internship slots to graduate students for the TRUST Summer Experience, Colloquium and Research in Information Technology (SECuR-IT) program coordinated by Stanford, San Jose State, and Berkeley. Additionally, the Center has received external funding and increased the base of knowledge transfer collaborators through the iCAST and AF-TRUST-GNC research programs. These programs provide an opportunity to leverage fundamental cyber security and critical infrastructure protection research being conducted in the Center and apply it to other areas.

5.7 Next Reporting Period External Partnership Plans

During the next reporting period, we hope to increase the number of companies participating in the Center's Industrial Research Partnership program and, in particular, further pursue opportunities for external industry funding to augment the government investment made in the Center. We feel that this effort will not only further grow the number of knowledge transfer opportunities for Center research results but it will also provide TRUST faculty and students more opportunities to collaborate with industry executives and professionals and apply their research to real-world problems.

We also hope to increase the center's global presence by identifying international partners with whom the Center can partner to broaden our research, education, and knowledge transfer impact. Initial discussions have taken place with cyber security researchers, government organizations, and commercial companies in the Belgium, Denmark, Finland, India, Northern Ireland, Taiwan, and the United Kingdom.

6 DIVERSITY

6.1 Goals and Objectives

In TRUST, our diversity efforts will take a "grass roots" approach by building strong partnerships with faculty and institutions that will help us achieve our goals of inclusion of women and underrepresented minorities (URM). These partnerships will help us to cultivate the role models and mentors necessary to meet the diversity goals and objectives of the Center. Our programs can be grouped under the following goals:

- Infuse the computer science and engineering pipeline with new, diverse, and talented individuals
- Retain those individuals within TRUST research areas
- Prepare those individuals for successful careers, especially as researchers and educators in academia

Our objectives are quantified by the level of participation of women and underrepresented minorities within the Center. We seek to achieve 30% women among the Center's participants (i.e., faculty, students, research scientists, and Center staff). We also seek to achieve 10% underrepresented minorities among the Center's participants. The Center conducts assessments to track our progress towards these objectives.

6.2 Performance and Management Indicators

TRUST diversity activities are periodically monitored for meeting the Center's overall diversity objectives. Periodic monitoring consists of meetings of the TRUST Executive Board where progress of each diversity activity (or sets of activities) is formally reviewed. The diversity evaluation metrics are outlined in the table below.

Goals	Objectives	Evaluation Criteria	Frequency
Minority Faculty Research	Guided Summer Program	Number of faculty, Exit Surveys, Tracking surveys of alumni	Every 3 Years
Immersion Institute	Attract more women students to TRUST and related fields	Exit surveys, Tracking surveys of alumnae, Module development	Every 3 Years
SUPERB-TRUST	Research opportunities for minority undergraduate students at non-partner institutions	Exit surveys, Tracking surveys of alumni, Graduate school applications	Every 3 Years
Community Outreach	Dialog with public about policy, privacy, and economics	Exit surveys	Every 2 Years

Recruitment of underrepresented minority groups and women is a high priority for TRUST. For example, announcements for the SECuR-IT program were distributed via email to the following organization and websites: The Computer Alliance of Hispanic Serving Intuitions (CAHSI), Historically Black Colleges and Universities (HBCU), Louis Stokes Alliance for Minority Participation (LSAMP), Alliances For

Graduate Education and the Professoriate (AGEP), Committee for the Status of Women in Computing Research (CRA-W), California State University Computer Science Department Chairs and EECS university department chairs, Quality Education for Minorities Network (QEM) and Integrative Graduate Education and Research Traineeship (IGERT) website program portal.

6.3 Current and Anticipated Problems

No significant problems were encountered during the reporting period. No significant problems are anticipated in the next reporting period.

6.4 Diversity Activities

The sections below describe some of the Center's activities which are contributing to the development of U.S. human resources in science and engineering at the postdoctoral, graduate, undergraduate, and pre-college levels—especially those aimed at attracting, increasing, and retaining the participation of women and underrepresented groups.

Summer Undergraduate Program for Engineering Research at Berkeley-TRUST (SUPERB-TRUST) – This program supports a cohort of URM undergraduate students for an eight week summer residential program at Berkeley. The program allows undergraduate students to work with TRUST faculty and graduate students in a TRUST-related research area, experience firsthand a rigorous academic research environment, participate in technical seminars, participate in professional development activities, and present the results of their research.

Student Transitional Alliance for Research in STEM (STARS) – This is a NSF sponsored program and the partnership is designed to provide faculty and students from minority serving institutions (MSIs) with increased access to undergraduate and graduate research opportunities at six NSF-sponsored STCs (those funded in FY 2005 and FY 2006). The goals of this program are: 1) To increase the number of students from MSIs completing graduate degrees on STC campuses, 2) To increase the number of students and faculty members from under-represented groups to obtain research experience at STC sites, 3) To increase the involvement of MSI researchers on STC projects, 4) To provide an expanded forum for STCs to share their education and knowledge transfer initiatives, and 5) To increase faculty and staff diversity at STCs.

Women's Institute in Summer Enrichment (WISE) – This is a one-week residential summer program that brings together graduate students, post-doctoral fellows, and professors from all disciplines that are interested in ubiquitous secure technology and the social, political, and economical ramifications that are associated with this technology. The third offering of this program was held at Cornell University (Summer 2008); there were 30 participants with nine speakers. The Institute emphasizes the inclusion of women and underrepresented graduate students, post-doctorates, and junior faculty.

6.5 Diversity Activity Impact

"[The] under-participation in CS [computer science] by large segments of our society represents a loss of opportunity for individuals, a loss of talent in the workforce, and a loss of creativity in shaping the future of technology. Not only is it a basic equity issue, but it threatens our global economic viability as a nation."

> P. A. Freeman and J. Cuny, "Common ground: A diverse CS community benefits all of us," Computing Research News, vol. 17, 2005.

TRUST seeks to address the grand challenge as described by Freeman and Cuny. Our efforts in computer science and engineering must make strides to diversify the workforce in order to meet the future demands

of our technical profession. To that end, TRUST faculty and staff are engaged in a number of diversity activities:

Women's Institute for Summer Enrichment (WISE): WISE has now become a signature program of TRUST to attract women researchers.

Summer Undergraduate Program for Engineering Research at Berkeley-TRUST (SUPERB-TRUST): During summer 2008, Berkeley led a research project for 10 undergraduate students sponsored by the annual SUPERB-TRUST program. Participants were from diverse backgrounds and cultures, including seven female students. Many of the students came from undergraduate institutions with a limited research program and thus this was an important opportunity for them to be exposed to the cutting edge of academic research. The students worked together, as a team, on a joint project examining fuzzing and symbolic execution as an approach to finding security vulnerabilities in software. The students' research was tightly integrated into research being undertaken by TRUST faculty and graduate students.

Information Assurance Capacity Building Program (IACBP): The IACBP is a capacity building program supporting faculty development and retention in minority serving intuitions. This program also creates opportunity for future collaboration between IACBP and TRUST faculty.

Curriculum Development in Security and Information Assurance (CDSIA): The CDSIA is a capacity building program with the objective to (1) reach out to the many universities of the California State University system and to other universities whose mission is focused on work-force preparation and undergraduate education, (2) to share with faculty members of these institutions material and support structures developed by the TRUST partners, (3) to strengthen the TRUST-related community of educators, and (4) to facilitate the education of members of underrepresented communities in the domain of secure technologies.

Community Outreach: Programs like the TRUST Speakers Series provide information and technology transfer to the community at large. In addition to providing on-campus presentations, TRUST archives speaker presentations on the TAO Portal. This program is learning exchange for professionals and academics in the security field.

Recruitment of URM Faculty: Cornell invited Hakim Weatherspoon, an African-American UC Berkeley Ph.D., to join its research team as a post-doctoral researcher. During a two-year stay, Weatherspoon completed a number of papers, publishing in prominent venues such as NSDI, Eurosys, and HotOS, with additional papers still in the "pipeline". Weatherspoon has now joined Cornell's faculty as a tenure-track Assistant Professor.

6.6 Diversity Metrics/Indicators

As stated previously, the Center has established the goals of 30% participation by women and 10% participation by members of underrepresented groups. Figure 1 and Figure 2 provide the historical participation within the Center by gender and by race/ethnicity respectively. In each case, the Center is at or near the goal for participation in each category. For a perspective in computer science and engineering, the Taulbee Survey reports approximately 20% participation by women and approximately 5% participation by underrepresented minorities.

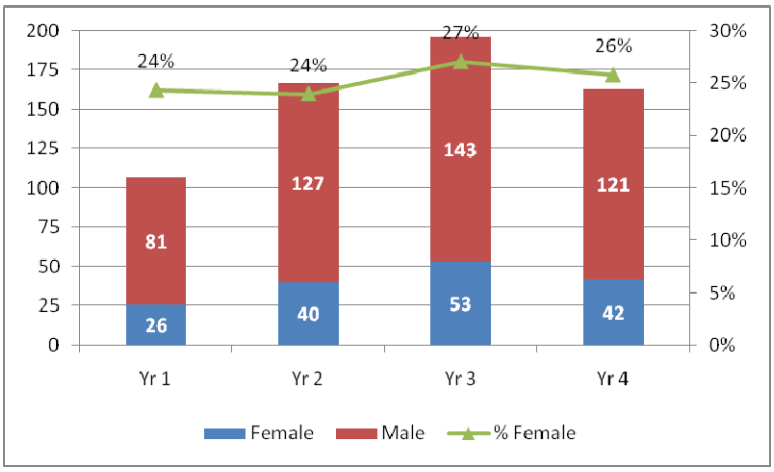

Figure 3: TRUST Participation by Gender

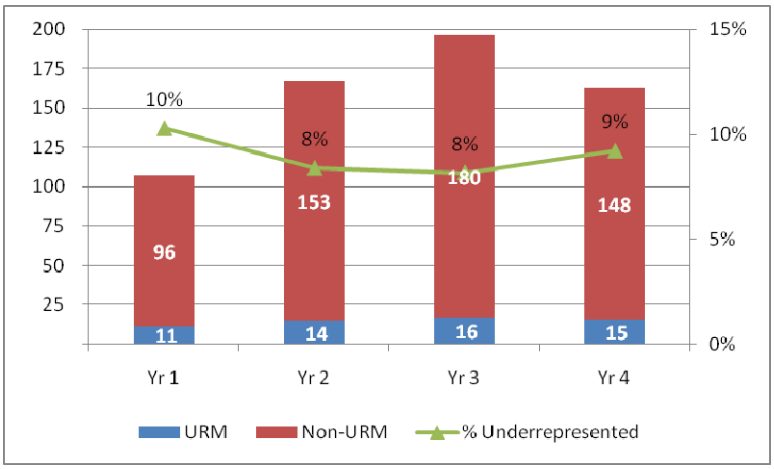

Figure 4: TRUST Participation by Race/Ethnicity

The tables below provide detail on the gender, race, and citizenship breakdown of TRUST participants in WISE, SECuR-IT, SUPERB-TRUST, and CDSIA programs during the June 1, 2008 to May 31, 2009 reporting period.

