TRUST
Team for Research in Ubiquitous Secure Technology

Annual Report (2007 – 2008)

June 17, 2008

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

TABLE OF CONTENTS

1 GENERAL INFORMATION

1.1 Summary

Date Submitted	March 28, 2008
Reporting Period	June 1, 2007 – May 31, 2008
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, if changed since last reporting period	
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 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 partner in the area of sensor networks 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 cybersecurity as it affects society. First, 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 are now used for controlling critical infrastructures for electricity, healthcare, finance, and medical networks. Second, and somewhat contradictorily, many such 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. Privacy remains poorly understood and poorly supported; 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. TRUST represents the strongest and most diverse engagement of the issue of trusted systems ever assembled. TRUST is the first to recognize the breadth of the problem and to combine fundamental science with a broader multidisciplinary focus on economic, social and legal considerations and a substantial educational mission. TRUST will enable dialog with stakeholders whose needs simply cannot be approached in a narrower and purely technical manner, or by any single research group. TRUST seeks to be an intermediary between the policy makers and society at large on the one hand, and the 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, industrial outreach and knowledge transfer, education, and diversity outreach.

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 third year was increasing content in the TRUST Academy Online (TAO) and redesigning the TAO portal as well as the center's Education Community Development efforts. Again, all these are reported below in the section on education.

In knowledge transfer, TRUST has continued a robust program of technology transition with industry (from bug reports of open source software to tools such as Spoofguard and 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 such as Intel, Microsoft, Sun, Telecom Italia, and United Technologies with whom we are beginning to engage in collaborations of mutual interest. More details are provided in the section on knowledge transfer.

In diversity, TRUST has an ambitious goal of reaching a diversity goal of 30% of women in its faculty and students, and 10% of researchers from underrepresented communities, and has been proactive in this regard. Several activities for enhancing diversity are reported in the corresponding section.

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 projects are organized into several research 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.

For this reporting period, there were six research thrust areas, listed below:

1. Electronic Medical Records
2. End User Security
3. Network Defenses
4. Policy
5. Secure Sensor Networks
6. Trustworthy Systems

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 Electronic Medical Records

Project Leader: *Janos Sztipanovits (Vanderbilt University)*

Computer technology, patient sensors, and networking are revolutionizing several aspects of healthcare and medical information processing. Small wireless sensors will free many patients from managed care facilities, while providing timely medical assistance when needed. At the other end of the spectrum, virtually all patients will soon gain greater control over their records and treatment options through web portals. The TRUST Electronic Medical Records (EMR) research thrust addresses the complex security and privacy issues emerging from the rapidly increasing use of electronic media for the archival and access of patient records. This change is driven and strongly influenced by the Health Insurance Portability and Accountability Act (HIPAA) of 1996. EMR has become an area where technology, public policy and individual interests intersect and conflict, making the development of information systems for EMR archiving and access a very challenging problem. There is clear evidence that without a detailed understanding of the relevant issues on all sides, an acceptable solution cannot and will not emerge.

The projects leverage a cooperative relationship established with the Informatics Institute and the Biomedical Informatics Department of the Vanderbilt University Medical Center. The MyHealth at Vanderbilt System – a functioning web-based patient portal – is a unique resource that serves as the basis for experimentation and interaction through real-life deployment scenarios. The EMR projects also utilize the on-going Information Technology for Assisted Living at Home project at the University of California, Berkeley to develop the tools necessary to produce high confidence and secure embedded software systems necessary to investigate the nature of automated and semi-automated sensor data inclusion into EMRs. This project develops smart sensing technologies that enable alert monitoring and gathering long-term out-patient biometric data. Decisions on how and when to include this data in an EMR and how to apply the methods to existing and new in-patient sensor systems are of primary importance. EMR researchers collaborate extensively with TRUST Secure Sensor Networks researchers, with a particular emphasis on the latter's Real-Time Patient Monitoring Project.

We have three areas that represent challenges for the TRUST research agenda and have direct relevance and applicability in EMR. They are:

- Architecture Modeling for Patient Portals
- Privacy Modeling and Analysis
- Integration of Real Time Sensor Data with Patient Portals.

The accomplishments of each project are discussed in more detail below.

Architecture Modeling for Patient Portals – The Vanderbilt team continued the development of a formalized design approach to Clinical Information Systems (CIS). The approach is based on a combination of standards-based design methodologies successfully applied individually in other domains previously:

- Service-Oriented Architecture (SOA) provides the framework to assemble loosely coupled software services that can be deployed on different platforms in complex distributed applications.
- Platform-Based Design (PBD) focuses on the creation of abstraction layers in the design flow and investigates the semantic properties of mapping across these layers.
- Model Integrated Computing (MIC) is predicated on the notion that models play an essential role in every part of a system's life-cycle, which commences with specification,

progresses through design, development, verification, integration, and concludes with upgrade and maintenance.

The combination of SOA, PBD, and MIC techniques enables the design of complex CIS to ensure reliability, performance, privacy and security beyond what can be achieved by current ad-hoc practices. During this reporting period, research focused on refining the modeling language, creating a model translator for BPEL4WS, designing a policy representation language and a corresponding model translator for XACML, and building a test deployment platform for policy-based workflow execution.

A new addition to the EMR area is Prof. Brad Malin from Vanderbilt. Prof. Malin's work has been in ways to extract patterns and partial workflows from access logs of existing EMR system. In various investigations, it has been shown that the access logs of electronic medical records can provide feedback to understand how healthcare systems are used and make them more efficient. However, the study of the information within the access logs for health information privacy and security has been neglected. The automated analysis of healthcare access logs has focused on "what" providers and patients view, but to build security policies it is also necessary to know "who" is viewing the health records of whom, and most importantly but often problematically, "why" they are viewing those health records. The size and detail of electronic medical record systems access logs provides a prime opportunity to leverage automated methods to learn from, as well as monitor, access logs. Prof. Malin's approach was to first process access logs into care provider user sessions based on login/logout, or timeout, to the system. Based on user sessions, he then used static and temporal data mining techniques to extract patterns and regularities regarding teams and workflows. The workflows that exist in a healthcare environment are diverse, complex, and difficult to specify. As opposed to having to model every workflow in the system, the automatic extraction technique will speed up the process and provide models that are potentially much closer to reality.

Privacy Modeling and Analysis – The Stanford team developed a general framework around actions that may use information in ways that are significant for privacy, and/or transmit information between a sender associated with one role in the organization and a receiver associated with a possibly different role in the organization. Each action, including messages (which could represent an actual message, a web form, the state of some data structure associated with a workflow process, etc.) carries some information and has a set of associated tags that identify the kind of information used. Initially using temporal logic as a starting point to express properties of sequences of actions, they identified certain kinds of formulas that define workflow, utility goals, and privacy policies. A second and overlapping thrust was to further study HIPAA and hospital privacy policies, expressing HIPAA in a formal logic of privacy, and conduct further basic research on the logic of privacy. A new direction was to collaborate with the Stanford Center for Computers and Law (CodeX) at Stanford, working with lawyers and law students on codifying privacy laws. In addition, they explored additional ways to work with CodeX to collaborate on legal and policy topics.

A new participant in this area, Prof. Johannes Gehrke from Cornell, explored the possibility of anonymizing the text data in medical records so that the result is useful to researchers but still preserves the privacy of the individuals associated with the medical records. Text data is rich in information and has the complex structure of natural languages but, at the same time, search engines have demonstrated that natural language has much redundancy and that the meaning of text can often be captured by a small set of keywords and phrases. In fact, the keywords that users type into the search box of a search engine create many of the same privacy issues that are present in the text data of medical records. In terms of privacy leakage, search log data is very similar to patient notes written by doctors, but due to the short nature of many search queries search log data is also simpler than text data in medical records. For this reason we proposed, as a first step, to develop techniques for anonymizing search logs and reasoning about their privacy leakage to highlight and help address many of the issues involved in the more complicated task of sanitizing the text in medical records which is the long-term goal.

Integration of Real Time Sensor Data with Patient Portals – The patient monitoring project is a collaborative effort between Berkeley, Cornell, and Vanderbilt. Built on top of the results in the past two years, a health monitoring system prototype has been built that integrates the previous research results into a coherent system design. The goal of this work was to design, implement, and validate a set of algorithms and protocols that support secure, reliable, and quality-of-service-assured sensor-based health monitoring service. In order for a continuous monitoring system to be most effective, it must have the capacity to dynamically notify multiple parties based on their availability. It must also consider the role of those parties, and do so based on the context of the event that occurs. What is necessary in this scenario is a distributed policy system. There were several challenges to address in development of a distributed policy scheme for a notification network and two types of policy descriptions from a user interface standpoint were supported: (1) An abstract level of information sharing which results in general characteristics of the resulting policy, suitable for a patient interface allowing them a level of control without the need for excessive decision making and repeated input and (2) A more precise description of policy resulting from detailed data about the role of various care givers, especially supporting frequent updates.

The synergistic research projects described above provided a comprehensive approach to the EMR thrust. Knowledge transfer activities were also broad. Of particular importance was the Model-Based Design of Trustworthy Health Information Systems (MOTHIS) workshop organized by several TRUST investigators (Sztipanovits from Vanderbilt, Mitchell from Stanford, and Bajcsy from Berkeley) and Ruth Breau from the University of Innsbruck. The workshop was held October 5, 2007 in conjunction with the prestigious Models 2007 conference in Nashville, TN. This presented a tremendous opportunity for TRUST researchers to present their work and interact with experts outside of the TRUST center.

The Vanderbilt team continued the close cooperation with the Vanderbilt University Medical Center including the MyHealth portal development team. The information exchange has been two way: the Vanderbilt team continued to learn about the EMR domain from the VUMC and the VUMC was the recipient of the technology transfer of development tools as well as information on vulnerabilities uncovered in the current operational version of the VUMC patient portal. Stanford continued its collaboration with TRUST industrial partner Tata Consultancy Services (TCS) on a TCS privacy product and methods for ensuring privacy in commercially deployed software for U.S. and European banks, hospitals, and others customers. The home patient monitoring team led by Berkeley began to work with several senior medical personnel, in particular Dr. Michael Aminoff of the University of California, San Francisco. They also worked with Vanderbilt Homecare Services, Inc and its managed assisted living/independent living facilities in Nashville, TN. These connections helped them better understand in-home patient care scenarios, get the first-hand experience in terms of appropriate target groups who will benefit the most from our patient monitoring system, identify the medical data that are critical for patient health, and prototype sensor devices that are feasible for deployment.

On the education front there is also substantial work. Maryanne McCormick and Ruzena Bajcsy from UC Berkeley developed an undergraduate course to be offered beginning in the fall 2008 semester on privacy and ethics in the use of medical technology. Yuan Xue from Vanderbilt taught the “Network Security” course at Vanderbilt in the fall 2007 semester which covered issues on wireless and sensor network security. She also participated in the 2007 WISE summer program and gave a lecture on privacy issues in the homecare environment. The Stanford team participated in the Computer-Science Undergraduate Research InternShip (CURIS) program, offering research opportunities that help undergraduates develop an interest in computer science research. Finally, all EMR teams participated in the development of course modules and contributed to TRUST educational efforts.

2.4.2 End User Security

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

This research area is concerned with issues that address questions of security for end users. The proposal studies three major areas: innovative new browser protection systems; transaction mechanisms for protecting systems from malicious interference, and development of forensic techniques for diagnosing attacked systems.

Each of these builds on earlier techniques. Browser protection systems provide underlying technology that can be used to inform anti-malware systems. These, in turn, leave an evidence trail that can be used to understand systems that have come under attack. A central aspect of this proposal will be the exchange of information across the three institutions as part of a coordinated effort to tackle End User Security problems at multiple levels.

A major focus of this work is the exploration of online identity theft and related threats that pose risks for millions of Americans using the World Wide Web on a daily basis. Online identity crimes involve multiple victims, result in large dollar losses, compromise privacy, are often used by organized criminal groups, and may be associated with other crimes such as illegal drugs, mail fraud, and terrorism. This research area was originally conceived around technology for preventing phishing, which uses fraudulent e-mail to deceive consumers into visiting fake replicas of familiar Web sites and disclosing sensitive information. While TRUST researcher developed and deployed various ways of mitigating phishing in prior years the problem still remains and the opportunity exists for greater TRUST impact through improved methods, outreach, and technology transfer. In addition, perpetrators are developing and deploying increasingly sophisticated and powerful methods, leveraging spyware, botnets, and related malware. These advancing threats pose new technical problems, and raise questions about legal status of organizations that produce and deploy software that facilities identity compromise.

This collaborative TRUST research area, involving faculty and students from computer science and law at the University of California, Berkeley, Carnegie Mellon University, and Stanford University, has studied the social and legal context of identity theft, developed improved technology to combat phishing, spyware, botnets, and related threats, pursued technology transfer opportunities, and studied the policy and legal implications of intrusive activities and possible defensive measures.

The main topics addressed during this reporting period were:

- Usable Web Authentication
- Protection of Web Content from Malicious Interference
- Computer Forensics and Privacy
- Education and Outreach.