WISE 2008

Constituency	Gender		Race					US Citizen		Total
	M	F	White	African American	Asian	Hispanic	Other	Y	N	
Faculty	0	6	2	0	4	0	0	6	0	6
Graduate Students	3	15	8	1	6	3	0	14	4	18
Research Scientists	1	0	0	0	1	0	0	0	1	1
Post Doctorates	0	1	1	0	0	0	0	0	1	1
TOTALS:	**4**	**22**	**11**	**1**	**11**	**3**	**0**	**20**	**6**	**26**

SECuR-IT 2008

Constituency	Gender		Race					US Citizen		Total
	M	F	White	African American	Asian	Hispanic	Other	Y	N	
PhD Students	1	0	1	0	0	0	0	1	0	1
MS Students	4	1	2	0	2	1	0	3	2	5
TOTALS	**5**	**1**	**3**	**0**	**2**	**1**	**0**	**4**	**2**	**6**

SUPERB-TRUST 2008

Constituency	Gender		Race					US Citizen		Total
	M	F	White	African American	Asian	Hispanic	Other	Y	N	
Undergraduates	3	7	3	1	4	2	0	10	0	10
TOTALS	**3**	**7**	**3**	**1**	**4**	**2**	**0**	**10**	**0**	**10**

CDSIA 2008

Constituency	Gender		Race					US Citizen		Total
	M	F	White	African American	Asian	Hispanic	Other	Y	N	
Faculty	28	9	16	1	20	0	0	37	0	37
TOTALS	**28**	**9**	**16**	**1**	**20**	**0**	**0**	**37**	**0**	**37**

6.7 Next Reporting Period Diversity Plans

We plan to continue our successful activities such as WISE, SECuR-IT, SUPERB-TRUST, and CDSIA. To that portfolio, we will add the Bridges to Underrepresented Institutions for Long-term Development in Information Technology (BUILD-IT) conference. We plan to hold this conference at Vanderbilt with support from all TRUST campuses. The conference will select faculty mentors from Historically Black Colleges and Universities (HBCUs) and Hispanic Serving Institutions (HSIs) and provide them with the opportunity to (a) learn about TRUST research thrusts, (b) meet TRUST faculty and graduate students, and (c) discuss the Center's diversity mission, objectives, and programs. Target institutions are among the largest producers of undergraduates in computer science and engineering from traditionally underrepresented groups. We will form relationships with faculty at schools who produce our targeted participants so those faculty feel confident in sending their best and brightest to TRUST Center partner institutions. We will also address key issues for the success of a diverse population such as the campus climate, community support structure, and availability of professional development.

We have identified the Computing Alliance for Hispanic-Serving Institutions (CAHSI) and the Association of Computer/Information Sciences and Engineering Departments at Minority Institutions (ADMI) as potential partners to aid us in reaching a new generation of researchers in the area of security. We also plan to work with the Coalition to Diversify Computing to support their programming efforts, such as participating in the Richard Tapia Celebration of Diversity in Computing Conference, which TRUST Diversity Director Prof. William Robinson and TRUST Education Director Dr. Kristen Gates attended in Portland, OR. In addition, we intend to leverage our affiliation with the Empowering Leadership Alliance to connect our students with a larger community of scholars and mentors.

7 MANAGEMENT

7.1 Organizational Strategy

TRUST is organized to support the Center's strategic goals and objectives and to provide an operational structure that enables collaboration and allows the Center's researchers to primarily focus on research. At the same time, the TRUST organization has the necessary management and leadership resources that allow such a large, diverse organization to effectively function.

The TRUST organization chart is shown in Appendix B. The Center is guided by the Director (and Principal Investigator) Prof. Shankar Sastry from Berkeley. Additional Center leadership and management is provided by the Chief Scientist, Prof. Fred Schneider from Cornell; the Executive Director, Larry Rohrbough, from Berkeley; the Education Director, Dr. Kristen Gates from Berkeley; the Outreach Director, Prof. William Robinson from Vanderbilt; the Policy Director, Prof. Deirdre Mulligan from Berkeley; the Program Manager, Gladys Khoury from Berkeley; and the Program Coordinator, Sally Alcala, from Berkeley.

The Executive Board manages and executes the overall administration of the Center. The Executive Committee consists of the Center Director, Chief Scientist, Executive Director, Education Director, Outreach Director, Policy Director, Program Manager, and university Principal Investigators.

7.2 Performance and Management Indicators

Effective operation and management of the Center depends on several key processes and agreements. One of which is the set of TRUST Center By-Laws. The By-Laws were drafted and accepted into practice in the first year of the Center and govern the operation and management of the Center.

The TRUST Center By-Laws are as follows:

1. The TRUST center will be administered by a board of directors with no more than nine directors and no fewer than five directors. The Board will have a Chairman.

2. The board will have as ex-officio members the co-PIs of the NSF STC TRUST proposal: that is, John Mitchell, Mike Reiter, Shankar Sastry, Janos Sztipanovits and Steve Wicker will be the Board members. Shankar Sastry will be the Chairman of the Board. The chairman of the board will be responsible for conducting the meetings, or delegating the conducting of the meeting to another board member.

3. Directors are elected to or removed from the board by 2/3 vote of the standing directors rounded up to the next integer (for example, if the board has 5, then 4 must vote in favor, if 4, then 3, and if 3, then 2).

4. A quorum for a directors meeting consists of 2/3 of the directors. Meetings will be scheduled at an average interval of once a month until modified by the directors.

5. Directors meetings can be scheduled by a 2/3 vote, and directors will be notified at least one week in advance.

6. A quorum for a directors meeting consists of 2/3 of the directors and decisions made at such a meeting are final. Participation by telephone at the meetings is fine.

7. Unless otherwise stated, any decision by the board is by majority vote (either a majority of the directors present at a meeting, or a majority of the standing directors if the decision is made without a meeting). Obtaining votes by email is acceptable.

8. Major TRUST activities including research, education and outreach directions will be reported to the board on a periodic basis, not to exceed three months, for concurrence.

9. A Secretary will be appointed by the board, and will be responsible for recording decisions made by the board and distributing a summary of the deliberations to any board members not present at a meeting.

10. A Treasurer will be appointed by the board, and will be responsible for reporting financial status to the board, including cash flow position and projections for all accounts that are part of the TRUST center.

11. The bylaws can be modified by a 2/3 vote of the standing board. Amendments will be logged in and kept current by the secretary of the Board.

7.3 Management Metrics/Indicators

During this reporting period, the Center leadership provided effective management and guidance. Center staff, Principal Investigators, and members of the Executive Board worked together to provide an operational structure that supported the research, education, and knowledge transfer goals of the Center as well as an infrastructure for running the day-to-day aspects of the Center.

7.4 Current and Anticipated Problems

No significant problems were encountered during the reporting period. No significant problems are anticipated in the next reporting period.

7.5 Management and Communications System

The TRUST management structure includes a number of systems and processes that foster communication within the Center. First, the TRUST website (www.truststc.org) is designed to be a comprehensive resource for obtaining TRUST-related material and communicating with TRUST researchers and staff. The TRUST website provides e-mail lists, collaborative workspaces, access to publications and presentations, news items, blogs, information on past and future TRUST events, and workshop/conference registration pages. Industrial, governmental and academic participants have individual accounts and membership in multiple workspaces via a secure login procedure. E-mail lists and newsgroups are linked to each other providing easy access to discussion threads. E-mail messages are archived and are searchable. Resources such as workgroups and publications have fine grained access control and the website provides workgroup web pages via participant supplied HTML and Wiki pages. There have been no problems with the website, despite that fact that its content has grown significantly as has the number of registered users and page views and its infrastructure has become the primary means by which information is communicated to TRUST researchers and the wider TRUST community.

In order to ensure regular dialogue and communication across partner institutions, the TRUST Executive Board holds standing monthly meetings to discuss the current status of projects, funding and resource allocation, and other management and operational issues. Ad hoc meetings are also arranged as necessary in addition to these regularly scheduled meetings and the frequency of the Executive Board meetings has

changed from monthly to bi-monthly to weekly as necessary to allow the group ample opportunities to confer and make timely decisions.

7.6 Center Advisory Personnel

TRUST receives outside advice, guidance, and counsel from two groups: the External Advisory Board (EAB) and the Industrial Advisory Board (IAB). Each group is described in more detail below.

External Advisory Board – The TRUST EAB is a distinguished group of experts in research, education, policy, and management whose guidance supplements the strategic planning by TRUST management and the TRUST Executive Board. The primary goal of the EAB is to offer an independent assessment of TRUST research, education, outreach, and diversity accomplishments, goals, and plans. EAB input plays a crucial role in the annual revision of the TRUST strategic plan.

The EAB's effectiveness is directly related to its ability to offer unbiased counsel; as such, self-governance is a guiding principle in the EAB's charter. EAB members are appointed for three year terms and the EAB is headed by a chairperson, who is also appointed for a term of three years.

NSF policies on conflict of interest govern the independence of the EAB and require that EAB members do not have financial interests or collaborations with faculty and staff being supported by TRUST funding. The EAB meets annually and performs the following functions:

- First, it reviews the TRUST strategic plan, project plans, and annual report on research, education, and outreach. Unfettered Q&A sessions during TRUST briefs facilitate collecting information on pivotal points.
- Second, the EAB conducts deliberations, which occur in closed session presided by the EAB chairperson.
- Third, the EAB produces a report and presents its findings to the TRUST Executive Board and the Vice Chancellor of Research at the TRUST lead institution, UC Berkeley.

EAB members and their affiliations are listed in the table below.

	Name	Affiliation
1	Alfred Aho	Columbia University
2	Annie Anton	North Carolina State University
3	Patricia Bellia	University of Notre Dame
4	Matthew Davis	University of California
5	Lee Burge	Tuskegee University
6	David Clark	Massachusetts Institute of Technology
7	George Cybenko	Dartmouth College
8	James Johnson	Howard University
9	Jay Lala	Raytheon
10	Carl Landwehr	University of Maryland
11	Teresa Lunt	Palo Alto Research Center
12	Dan Manson	California State Polytechnic University
13	Andrew Odlyzko	University of Minnesota
14	William Sanders	University of Illinois at Urbana-Champaign
15	Joseph Sifakis	CNRS, Verimag
16	Gene Spafford	Purdue University

Industrial Advisory Board – The TRUST IAB consists of senior executives and thought leaders from industry, academia, and government and commercial research laboratories. The primary goal of the SAB is to engage the TRUST Executive Board to communicate industry's perspective and research needs and help the Executive Board develop and execute a successful Center/Industry partnership model.

IAB members and their affiliations are listed in the table below.

	Name	Affiliation
1	Andrew Chien	Intel
2	Jean Colpin	United Technologies Research Center
3	Phil Edholm	Nortel Networks
4	Pieroguido Iezzi	Perelli
5	Wayne Johnson	HP Laboratories
6	William Mark	SRI International
7	John W. Noerenberg	Qualcomm
8	Giovanni Penna	Telecom Italia
9	Emil Sarpa	Sun Microsystems
10	Steve Trilling	Symantec

7.7 Center Strategic Plan Changes

Changes to the TRUST Strategic Plan are indicated within that document. The TRUST Strategic Plan was last updated September 18, 2008.

8 CENTER-WIDE OUTPUTS AND ISSUES

8.1 Center Publications

The following sections provide lists of various TRUST Center publications produced during this reporting period. Publications are listed in reverse chronological order and are grouped into the following categories based on their publication type: Peer Reviewed Publications, Journal Articles, Books and Book Chapters, and Non-Peer Reviewed Publications. For each publication, a link to the TRUST publications database is provided as reference.

8.1.1 Peer Reviewed Publication

- Rethinking Reliable Transport for the Datacenter, Mahesh Balakrishnan, Joe Hoffert, Ken Birman, Douglas Schmidt, Proceedings of the Large-Scale Distributed Systems and Middleware Workshop (LADIS 2008), September, 2009

- Automatic Dimension Inference and Checking for Object-Oriented Programs, Sudheendra Hangal and Monica S. Lam, 31st International Conference on Software Engineering, May, 2009

- Reducing Power Consumption with Relaxed Quasi Delay-Insensitive Circuits, Christopher LaFrieda, Rajit Manohar, IEEE International Symposium on Asynchronous Circuits and Systems, May, 2009

- Smoke and Mirrors: Reflecting Files at a Geographically Remote Location Without Loss of Performance, Hakim Weatherspoon, Lakshmi Ganesh, Tudor Marian, Mahesh Balakrishnan, Ken Birman, 7th USENIX Conference on File and Storage Technologies, February, 2009

- Blue versus Red: Towards a model of distributed security attacks, Neal Fultz and Jens Grossklags, Proceedings of the Thirteenth International Conference Financial Cryptography and Data Security, February, 2009

- Conditioned-safe Ceremonies and a User Study of an Application to Web Authentication, Chris Karlof, Doug Tygar, David Wagner, Sixteenth Annual Network and Distributed Systems Security Symposium, 2009

- Device driver safety through a reference validation mechanism, Fred Schneider, Dan Williams, Patrick Reynolds, Kevin Walsh & Emin Gun Sirer, Proceedings of the 8th USENIX Symposium on Operating Systems Design and Implementation OSDI '08, December, 2008

- LQG Control Over Multi-Channel TCP-like Erasure Networks with Probabilistic Packet Acknowledgements, E. Garone, B. Sinopoli, A. Casavola, Proceedings of IEEE Conference on Decision and Control, December, 2008

- Characterization of the Critical Value for Kalman Filtering with Intermittent Observations, Y.Mo, B. Sinopoli, Proceedings of IEEE Conference on Decision and Control, December, 2008

- CyberRadar: A Regression Analysis Approach to the Identification of Cyber-Physical Mappings in Process Control Systems, Julian L. Rrushi, Kyoung-Don Kang, Proceedings of the IEEE/ACM Workshop on Embedded Systems Security, Catherine Gebotys, Dimitrios Serpanos, October, 2008

- An Outer Bound to the Admissible Source Region of Discrete Memoryless Arbitrarily Varying General Broadcast Channels, Amin Aminzadeh Gohari, Venkatachalam Anantharam, "Proceedings of the 46th Annual Allerton Conference on Communications, Control and Computing", Urbana, Illinois, September 23 -26, 2008, 2008

- Supporting Large-scale Continuous Stream Datacenters via Pub/Sub, Joe Hoffert, Douglas Schmidt, Mahesh Balakrishnan, Ken Birman, Proceedings of the Large-Scale Distributed Systems and Middleware Workshop (LADIS 2008), September, 2008

- Integration of Clinical Workflows with Privacy Policies on a Common Semantic Platform, Jan Werner, Bradley Malin, Yonghwan Lee, Akos Ledeczi, Janos Sztipanovits, 2nd International Workshop on Model-Based Design of Trustworthy Health Information Systems, September, 2008

- LQG Control over Lossy TCP-like Networks with Probabilistic Packet Acknowledgements, E. Garone, B. Sinopoli, A. Casavola, The Forty-Sixth Annual Allerton Conference on Communication, Control, and Computing, September, 2008

- Optimal Sensor Density for Remote Estimation Over Wireless Sensor Networks, R. Ambrosino, B. Sinopoli, K. Poolla, S. Sastry, Proceedings of the Forty-Sixth Annual Allerton Conference on Communication, Control, and Computing, September, 2008

- Distributed Sensor Localization in Euclidean Spaces: Dynamic Environments, U. A. Khan, S. Kar, B. Sinopoli, J. M. F. Moura, Proceedings of the Forty-Sixth Annual Allerton Conference on Communication, Control, and Computing, September, 2008

- An Algorithmic Approach to Authorization Rules Conflict Resolution in Software Security, Weider D. Yu, Ellora Nayak, Proceedings of the Thirty-second IEEE International Computer Software & Applications Conference(COMPSAC), IEEE Computer Society, 32-35, July, 2008

- A Web-Based Wireless Mobile System Design of Security and Privacy Framework for u-Healthcare, Weider D. Yu, Sriram Mudumbi, Roopa Gummadikayala, Proceedings of the 10th IEEE International Conference on e-Health Networking, Applications and Services(HEALTHCOM), IEEE Computer Society, July, 2008

- Security, Safety and Privacy – Pervasive Themes for Engineering Education, Sigurd Meldal, Kristen Gates, Russell Smith, Xiao Su, ICEE 2008, Geza Varady, iNEER, July, 2008

- Automatic Generation of XSS and SQL Injection Attacks with Goal-Directed Model Checking., Michael Martin and Monica S. Lam., 17th USENIX Security Symposium, July, 2008

- Security and Insurance Management in Networks with Heterogeneous Agents, J. Grossklags, N. Christin, J. Chuang, Proceedings of ACM E-Commerce Conference, July, 2008

- Sound, Complete and Scalable Path-Sensitive Analysis, I. Dillig, T. Dillig, A. Aiken, Conference on Programming Language Design and Implementation, 270-280, June, 2008

- End-to-End Enforcement of Erasure and Declassification, Stephen Chong, Andrew C. Myers, Proceedings of the IEEE Computer Security Foundations Symposium, 98–111, June, 2008

- Perspectives: Improving SSH-style Host Authentication with Multi-path Network Probing, Dan Wendlandt, Dave Andersen, Adrian Perrig, USENIX Annual Technical Conference, June, 2008

- Secure Control: Towards Survivable Cyber-Physical Systems, Alvaro Cardenas, Saurabh Amin, S. Shankar Sastry, First International Workshop on Cyber-Physical Systems (WCPS2008), IEEE, June, 2008
- Mechanisms to Provide Integrity in SCADA and PCS Devices, A. Shah, A. Perrig, B. Sinopoli, Proceedings of the International Workshop on Cyber-Physical Systems - Challenges and Applications (CPS-CA '08)., June, 2008
- Securing Nonintrusive Web Encryption through Information Flow, Lantian Zheng, Andrew C. Myers, June, 2008
- CareNet: An Integrated Wireless Sensor Networking Environment for Remote Healthcare, Shanshan Jiang, Yanchuan Cao, Sameer Iyengar, Philip Kuryloski, Roozbeh Jafari, Yuan Xue, Ruzena Bajcsy, Stephen Wicker, BodyNets, 2008
- On the Asymptotic Behavior of Selfish Transmitters Sharing a Common Wireless Communication Channel, Hazer Inaltekin, Mung Chiang, H. Vincent Poor, Stephen Wicker, IEEE ISIT 2008, 2008
- New Bounds on the Information-Theoretic Key Agreement of Multiple Terminals, Amin Aminzadeh Gohari, Venkatachalam Anantharam, To appear in ``Proceedings of the IEEE International Symposium on Information Theory'', Toronto, Canada, 2008
- A Logic for Reasoning about Networked Secure Systems, Deepak Garg, Jason Franklin, Dilsun Kaynar, Anupam Datta, Joint Workshop FCS-ARSPA-WITS, 2008
- Taking Advantage of Data Correlation to Control the Topology of Wireless Sensor Networks, Sergio Bermudez, Stephen Wicker, International Conference on Telecommunications, 2008
- Automatic Inference of Stationary Fields: a Generalization of Java's Final Fields, C. Unkel, M. Lam, 35th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, (San Francisco, CA) ACM, 2008
- Link Privacy in Social Networks., A. Korolova, S. Nabar, Y. Xu, R. Motwani, Proceedings of the 21st International Conference on Data Engineering (ICDE), 2008
- Proceedings of the 17th Usenix Security Symposium, Michael Martin, Monica Lam, Automatic generation of XSS and SQL injection attacks with goal-directed model checking., 2008
- QuickPay online payment protocol, Mark Stamp, Jian Dai, Proceedings of SEKE '08, 2008
- On the connectivity of finite wireless networks with multiple base stations, Sergio Bermudez, Stephen Wicker, International Conference on Computer Communications and Networks, 2008
- Gossip-based Distribution Estimation in, Maya Haridasan, Robbert van Renesse, 7th International Workshop on Peer-to-Peer Systems (IPTPS '08, 2008
- Programming with Live Distributed Objects, Krzysztof Ostrowski, Ken Birman, Danny Dolev, Jong Hoon Ahnn, 22nd European Conference on Object-Oriented Programming (ECOOP 2008, 2008
- Quicksilver Scalable Multicast (QSM), Krzysztof Ostrowski, Ken Birman, Danny Dolev, 7th IEEE International Symposium on Network Computing and Applications (IEEE NCA 2008, 2008

- Exploiting Interference Diversity for Event-Based Spectrum Sensing, Arash Parsa, Amin Aminzadeh Gohari, Anant Sahai, 2008 IEEE Symposium on Dynamic Spectrum Access Networks (DySpAN)., 2008
- Oblivious Routing for Wireless Mesh Networks, Jonathan Wellons, Yuan Xue, IEEE International Conference on Communications (ICC), 2008
- Automated Whitebox Fuzz Testing, Patrice Godefroid, Michael Y. Levin, David A Molnar, Network Distributed Security Symposium (NDSS), Internet Society, 2008
- Active Property Checking, Patrice Godefroid, Michael Y. Levin, David A Molnar, EMSOFT, Association for Computing Machinery, 2008
- The Building Blocks of Consensus, Yee Jiun Song, Robbert van Renesse, Fred Schneider, Danny Dolev, 9th International Conference on Distributed Computing and Networking (ICDCN '08, 2008
- Enforcing Fairness in a Live-Streaming System, Robbert van Renesse, Maya Haridasan, Ingrid Jansch-Porto, Multimedia Computing and Networking (MMCN 2008, 2008
- DQML: A Modeling Language for Configuring Distributed Publish/Subscribe Quality of Service Policies, Joe Hoffert, Douglas Schmidt, Aniruddha Gokhale, Proceedings of the 10th International Symposium on Distributed Objects, Middleware, and Applications, 2008
- Minimum Disruption Service Composition and Recovery in Mobile Ad Hoc Networks, Shanshan Jiang, Yuan Xue, Douglas Schmidt, Computer Network Journal, Special Issue on Autonomic and Self-Organizing Systems, 2008
- Predictive or Oblivious: A Comparative Study of Routing Strategies for Wireless Mesh Networks Under Uncertain Demand, Jonathan Wellons, Liang Dai, Yi Cui, Yuan Xue, The Fifth Annual IEEE Communications Society Conference on Sensor, Mesh, and Ad Hoc Communications and Networks, 2008
- Integrating Traffic Estimation and Routing Optimization for Multi-Radio Multi-Channel Wireless Mesh Networks, Liang Dai, Yuan Xue, Bin Chang, Yi Cui, IEEE INFOCOM, 2008
- A Model-Integrated Approach to Implementing Individualized Patient Care Plans Based on Guideline-Driven Clinical Decision Support and Process Management - A Progress Report, Jason Martin, Janos Laszlo Mathe, Peter Miller, Akos Ledeczi, Liza Weavind, Anne Miller, David Maron, Andras Nadas, Janos Sztipanovits, 2nd International Workshop on Model-Based Design of Trustworthy Health Information Systems MOTHIS 2008, 2008
- On The Impossibility of Basing Identity Based Encryption on Trapdoor Permutations, D. Boneh, A. Papakonstantinou, C. Rackoff, Y. Vahlis, and B. Waters, FOCS, 2008
- Overshadow: A Virtualization-Based Approach to Retrofitting Protection in Commodity Operating Systems, M. Chen, P. Subrahmanyam, C. Waldspurger, E. C. Lewis, Tal Garfinkel, D. Boneh, D. Ports, and J. Dwoskin, ASPLOS, 2008