The sections below provide more details on the key accomplishments and outcomes in each End User Authentication topic.

Usable Web Authentication

There are many proposals for how to improve web browsers and/or web sites to increase resistance against phishing attacks. Unfortunately, all known schemes have substantial limitations in their ability to defend against phishing. Moreover, pharming has so far received considerably less attention from researchers. We have identified a sophisticated new type of pharming attack, which we call *dynamic pharming*. Because dynamic pharming hijacks a web session after it has already been authenticated, dynamic pharming can be used to attack all known methods for web authentication, including password-, cookie-, and SSL-based schemes.

One key insight is that if authentication credentials are *disclosure-resistant* – namely, if a user cannot reveal her authentication credentials to others even if she wants to – then it will be hard for a phisher to steal the user's credentials. Existing solutions have been based on using cookies and SSL. These solutions provide excellent usability, are incrementally deployable, and are inherently resistant to phishing attacks and to conventional pharming attacks. However, they are still vulnerable to dynamic pharming attacks.

In this project, we have refined our methods for web authentication. In particular, the goal is to prevent web site authentication credentials from being stolen by phishers, pharmers, and other attackers. We developed countermeasures to these attacks and develop new web authentication techniques that resist dynamic pharming and other advanced attacks. Our work is based on refining the *same-origin policy*, a core component of web browser's security architecture, which is unfortunately flawed in its current form. We will show to revise the same-origin policy in a way that defends against pharming and provides a basis for highly secure web authentication using cookies or other persistent web objects.

Second, we have developed a new methodology for user studies of web security schemes. Today, standard practice for evaluating anti-phishing tools and other web security schemes was to perform a user study where users are asked to play a fictional role, pretending to be someone who must protect sensitive data. Recent work by Schechter et al. has shown that role-playing introduces significant biases. The implication is that conventional user study methodologies are flawed. We have developed new methods for designing user studies that avoid this methodological bias. We have examined approaches based on use of subterfuge to prevent users from realizing that they are in a security-related user study and to better simulate real-world conditions; as well as exploring ethical implications of such approaches. Protection of human subjects plays a large role in our research.

We have evaluated the usability and security of our web authentication schemes using our new user study methodology in a study involving several hundred users.

We also studied an active underground economy which specializes in the commoditization of activities such as credit card fraud, identity theft, spamming, phishing, online credential theft, and the sale of compromised hosts. Using a seven month trace of logs collected from an active underground market operating on public Internet chat networks, we measure dhow the shift from "hacking for fun" to "hacking for profit" has given birth to a societal substrate mature enough to steal wealth into the millions of dollars in less than one year.

Computer security is a field that lives in co-dependence with an adversary. The motivation for security research is ever to stymie the goals of some hypothetical miscreant determined to violate one of our security policies. Typically, we abstract away their motivations and consider the adversary solely in terms of their capabilities. There is good reason for this since the threat model for any security mechanism is generally driven entirely by the adversary's abilities. Moreover, reasoning about any individual's state of mind, let alone predicting their behavior, is inherently prone to error. That said, the nature of Internet-based threats has changed over the last decade in ways that make it compelling to attempt a better understanding of today's adversaries and the mechanisms by which they are driven. First and foremost among these changes is the widespread observation that Internet-based criminal activity has been transformed from a reputation economy (i.e., receiving "street cred" for defacing Web sites or authoring viruses) to a cash economy (e.g., via SPAM, phishing, DDoS extortion). Indeed, even legal activities such as vulnerability research has been pulled by the gravity of a cash economy and today new vulnerabilities are routinely bought and sold by public companies and underground organizations alike. Thus, there is a large fraction of Internet-based crime that is now fundamentally profit driven and can be modeled roughly as rational behavior. Second, and more importantly, the nature of this activity has expanded and evolved to the point where it exceeds the capacity of a closed group. In fact, there is an active and diverse on-line market economy that trades in illicit digital goods and services in the support of criminal activities. Thus, while any individual miscreant may be difficult to analyze, analyzing the overall market behavior and the forces acting on it is far more feasible.

We performed a first exploration into measuring and analyzing this market economy. Using a dataset collected over seven months and comprising over 13 million messages, we documented a large illicit market, categorized the participants and explored the goods and services offered.

Protection of Web Content from Malicious Interference

We have developed tools for enhancing the security of web transactions and for defending against upcoming attacks. Current phishing attacks focus primarily on stealing user credentials such as passwords. Web sites, in response, are deploying stronger authentication and backend analytics systems that make it harder for phishers to extract value from stolen passwords. We anticipate that cyber criminals will soon adapt. In particular, we expect to see huge growth in the use of a new type of malware called a Transaction Generator (TG). A TG waits for the user to log in to his account at a site and then issues transactions on behalf of the user. TGs can be stealthy and issue transactions without leaving any trail. Consequently, we claim that strong identity systems are insufficient for securing web transactions. One must, in addition, deploy a transaction confirmation system that enables users to confirm every transaction they make.

In this project we have developed a tool called SpyBlock to defend against the upcoming threat of Transaction Generators. SpyBlock is a secure transaction confirmation system isolated from malware, a VMM-based approach for this purpose that will provide strong isolation in commodity Oses, integrated with existing identity systems such as OpenID and CardSpace so as to be easy to deploy and use.

Spyblock is a system that protects web passwords from malicious spyware and keyloggers. The system consists of two components: a browser extension that runs in an untrusted environment with the browser and other applications, and an authentication agent that runs in an environment that is protected from spyware. Using a virtual machine monitor, the trusted and untrusted components can both run on the same physical machine. The user only interacts with the trusted component during authentication; all other web browsing activity can be conducted using the untrusted application environment. The SpyBlock system protects user passwords from Keyloggers and Transaction Generators on the user's machine. All user passwords are kept hidden from the VM and any spyware running inside the VM. Instead, users enter passwords into the SpyBlock agent running on the host OS. The agent embeds (hashed) passwords in outgoing HTTP login requests. As a result the VM never sees user passwords. The agent enables users to confirm transactions so that a malicious transaction generator cannot fake user requests. Deploying the system on a large scale can now be done free of charge thanks to freely available virtual machine monitors.

Of related interest is the protection of individual operating system kernels. Computing platforms are encompassing an ever-growing range of applications, supporting an ever-growing range of hardware, and providing tremendous functionality. Consequently, the complexity of our computing platforms is steadily increasing, resulting in kernel code sizes of 4.3 million lines of code for Linux 2.6 and 40 million lines of code for Windows XP.

The increased complexity of OSes, unfortunately, increases the number of vulnerabilities and thus makes them more vulnerable to attacks. This is compounded by the fact that, despite many efforts to make OSes modular, most OSes in common use today are monolithic in their design. A compromise of any part of a monolithic OS can potentially compromise the entire OS. Given that an OS occupies a privileged position in the software stack of a computer system, compromising the OS gives an attacker complete control over a computer system.

In view of the importance of the security of the OS to the security of a system, securing existing operating systems as they exist today is of critical importance. In other words, it is preferable to propose approaches that do not mandate large-scale design changes to existing OSes or call for building new OSes. In this area, we have taken a first step in that direction by developing SecVisor, which prevents an adversary from injecting code into an OS (i.e., SecVisor can guarantee integrity of the code executing in the OS). We can achieve this guarantee even in the presence of an attacker with complete control over the computer system except the CPU, memory controller, and the memory bus.

SecVisor is a tiny hypervisor that virtualizes the CPU's Memory Management Unit (MMU) and the IO Memory Management Unit (IOMMU) thereby ensuring that the OS does not have control of the CPU and DMA protections over its own addresses. Since SecVisor does not support multiple Virtual Machines (VMs), its MMU and IOMMU virtualization code is smaller than that found in Virtual Machine Monitors (VMMs). SecVisor sets memory protections over OS addresses to ensure that only code that it approves of can execute with kernel privilege. It also ensures that the approved code cannot be modified. The approval policy is flexible and can be set to whatever the user wants.

SecVisor prevents numerous attacks against current OS kernels. For example, there are at least three ways in which an attacker can inject code into the OS kernel. First, the attacker can misuse the modularization support that is part of many current OSes. Modularization support allows privileged users to add code to the running kernel. An attacker can employ a privilege escalation attack to attain sufficient privileges to load a module into the OS. Privilege escalation is a common feature observed in many of the current attacks. Second, the attacker can locally or remotely exploit software vulnerabilities in the OS kernel code and, for example, inject code by exploiting a kernel-level buffer overrun. The NIST National Vulnerability Database shows that both the Linux Kernel and Windows XP SP2 were found to have 81 and 31 such vulnerabilities, respectively. Third, DMA-capable peripheral devices can corrupt kernel memory via DMA writes. Since SecVisor requires very minimal changes to the OS, it can be deployed to protect legacy Oses against these attacks.

We implemented SecVisor on a computer with an AMD CPU, running Linux. Our implementation uses AMD's Secure Virtual Machine (SVM) technology to virtualize the MMU. SecVisor also uses the Device Exclusion Vector (DEV) protections that are part of the SVM technology to protect memory from DMA-writes by peripherals. The virtualization features of SVM technology enabled us to greatly minimize the size and complexity of SecVisor. Virtualization support is now widely available on both Intel and AMD CPUs, making it practical to widely deploy SecVisor.

Computer Forensics and Privacy

We developed three major systems for forensic analysis: the Rapidly Reconfigurable Security Scanner (developed with Taiwan's Institute for Information Industry and now deployed to protect real systems); SWOON, an extension to the DETER testbed that allows it to fully simulate wireless systems; and SuperTED, a testbed for the security of sensor networks. These systems allow us to detect potential problems in deployed systems and to recreate events that could have lead to particular evidence trails of attacks. This testbed event recreation allows for more accurate diagnosis of real attacks.

Leveraging the DETER architecture, users can make their own wireless topologies on SWOON. Since DETER is a testbed for wired network, two experiment nodes are used to present one wireless node in SWOON: one presents the application node while the other, the shadow node, behaves as the wireless network interface of the application node. The application node may be an access point (AP) or a wireless station (STA). The shadow node is responsible for broadcasting wireless packets to other nodes on the wireless LAN. According to configuration from users, a shadow node can adjust the loss rate, bandwidth and latency for its application node. Thus, users can execute their systems in the application node and simulate the network behaviors in the shadow node.

Education and Outreach

Education and outreach are central goals of this research area. Identity theft is a real problem, and we continued to work with law enforcement groups such as the U.S. Secret Service, the Federal Bureau of Investigation, Infragard, the Department of Homeland Security Identity Theft Technology Council, the anti-phishing working group (www.antiphishing.org), and industry to get our anti-phishing information, our vulnerability testing information, and our legal and policy analyses as widely disseminated as possible. Our ultimate goal is to effect real change that will touch everyone who uses the Web. Several technology transition partners have been identified and engaged to varying degrees, including PassMark Security, RSA Security, and divisions of Microsoft. In order to increase public awareness of this problem and its potential solutions, we also continue to talk with the press.

We have also worked to influence the teaching of security by developing study units on topics in End User Security. During this reporting period, we expect to add a number of such units addressing topics such as web authentication, basics of malicious web content, statistical learning theory in computer security, usable security, and forensic techniques.

2.4.3 Network Defenses

Project Leader: *Adrian Perrig (Carnegie Mellon University)*

Computer networks are, arguably, one of the key technical developments of our era. They have enabled us to construct powerful systems of tremendous scope and complexity. But with this scope and complexity they also bring exposures to failures, concurrency-related bugs, poor management, and outright misuse. Given their design assuming a trustworthy environment, modern networks have become exceedingly hard to defend against mishap, whether accidental or deliberate, and this observation has made research into network defense, broadly construed, an obvious and central area for investigation by members of the TRUST team.

TRUST researchers are pursuing a gamut of innovative topics in the area of computer networks, which we classify roughly into the area of "network defenses" techniques. During the 2007-2008 reporting period, Network Defenses activities unified several closely related internal TRUST projects. Most of these activities involved multiple institutions and all had an "organic" need for dialog, sharing of ideas, and other forms of participation by multiple organizations and multifaceted research teams capable of looking at a spectrum of issues that range from social and pragmatic to highly technical. The remaining activities fall into categories in which TRUST researchers are proposing work complementary to the primary, more collaborative, activities. In aggregate, the work includes efforts from essentially every facet of the TRUST Center.

During 2007-2008, progress was made in all areas. The DETER testbed has grown to encompass more TRUST partners, with Cornell University and the University of California, Berkeley also extending the basic concept to explore questions associated with long-haul, high latency links. The shared testbed is an exceptionally powerful resource, not just for the original purpose of studying virus outbreaks but also for exploring new concepts that might be developed further in the context of the National Science Foundation Global Environment for Network Innovations (GENI) initiative.

At the University of California, Berkeley, the DETER testbed is also being used to explore collaboration-based attacks against network switch-based worm and virus detection. The work is joint with researchers from HP Labs in Bristol UK and HP in Roseville, California and has already yielded new insights into how steathly worms can collaborate to quickly propagate in a nearly-undetectable manner, even in the presence of switch-based throttling. The work has also yielded preliminary results on effective, efficient defenses against attacks relying on collaborating participants.

Another DETER testbed project has developed a framework for using machine learning in the presence of adversaries for security applications (e.g., worm/virus detection, intrusion detection, etc.). The framework can be used both as a tool for evaluating the vulnerabilities of a learning-based security application, and as a tool for designing new applications.

Finally, at the application level, there is an opportunity to better understand the network challenges when meeting the requirements of applications with strong real-time requirements and in the presence of on-going network attacks.

During the 2007-2008 reporting period, Cornell University has made significant progress on epidemic protocols, which are an important direction in robust networking and are increasingly deployed, particularly within datacenter applications. Some epidemic protocols pertain only to tolerating benign failures, but overall they have been focusing much attention on malicious failures and Byzantine failures in general.

While the robustness of epidemic protocols appears universally accepted, we have identified various weaknesses in the work by Alvisi et al. published in ACM SIGOPS Operating Systems Review, Special Issue on Gossip-Based Networking, October 2007. The Cornell Fireflies system, an epidemic-based scalable intrusion-tolerant group management system, addresses many of the issues, and recently, in cooperation with researchers at the Hebrew University of Jerusalem, we have extended Fireflies and made it self-stabilizing. Thus even if at times the system inadvertently deviates from its specification, it converges back to a correct state. Doing so provides a significant added level of protection.

On top of Fireflies we have developed SecureStream, an intrusion-tolerant live-streaming system. With help from Ingrid Jansch-Porto of the Federal University of Rio Grande do Sul, Brazil, we have extended SecureStream to deal with heterogeneity in the user population, moving low-capacity peers, as well as freeloading peers, to the edges of the dissemination graph, while high-capacity peers form the core of the graph.

Related to this effort is the NightWatch system, a system based on an epidemic protocol that generates a synopsis of a probability distribution of some monitored value. NightWatch can be used to classify peers (low-capacity or high-capacity) in a system like SecureStream, and can be used by self-organizing systems in general.

We are also investigating how we can improve the robustness of more traditional protocols. We have formulated a general framework for consensus, a project that we are developing still further. It forms the basis for developing one-size-fits-all solutions to a variety of replication problems, both in the realm of crash failures and more general failures such as crash failures. The framework exploits opaque quorum systems, and a direction for future research is to use the probabilistic opaque quorum systems of Mike Reiter at the University of North Carolina, Chapel Hill to build a probabilistic consensus protocol, somewhat akin to our Fireflies system but using a very different approach. Interestingly, though, we have discovered that in general it is unnecessary to use Byzantine consensus for tolerating Byzantine failures. Again with cooperation with people at the Hebrew University of Jerusalem we designed a general transformation of crash-tolerant protocols to Byzantine-tolerant protocols for asynchronous systems. Up to now such translations were only known for a limited class of protocols or for synchronous systems. The initial transformation technique only works for small scale systems, but we have since developed a scalable technique called Nysiad. Nysiad could be applied to protocols such as BGP or DNS. This would go beyond the current secure version of such system in that it would not only avoid compromise, but indeed tolerate (limited) compromise.

Carnegie Mellon University has made important progress in making networks more available by making them resilient against large-scale distributed denial-of-service attacks. The two major accomplishments are the Portcullis and SNAPP systems, which remove the major obstacles that have plagued network capability systems.

In a Distributed Denial-of-Service (DDoS) attack, an adversary, sometimes controlling tens of thousands of hosts, sends traffic to a victim to exhaust a limited resource, e.g., network capacity or computation. The victim of a network DDoS attack can often identify legitimate traffic flows but lacks the ability to give these flows prioritized access to the bottleneck link; in contrast, routers have the power to prioritize traffic, but cannot effectively identify legitimate packets without input from the receiver.

Network capabilities enable a receiver to inform routers of its desire to prioritize particular flows, offering a promising DDoS defense. To set up a network capability, the source sends a capability request packet to the destination, and routers on the path add cryptographic markings to the packet header. When the request packet arrives at the receiver, the accumulated markings represent the capability. The receiver permits a flow by returning the capability to the sender, who includes the capability in subsequent packets to receive prioritized service from the routers.

Current proposals for capability-based systems treat prioritized traffic (i.e., packets with a valid capability) preferentially over non-prioritized traffic. However, capability-based systems still suffer from a critical weakness: they cannot protect the initial capability request, because that request is sent unprotected as non-prioritized traffic. An attacker can flood the capability-setup channel, thus preventing a legitimate sender from establishing a new capability-protected channel. This attack, referred to as Denial-of-Capability (DoC) by Argyraki and Cheriton, is the Achilles heel of current capability proposals. Agryraki and Cheriton show that several thousand attackers can easily saturate the request channel of a typical network service, preventing legitimate senders from acquiring capabilities.

Our Portcullis system addresses these drawbacks and strictly bounds the amount of delay a collection of attacking nodes can create for any client. With realistic Internet-scale simulations, we show the strong fairness Portcullis' computational puzzles provide. Portcullis introduces a powerful mechanism for providing DDoS resistance.

The central contribution of the Stateless Network-Authenticated Path Pinning (SNAPP) project is a stateless, secure path-pinning building block that provides numerous benefits without the drawbacks of prior approaches. SNAPP adds a small amount of information to packet headers as they pass through the network, akin to the techniques used by network capabilities. Once a packet has traversed a path, the sender and receiver can send additional packets that are forwarded based upon the information encoded in the packet.

In considering architectural primitives for designing a network architecture, whether for an overlay network or a next-generation network core, we find that SNAPP represents a versatile building block for achieving a number of useful properties. SNAPP provides sufficient flexibility to: 1) decouple forwarding from routing to enhance the availability of paths in the face of routing disturbances, 2) provide route-selection control to the sender (to request multiple routes and select among them), 3) enable applications with expensive route lookups, 4) provide capability-based systems with stable paths despite routing changes, 5) enable load balancing at the sender, 6) provide sender anonymity at the network layer, and 7) provide sender accountability. SNAPP also provides additional flexibility for implementing routers and other forwarding devices; for example we can envision a system where high-speed switches perform the packet forwarding, and separate servers are used to aid in path setup. This may lead to an approach for optical networks, where switching may be feasible in the all-optical domain, whereas the more complex routing decisions occur in traditional hardware.

2.4.4 Policy

Project Leader: *Deirdre Mulligan (UC Berkeley)*

The TRUST research agenda includes a robust, interdisciplinary policy component. This research is aimed at contributing to the creation of secure, private and trustworthy systems by structuring incentives for research, investment, policies and procedures directed towards privacy and security enhancing technology.

Trustworthy systems are achieved through a mix of component parts, some technical, some procedural, some informed by economics and others by legal obligations. To create secure, private and trustworthy technology and systems requires an understanding of the relationship between the component parts and an active consideration of how one domain interacts with the other. Technology deployment decisions made without an understanding of how the decisions relate to policy, and policy decisions made without an understanding of the existing assumptions of the security architecture often yield problematic results. In the absence of a holistic approach to considering how to embed values in technical systems a range of failure modes appear. The research conducted under the Policy thrust seeks to understand organizational and individual roles in making security and privacy decisions for trustworthy systems, the barriers to implementing security and privacy policies, and potential policy mechanisms that can improve better privacy, security and compliance with privacy and security mandates.

The TRUST policy work has four primary objectives.

- Understand the manner in which law and other external and internal forces influence organizational strategies, policies and practices for protecting privacy and security; identify the extent to which different forms of legal intervention create institutional engagement and response; and, identify principles to guide future efforts;

- Understand how the process of translating legal objectives into policies managed by information technology shapes and skews the "on the ground" meaning of each the individual goals reflected in the law;

- Understand the manner in which policies and mechanisms for improving information and computer security conflict with and can be reconciled (or not) with other values through law, technology, or other mechanisms; and

- Develop approaches for modeling systems that handle sensitive information, languages for specifying privacy policies, and algorithms for their enforcement.

The following sections describe TRUST research during this reporting period to advance these objectives.

Understanding and Improving Organizational Behavior – CSOs & CPOs

- In order to understand how best to create incentives for optimal privacy and security investments in trustworthy systems, we must understand how and why organizations and individuals make security and privacy decisions. We must understand the barriers to implementing chosen policies (whether externally or internally driven) and identify potential avenues for achieving better privacy, security and compliance.

  Interviews with 10 leading Chief Privacy Officers were completed in the Fall of 2007. A series of papers discussing the effect of various internal and external forces on organizational behavior with respect to the implementation of privacy have resulted from this initiative. The first, "Privacy Decisionmaking in Administrative Agencies," was presented at the Conference on Surveillance at the University of Chicago Law School June 15–16, 2007, and subsequently published in the Chicago Law Review. The second, "Catalyzing Privacy: Corporate Privacy Practices Under Fragmented Law," was presented at Center for the Study of Law and Society, at UCB, on December 3, 2007, and will be placed in a law review during Fall 2008. The researchers are also preparing short articles for publications aimed at security and privacy professionals (lawyers, technologists, business) to translate their findings for these audiences.

  Interviews with Chief Security Officers began this spring. Fifteen CSO/CISO from a range of sectors (infrastructure (information and physical), financial services, retail, and healthcare) have been contacted and we are finalizing the interview schedule. The initial set of interviews have revealed interesting and nuanced perspectives on the utility of law in effecting positive changes in security.

  A literature review on the economics of security, existing security laws and regulations, and self regulatory security standards, and information and security experts perspectives on security has been conducted.

Case Studies of RFID in Public Identification

- In 2007, both federal and state agencies proposed the use of RF technology in types of state ID, from the Dept. of State's PASS card to allow US citizens to travel between selected countries in the Western Hemisphere without a passport, to the State of Washington's proposal to create a PASS-like card that also doubles as a driver's license.

  During fall of 2007 Jennifer King conducted exploratory research with nine subjects to test methods and study protocols to elicit end user mental models of RFID technology. We focused on three RF-enabled objects: BART EZ-Rider transit cards, credit cards, and the US e-Passport. Initial results from this research were then used to refine the study; a short paper outlining theory, related research, methods and results was submitted and accepted to the USENIX Psychology and Security workshop in April 2008. The remaining user research will be conducted during Summer 2008 and final results should be available in Fall 2008.

Privacy, Compliance, and Risk Management

- This work builds on early TRUST work to develop approaches for modeling systems that handle sensitive information, languages for specifying privacy policies, and algorithms for their enforcement. Stanford and CMU have collaborated on privacy research motivated by interactions with Vanderbilt and Vanderbilt University Medical center. Our goals include producing web portal concepts, foundations, and technology that can be applied in the MyHealth@Vanderbilt patient portal, and also for other systems that have similar structure and goals. Through a consideration of patient portal workflow we have identified key questions to ask of systems with respect to measuring privacy and developed auditing techniques for detecting potential policy violations. "Privacy and Utility in Business Processes," by A. Barth, J.C. Mitchell, A. Datta, and S. Sundaram, was published in the proceedings of the 20th IEEE Computer Security Foundations Symposium (CSF 20), held in Venice, Italy in July, 2007.

  This paper considers agents in an organization whose actions transmit personal information from a sender in one role to a receiver in some other role. Privacy and utility goals are expressed using a Logic of Privacy and Utility (LPU) developed in the paper, with workflow, expressed as a combined responsibility of a mechanical workflow engine and organizational roles of agents. In the Vanderbilt patient portal workflow, for example, doctors are responsible for answering health questions and secretaries are responsible for scheduling appointments.

  With this mode, the paper formulates and addresses the key question: Does a given workflow achieve its privacy and performance goals? We have also developed algorithms for examining audit logs for policy violations, and determining the set of possible agents who may have acted improperly relative to their business process responsibilities.

- Demonstration System: In summer 2007, a team at Stanford including one MS student, one PhD student and two undergraduate students built a web portal demonstration system implementing portions of the privacy framework described in the 2007 publication. This system combined a web server front end with a relational database backed and a Prolog-based policy engine to interpret and enforce declarative privacy policy. This demonstration system allowed Stanford to demonstrate automated policy enforcement, and will also provide a demonstration platform for further policy formulations as they develop. In addition to demonstrating how research could be applied to the Vanderbilt portal, this project raised several research and design questions that were addressed during the summer project. The summer 2007 web portal project was carried out in connection with the Stanford CURIS summer undergraduate research opportunities program.

- HIPPA Formulation: Stanford and Carnegie Mellon continue 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 the regulation that can be used by other organizations as the basis of automated operational compliance, thus addressing the need established through visits and interaction with Vanderbilt.