8.1.2 *Journal Articles*

- The Monoculture Risk Put into Context, Fred Schneider, Ken Birman, IEEE Security & Privacy Magazine, 2009
- Quantifying Information Flow with Beliefs, Andrew C. Myers, Michael R. Clarkson, Fred Schneider, Journal of Computer Security, 2009
- Trustworthiness as a Limitation on Network Neutrality, Fred Schneider, Aaron Burstein, Federal Communications Law Journal, 61, 2009
- A One-shot Random Access Game for Wireless Networks – Behavior of Nodes at Nash Equilibria, Hazer Inaltekin, IEEE/ACM Transactions on Networking, December, 2008
- Privacy Decision Making in Administrative Agencies, Kenneth A. Bamberger, Deirdre Mulligan, Chicago Law Review, 75, 1, 75, January, 2008
- A Combined Localization and Geographic Routing Algorithm for Rapidly-Deployed Wireless Sensor Networks, Hui Qu, Stephen Wicker, International Journal of Distributed Sensor Networks, 4, 1, 44-63, 2008
- Confidentiality in Sensor Networks: Transactional Information, Sameer Pai, Marci Meingast, Tanya Roosta, Sergio Bermudez, Stephen Wicker, Deirdre K. Mulligan, S. Shankar Sastry, IEEE Security and Privacy Magazine, 2008
- Co-designed anchor-free localization and location-based routing algorithm for rapidly-deployed wireless sensor networks, Hui Qu, Stephen Wicker, Information Fusion, 2008
- Information-Theoretic Key Agreement of Multiple Terminals - Part II: Channel Model, Amin Aminzadeh Gohari, Venkatachalam Anantharam, Submitted to ``IEEE Transactions on Information Theory", 2008
- Model-based design of clinical information systems., Janos Laszlo Mathe, Jan Werner, Yonghwan Lee, Bradley Malin, Akos Ledeczi, Methods of Information in Medicine, 47, 5, 399-408, 2008
- Model-Based Design of Trustworthy Health Information Systems, Ruth Breu, Janos Sztipanovits, Elske Ammenwerth, Methods of Information in Medicine, 2008
- Formalizing the Structural Semantics of Domain-Specific Modeling Languages, Ethan Jackson, Janos Sztipanovits, Journal of Software and Systems Modeling, 2008
- A Survey of Worm Detection and Containment, Pele Li, Mehdi Salour, Xiao Su, IEEE Communications Surveys and Tutorials, 10, 20-35, 2008
- Distributed Intrusion Detection System for Resource-Constrained Devices in Ad Hoc Networks, Adrian Lauf, William H. Robinson, Elsevier Journal of Ad Hoc Networks, 2008
- Information-Theoretic Key Agreement of Multiple Terminals - Part I: Source Model, Amin Aminzadeh Gohari, Venkatachalam Anantharam, Submitted to ``IEEE Transactions on Information Theory", 2008

- Profile hidden Markov models and metamorphic virus detection, Mark Stamp, Srilatha Attaluri, Scott McGhee, Journal in Computer Virology, 2008
- An agent-based privacy enhancing model, M. Stamp, H.-H. Lee, Information Management & Computer Security, 2008

8.1.3 Books and Book Chapters

- Memory Corruption Attacks, Defenses, and Evasions, Carlo Bellettini, Julian L. Rrushi, Jatinder N. D. Gupta, Sushil K. Sharma, 12, 139-151, 1st, Information Science, 2008
- P2PTunes: A peer-to-peer digital rights management system, Mark Stamp, Ramya Venkataramu, IGI Global, 2008
- Digital rights management for untrusted peer-to-peer networks, Mark Stamp, Pallavi Priyadarshini, IGI Global, 2008
- Digital rights management for streaming media, Mark Stamp, Deepali Brahmbhatt, IGI Global, 2008
- Secrecy Analysis in Protocol Composition Logic, A. Roy, A. Datta, A. Derek, J. C. Mitchell, J.P. Seifert, Formal Logical Methods for System Security and Correctness, IOS Press, 2008

8.1.4 Non-peer Reviewed Publications

- TRUST Annual Report 2007-2008, Faculty and Staff, Team for Research in Ubiquitous Secure Technology (TRUST), an NSF Science and Technology Center., June, 2009
- SmartFuzz and MetaFuzz, David A Molnar, David Wagner, University of California Berkeley, 2009
- Hyperproperties, Fred Schneider, Michael R. Clarkson, Journal of Computer Security, December, 2008
- TRUST 2nd 5-Year Strategic and Implementation Plan, Ruzena Bajcsy, Kristen Gates, Sigurd Meldal, John C. Mitchell, Deirdre Mulligan, Adrian Perrig, William H. Robinson, Larry Rohrbough, S. Shankar Sastry, Fred Schneider, Janos Sztipanovits, Doug Tygar, Stephen Wicker, TRUST Science & Technology Center, September, 2008
- A New Outer Bound to the Capacity Region of Deterministic-Code Discrete Memoryless Arbitrary Varying General Broadcast Channel, Amin Aminzadeh Gohari, Venkatachalam Anantharam, UC Berkeley, UCB/EECS-2008-81, June, 2008

8.2 Conference Presentations

The following is a list of conference presentations made by TRUST Center personnel during this reporting period. For each presentation, a link to the TRUST publications database is provided as reference.

- A Model-Integrated Approach to Implementing Individualized Patient Care Plans Based on Guideline-Driven Clinical Decision Support and Process Management, Janos Laszlo Mathe,

Jason Martin, Peter Miller, Liza Weavind, David Maron, Akos Ledeczi, Anne Miller, Andras Nadas, Janos Sztipanovits, November, 2009

- Co-design Environment for Secure Embedded Systems, Matthew Eby, Janos Laszlo Mathe, Jan Werner, Gabor Karsai, Sandeep Neema, Janos Sztipanovits, Yuan Xue, February, 2009
- MD5 considered harmful today: Creating a rogue CA certificate, Alex Sotirov, Marc Stevens, Jacob Appelbaum, Arjen Lenstra, David A Molnar, Dag Arne Osvik, Benne de Weger, 30, December, 2008
- Bootstrapping Trust in a 'Trusted' Platform, Bryan Parno, 11, November, 2008
- Secure Control and the Analysis of Denial of Service Attacks, Saurabh Amin, Alvaro Cardenas, Alex Bayen, S. Shankar Sastry, 11, November, 2008
- Detecting Forged TCP Reset Packets, Nicholas Weaver, Robin Sommer, Vern Paxson, 11, November, 2008
- Expressing and Enforcing Flow-Based Network Security Policies, Tim Hinrichs, Natasha Gude, Martin Casado, John C. Mitchell, Scott Shenker, 11, November, 2008
- Open Problems in the Security of Learning, Marco Barreno, Peter L. Bartlett, Fuching Jack Chi, Anthony Joseph, Blaine Nelson, Benjamin I. Rubinstein, Udam Saini, Doug Tygar, 11, November, 2008
- Automatic Patch-Based Exploit Generation is Possible: Techniques and Implications, David Brumley, Pongsin Poosankam, Dawn Song, Jiang Zheng, 12, November, 2008
- Programming with Live Distributed Objects, Krzysztof Ostrowski, Ken Birman, Danny Dolev, Jong Hoon Ahnn, 12, November, 2008
- Smoke and Mirrors: Shadowing Files at a Geographically Remote Location Without Loss of Performance, Hakim Weatherspoon, Lakshmi Ganesh, Tudor Marian, Mahesh Balakrishnan, Ken Birman, 12, November, 2008
- Online Information Security Education through Anchored Instruction, Eric Imsand, Larry Howard, Ken Pence, Mike Byers, Dipankar Dasgupta, 11, November, 2008
- A Security Standard for Smart Power Meters, Coalton Bennett, Darren Highfill, Stephen Wicker, 11, November, 2008
- The FBI and Emerging Threats of Computer Intrusions and Cyber Crime, Scott E. Augenbaum, 11, November, 2008
- A Model-Integrated Approach to Implementing Individualized Patient Care Plans Based on Guideline-Driven Clinical Decision Support and Process Management, Jason B. Martin, Janos L. Mathe, Peter Miller, Akos Ledeczi, Liza Weavind, Ann Miller, David J. Maron, Andras Nadas, Janos Sztipanovits, 11, November, 2008
- Integration of Clinical Workflows with Privacy Policies on a Common Semantic Domain, Jan Werner, Bradley Malin, Yonghwan Lee, Akos Ledeczi, Janos Sztipanovits, 11, November, 2008

- Automatic Detection of Policies from Electronic Medical Record Access Logs, John Paulett, Bradley Malin, 11, November, 2008
- Fault Tolerant Sensor Network Routing for Patient Monitoring, Shanshan Jiang, Annarita Giani, Allen Yang, Yuan Xue, Ruzena Bajcsy, 11, November, 2008
- Redundancy Minimizing Techniques for Robust Transmission in Wireless Networks, Anna Kacewicz, Stephen Wicker, 11, November, 2008
- DexterNet: An Open Platform for Heterogeneous Body Sensor Networks and Its Applications, Philip Kuryloski, Annarita Giani, Roberta Giannantonio, Katherine S. Gilani, Ville-Pekka Seppa , Edmund Seto , Raffaele Gravina , Victor Shia, Curtis Wang, Posu Yan, Allen Yang, Jari Hyttinen, S. Shankar Sastry, Stephen Wicker, Ruzena Bajcsy, 11, November, 2008
- An Intrusion Detection System for Wireless Process Control Systems, Adrian Lauf, Jon Wiley, Tanya Roosta, William H. Robinson, Gabor Karsai, 11, November, 2008
- On the Connectivity of Finite Wireless Networks with Multiple Base Stations, Sergio Bermudez, Stephen Wicker, 11, November, 2008
- Quantitative Information Flow as Network Flow Capacity, Stephen McCamant, Michael D. Ernst, 12, November, 2008
- Comparison of Blackbox and Whitebox Fuzzers in Finding Software Bugs, Marjan Aslani, Nga Chung, Jason Doherty, Nichole Stockman, William Quach, 12, November, 2008
- Applying policy-based intrusion detection to scada networks, Adrian Lauf, William H. Robinson, Gabor Karsai, Jon Wiley, Tanya Roosta, 11, November, 2008
- An Algorithmic Approach to Authorization Rules Conflict Resolution in Software Security, Weider D. Yu, Ellora Nayak, 12, November, 2008
- Verifying the Safety of User Pointer Dereferences, Suhabe Bugrara, Alex Aiken, 12, November, 2008
- Verifiable Functional Purity in Java, Matthew Finifter, Adrian Mettler, Naveen Sastry, David Wagner, 12, November, 2008
- The TRUST-SCADA Experimental Testbed: Design and Experiments, Annarita Giani, Gabor Karsai, Aakash Shah, Bruno Sinopoli, Jon Wiley, 12, November, 2008
- Towards a Scalable System for Distributed Management of Private Information, Michael Siegenthaler, Ken Birman, 11, November, 2008
- Poster - Applying Policy-Based Intrusion Detection to SCADA Networks, Adrian Lauf, Tanya Roosta, Jon Wiley, William H. Robinson, Gabor Karsai, 11, November, 2008
- STEEP - A Model-Integrated Clinical Information System Application, Janos Laszlo Mathe, Jason Martin, Peter Miller, Akos Ledeczi, Liza Weavind, Anne Miller, David Maron, Andras Nadas, Janos Sztipanovits, November, 2008
- Integration of Clinical Workflows with Privacy Policies on a Common Semantic Platform, Jan Werner, Bradley Malin, Yonghwan Lee, Akos Ledeczi, Janos Sztipanovits, 30, September, 2008

- A Model-Integrated Approach to Implementing Individualized Patient Care Plans Based on Guideline-Driven Clinical Decision Support and Process Management - A Progress Report, Jason Martin, Janos Laszlo Mathe, Peter Miller, Akos Ledeczi, Liza Weavind, Anne Miller, David Maron, Andras Nadas, Janos Sztipanovits, 30, September, 2008

- Supporting Large-scale Continuous Stream Datacenters via Pub/Sub Middleware and Adaptive Transport Protocols, Joe Hoffert, Douglas Schmidt, Mahesh Balakrishnan, Ken Birman, September, 2008

- Supporting Scalability and Adaptability via ADAptive Middleware And Network Transports (ADAMANT), Joe Hoffert, Douglas Schmidt, Mahesh Balakrishnan, Ken Birman, July, 2008

- TRUST: A Collaborative Approach to Advancing Cyber Security Research and Development, Larry Rohrbough, 12, June, 2008

8.3 Other Dissemination Activities

The following is a list of other dissemination activities associated with TRUST Center personnel during this reporting period that are not covered elsewhere in this report.

- July 2008 Jennifer King presented the preliminary findings of her RFID work at a Federal Trade Commission workshop on RFID-based wireless payment systems in Seattle, WA. This project was also highlighted in a poster at the Symposium on Usable Privacy and Security at Carnegie Mellon University.

- September 4-5, 2008 Professor Deirdre K. Mulligan presented at the Stanford Energy & Feedback Workshop: End-use Energy Reductions through Monitoring, Feedback, and Behavior Modification at the Precourt Institute for Energy Efficiency. She discussed the privacy and security issues arising in smart grids and mechanisms to address them.