Security Considered with Other Values

- TRUST fellow Aaron Burstein completed two projects in this area. The first was an examination of legal and ethical issues surrounding cybersecurity research, with a particular focus on obtaining access to network data. This work benefited from collaboration with Vern Paxson and has resulted in two publications: a conference paper in the USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET), and a law review article, which contains an extensive analysis of the legal issues and recommends a security research exception to federal communications privacy laws. The second activity was a collaboration with Fred Schneider on the interaction between Internet trustworthiness and other Internet policy goals. For a specific example of this tension, Burstein and Schneider analyzed the relationship between trustworthiness and network neutrality, a much-debated notion of prohibiting broadband service providers from degrading network traffic, applications, or protocols. This collaboration resulted in a paper presented at the 35th Telecommunications Policy Research Conference (TPRC) as well as a paper that is currently under review by law journals.

Technologies and Compliance

- Authentication and Identity Theft: Fair and Accurate Credit Transactions Act of 2003 (FACTA) Access Study seeks to study the procedures and policies in place at credit card issuers that are the first line of protection against identity theft. Under the FACTA, victims of identity theft may obtain business records associated with the crime from the company that created an account for the impostor in the victim's name. Through obtaining the business records in identity theft cases, we will be able to evaluate business practices and defenses to identity theft, make policy recommendations, and perhaps reduce the amount of "new account" identity theft that causes enormous damage to the economy. Ultimately, the goal of the FACTA Access Study is to promote procedures, through technical recommendations or changes in the law, that will create more trustworthy authentication systems in the credit granting context. Work began on the FACTA Access Study last summer, when an application was made to the Committee for Protection of Human Subjects at UC-Berkeley in August 2007. The Committee approved the protocol on November 20, 2007. This semester, we have commenced work on the actual study, by designing a questionnaire and developing a strategy to recruit subjects. We intend to start interviews with identity theft victims and make FACTA requests in April 2008.

- Technologies of Compliance: This project considers the ways in which information technology systems geared to compliance with a number of different end goals – reducing information privacy risks and data security risks and financial control risk and enterprise risk more generally – might shape (and skew) the “on the ground” meaning of each one of these individual goals. The team has begun research on how technology is being used by firms in compliance activities. Several JD students and a PhD student at the UC Berkeley School of Information have conducted a literature review and the lead researcher has participated in several meetings and conferences about information technology and compliance, including a meeting with John Mitchell at Stanford and Tata Consultancy Services (Stanford collaborators in India), to discuss privacy and compliance issues. The lead researcher has also created links with the Information Systems Compliance Consortium, based at SUNY Stony Brook.

The Policy work has resulted in substantial multi-institutional collaboration that is beginning to bear fruit in the form of collaborative publications and proposals. Several conference papers and a journal article have already been accepted (see list of publications elsewhere in this report).

2.4.5 Secure Sensor Networks

Thrust Leaders: *Steve Wicker (Cornell University), Deirdre Mulligan (UC Berkeley)*

The TRUST Secure Sensor Networks initiative focuses on the development and use of secure embedded sensor networks in a variety of large-scale applications. Applications to be emphasized include the protection and monitoring of critical infrastructure, rapid response systems for homeland defense, and the remote monitoring of individuals for clinical purposes, whether living at home or in group facilities. Recent developments in the field of sensor and networking technology have made such networks possible; this initiative will consider the further development of the requisite deployment, network configuration, data dissemination and query generation and response, and security technologies. This initiative also considers the privacy and security issues arising from the use of sensor networks, and the ways in which embedded sensor networks affect the expectations, experiences and activities of individuals in public and private spaces. An emphasis is currently being placed on developing privacy metrics, and limiting the acuity of sensing technologies to the minimum required to meet mission objectives. As well as considering the questions raised about the relationship between citizens and government by the possibility of constant monitoring enabled by widespread deployment of visual and other sensors in public spaces. A significant educational and outreach component has been developed with the joint objective of increased diversity in the ongoing development of these technologies and an increase in public awareness of the surrounding technical, legal, economic, and social issues.

The TRUST sensor networking team has four primary objectives.

- Develop technologies that facilitate the use of large-scale embedded sensor networks in applications that are critical to the nation’s economy, security, and health.
- Demonstrate these technologies through the use of realistic testbeds, enhancing technical development while enabling ties to our corporate sponsors.
- Examine the legal, economic and societal issues that emerge from the use of these technologies in public and private places, and develop policies that guide their design, development, deployment, use, and regulation to protect the privacy, security, and economic and societal interests of the public.
- Develop security technologies that limit and characterize the potential threat from passive and active network intruders.
- Develop mechanisms for increasing diversity among the practitioners of sensor networking technology and social sciences, while building teaching tools that increase awareness of the capabilities of this technology.

In the past year we made extensive progress in establishing our testbeds and further developing cross-institutional research teams. Our research has coalesced around three main facilities: a *Camera Networks Testbed*, a *Microgrid/SCADA Testbed*, and a *Medical Sensing Testbed*. Additionally, several PhD students have been exchanged between Cornell, Vanderbilt, and Berkeley. This effort has resulted in substantial multi-institutional collaboration that is beginning to bear fruit in the form of collaborative publications and proposals. Several conference papers and a journal article have already been accepted (see list of publications elsewhere in this report).

In the remainder of this section, we provide a description of each testbed as well as introductions to research projects and specific results that have emerged during this reporting period from the use of the testbeds.

Camera Network Testbed – The camera network testbed is based at Berkeley and has involved students and faculty from Berkeley, Cornell, Stanford, and Vanderbilt. Key research topics include the following:

- Extremely-low power computational platform
- Camera localization strategies
- Privacy-aware image capture and transfer
- Multi-level, policy-directed sensing

Projects and research results that leveraged the Camera Network Testbed area as follows:

- *Camera Network Development*: The use of high-bandwidth, wireless, mesh, networked camera systems for surveillance of public spaces has increased dramatically during the past ten years. This increase has been due, in part, to the development of inexpensive, simple, yet powerful camera technology. With the increased prevalence of these systems, there have been a number of privacy concerns raised. Many of these concerns pertain to the amount of detailed information that is provided by the images captured by the cameras and the near-real time operation and monitoring of the cameras. In response to the concerns, policies and best practices have been suggested. To complement this effort, during the last year, we examined current systems and their policies for deployment, and we suggested novel system-level technical solutions to help preserve privacy and to enhance privacy awareness. To validate our work, we designed and tested proof-of-concept, privacy-enabling, software for the camera systems. We also designed a wireless, mesh, camera network testbed incorporating some of our suggested technical solutions and deployed it in the public spaces of Cornell's Johnson Art Museum.

  In particular, to complement policies and best practices for camera system deployment, we suggested the use of system design measures, such as online notices and public-input surveys. We designed and suggested the use of a validation code based automatic shut-off software to help enforce policies that demand a fixed period of time during which the camera system can be in use. To help systems adhere to the policies concerning minimal collection of personally identifying information, we also suggested that only certain individuals or objects of interest be tracked or identified (e.g. a person carrying a gun) using real-time computer vision data abstraction techniques. Specifically, we designed vision software capable of tracking objects or humans of interest using Kalman-filter based predictive searching, robust background subtraction and object segmentation (using quick erosion, closing, and region growing algorithms), and intelligent region selection (based on dimensions of the region).

  To further validate our suggested solutions, we constructed a wireless mesh network testbed system comprised of four off-the-shelf AXIS IP cameras, two Linksys 802.11

routers, and a wireless sound actuation component comprised of distributed, wireless Jornada PDAs and deployed it in the floor-space of the Johnson Art Museum. The system was used to track multiple individuals around the Museum's floor-space (without collecting any Personally Identifiable Information) and to generate distributed sounds based on individual presence or absence in particular portions of the space. While the system accomplished this task, the data abstraction techniques that we employed ensured that we achieved the goal of preserving the privacy of the individuals captured in the image data from the cameras. The system was also capable of targeting and revealing particular objects of interest based on key features (e.g. size and shape) and sounding an alarm if these objects were detected in the field of view of the cameras.

We ensured that performance and cost were not compromised when designing our testbed system with our stated privacy goals in mind. Using the latest off-the-shelf wireless routers and simple software engineering techniques, such as distributed programming with multi-threading, and a fast interpreted programming language codebase, we were able to keep the latency of network data aggregation and processing at a minimum. This was also in spite overhead of using real-time data abstraction techniques. We compared the frame rate of our testbed with that of a similar surveillance camera deployment contained on one block of San Francisco. We found that we had approximately a factor of 6 increase in frame rate with a comparable distance between routers in both deployments and despite the overhead of performing image processing to meet our privacy goals. The cost was also kept to a minimum (~$300/camera for our testbed system versus ~$10,000/camera in the San Francisco system).

In the coming months, leveraging our accumulated knowledge of privacy in the camera networking domain, we will survey the current policies for the use of San Francisco's camera networks and we will suggest additional solutions that may be necessary to provide additional privacy protections. The results of this work will be turned into reports for the San Francisco Mayor's Office.

- *Camera Network Policy:* The emphasis thus far in this area has been an exploration of the privacy issues and broader questions about policing and democracy posed by the potential capacity of the state to engage in persistent monitoring of public places. Data about the configurations and use of deployed systems and related policies has been sought for analysis. The ability of current privacy law to respond to the qualitative and quantitative changes in the use of camera networks by the government has been examined and found insufficient. The use of theories considering the relationship between policing and democracy are currently being explored as vehicles to more successfully frame the conversation about the policies that should guide the development, deployment, and use of permanent public video surveillance systems. We have developed impact assessment tools, guidelines, and model legislation to create opportunities for public input into system decisions and substantive policies designed to protect the privacy, associational, expressive, and equality interests of society in camera networks. We are exploring the ability of technology to affirmatively protect privacy or other values in camera networks including technical options for de-identification, abstraction, and triggering that reduce the collection of data in ways that respond to articulated privacy concerns. We are also considering potential attacks on these networks and creating technical counter-measures and design options to diminish the attack surface. This area has included outreach with relevant state and federal agencies as well as localities considering camera networks.

Microgrid/SCADA Testbed – The microgrid/SCADA testbed is based at Cornell and involves faculty and students from Cornell, Smith, and Berkeley. Key research topics include the following:

- Secure data transport mechanisms based on public 3G/4G cellular
- Multi-layer security against insider attacks
- Info-Theoretic characterization of privacy content/privacy filtering
- Demand/Response system performance
- Microgrid control performance
- Main grid SCADA integration
- Secure SCADA data and signaling transport
- Security Co-design strategy using Vanderbilt design tool

Projects and research results that leveraged the Microgrid/SCADA Testbed area as follows:

- *Power Sensing Systems and Privacy:* The next decades will see a transformation of our nation's power distribution systems. Next generation Supervisory Control and Data Acquisition (NG-SCADA) architectures will precipitate an exponential increase in both the data and control available to consumers and utilities. Utilities are increasingly adopting automated metering, advanced demand response architectures, microgrids, and other systems which will provide cost savings in power generation, increase grid reliability, and create new modes of consumer-utility interaction. This transformation is already well underway. Recent years have seen several pilot microgrid projects, as well as increased deployment of Advanced Metering Infrastructure (AMI) systems by major utilities across the US. AMI systems in particular have been deployed on a large scale by entities such as California Public Utilities Commission. According to a 2006 Federal Energy Regulatory Commission staff report, six percent of meters installed in the US are `smart' meters supporting some advanced metering project, and the number continues to steadily increase.

  In a typical Advanced Metering setup, the customer is equipped with solid state electronic meters that collect time-based data at daily, hourly or sub-hourly intervals. The types of available devices differ from project to project, but may include electricity, gas, and water meters. These meters have the ability to transmit the collected data through commonly available fixed networks such as Broadband over Power Line (BPL), Power Line Communications (PLC), and public networks (e.g., landline, cellular, paging). The meter data are received by the AMI host system and sent to the Meter Data Management System (MDMS) that manages data storage and analysis, shaping the information into a form useful for the utility.

  A NILM system collects data much like its AMI counterpart, but goes a step further by processing the data to determine the operating schedules of individual electrical loads. This is typically done by disaggregating the collected data stream into individual load signatures and matching each signature with reference signatures stored in a database. For private residences, these loads are usually appliances such as the refrigerator, air conditioner, or water heater. Several NILM systems of varying capabilities exist, including a commercially available system which can distinguish between major appliances, a system, based on a genetic algorithm, that does not need training data, and various experimental high-capability systems developed at MIT which identify complex loads and even pinpoint malfunctioning appliances.

Although it is known that first-stage parameters such as appliance usage may be accurately estimated, to our knowledge no one had ever tried to extrapolate activity from power consumption data. In this past year we chose to prove that activity extrapolation is feasible, thus lending credibility to our thesis and providing an experimental precedent which others can cite in future efforts to guide policy development. To do this, we conducted a small-scale monitoring experiment on a private residence.

We conducted our experiment in a typical Cornell University student residence. For data gathering, we used the Brultech EML energy usage monitor. The energy monitor was attached to the residence's breaker panel and sent real-time power usage information to a workstation responsible for data collection. The station recorded power usage at intervals of 1 or 15 second(s) and with a resolution of 1 Watt. The same workstation then ran the NILM and behavior extraction algorithms. To evaluate the system's performance, we placed a network of cameras around the residence. We elected to use the Axis 206 network camera, which we connected to a workstation using an Ethernet switch. The workstation ran the AXIS Camera Station software and recorded motion events for later processing.

Generally, our activity extrapolation algorithm performed quite well in determining presence and sleep cycles. In both cases, over 90% of the total interval length was correctly classified, for both training and experimental data. We believe this is due to our success in identifying the refrigerator load, the small number of autonomous appliances in the residence, and the consequent simplicity of presence / sleep-wake heuristics. We expect that this clear connection of the collection of power consumption information to the revelation of activities within the home will guide lawmakers to establish clear guidelines as to the collection, re-use, and/or sale of such data.

- *Power Grid Policy:* This area is building off work initially conducted under a grant from the California Energy Commission examining the potential privacy and security issues raised by the move to a demand response energy infrastructure in which two way communication between utilities and residences is the norm, data about energy consumption is collected in fifteen to thirty minute increments, and appliances, other energy consuming devices within the home, and sensors are in communication with programmable, computable thermostats within the home. The security and privacy challenges presented by the in-home sensor networks, the increasingly detailed data flowing out of the home, the introduction of additional players into home energy management, and the ability to remotely control devices within the home are formidable. Our current work focuses on influencing standards, regulations, and rules around demand response energy systems to ensure that the heightened ability of detailed energy data to reveal personal in-home activities is addressed. Understanding the extent to which the increased frequency of energy readings and the information from in-home appliances, devices, and sensors alter the privacy concerns around utility records and sensor readings is essential to identifying appropriate policy and technology options. Legal and technical analysis, work with regulators and industry, as well as theoretical exploration of the relationship between in-home sensor networks and existing privacy and computer security statutes and laws are among the contributions in this area.

- *SCADA Testbed Architecture:* We have developed a SCADA testbed architecture that is based on a layered model, ranging from the low-level sensors and actuators, through the Remote Terminal Units (RTU-s), to the SCADA hosts, and connected up to the corporate network. The corporate network represents the business end of a utility. This network is typical of an enterprise with a LAN/WAN connected to the Internet. However, in the case of utilities and industrial plants, the corporate network is often connected to the SCADA network in order to simplify business processes by allowing network access to critical data on SCADA servers. This is one of the biggest information assurance concerns related to SCADA systems as an attacker can now connect to the SCADA network via the Internet by compromising nodes on the corporate network. The SCADA master station consists of the SCADA master servers and the HMI. The master station is located in a central control center from where operators can monitor the performance of the entire system. SCADA master servers run the server side applications that communicate with the RTUs. The SCADA master servers poll the RTUs for data and send control messages to supervise and control the utility's physical infrastructure. Backup servers are used to increase fault-tolerance of the system.

  We envision (at least) three different realizations of the reference architecture: single simulation-based, federated simulation-based, and emulation- and implementation-based. The single simulation-based instantiation has all elements implemented using a simulation framework and language, like Simulink/Stateflow from Mathworks. We envision that the individual components of the architecture are implemented as Simulink subsystems that include the plant simulation, sensor simulations, simulations for the data acquisition and control activities on the RTUs, simulation of the computations performed on the SCADA servers, etc. For high-fidelity simulations we will model and simulate the implementation platforms as well: the OS schedulers and the networking mechanisms. The TrueTime toolsuite provides a good example for doing this in the Simulink framework. For some, e.g. network attack scenarios these models will be extended to faithfully simulate the dynamic behavior of the network under attack. The federated simulation-based instantiation uses several dedicated, coordinated simulation engines that simulate the various architectural elements. Here, the key is that the individual simulation engines work with high-fidelity, industrial-grade models, possibly using off-the-shelf, commercial products. The same architectural elements are instantiated with a different technology, for example Speedup for plant simulations, Omnet++ for network simulation, and DEVS for simulating software modules, etc. In this case the problem is the timed coordination across these simulation engines, but DoD's High-Level Architecture (HLA) offers a platform to solve this problem. HLA provides services for simulation time coordination and data interchange during the simulation process, and several simulation engines have HLA interfaces implemented. The emulation- and implementation-based instantiation uses actual commercial SCADA devices along with implementations of the software modules performing the data processing (running on realistic hardware), emulations of the network (running on a network emulator like EmuLab), and real-time simulations for the plant (running on dedicated, high-performance hardware). We believe such an emulation/implementation-based realization is feasible and could be made highly realistic and scalable. Attacks on the network and computing nodes could be analyzed in a contained laboratory environment, which is safely decoupled from the "real network", yet provides a highly realistic environment (e.g., like the DETER testbed).

Currently, we are in the process of working out the details (interfaces, protocols, data structures) of the above architecture. We have also started prototyping the emulation/implementation-based variant on small networked embedded computers that are very cost-effective for building up the SCADA testbed hardware.

- *Secure Wireless Sensing:* A taxonomy of wireless sensor network attacks was developed at Berkeley. The taxonomy details threats in terms of the OSI layer and the technology and knowledge available to the attacker. This tool has proven important in our later work on security design, as it allows threat analysts to identify and focus on threats that are specific to a given context.

  In related work at Vanderbilt, faculty and students have developed security co-design tools that couple security with the initial design stages of sensor networks. The basis 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.

<u>Medical Sensing Testbed</u> – The medical sensing testbed is based at Vanderbilt and includes students and faculty form Vanderbilt, Cornell, and Berkeley. Key research topics include the following.

- Live patient demonstration through Nashville Home Care facility
- Secure data transport mechanisms based on public 3G/4G cellular
- Variable QoS strategies for single transport mechanism
- Multi-layer security against insider attacks
- Security Co-design strategy using Vanderbilt design tool
- Tie-in to multi-level authorization/confidentiality EMR project

Projects and research results that leveraged the Medical Sensing Testbed area as follows:

- *Medical Sensing Systems:* Sensing devices such as body temperature monitors, blood pressure measurement devices, glucometers, accelerometers, acoustic sensors and video cameras are playing an increasingly prominent role in health care. Such devices have become increasingly integrated and networked within the confines of modern medical centers – the electrocardiograph in emergency rooms immediately dispatch their measurements, through wireless networks, to staff cardiologists who may begin to evaluate a patient within seconds of the test, regardless of the patient's placement within the building. Physicians may also download CAT scans and other tests onto laptops as they move from patient to patient in a typical care facility.

  We are extending the scope and reach of these technologies so that they can support remote monitoring of patients. The goal is to facilitate the movement of patients into medium-care facilities or their own homes while still allowing frequent monitoring of their condition by a physician, as well as rapid detection medical events that require rapid care.

During the past three years we have designed and built CareNet -- An Integrated Wireless Sensor Networking Environment for Remote Healthcare as a small-scale medical sensing environment and conducted preliminary experiment studies at several patients' homes.

Our system has the following features.

- High Reliability and Performance – Our wireless networking infrastructure features a two-tier design. A high-performance IEEE 802.11 wireless network is used as the backbone structure to provide local area communication coverage, while the wearable body sensors communicate with the base-stations of the backbone structure directly. Compared with sensor networks in which wireless communications are solely based on IEEE 802.15.4 standard, this design greatly improves system reliability and performance. Further, the backbone structure can also be equipped with audio and video sensors which support additional multimedia communication needs in the homecare services.

- Good Scalability and Extensibility – Our networking and security software at the backbone network is built on top of the ACE environment, which is commonly used to build extensible concurrent and networking applications. Using a backbone structure, our hybrid network design scales much better than a pure IEEE 802.15.4-based sensor network.

- Privacy Aware Data Confidentiality Protection – Privacy and data confidentiality is a primary consideration in our system design. Our system design features built-in secure communication components which are adaptively implemented for different networking environments and used at all communication phases of the system.

- Integration with Web-Based Patient Portal – In our system, the sensor data collection does not mark the end, but simply the beginning of data transport. Data from the patient centric sensor network will be collected in a web-based patient portal that is under development at Vanderbilt Medical School. The portal allows care givers to efficiently access sensor network data through a unified medical record system. The patient portal also allows for arbitration of access permissions for different parties (e.g., families, physicians, nurses, insurance companies).

In the past year we collaborated with Vanderbilt Homecare Services, identifying four senior volunteers to participate the experiment of our system prototype. In our experiment, five sensor motes were mounted on each volunteer, two on the wrists, two on the ankles, and one on the waist. Each sensor mote was capable of recording accelerations in three dimensions (x, y and z axis) as well as gyroscope in two dimensions (x and y axis).

The experiment process involved a controlled experiment, where volunteers were required to perform a set of designed movements, and an uncontrolled experiment, where volunteers performed their daily physical activities. The designed movement set included (1) Vertical stretching and grabbing of each arm; (2) Vertical stretching and grabbing of both arms; (3) Fetching cup and drinking water; (4) Sit-to-stand and stand-to-sit combination; (5) Raising each leg in sitting position; (6) Raising both legs in sitting position.

Each set of movements was repeated five times in the experiment. While sensor motes were recording the movement data, a video camera simultaneously recorded the volunteer's motion images. Both sensor movement data and video images were sent back and stored in the home healthcare gateway. We synthesized movement and video data based on their timestamps, then played back and compared their synchronicity.

The proposed project for 2008 is built on top of our previous research results and addresses the unique challenges in integrating our previous research results and improving our system for practical deployment. In particular, we will enhance the system from the following perspectives:

- <u>Scalability and Extensibility</u> – Our current monitoring system only consists of five movement sensors, two video sensors, and four gateway routers, which could cover a monitoring area up to a single family house and support monitoring service for two patients simultaneously. To enable deployment in a large facility such as the assisted living village and enable simultaneous monitoring for hundreds of people, we need to scale up our system design. To achieve this goal, we will study scalable packet routing and forwarding mechanisms, and efficient data compression and reduction algorithms. We will also extend our system to handle other medical sensors such as pulse oximeter.

- <u>Reliability and QoS-Awareness</u> – While scaling up our system to more heterogeneous sensors, we face diverse QoS requirements from these devices. This poses a challenge for our system design in the resource-constrained wireless sensor networking environment. To enable reliable and timely delivery of the life-critical sensory data, we will design situation-aware service differentiation mechanisms, where sensory data packets will be classified into different priorities dynamically in the system based on the patients' medical situations. Packets with different priorities will be handled differently by the packet scheduler and queue management in the network, so that critical packets with high priorities will be delivered with reliability assurance.

- <u>Privacy and Security Considerations</u> – 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 systematically evaluate these security measures and analyze their performance tradeoffs. Further we will investigate the detection mechanisms to identify anomaly situations caused by malicious attacks or device malfunctions. We will develop data and system behavior models for the sensory data and the monitoring system and apply classification algorithms to detect data out of range and system anomalies.

- Integration with Web-Based Patient Portal – In our system, the sensor data collection does not mark the end, but simply the beginning of data transport. Data from the patient monitoring sensor network will be collected to an EMR system and accessed from a web-based patient portal. To achieve this goal, we will connect our remote healthcare system with the EMR/patient portal prototype developed at Vanderbilt University based on a service-oriented design. The portal system allows care givers to efficiently access sensing data through a unified medical record system and allows for arbitration of access permissions for different parties (e.g. families, physicians, nurses, and insurance companies).

2.4.6 Trustworthy Systems

Thrust Leaders: *Alex Aiken (Stanford University, Mike Reiter (University of North Carolina, Chapel Hill), David Wagner (UC Berkeley)*

The Trustworthy Systems area of the TRUST center encompasses research addressing the full range of issues in trustworthy computing via securing software, securing hardware, and ensuring survivability of critical systems. During this reporting period, Trustworthy System research projects were focused in the following areas:

- Robust Software
- Security Policies
- Trustworthiness by Construction
- Platform Integrity
- Intrusion-Tolerant Systems

The activities of each project are discussed in more detail below.