- September 9, 2008 Professor Deirdre K. Mulligan participated in a press dinner on the Future of privacy sponsored by Microsoft.

- September 15, 2008 Professor Deirdre K. Mulligan was the Constitution Day Speaker at UNC-Charlotte. Mulligan discussed changes in legal rules about government access to personal information and the use of technology in the war on terrorism.

- September 16, 2008 Professor Deirdre K. Mulligan presented at the Second Annual IT Ethics Luncheon/Workshop, "Privacy's Very Uneasy Relationship With Technology," at UNC Charlotte Center for Professional & Applied Ethics.

- September 24-25, 2008 Professor Deirdre K. Mulligan participated in a NSF workshop on Network Design and Societal Values and assisted in writing the report.

- October 9, 2008 Professor Deirdre K. Mulligan participated in a research discussion roundtable with Craig Mundie, Chief Research and Strategy Officer of Microsoft.

- October 20, 2008: Presented research on identity theft to the Cybercrime Studies Center at John Jay College of Criminal Justice, City University of New York (CUNY). This presentation

focused upon methods of identity theft and the need for training in the criminal justice field to interdict the crime.

- November 18-19, 2008 Professor Deirdre K. Mulligan participated in a Harvard workshop on Protecting Subject Privacy for Large Scale Experimentation and Research.

- Fall 2008: Published "Toward a Market for Bank Safety," in 21 Loy. Consumer L. Rev. 101 (Fall 2008), an analysis of Federal Trade Commission identity theft statistics. This article argued that consumers have no ability to choose financial institutions based upon their effectiveness in preventing identity fraud. It proposes metrics for measuring and comparing the prevalence of fraud at financial institutions.

- January 7, 2009 Professor Deirdre K. Mulligan and researcher Jennifer King presented the San Francisco Community Safety Camera Study to the President of the SF Police Commission.

- January 14, 2009 Professor Deirdre K. Mulligan, researcher Jennifer King, and Professor Stephen Rafael presented the San Francisco Community Safety Camera Study to the SF Police Commission.

- March 5, 2009 Professor Deirdre K. Mulligan chaired the second Workshop on the Economics of the Information Infrastructure. The workshop brought together economists, lawyers, computer scientists and others to discuss specific proposals to alter the technical architecture and policy framework of the internet.

- March 6, 2009 Co-chaired Berkeley Center for Law & Technology and Berkeley Technology Law Journal Symposium on security breach notification. This day-long event brought together experts in economics, law, and information security to discuss the problem of security breaches, their prevention, and remediation.

- March 31, 2009 Professor Deirdre K. Mulligan gave a presentation on the privacy and civil liberties concerns in public safety surveillance cameras for the International Association of Chiefs of Police Public Safety Technology and Policy Symposium. Mulligan discussed the findings of the San Francisco Community Camera Report.

- April 21, 2009 Professor Deirdre K. Mulligan participated in a panel on Surveillance: Security, Privacy and Risk at RSA.

- May 4, 2009, Professor Deirdre K. Mulligan organized and moderated a panel about human rights and corporate responsibility and complicity in the telecommunications, communications and information sector at the Soul of the New Machine Conference at UC Berkeley.

8.4 Awards and Honors

The following table describes awards and honors received by TRUST Center personnel during this reporting period.

	Recipient	Reason for Award	Award Name and Sponsor	Date	Award Type
	Ruzena Bajcsy	Benjamin Franklin Medal in Computer and Cognitive Science	"Best of the Best" in Science and Technology, The Franklin Institute	February 18, 2009	Election

8.5 Graduates

During this reporting period, the following undergraduate, graduate, and Ph.D. students from across all TRUST universities graduated. Students are listed alphabetically by last name along with their institution name and degree.

#	Student Name	Degree(s)
1	Barreno, Marco (Berkeley)	Ph.D.
2	Barth, Adam (Stanford)	Ph.D.
3	Bhatia, Nisha (San Jose State)	M.S.
4	Brumley, David (Carnegie Mellon)	Ph.D.
5	Cadar, Cristian (Stanford)	Ph.D.
6	Hartwig, Cody (Carnegie Mellon)	M.S.
7	Karlof, Chris (Berkeley)	Ph.D.
8	Khera, Ashira (San Jose State)	M.S.
9	Martin, Michael (Stanford)	Ph.D.
10	Meingast, Marci (Berkeley)	Ph.D.
11	Merideth, Mike (Carnegie Mellon)	Ph.D.
12	Misra, Saswat (Cornell)	Ph.D.
13	Molnar, David (Berkeley)	Ph.D.
14	Qu, Hui (Cornell)	Ph.D.
15	Schiff, Jeremy (Berkeley)	M.S. and Ph.D.
16	Segura, Eduardo (San Jose State)	M.S.
17	Wiley, Jonathan (Vanderbilt)	M.S.
18	Zeldovich, Nikolai (Stanford)	Ph.D.

8.6 General Knowledge Transfer Outputs

Details of knowledge transfer outputs are provided in Section 4.

8.7 Institutional Partners

The following table lists all TRUST Center research, education, knowledge transfer, and other institutional partners.

	Org. Name	Org. Type	Address	Contact Name	Type of Partner	160+ Hrs?
1	Academia Sinica	Other	Taipei, Taiwan	D.T. Lee	Research	Y
2	Air Force Office of Scientific Research	Federal Government	Arlington, VA	Bob Bonneau	Research	Y
3	Air Force Research Laboratory	Federal Government	Rome, NY	Rick Metzger	Research	Y
4	Cisco Systems	Company	San Jose, CA	Ken Watson	Research	N

					Knowledge Transfer	
5	Cyber Security Industry Alliance	Non-Profit	Arlington, VA	Liz Glasser	Education	Y
6	Deloitte & Touche LLP	Company	San Jose, CA	Dennis Kushner	Education	Y
7	DoCoMo USA Labs	Company	Palo Alto, CA	Svetlana Radosavac	Research	Y
8	EADS	Company	Paris, France	Cedric Blancher	Research	Y
9	eBay	Company	San Jose, CA	Dave Cullinane	Education	Y
10	General Electrical Capital	Company	McKinney, TX	James Beeson	Education	N
11	Greater Bay Bank	Company	Palo Alto, CA	Jason Hoffman	Education	Y
12	Hewlett-Packard	Company	Palo Alto, CA	Rich McGeer	Research Knowledge Transfer	N
13	ING	Company		Robert Weaver	Education	N
14	Intel	Company	Santa Clara, CA	Anand Rajan	Research Knowledge Transfer	N
15	Jefferson Wells	Company	Brookfield, WI	Jeffrey Camiel	Education	N
16	Microsoft Research	Company	Redmond, WA	Mike Schroeder	Research	N
17	Oracle	Company	Redwood Shores, CA	Mary Ann Davidson	Knowledge Transfer	N
18	Pirelli Research Laboratory	Company	Berkeley, CA	Marco Sgroi	Research Knowledge Transfer	N
19	Rapport, Inc.	Company	Redwood City, CA	Andrew Singer	Education	Y
20	Silicon Valley Bank	Company	Santa Clara, CA	Andrew Neilson	Education	Y
21	Sun Microsystems	Company	Menlo Park, CA	Emil Sarpa	Research Education	Y
22	Symantec	Company	Santa Monica, CA	Ken Baylor	Research Knowledge Transfer	N
23	Tata Consultancy Services	Company	Chennai, India	Sanjay Bahl	Education	N
24	United Technologies	Company	East Hartford, CT	Clas Jacobson	Research Knowledge Transfer	N
25	Visa International	Company	San Francisco, CA	George Sullivan	Research Knowledge Transfer	N
26	Yahoo Inc.	Company	Sunnyvale, CA	Mark Seiden	Education	Y
27	Xilinx, Inc.	Company	San Jose, CA	Abe Smith	Research Knowledge Transfer	Y

9 INDIRECT/OTHER IMPACTS

9.1 International Activities

As part of TRUST's goals of disseminating results, we are eager to establish relationships with international programs where mutually beneficial opportunities exist. Our first large effort in this area is with Taiwan. The TRUST Center has received significant attention from Taiwan, and funds for cooperating with TRUST have been approved the National Legislature (the Legislative Yuan) and a member of the Taiwanese Cabinet at the level of Minister of State has been assigned to oversee the program: The International Collaboration for Advancing Security Technology (iCAST).

Taiwan is a leading player in the world of electronics and IT. Taiwan has been expanding its scope from more narrowly focused areas in manufacturing and integrated circuit design to become an aggressive player in the world of IT services. Taiwan by most accounts has the second or third largest penetration of broadband services (as of July 2005, with 10.5 million broadband users and 14.6 Internet users out of a total population of 22.8 million.) Taiwan also faces unique challenges because of its relationship with mainland China, and both public and private institutions in Taiwan are under constant attack from mainland Chinese sources. Some of these are believed to be government sponsored.

Based on TRUST, Taiwan has set up an inter-university institute called the Taiwan Information Security Center (TWISC) and has adopted an international collaboration center for research in computer security, directed by Dr. D. T. Lee, a former NSF program officer. TWISC is overseen by the cabinet level Science and Technology Advisory Group (run by a Minister of State). Major members include the National Science Council (NSC, the "Taiwanese NSF"); the Institute for Information Industry (III, a public/private software industry coordinating group); the Industrial Technology Research Institute (ITRI); major infrastructure groups (e.g., telecommunication companies); and government representatives from public safety and law enforcement.

Funding has been provided to TRUST and partner institutions Carnegie Mellon University and the University of California, Berkeley at approximately US$2M per year. The Center is very excited about this collaboration because of the outstanding quality of our Taiwanese research counterparts, their impact in the IT area, and the chance to observe and address the emerging patterns of cyber attack within Asia (and particularly emerging from mainland China) firsthand.

Please see Section 5.4 for additional information on iCAST and TRUST.

9.2 Other Outputs, Impacts, and Influences

None to report.

10 ATTACHMENTS

Appendix A: Biographical Information of New Faculty

Cornell University:

Edward Suh – Suh is an Assistant Professor of Electrical and Computer Engineering at Cornell University, where I am a member of the Computer Systems Laboratory. He received a Ph.D. degree in Electrical Engineering and Computer Science from the Massachusetts Institute of Technology (MIT) with work on a single-chip secure processor. Following the graduate school, he spent a year at PUFCO Inc., where he led the development of unclonable RFIDs and secure embedded processors. He joined the faculty of the School of Electrical and Computer Engineering at Cornell in 2007. Suh's research interests include computer systems in general with particular focus on computer architecture. He is interested in combining architectural techniques with low-level software to enhance various aspects of computing systems such as performance, security, and reliability. Currently, he is mainly working on secure embedded processors, reconfigurable multi-core architecture for security and new programmable architecture (such as FPGAs) for simplified synthesis.

Hakim Weatherspoon – Weatherspoon received his Ph.D. in 2006 from the University of California, Berkeley in the area of secure and fault-tolerant distributed wide-area storage systems (e.g. Antiquity, OceanStore, etc.). He received a B.S. in Computer Engineering from the University of Washington in 1999. Weatherspoon's research interests cover various aspects of information systems, distributed systems, network systems, and peer-to-peer systems. In these areas, he is particularly focus on fault-tolerance, reliability, security, and performance of Internet-scale systems with decentralized—autonomous, federated, multi-organizational, and cooperative—control.

San Jose State University:

Simon Shim –Shim is a Professor in Computer Engineering department. He received and M.S. from Rensselaer Polytechnic Institute, and Ph.D. from University of Minnesota both in Computer Science. His expertise includes web technologies, media and network security, database, data mining, and SAN. He is a co-director of the Internet Technology Laboratory which is supported by grants from Intel, Microsoft, Wytec, and Informix Corporation. He has authored and co-authored more than 35 technical publications in IEEE Computer, IEEE Transaction on Knowledge and Data Engineering, IEEE Concurrency, ACM MONET, and Journal of Multimedia Tools and Applications (Kluwer Academic Publishers). His work was recognized by IBM Rochester as a significant Research Contribution in the study of SAN with four colleagues. He has served as the workshop chair of IEEE E-Commerce conference, co-chair of IEEE workshop on Mobile Commerce and Services, and co-chair of IEEE workshop on Data Engineering Issues in E-commerce. He served as the lead guest editor of the special issue on high-speed security for IEEE Computer, June 2004. He has served in many international conferences/workshops as a technical committee member.