Robust Software – Software robustness is a central problem in the construction of trustworthy systems. These projects address ways to eliminate software vulnerabilities to better enforce least privilege in software programs, and to compose software systems to provide robustness. Projects were focused in several areas:

- *Eliminating Software Vulnerabilities*: It is well-known that software errors are the source of numerous vulnerabilities. This area of research seeks to eliminate the errors and/or vulnerabilities through automated means. One project focused on formally verifying a number of properties of a large, real-world software system, the Linux kernel, which together would imply that the system cannot under any circumstances commit a class of undefined behaviors that would result in either a security hole or system crash. These properties include such things as verifying the absence of buffer overruns, null pointer dereferences, use of un-initialized variables, misuse of user/kernel pointers (this one is specific to Linux), and the absence of certain integer overflows, among others. In contrast to safe C compiler projects (e.g., George Necula's CCured) where the goal is to ensure that any undefined behavior is detected, the goal here was to prove that such errors do not arise in the first place. The same basic techniques could be used, for example, to show that uncaught exceptions do not arise in programs written in safe languages such as Java. In a second project, TRUST researchers developed new methods for detecting security vulnerabilities in commercial software using a technique known as symbolic execution. One TRUST researcher collaborated with researchers at

Microsoft Research to build a software testing tool called SAGE that proved effective at finding security vulnerabilities: it found dozens of important security vulnerabilities in several Microsoft products and is reportedly now internally used as part of the weekly build and testing process for at least one Microsoft product. A third project in this area researched methods of automating and performing vulnerability and exploiting analysis and defense in commercial off-the-shelf software, where source-code may not be available. In particular, this effort designed and developed novel techniques by employing program slicing, model checking, and other program analysis techniques to automatically identify whether a potentially vulnerable point can be reachable by un-trusted inputs, and then to automatically generate input-based filters to filter out malicious attacks and hardening mechanisms to protect vulnerable software from malicious incoming attacks. By enabling reachability analysis of un-trusted inputs, this effort can determine whether there exists an un-trusted input capable of exploiting a potential vulnerability in the software, and by automatically identifying the conditions under which a vulnerability can be exploited, this effort can automatically generate input-based filters that can filter out attack packets even for polymorphic worms. A fourth project developed more advanced static and dynamic techniques for finding security vulnerabilities in Java web applications. The project used existing model checking techniques such as Java Pathfinder to design, develop, and evaluate new algorithms and apply them to a large number of open-source Java web applications. The software will be made freely available so others can use the results as well as build upon our work. TRUST researchers continue to advance the state of the art in these areas and are working closely with industry partners to transition these ideas to commercial practice.

- *Enforcing Least Privilege*: A system satisfies the principle of least privilege if it possesses only the permissions it requires to perform its tasks. Unfortunately, today's systems do a poor job of supporting and implementing least privilege. For instance, when you run a mail client program, it inherits the power to read and write all files in your user account; this is far more than the mail client legitimately needs, and it means that an email worm can destroy all your files. One project developed a language, called Joe-E that will be familiar and accessible to programmers but that helps improve least privilege. To make Joe-E accessible, Joe-E is chosen to be a subset of Java. Joe-E builds on prior work on object capabilities (i.e., where a reference to an object represents a capability to affect that object) and the system is built so that this is the only way that code can get any kind of privilege. The goal of Joe-E is to bring object capabilities to a mainstream language, eliminating the need for programmers to learn a new language and thereby reducing barriers to adoption. In addition, Joe-E is intended to enable programmers to reason about the flow of privilege in the program, thereby enabling composition of modules into a larger system without putting security at risk. That is, a key goal is to support modular reasoning, so that a programmer who examines one module in isolation (along with the interface to all other modules that it calls) can reason about the set of privileges available to that module and about the trust relationships it has with other modules. The Joe-E project made significant advances over the past year. The researchers released version 2.0 of the language to the public under an open source license. This project is having impact outside of TRUST. For example, software developers at HP Laboratories have adopted Joe-E in one of their software projects. As a result, we have gained considerable experience with the strengths and weaknesses of this approach to securing software systems. Meanwhile, Google has launched a new project for web security that builds upon similar ideas.

- *Live Distributed Objects and their Robust Communication Infrastructure*: At Cornell University and Vanderbilt University, work is underway to develop a new way of building powerful edge-computing options that enable non-programmers to implement sophisticated distributed applications. A live object is a new lightweight component programming technology developed by Cornell for encapsulating event-driven applications in a simple way that enables composition with basic forms of type checking, reflection, and protection. We're using live objects to encapsulate all sorts of distributed functionality, notably data replication technologies needed for building collaboration applications, but also basic graphics containers such as the Windows XNA interfaces and basic office automation technologies such as databases and spreadsheets. With live objects, anyone who understands the drag-and-drop mindset can rapidly construct powerful collaboration tools and other forms of distributed functionality. For most purposes, no programming is needed at all—the skill level is similar to that needed to build web pages or documents. The connection to robustness and security arises because, once applications are built in this manner, we can automatically endow them with desired robustness properties. For example, we've constructed a powerful new fault-tolerance layer (Quicksilver) that can offer any of a range of consistency properties and encapsulated it as a live object. An end user who adopts the platform for its ease of use and productivity benefits thus gains fault-tolerance and security properties without even realizing that our platform automatically offers these kinds of guarantees. In joint work between Cornell and Vanderbilt, TRUST researchers are adding real-time properties to the platform. Live objects have attracted keen interest from corporate partners such as Microsoft, Intel, Red Hat, JP Morgan, and Raytheon and have also been briefed to high level executives at the U.S. Air Force. A number of papers have been written on this work; we're particularly excited about one that was conditionally accepted to appear at ECOOP 2008.

- *Enterprise Data Recovery and Business Continuity Solutions for the Financial Sector*: Cornell University is also pursuing work on new technologies for ensuring that banks and other critical financial infrastructure providers can ride out disruptions such as major power outages or environmental disasters, which often require that operations be shifted to a remote data center that wasn't impacted by the event. Our solutions include a new network appliance technology, Maelstrom (to be presented at NSDI 2008), that masks the latency associated with wide-area optical links in a manner that dramatically improves the performance of TCP when packet loss occurs. We're also working to transition a previously developed real-time event notification protocol, Ricochet (NSDI 2007), into the Red Hat community. We've used our own solutions to implement the Smoke and Mirrors File System, which continuously maintains a backup at a distance over a Maelstrom-equipped long-haul optical links. This work is demonstrating dramatically improved robustness to network faults and congestion and, for the first time, enabling unmodified applications to maintain backup data at a safe geographic distance. There has been tremendous interest in all three solutions by the financial industry and we are working with the Financial Services Technology Consortium (FSTC) and the Financial Services Sector Coordinating Council (FSSCC) to explore technology transition options.

Security Policies – To a first approximation, a trustworthy system is one that enforces desired security policies, and so security policy research is central to the trustworthy systems agenda. These projects distill and enforce security policies in a variety of settings. One project uses information flow to derive the access-control policy implemented by a program. Currently, most applications make use of access-control checks spread through out the code. The goal in this project is to develop a tool that aggregates such checks together into an access-control policy that could ease the transition to using a centralized policy. Users could examine the extracted policy and analysis engines could answer queries about it. Such tools could check if the extracted policy matches a specified policy. Even in the absence of a formal specification, change-impact analysis could be possible: given application code before and after edits, one could compare the extracted policies to ensure that no new security holes were introduced. A second project seeks to define, validate, and optimize a unified framework for QoS (including access-control) policy management that enables the predictability and resource control required by information management systems, while preserving the modularity, scalability, and robustness that's the hallmark of Service-Oriented Architecture (SOA) platform technologies. This includes approaches for converting user intent - in conjunction with a static/dynamic runtime environment - into QoS policies and building technologies that (1) enable the decentralized creation of access control policies for distributed resources and (2) exercise that authority efficiently when resources need to be accessed. A third project seeks to ensure that an authenticated user has access to only those services for which he/she has authorization. Web based resources available via Web Services are typically dynamic and distributed in nature and hence require adaptive authorization models that can keep pace with the dynamically changing security requirements of the target enterprise. The goal of this project is to develop an approach aimed at a more generalized and reusable solution which provides the flexibility to handle authorization rule updates in real time. A fourth project in this area is the Civitas secure voting system (Oakland '08), arguably the most secure electronic voting system yet constructed. Civitas provides both universal verifiability of election results and coercion resistance, yet does not require a single trusted voting terminal supplier—in fact, voters can vote with their home machines. The Civitas software is implemented in the Jif programming language, which was extended with two new kinds of information security policies: declassification policies, permitted controlled information release, and erasure policies, requiring removal of information from the system. These two kinds of policies enable the compiler to check important aspects of the security of the resulting system, such as the mandatory deletion of key shares.

Trustworthiness by Construction – The Swift system (SOSP'07) supports building web applications that are secure by construction: explicit security policies are used to drive the construction of web applications. In modern web applications, some application functionality is usually implemented as client-side code written in JavaScript. Moving code and data to the client can create security vulnerabilities, but currently there are no good methods for deciding when it is secure to do so. Swift automatically partitions application code while providing assurance that the resulting placement is secure and efficient. Application code is written in the Jif programming language, which supports Java-like code annotated with information flow policies that specify the confidentiality and integrity of web application information. The compiler uses these policies to automatically partition the program into JavaScript code running in the browser and Java code running on the server. To improve interactive performance, code and data are placed on the client side. However, security-critical code and data are always placed on the server. Code and data can also be replicated across the client and server to obtain both security and performance.

Platform Integrity – Platform integrity refers to techniques to validate a computing platform or to limit users' dependencies to those properties that can be validated. This includes validating the software platform (or its properties) running on a host (also often referred to as "trusted computing"), or validating that a component encountered in an unfamiliar environment can reasonably be trusted for a limited purpose. One project in this space seeks to build system infrastructure for trustworthy computing spanning basic research in operating systems, cryptography, and distributed systems. This project is centered on the construction of a new operating system called Nexus that will provide new abstractions and mechanisms for trusted computing. The Nexus will provide strong isolation, reduce application TCB, and support the principle of least privilege. It will also provide higher-layer programming abstractions that virtualize the primitives offered by the secure coprocessor. A second component of this project is to integrate privacy-preserving attestation into Nexus. This type of attestation provides the same assurances as traditional hash-based attestation with signatures but without revealing the identities of the hosts and without enabling a third party to link together independent activities performed by a given node. A third component of this effort involves using Nexus to wrap a process inside another, track all inputs to and outputs from the encapsulated process and ensure via active attestation through a reference monitor that the process is behaving correctly (i.e., that outputs are legitimate given the set of inputs). Finally, this project seeks to develop an application-oriented security policy language and enforcement structure to capture higher-level security policies and ensure that they are correctly mapped to the available primitives. Included in this are uses of attestation in connection with data collection and provenance (e.g., so that data can be reliably "timestamped" upon its collection and its credibility can be evaluated based on what influenced it). A second project in this area focuses on increasing the security of mobile computing environments, focusing on two specific challenges: (1) simple and secure trust establishment in local environments, and (2) execution of un-trusted components in isolated execution environments. The first of these refers to developing techniques to help users identify what components (e.g., base stations, printers) in an unfamiliar environment should be trusted. The second involves mechanisms to limit the effects of using components (and, e.g., the drivers they require) when their trustworthiness cannot be established, using virtualization and isolation technologies. A third project in this area seeks to develop a System-on-a-Programmable-Chip (SoPC) implementation of a trustworthy hardware platform that provides software protection against malicious attacks. With the programmable nature of FPGAs, several techniques can be evaluated in isolation or in combination for tunable levels of security (e.g., watermarking, cryptographic algorithms). The programmability also allows implementation of stronger encryption techniques in future systems. In addition, processor cores from different vendors (e.g., Nios II from Altera, MicroBlaze from Xilinx) are being investigated to evaluate the performance impact of various levels of security. Other tradeoffs being studied are the choice between hard microprocessor cores and soft microprocessor cores and the use of multiple processor cores. This project benefits from collaboration with other TRUST members to provide contextual applications that quantify the security benefits of FPGAs. A fourth project constructs a more robust, secure and flexible operating system by "deconstructing" a modern operating system using micro-kernel principles. Instead of using the traditional approach to micro-kernel construction of designing a small, elegant micro-kernel and constructing an operating system out of multiple protected subsystems, the project starts with a trusted virtual machine monitor capable of running a modern operating system and then rips major subsystems out of the operating system to run in specialized virtual machines on the same platform. The project extracts the OS components responsible for external communication including the file system, networking stacks, and user interface. These changes will result in an architecture similar to a 1980s micro-kernel, but one that is compatible with today's software environments.

Intrusion-Tolerant Systems – "Intrusion tolerance" refers to utilizing cryptography and/or distribution in the implementation of a service so that the service will retain desired properties despite the hostile corruption of components implementing the service. This area is itself very broad, including work in, for example, secure multiparty computation and Byzantine fault-tolerant protocols. These techniques have been used to construct experimental services implementing secure key distribution and certification, secure DNS, secure file systems, and even secure electronic voting systems. During this reporting period, our progress in this area focused on the development of new, more scalable approaches to the design and implementation of intrusion-tolerant services. Intrinsic to most such services is accessing subsets ("quorums") of servers in the course of issuing requests to the service, and scalability is improved by shrinking the sizes of those quorums. At the same time, however, decreasing quorum sizes tends to also decrease the intrusion-tolerance of these systems. In order to break this tension, we explored probabilistic quorum systems in which quorums of servers have needed intersection properties with high probability (versus with certainty, as is achieved in traditional approaches). By permitting a small and quantifiable possibility of error, we have shown that we can dramatically improve the scalability and fault-tolerance of intrusion-tolerant services simultaneously. Our work this past year has been foundational, focusing on the theory underlying this technique. We plan to transition to empirical studies of this approach during this coming year.

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.

To increase the chances that TRUST research results are relevant and to maximize the Center's impact, beginning next reporting period we will reorganize the Center research activities around three target application areas. These application areas were selected because each emphasizes fundamentally different kinds of trustworthiness problems from the others and each is poised to make increasing use of networked computing, but trustworthiness issues could be a substantial impediment to success if not adequately addressed.

For each research area, there are current research topics that support one or more of the three areas (e.g., trusted operating systems, reliable computing, languages and tool support for writing secure code) , and TRUST will continue in our investigations of these. As such, the new TRUST research areas build on prior achievements but will introduce new topics to encourage innovative, novel project ideas from the TRUST research community and knowledge transfer opportunities from external TRUST partners and collaborators.

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)*

Ordinary people are more and more conducting business and managing their assets using the Internet. These applications are largely organized as client-server systems: there's some server-farm run by a financial organization, and individuals use desktop computers (typically at home) to access these servers via the web. The web browser provides the lingua franca for these interactions, with email a distant second.

One taxonomy of security enforcement mechanisms speaks in terms of the "gold standard" because the first letters of each element are the atomic symbol for gold (Au): Authentication, Authorization, and Audit. Their relative importance varies, depending on the application. Inadequate means of Authentication is responsible for much of the identity fraud and theft seen today. Not only must the financial institution authenticate the customer, but the customer must also authenticate the financial institution to avoid phishing, pharming, etc. This is a rich problem space where TRUST researchers have already made some progress—but there is more to be done. We expect the client-server structure (with its asymmetry in processing and with its needs to support scalability over numbers of clients) will enable certain kinds of solutions but preclude others. Also, studying this problem area invite consideration of physically secured machine-room computing; it's a very different world from large networks of badly managed desktop computers.

In this application area, integrity and confidentiality are primary requirements, though availability and the need to store data in ways that resist catastrophic physical disruption and attacks are also quite important. In terms of specific research topics, this area will focus on improving the design of core systems applicable to financial infrastructure with scalable intrusion-tolerant distributed systems and reliable, fast transaction processing and event notification; developing design and construction principles for secure web systems that protect web content from malicious interference and address issues of secure human computer interfaces; developing principles for secure and reliable network infrastructure including exploration of trusted computing platforms and secure network enforcement and analyzing the security of network protocols; development of algorithms and tools for code analysis, monitoring and malware detection including automated error detection, symbolic execution, intelligent fuzzing, and botnet detection and mitigation; and conducting public policy studies and users studies in the area of computer security to examine risk management of computer security, including analysis of security breach notices, measuring user perception and personal information, and studying rationality, risk, and interdependent security.

2.6.2 Health Infrastructures

Thrust Leader: *Janos Sztipanovits (Vanderbilt University)*

It's common practice today for billing and other financial records for medical procedures to be automated. The next step will be computerization of medical histories—test results, imaging, and consultations with medical professionals. Ideally, this medical history would be available when and where it is needed and presented to medical practitioners in an integrated way.

Because medical research benefits from access to real medical records, there is much to be gained by supporting data mining of these medical records. To do so, however, requires protecting patient privacy—and there area a number of fundamental research challenges that need to be addressed to achieve this. The implementation of authorization is particularly challenging in the medical records setting. For example, whether a medical practitioner is authorized to see an individual's records could depend on:

- Area of expertise (e.g., a podiatrist might be denied access to a patient's psychological records)
- Previously established long-standing trust agreements (e.g., a patient's primary doctor can see his records but a random internist can not)
- Ad-hoc, short-term trust agreements (e.g., granting a specialist access to those part of a patient's records relevant to a particular malady)
- Expedience (an ER gets access to a patient's records whenever they are treating him).

These and other scenarios illustrate that even rich authorization structures such as "role based access control" are inadequate. Additionally, privacy is a paramount concern in this space. Moreover, there is much to understand vis-à-vis interactions between the law and technology because there is so much law that applies to medical information.

Finally, there are challenging research problems due to physical distribution and heterogeneity. Medical records might not be stored in a single computer but rather each record might be stored on behalf who has created it (e.g., x-rays at the hospital, consultation summary at the doctor's office) and each record is most likely in a different format, which creates interoperability and engineering challenges.

2.6.3 Physical Infrastructures

Thrust Leaders: *Steve Wicker (Cornell University)*

The focus of this area is Supervisory Control and Data Acquisition (SCADA) and other networked systems that control existing physical infrastructures (e.g., the power grid, natural gas distribution, automated railroad control) and more futuristic enterprises such as "smart buildings" and "smart structures" (e.g., active-bridges whose structural integrity depends on dynamic control or actuators to replace static brute-force physical structure).

Unlike the other research areas, there is virtually no legacy system problem to overcome with physical infrastructures. Thus, we are not constrained by existing deployed systems and, when we are successful, it should be easy to transition our work into real deployments. We therefore are free to contemplate new system software and new network architectures and protocols.

Security requirements are traditionally enumerated in terms of confidentiality, integrity, and availability. In this applications area, confidentiality is not nearly as important as the other two dimensions. Moreover, availability is probably too weak—real-time bounds must be satisfied which, for example, changes the solution space for defending against denial of service attacks.

Privacy issues also arise in this area, but in subtle ways. Most people don't think about what could be inferred from their uses of infrastructure since information is revealed in indirect ways (e.g., increased power draw implies somebody is at home). Moreover, when distributed networks of sensors are widely deployed, then opportunities for privacy abuse arise—again, through abuse of information that is being collected (presumably) for other reasons. These will present a number of challenges for TRUST researchers but the Center has established a number of multi-disciplinary, cross-campus teams that are poised to address these issues.

3 EDUCATION

3.1 Goals and Objectives

One of the drivers of this Center is the view that concerns regarding security must be consciously engineered into new and legacy critical infrastructure systems, and that to do so requires a rethinking of every component level of the system. To ensure that these concerns are shared and addressed by the next generation of computer scientists, engineers and social scientists, TRUST researchers will incorporate their findings and methods wherever possible into the standard. Thus, this project will result in a broad curriculum reform of existing computer science and engineering courses. We will develop a whole set of courses from the lower division to the advanced graduate level as the research on trust matures.

The center has distinct education constituencies – both undergraduate and graduate programs – for which there are distinct mechanisms for knowledge dissemination. For undergraduates, the center has adopted a two-pronged approach. On the one hand, the center will have activities concerned with diffusing ideas of trustworthiness throughout the entire undergraduate curriculum. On the other hand, the center needs is working towards defining a modern “standard” computer security course at the undergraduate level. For graduate students, the center finds that a series of summer workshops or seminars on specific disciplines is where a significant impact can be made, in addition, of course, to developing topic specific customized courses. The summer seminars are to be 1-week courses, where research leaders provide intensive short courses in areas of current research interest.

Beyond the above partition, the realization that **TRUST solutions = policy options + technology options** requires TRUST to bring together two communities of researchers: *technology* researchers and *policy* researchers. Technology done independent of policy risks irrelevance; policy done independent of the technology risks obsolescence or suppresses options.

From the marriage of policy and technology arises some horizontal partitions in addition to the ones by education level, and the TRUST center will engage the educational community to work towards:

- A broader understanding of TRUST *technology* options as such among (future) *technologists*
- A broader understanding of TRUST *technology* options as such among (future) *policy shapers*
- A broader understanding of TRUST *policy* options as such among (future) *policy shapers*
- A broader understanding of TRUST *policy* options as such among (future) *technologists*.

The center strategy for achieving this broad influence is through a combination of *push* and *pull* tactics: to *generate learning material* (such as learning modules, course syllabi, textbooks, broader curricula), provide *effective dissemination structures* (such as on-line repositories, internet delivery mechanisms, summer seminars and workshops, center-wide seminar series), and establishing *broad educator communities* (such as summer schools, education conference participation) that engage with the center in adopting and adapting the results of the center to their instructional context.

Specifically, the TRUST Objectives in Research are to establish the following:

a) Learning Technology Infrastructure
b) Undergraduate Programs: generate best-practices material for computer science courses, security modules for other engineering programs and the social sciences, create a signature new undergraduate trusted system course, capstone experience for undergraduates
c) Graduate programs: specialized material for both engineering and policy
d) TRUST Summer seminars for Students, for Industry, for Instructors and for Researchers
e) A recurring and significant presence at key education conferences
f) A series of TRUST domain workshops.

During this past year the center-wide activities in the education area have focused on c, d, e and f with a ramping up of the efforts related to a and b: on establishing the infrastructure for the learning modules repository, and on establishing a set of pilot course modules within this repository, bringing together material from the various TRUST partner institutions in an integrative learning material generation exercise. Summer programs offered during 2007 included: SECuR-IT at Stanford, WISE at UC Berkeley, SUPERB-IT at Berkeley, Information

Assurance Capacity Building Program (IACBP) and the Curriculum Development in Security and Information Assurance (CDSIA) were held at San Jose State University.

3.2 Performance and Management Indicators

During the second-half of 2007, a compressive review and evaluation of TRUST Academy Online (TAO) was performed. This review was fundamental to the portals 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 thrust and course materials. Additional user feedback with surveys will help refine the portal's technology and user functions.

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

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 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. These materials are 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. This approach was pioneered by the NSF VaNTH Engineering Research Center for Bioengineering Educational Technologies, and the TAO Portal is a reuse and refinement of technology originally developed for the VaNTH Portal using a public domain portal framework called Plone.

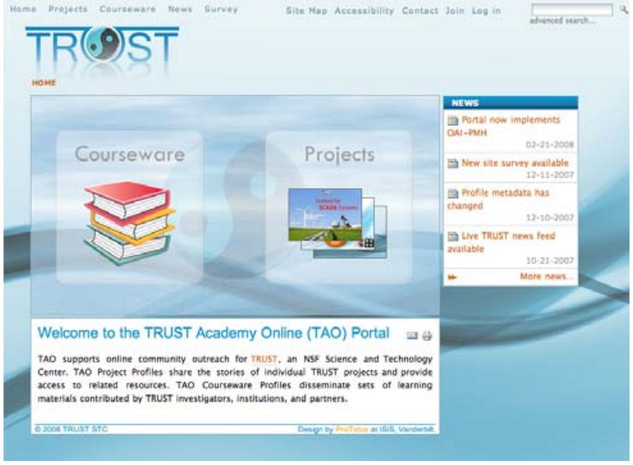

Portal now implements OAI-PMH 02-21-2008

New site survey available 12-11-2007

Profile metadata has changed 12-10-2007

Live TRUST news feed available 10-21-2007

Welcome to the TRUST Academy Online (TAO) Portal

TAO supports online community outreach for TRUST, an NSF Science and Technology Center. TAO Project Profiles share the stories of individual TRUST projects and provide access to related resources. TAO Courseware Profiles disseminate sets of learning materials contributed by TRUST investigators, institutions, and partners.

Figure 1: The TAO Portal Front Page

An important expansion of outreach objectives for the TAO was undertaken during this year. It was motivated by recognizing that while TAO's provision of learning materials is strategic, since it targets a community (educators) with high potential impact, the actual impact will likely be felt gradually as attitudes among educators and administrators admit greater roles for reused or adapted units and materials in curriculum designs. To broaden the impact of TAO, and to make it more immediate, we saw an opportunity to extend our view of potential “consumers” to include all those interested in the center’s individual projects and their particular emphases, technical strategies, and outputs.

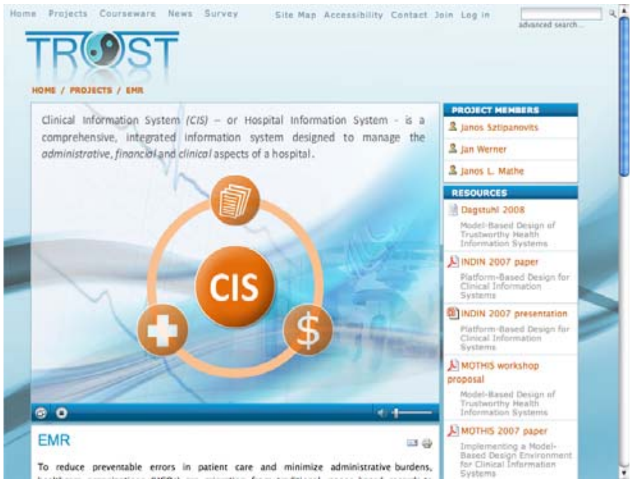
Home
Projects
Courseware
News
Survey
Site Map
Accessibility
Contact
Join
Log in
advanced search...
TRUST
HOME / PROJECTS / EMR
PROJECT MEMBERS
Janos Sztipanovits
Jan Werner
Janos L. Mathe
RESOURCES
Dagstuhl 2008
Model-Based Design of Trustworthy Health Information Systems
INDIN 2007 paper
Platform-Based Design for Clinical Information Systems
INDIN 2007 presentation
Platform-Based Design for Clinical Information Systems
MOTHIS workshop proposal
Model-Based Design of Trustworthy Health Information Systems
MOTHIS 2007 paper
Implementing a Model-Based Design Environment for Clinical Information Systems
CIS
$
EMR

Clinical Information System *(CIS)* – or Hospital Information System - is a comprehensive, integrated information system designed to manage the *administrative, financial* and *clinical* aspects of a hospital.