Stanford University:

Nick Bambos –Bambos is a Professor at Stanford University, having a joint appointment in the Department of Electrical Engineering and the Department of Management Science & Engineering. He heads the Network Architecture and Performance Engineering research group at Stanford, conducting research in wireless network architectures, the Internet infrastructure, packet switching, network management and information service engineering, engaged in various projects of his Network Architecture Laboratory (NetLab). He has graduated over 20 Ph.D. students, who are now at leadership

positions in academia (Stanford, CalTech, Michigan, GaTech, NYU, UBC, etc.) and the information technology industry (Cisco, Broadcom, IBM Labs, Qualcomm, Nokia, MITRE, Sun Labs, ST Micro, Intel, Samsung, TI, etc.) or have become successful entrepreneurs. From 1999 to 2005 he served as the director of the Stanford Networking Research Center, a major partnership/consortium between Stanford and information technology industries, involving tens of corporate members, faculty and doctoral students. He is now heading a new research initiative at Stanford on Networked Information Service Engineering.

He received his Ph.D. in Electrical Engineering and Computer Sciences (EECS) from the University of California at Berkeley (1989), as well as the M.S. in EECS (1987) and the M.A. in Mathematics (1989) from the same University. He graduated in Electrical Engineering from the National Technical University of Athens-Greece (1984) with first class honors. Before joining Stanford as an Associate Professor in 1996, he served as Assistant (1990-95) and tenured Associate Professor (1995-96) in the Electrical Engineering Department of the University of California at Los Angeles (UCLA).

Bambos has held the Cisco Systems Faculty Development Chair (1999-2003) in computer networking at Stanford and has won the IBM Faculty Award (2002) for high-impact research in performance engineering of computer systems and networks, as well as the Griffin Award (1997). He has been the David Morgenthaler Faculty Scholar (1996-99) at Stanford, and has received the National Young Investigator Award (1992) from the National Science Foundation (NSF) for research in computer networks and distributed computing architectures, as well as the NSF Research Initiation Award (1990) for studies in performance modeling of computer systems. He has also been a U.C. Regents Fellow, a David Gale Fellow, and an Earl Anthony Fellow. He is on the Editorial Boards of several research journals and serves on various international technical committees and review panels for networking research and information technologies. He has been serving on the boards of various start-up companies in the Silicon Valley, consults on high technology development and management matters, and has served as lead expert witness in high-profile patent litigation cases in networking and computing.

David Dill –Dill is a Professor of Computer Science and, by courtesy, Electrical Engineering at Stanford University. He has been on the faculty at Stanford since 1987. He has an S.B. in Electrical Engineering and Computer Science from Massachusetts Institute of Technology (1979) and an M.S and Ph.D. from Carnegie Mellon University (1982 and 1987). Dill has research interests in a variety of areas, including computational systems biology and the theory and application of formal verification techniques to system designs, including hardware, protocols, and software. He has also done research in asynchronous circuit verification and synthesis, and in verification methods for hard real-time systems. He was the Chair of the Computer-Aided Verification Conference held at Stanford University in 1994. From July 1995 to September 1996, he was Chief Scientist at 0-In Design Automation.

Dill's Ph.D. thesis, “Trace Theory for Automatic Hierarchical Verification of Speed Independent Circuits” was named as a Distinguished Dissertation by the Association for Computing Machinery (ACM), and published as such by M.I.T. Press in 1988. He was the recipient of a Presidential Young Investigator award from the National Science Foundation in 1988, and a Young Investigator award from the Office of Naval Research in 1991. He has received Best Paper awards at International Conference on Computer Design in 1991 and the Design Automation Conference in 1993 and 1998. He was named a Fellow of the Institute of Electrical and Electronics Engineers (IEEE) in 2001 for his contributions to verification of circuits and systems, and a Fellow of the ACM in 2005 for contributions to system verification and for leadership in the development of verifiable voting systems. In 2008, he received the first “Computer-Aided Verification” award with Rajeev Alur, for fundamental contributions to the theory of real-time systems verification.

Dill has been working actively on policy issues in voting technology since 2003. He is the author of the "Resolution on Electronic Voting", which calls for a voter-verifiable audit trail on all voting equipment, and which has been endorsed by thousands of people, including many of the top computer scientists in the U.S. He has served on the California Secretary of State's Ad Hoc Task Force on Touch-Screen voting, the Citizens DRE Oversight Board of the Santa Clara County Registrar of Voters, and on the IEEE P1583 Voting Equipment Standards Committee. He has testified on electronic voting before the U.S. Senate and the Commission on Federal Election Reform, co-chaired by Jimmy Carter and James Baker III. He is the founder of the Verified Voting Foundation and VerifiedVoting.org and is on the board of those organizations. In 2004, he received the Electronic Frontier Foundation's "Pioneer Award" for "spearheading and nurturing the popular movement for integrity and transparency in modern elections."

Phil Levis –Levis is an assistant professor in the Computer Science and Electrical Engineering departments of Stanford University. He received his Sc.B. from Brown University, his M.S. from the University of Colorado at Boulder, and his Ph.D. from the University of California, Berkeley. He researches wireless sensor networks, particularly software systems and networking. The results of his research have been adopted by thousands of users and researchers worldwide. He is the chair of the TinyOS Core Working Group and is a member of both the TinyOS Network Working Group as well as the TinyOS Alliance Working Group. His prior work includes the nesC language for network embedded systems, software design patterns for static embedded programming, using application-specific virtual machines to enable safe, high-level programming of sensor nodes, and the Trickle network algorithm for rapid yet efficient data dissemination.

UC Berkeley:

John Chuang –Chuang is Associate Professor of Information Management and Systems at the University of California, Berkeley. He received a B.S. and M.S. in Electrical Engineering from the University at Southern California and Stanford University respectively, and a Ph.D. in Engineering and Public Policy from Carnegie Mellon University. His research focus is on economics-informed design of computer networks and distributed systems, including incentive mechanisms for peer-to-peer networks and next-generation internet architecture design.

Appendix B: Center Organizational Chart

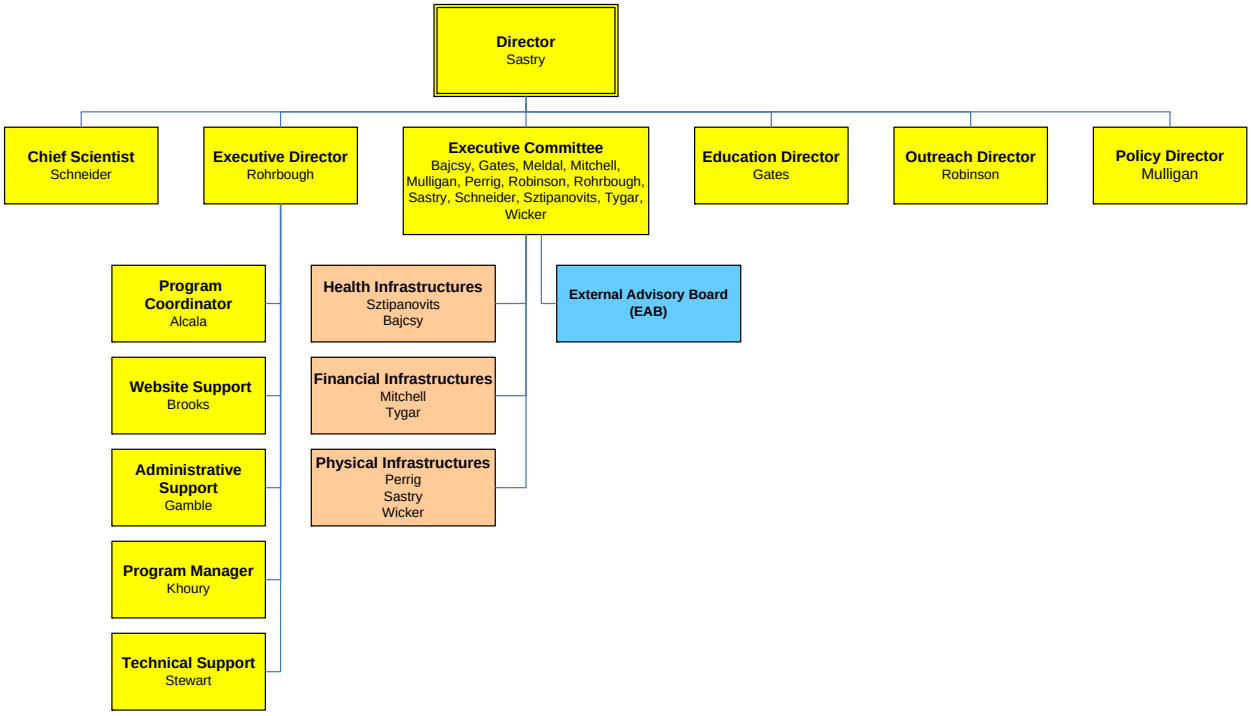

Appendix C: Minutes of External Advisory Committee Meetings

Minutes are in the form of Power Point slides. The slides below are from the October 11-12, 2007 External Advisor Board meeting in Ithaca, NY. There are a total of nine (9) slides for that meeting.


TRUST
External Advisory Board (EAB)
Out-Brief
Prepared for
Vice Chancellor Burnside &
TRUST Exec Comm
Ithaca NY
12 Oct 2007
Jay Lala & Carl Landwehr
EAB Co-Chairs

Distinguished External Advisory Board Members

	First Name	Last Name	Affiliation
1	Alfred	Aho	Columbia University
2	Annie	Anton	North Carolina State University
3	Patricia	Bellia	Notre Dame University
4	Matthew	Bishop	University of California, Davis
5	Lee	Burge	Tuskegee University
6	David	Clark	MIT
7	George	Cybenko	Dartmouth College
8	James	Johnson	Howard University
9	Jay	Lala	Raytheon
10	Carl	Landwehr	University of Maryland
11	Teresa	Lunt	Palo Alto Research Center
12	Dan	Manson	California State Polytechnic University
13	Andrew	Odlyzko	University of Minnesota
14	William	Sanders	UIUC
15	Joseph	Sifakis	CNRS, Verimag
16	Gene	Spafford	Purdue University

12 Oct 2007

2


EAB Members Present
Annie Anton
Al Aho
George Cybenko
David Clark
James Johnson
Jay Lala
Carl Landwehr
Gene Spafford
12 Oct 2007
3

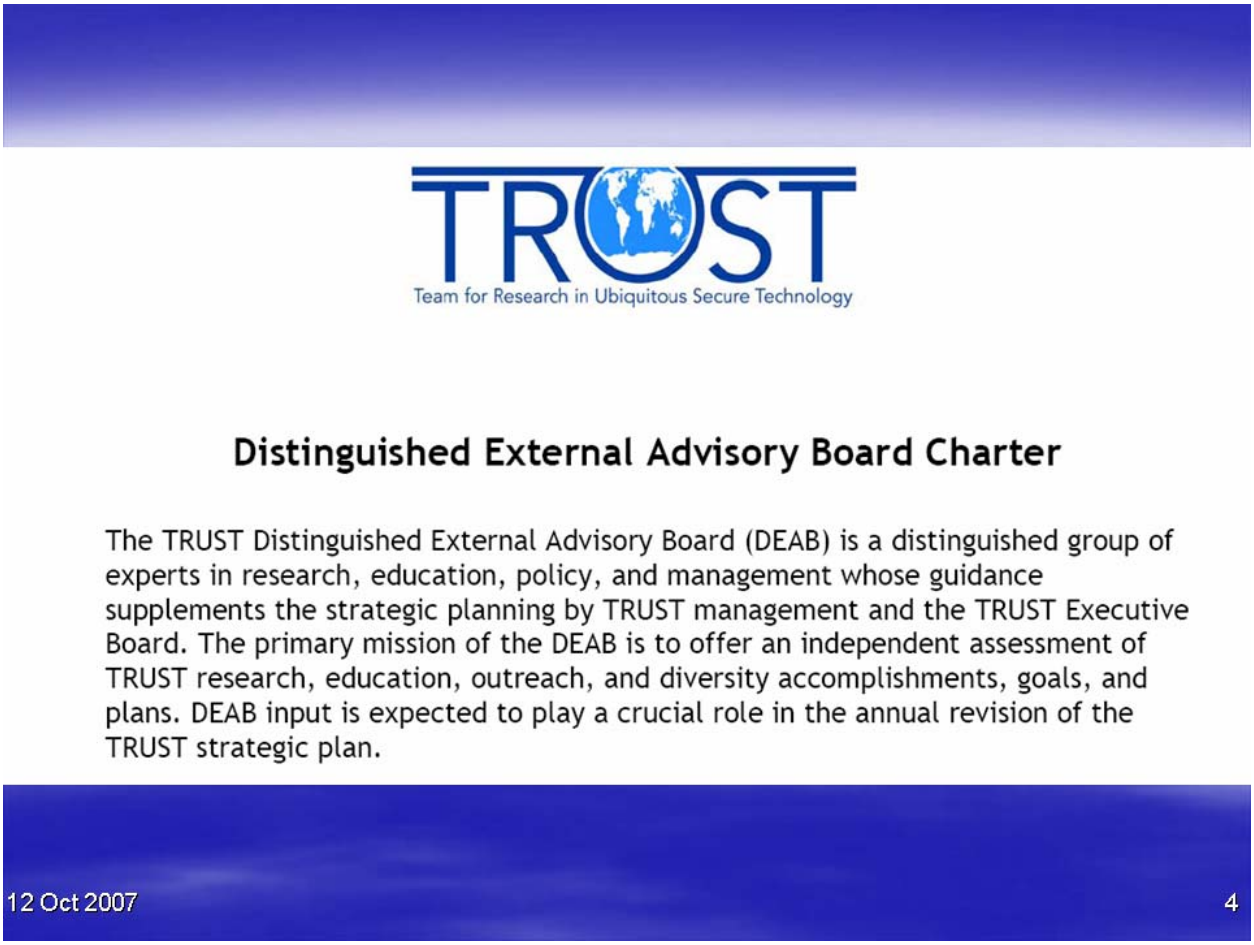
TRUST
Team for Research in Ubiquitous Secure Technology
Distinguished External Advisory Board Charter
The TRUST Distinguished External Advisory Board (DEAB) is a distinguished group of experts in research, education, policy, and management whose guidance supplements the strategic planning by TRUST management and the TRUST Executive Board. The primary mission of the DEAB is to offer an independent assessment of TRUST research, education, outreach, and diversity accomplishments, goals, and plans. DEAB input is expected to play a crucial role in the annual revision of the TRUST strategic plan.
12 Oct 2007
4

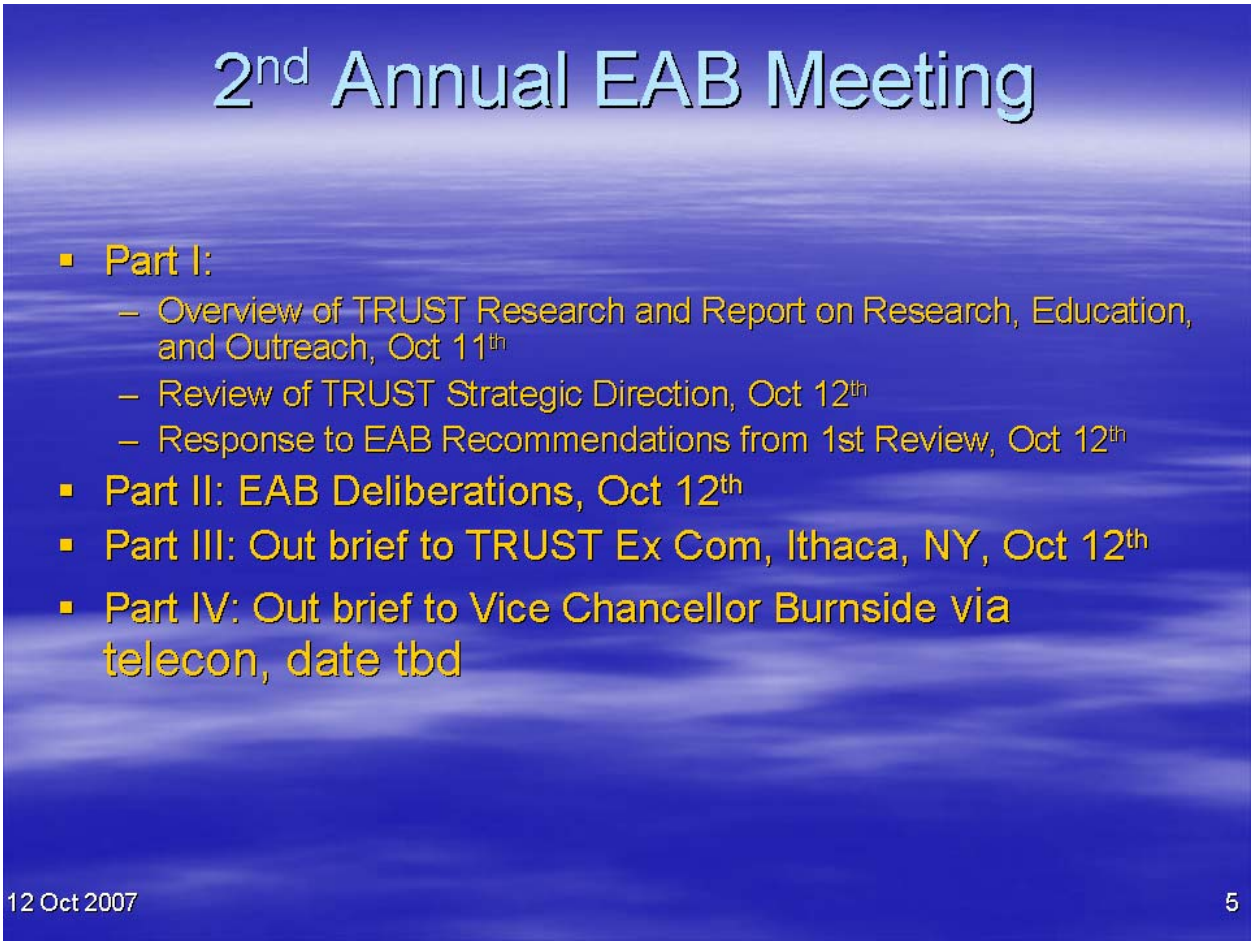
2nd Annual EAB Meeting
Part I:
Overview of TRUST Research and Report on Research, Education, and Outreach, Oct 11th
Review of TRUST Strategic Direction, Oct 12th
Response to EAB Recommendations from 1st Review, Oct 12th
Part II: EAB Deliberations, Oct 12th
Part III: Out brief to TRUST Ex Com, Ithaca, NY, Oct 12th
Part IV: Out brief to Vice Chancellor Burnside via telecon, date tbd
12 Oct 2007
5


Out-Brief Topics
TRUST Strengths
EAB Recommendations
12 Oct 2007
6


TRUST Strengths
Overall good progress
World class team of researchers
Passionate & committed leadership
Continued good evidence of cross-university collaboration
Heightened combination of technology and policy
Overachievement in education
Some evidence of maturation of overall strategy
12 Oct 2007
7

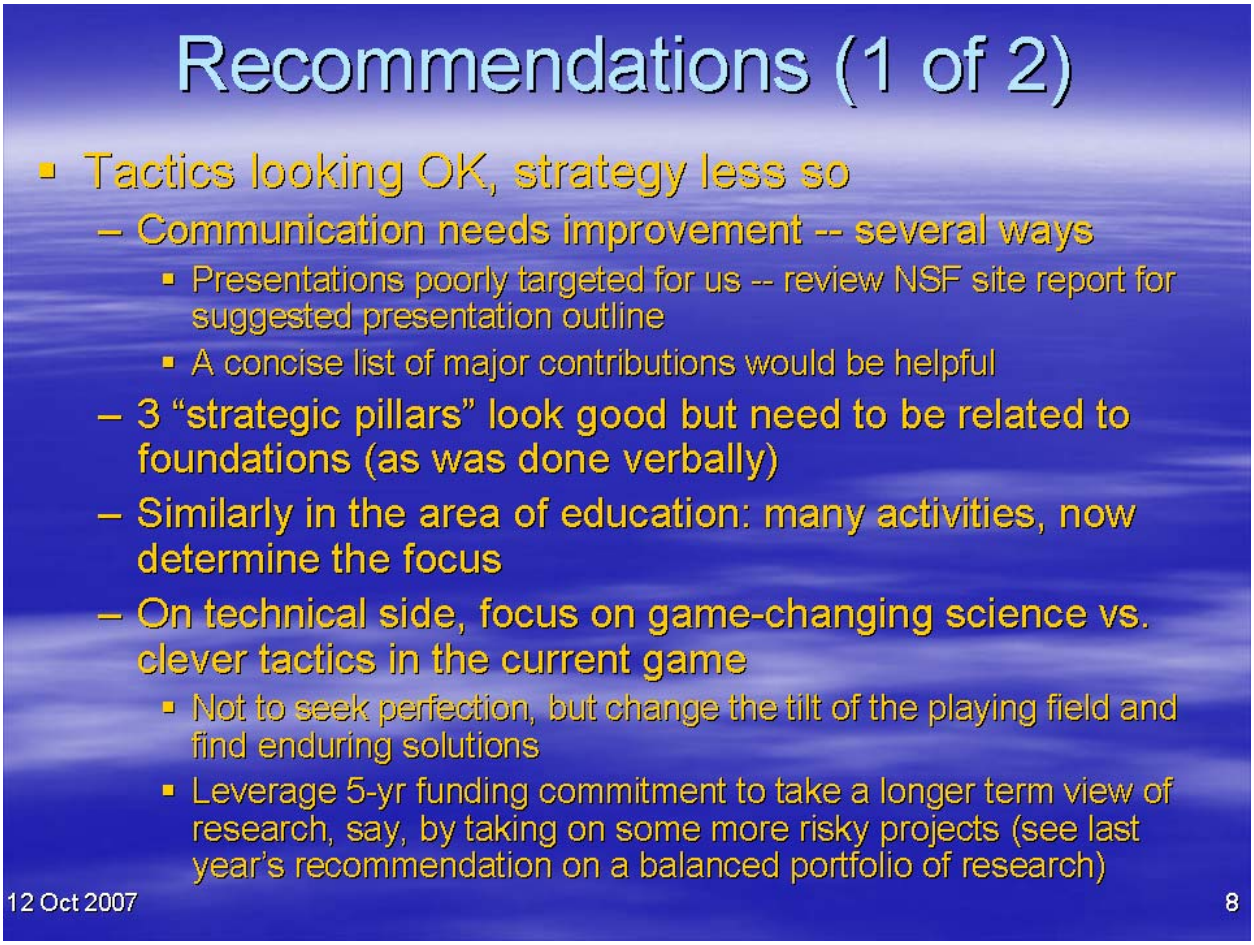
Recommendations (1 of 2)
Tactics looking OK, strategy less so
Communication needs improvement -- several ways
Presentations poorly targeted for us -- review NSF site report for suggested presentation outline
A concise list of major contributions would be helpful
3 "strategic pillars" look good but need to be related to foundations (as was done verbally)
Similarly in the area of education: many activities, now determine the focus
On technical side, focus on game-changing science vs. clever tactics in the current game
Not to seek perfection, but change the tilt of the playing field and find enduring solutions
Leverage 5-yr funding commitment to take a longer term view of research, say, by taking on some more risky projects (see last year's recommendation on a balanced portfolio of research)
12 Oct 2007
8

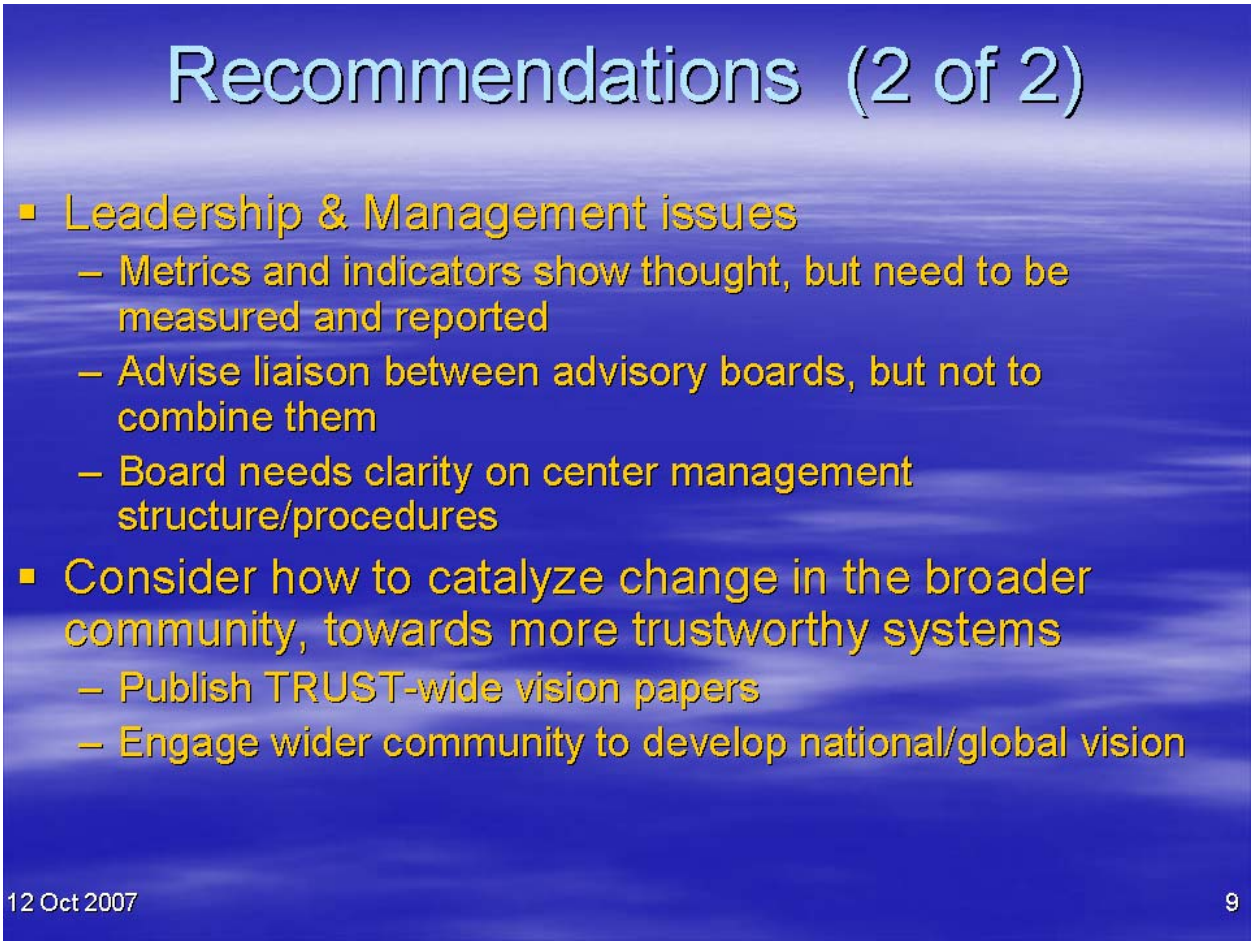
Recommendations (2 of 2)
Leadership & Management issues
– Metrics and indicators show thought, but need to be measured and reported
– Advise liaison between advisory boards, but not to combine them
– Board needs clarity on center management structure/procedures
Consider how to catalyze change in the broader community, towards more trustworthy systems
– Publish TRUST-wide vision papers
– Engage wider community to develop national/global vision
12 Oct 2007
9

Appendix D: Media Publicity Materials (if any)

Flyers for three TRUST Education and Outreach programs conducted in the summer 2008 are included on the following pages. Program flyers are for SECuR-IT, SUPERB-IT, and WISE.

SECuR-IT: Summer Experience, Colloquium and Research in Information Technology at Stanford University and San Jose State University

SUMMER 2008: Graduate Student Academic Emersion with Internship Program June 2 through August 8, 2008

The Team for Ubiquitous Secure Technology (TRUST) is proud to announce the Summer Experience, Colloquium and Research in Information Technology (SECuR-IT). This is a ten-week residential program with paid internship co-located at Stanford University and San Jose State University.

Program Overview

- Paid internship at a Silicon Valley technology company
- Learning cohort of 30 graduate students
- Seminars and presentations in security related topics at Stanford University
- Residential summer housing at San Jose State University
- College units for summer educational program

Graduate student internship opportunities available in: Security Architecture • Security Awareness and Security Management • Host and OS Security • Application Security • Network Security • Secure Software Engineering • Risk Management • Policy and Legal Compliance

PROPOSED Participating Technology Companies: Cisco Systems • eBay • Google • Intel • Yahoo • Sun Microsystems • Symantec Corporation • SalesForce • Raport Inc.

Program Structure

In addition to working with an industry mentor over the ten-week program, scholars participate in the following programmatic components:

- Seminars conducted by faculty and industry experts that expose students to a wide range of information technology and computer security research instruction;
- Faculty participation from: Stanford University, University of California, Berkeley and San Jose Sate University
- Informal social gatherings that provide a relaxed setting for students and faculty to exchange ideas and share experiences;
- Residential housing at San Jose State University;
- Ten week, paid 40-hour per week internship.

San Jose State University On-Campus Housing

Housing will be available at San Jose State University (SJSU) San Jose, California. Cost of housing and meals at SJSU will be the responsibility of program participants.

Application Process

SECuR-IT participation is open to graduate students (M.S. & Ph.D). Participation is limited to 30 people and will be selected from a nationwide pool of applicants, who have demonstrated, outstanding academic talent.

- **On-line application** only at **http://www.truststc.org/securit/apply**
  Application deadline: **January 31, 2008 at 5pm PST**
- Women and historically underrepresented ethnic minority groups will be given strong consideration although everyone is encouraged to apply.

Contact Information

Dr. Kristen Gates, Executive Director of Education
Team for Research in Ubiquitous Secure Technology (TRUST)
295 Hearst Memorial Mining Building
University of California, Berkeley CA 94720
(510) 642-3737 :: email: kgates@eecs.berkeley.edu :: URL: www.truststc.org

---

Team for Research in Ubiquitous Secure Technology (TRUST) :: URL: www.truststc.org
National Science Foundation Cooperative Agreement No. 0424422
University of California, Berkeley :: Carnegie Mellon University :: Cornell University :: Mills College
San Jose State University :: Smith College :: Stanford University :: Vanderbilt University

UNIVERSITY OF CALIFORNIA, BERKELEY

SUPERB-Information Technology 2008

Summer **U**ndergraduate **P**rogram in **E**ngineering **R**esearch at **B**erkeley

Team for Research in Ubiquitous Secure Technologies (TRUST)

Computing technologies are part of our nation's critical infrastructure. They form a part of everything from financial systems and the energy grid to telecommunication and transportation systems. Enhancing cybersecurity and computer trustworthiness is therefore of increasing importance as a scientific, economic, and social problem.

Program Description

- 8-week research experience: **June 9 – August 1, 2008**
- Research guided by faculty mentors and graduate students
- Educational activities include lab tours and industry field trips
- Graduate school advising and subsidized GRE prep course
- **$4,000 Stipend**
- **Travel Allowance up to $600**
- **Room and board provided at International House**

Exciting cutting edge TRUST research opportunities in:

- Physical Infrastructures: Control, Security, and Privacy
- Personal Health Records
- Phishing and ID Theft
- Financial Infrastructures
- Senor Networks
- and much more!

TRUST
Team for Research in Ubiquitous Secure Technology

Application Process

On-line application only (available, December 1, 2007):
http://www.truststc.org/superb/apply

- Application deadline: **Thursday, January 31, 2008, 5PM (Pacific Time)**
- Underrepresented students are encouraged to apply
- Must be US Citizen or Permanent Resident
- A minimum overall GPA of 3.0 is required with upward trends in grades being preferable

Contact Information:

Dr. Kristen Gates, Executive Director of Education
Team for Research in Ubiquitous Secure Technology (TRUST)
392 Cory Hall :: University of California, Berkeley, CA :: 94720
(510) 642-3737 :: email: kgates@eecs.berkeley.edu :: URL: www.truststc.org

National Science Foundation Cooperative Agreement No. 0424422

Berkeley Carnegie Mellon Cornell University MILLS
San José State University Smith College STANFORD University Vanderbilt University

WISE 2008: Women's Institute in Summer Enrichment

Sponsored by the Team for Research in Ubiquitous Secure Technology (TRUST)

June 8th through 13th, 2008: Ithaca, New York

Program Description

WISE is a 1-week residential summer program on the Cornell University campus that brings together graduate students, post-doctoral fellows, and professors from all disciplines that are interested in Ubiquitous Secure Technology and the social, political, and economical ramifications that are associated with this technology.

We are inviting scholars who are willing to share their knowledge, experience and skills with women faculty and graduate students in the various computer science, electrical engineering, and civil engineering disciplines associated with sensing systems for critical infrastructure, with an emphasis on the security and privacy issues that arise from the use of sensing systems in public places.

Topics may include but are not limited to:
* Large-Scale Sensing Systems * Secure Sensor Networking * Sensor Information Processing
* Engineering, Maintenance, and Security of Critical Infrastructure * Public Surveillance, Privacy, and the 4th Amendment * Rights and responsibilities of data, data owners and data users

Seminar Speakers

- Annie Anton: Computer Science, NC State
- Ruzena Bajcsy: Electrical Engineering and Computer Science, UC Berkeley– TRUST
- Judy Cardell: Computer Science, Smith College
- Deirdre Mulligan/Maryanne McCormick, Berkeley Law, UC Berkeley – TRUST
- Wendi Heinzelman: Electrical and Computer Engineering, University of Rochester
- Sheila Hemami: Electrical and Computer Engineering, Cornell University
- Susan Landau – Sun Microsystems Laboratories
- Christine Shoemaker: Environmental and Civil Engineering, Cornell University

WISE 2008 at Cornell

The seminar will be held on the campus of the Cornell University. The seminar will last one week and begin on June 9, 2008 with a welcome reception and includes lodging and meals.

WISE Tuition

Tuition for WISE 2008 is $2,500; however, NSF-TRUST fellowships are available to US professors, post-doctoral fellows, and Ph.D. candidates studying at US universities. There is a maximum of 20 fellowships for Ph.D. candidates, post-doctoral fellows, and professors of all levels for the Institute.

Application Process

WISE participation is open to US professors and post-doctoral fellows, and Ph.D. candidates studying at US universities. Participation is limited to 30 people and will be selected from a nationwide pool of applicants, who have demonstrated, outstanding academic talent.

- **On-line application** only (available December 1, 2007) at **http://www.truststc.org/wise/apply**
- Application deadline: **March 31, 2008 at 5 PM**
- Women will be given strong consideration although everyone is encouraged to apply

Contact Information

Dr. Kristen Gates, Executive Director of Education
Team for Research in Ubiquitous Secure Technology (TRUST)
295 Hearst Memorial Mining Building
University of California, Berkeley CA 94720
(510) 642-3737 :: email: kgates@eecs.berkeley.edu :: URL: www.truststc.org

Team for Research in Ubiquitous Secure Technology (TRUST) :: URL: www.truststc.org
National Science Foundation Cooperative Agreement No. 0424422
University of California, Berkeley :: Carnegie Mellon University :: Cornell University :: Mills College
San Jose State University :: Smith College :: Stanford University :: Vanderbilt University