To reduce preventable errors in patient care and minimize administrative burdens,

Figure 2: A TAO Project Profile

To enact this expansion strategy, we have added a new content type, called “project profiles,” to the TAO. These profiles aim to tell the stories of TRUST research projects and to provide access to associated resources such as papers, presentations, and posters. By “stories” we mean: what problems are being addressed, why these problems are important, how the problems are being pursued, what is the intended impact, who would benefit, how they would benefit, etc.

We have chosen “visual storytelling” as the vehicle for this communication. In project profiles, lightweight multimedia shorts are being used to quickly present the essential details of a project’s work in a way that is accessible to a broad audience and enjoyable. During the year we have used a small group of TRUST research projects at Vanderbilt to “prototype” and refine this concept. TAO media designers have collaborated with project graduate students to identify story elements and then produce these multimedia resources. The profiles can then be established and populated by the projects themselves using the same portal features that allow projects to populate courseware profiles. Given that TRUST projects comprise a fairly stable portfolio, we consider this strategy is scalable to incrementally include all TRUST projects, resulting in a rich information flow.

Accompanying this extension in audience, we have undertaken during the year to enhance the user experience on TAO. A keystone element in our strategy is 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. We have retained the tabular, text-based browser of the courseware profiles as a navigation alternative.

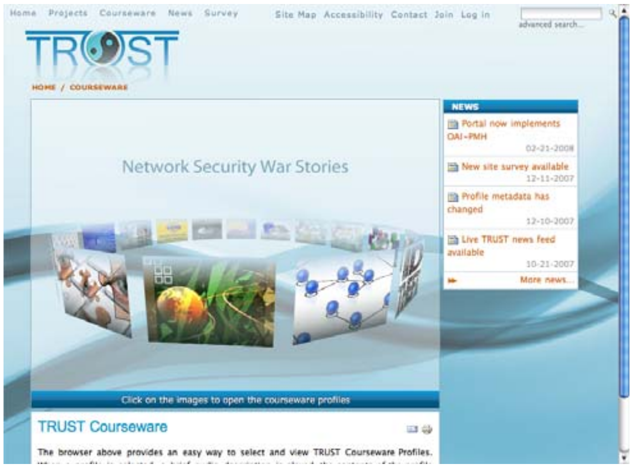
Home
Projects
Courseware
News
Survey
Site Map
Accessibility
Contact
Join
Log in
advanced search...
TRUST
HOME / COURSEWARE
NEWS
Portal now implements OAI-PMH
02-21-2008
New site survey available
12-11-2007
Profile metadata has changed
12-10-2007
Live TRUST news feed available
10-21-2007
More news...
Network Security War Stories
Click on the images to open the courseware profiles
TRUST Courseware

The browser above provides an easy way to select and view TRUST Courseware Profiles.

Figure 3: The Visual Browser for TAO Courseware Profiles

Many other features addressing usability and feedback have been introduced to the portal this year. Profile resources, previously incorporated into the bodies of profiles, are now externalized in a sidebar listing for easy access. Profile contributors and participants are similarly presented, and additional information about these members is provided. Profiles now elicit feedback from visitors, as does the entire site. We have improved the portal news service and have incorporated access to news from the TRUST web site via RSS. Overall the site has been given a fresh "look and feel" making extensive use of adaptability features of the Plone framework, demonstrating its versatility.

Finally, we have acted on recommendations provided to us by last year's Site Visit Team. First, we have simplified and adapted the classification metadata employed for the portal's courseware profiles. Feedback from visitors will be important to continuing this process, which was among the motivations for incorporating new feedback mechanisms. Second, we have made metadata from the portal accessible to archive cataloging services by implementing the OAI-PMH protocol. We anticipate that the reflection of such catalogs back to popular search engines will further improve detection of portal contents. Together these many extensions and improvements underscore our commitment to making the TAO a rich and effective dissemination vehicle for TRUST.

Activity Name	TRUST Learning Modules
Led by	Kristen Gates (UC Berkeley)
Intended Audience	TRUST portal users: students, faculty and Industry Professionals
Approx Number of Attendees (if appl.)	N/A

The TRUST Academy Online (TAO) is an online repository for TRUST Learning Modules. Accessible by the public, the TAO contains leading-edge learning materials available at no cost. By using these modules, educators have access to leading-edge research and teaching materials specific to trusted systems technology and policy issues.

The purpose of the modules is to create learning materials that are assessable via the TAO portal and used by teaching faculty as course content, lecture material, support materials, for a computer science or related higher education courses.

The modules represent a variety of learning materials and include: PowerPoint slide decks, lecture notes, case studies, assignments, related web site links and video clips.

Building on our second-year inventory of 17 modules, the TAO gained 25 learning modules during Year-3 in the following research thrust.

Year-3 Learning Modules	Category
Learning Care Provider Teams and Workflows from Electronic Medical Records Access Logs	EMR
Model-Integrated Clinical Information Systems	EMR
End User Security: From the Browser to Forensics	End User Security
Scaling and Evaluating Cluster Bro using DETER	Network Defense
RFID in Public Identification	Policy
Fair and Accurate Credit Transactions Act of 2003 (FACTA) Access Study	Policy
Is Increased Cyber Security Compatible with Other Policy Values?	Policy
Privacy, Compliance, and Risk Management	Policy
Technologies of Compliance	Policy

Year-3 Learning Modules	Category
Understanding and Improving Organizational and Individual Approaches to privacy and security for trustworthy systems	Policy
Towards a Workable Liability Framework for Honeyfarm and Botnet Research	Policy
Secure SCADA through Robust Estimation, Control, and Detection	Secure Sensor Networks
Security and Privacy in Microgrid SCADA	Secure Sensor Networks
Foundations of Intrusion-Tolerant Services	Trustworthy Systems
A Unified Framework for Trustworthy QoS Policy Management	Trustworthy Systems
Distributed Information Flow using Trusted Network Interface Cards	Trustworthy Systems
Symbolic Execution for Security	Trustworthy Systems
Joe-E: A Capability-based Programming Language for Security	Trustworthy Systems
Static Analysis of HiStar	Trustworthy Systems
Information Flow Inference and Visualization	Trustworthy Systems
Automatic Exploit Generation	Trustworthy Systems

Table 1: Learning Modules Inventory

Learning modules were created for the six research trust in addition to Education and Outreach and identified as 1) Electronic Medical Records (EMR); 2) End User Security; 3) Network Defenses; 4) Policy; 5) Secure Sensor Networks; and 6) Trustworthy Systems.

Year-3 TAO Modules	Modules
Education	4
Electronic Medical Records	2
End User Security	1
Network Defense	1
Policy	7
Secure Sensor Networks	2
Trustworthy Systems	8
TOTAL	25

Table 2: Learning Module by Category

Activity Name	Women's Institute in Summer Enrichment (WISE)
Led by	Kristen Gates (UC Berkeley)
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. Focused recruitment effort toward underrepresented minority groups and women.
Approx Number of Attendees (if appl.)	23 participants with 12 speakers

WISE is a 1-week residential summer program on the University of California, Berkeley 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. The program date was June 10-15, 2007.

Summer 2007, the program topic was sensor networks with a healthcare and policy emphasis and topics included:

- 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 2007 Seminar Speakers were:

- Terry Benzel, USC – Information Science Institute
- Ruzena Bajcsy, UC Berkeley – TRUST
- Deborah Estrin, UCLA – Center for Network Sensing
- Stephanie Forrest, University of New Mexico
- Jennifer Hou, University of Illinois
- Jennifer King, UC Berkeley – TRUST
- Maryanne McCormick, UC Berkeley – TRUST
- Deirdre Mulligan, UC Berkeley – TRUST
- Priya Narasimhan, Carnegie Mellon University
- Diana Smetters, Palo Alto Research Center (PARC)
- Dawn Song, Carnegie Mellon University – TRUST
- Yuan Xue, Vanderbilt University – TRUST

Tuition for WISE 2007 is $2,500; however, NSF-TRUST fellowships are available to US professors, post-doctoral fellows, and Ph.D. candidates studying at US universities. A maximum of 20 fellowships with travel stipend will be awarded.

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. 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.

Program Evaluation: Each WISE fellow completed a program evaluation. WISE participants will be tacked over a several year period to evaluate the programs' impact on educational, professional development, job placement and retention.

This was the second-year WISE has been hosted at the UC Berkeley campus; WISE 2008 will be hosted by Cornell University. An evaluation of first-year WISE participants was conducted with a follow-up survey scheduled for years one, three and five. Recommendations from the 2006 survey were put into place for the 2007 program. The WISE 2007 cohort was surveyed at the end of the program and at years one, three and five. Tracking of the WISE cohorts will 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 participants.

Activity Name	Summer Undergraduate Program in Engineering Research at Berkeley-Information Technology (SUPERB-IT)
Led by	Shankar Sastry (UC Berkeley)
Intended Audience	Undergraduate students, underrepresented minority groups and women
Approx Number of Attendees (if appl.)	2

The Summer Undergraduate Program in Engineering Research at Berkeley - Information Technology (SUPERB-IT) in the Electrical Engineering and Computer Sciences (EECS) Department 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 by affirming students' motivation for graduate study and strengthening their qualifications. The program provides students with the opportunity to gain research experience by participating in research projects with engineering faculty and graduate students. Upon completion of this program 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 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.

The SUPER-IT program is a nine week program—the program dates were June 10 – August 3, 2007. In 2007, SUPERB-IT had two students participating in TRUST related research topics. Each student was given a $3,750 stipend for the period, travel allowance, and provided on-campus housing. In addition to the undergraduate research experience, SUPERB-IT students participate in educational activities including lab tours and industry field trips. Graduate school advising and subsidized GRE prep course is also included.

Program Evaluation: The students are evaluated at midterm and at the end. They also report on their progress at the regular weekly meetings. They receive feedback on their work at the weekly meetings and after the midterm from the faculty advisors.

The students evaluate the program at the end of the program, using a questionnaire. The results of this survey are distributed to participating faculty and graduate students and used as feedback for program development. SUPERB-IT participants are tacked overtime to identify those students considering graduate school and those that have been accepted into graduate school programs.

For the Summer 2008 program, TRUST has six students participating in the SUPERB-IT program. Summer 2008 will also take a novel and innovative approach by creating a thematic cohort were students will be building a research prototype of a tool for finding security bugs in desktop applications.

Activity Name	Information Technology in Society "Trustworthy Systems: the societal/ethical impact of information technology applications"
Led by	Maryanne McCormick (UC Berkeley) and Ruzena Bajcsy (UC Berkeley)
Intended Audience	Undergraduate majors in computer science and engineering
Approx Number of Attendees (if appl.)	24 per semester

New Computer Science Course

University of California, Berkeley approved the new course called "Trustworthy Systems: the societal/ethical impact of information technology applications" and will be offered Fall 2008.

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.

Activity Name	Network Security CS285
Led by	Yuan Xue (Vanderbilt University)
Intended Audience	Undergraduate majors in computer science and engineering
Approx Number of Attendees (if appl.)	16 Fall 2007 semester

New Computer Science Course

Vanderbilt approved the new course "Network Security: CS285".

CS285 is a course for senior undergraduate and graduate student. This course was first developed in Fall 2005 as a special topic class (numbered CS291) and offered again in Fall

2006. Based on its positive feedbacks and class enrollment, the EECS department has recommended it to be changed into a regular course in Fall 2007. 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, authentication), practice of securing systems (PGP, IPSec, SSL), and recent research topics in network security. This course extends the student's knowledge on computer networks and improves their problem solving of security issues with programming.

Activity Name	TRUST Seminar Series
Led by	Kristen Gates (UC Berkeley), Annarita Giani (UC Berkeley) and 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 fall 2007. The program is a weekly event on the University of California, Berkeley campus. The fall 2007 series hosted 13 speakers with a total attendance of 715 participants. The spring 2008 series will host 14 speakers with a projected attendance of 640. The event is focused toward graduate students in computer science, industry professionals and campus community at large.

3.5 Professional Development Activities

TRUST students are active 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.

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

- Cornell University Research Exchange
- International Association of Privacy Professionals (IAPP), Certified Information Privacy Professional (CIPP) Training
- Internship at Google as a member of the application security team.
- Samuelson Law clinic–privacy-related research
- SECuR-IT Internship, Sun Microsystems, San Mateo, CA
- Tisdale Fellowship at Dell's Government Relations office in Washington, DC

TRUST students have membership in the following organizations:

- ACM: Association for Computing Machinery
- California State Bar Association
- HKN: Eta Kappa Nu National Electrical Engineering honor society
- IEEE: Computer Society
- IEEE: Institute of Electrical and Electronics Engineers
- WICSE: Women in Computer Science and Electrical Engineering

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

- ACM Computer and Communications Security, Alexandria, VA

- ACM Conference on Wireless Network Security (WiSec), Alexandria, VA
- ACM Symposium on Information, Computer and Communications Security, Tokyo, Japan
- ACM/IEEE International Conference on Model Driven Engineering Languages and Systems (MODELS 2007), Nashville, TN
- Advances in Neural Information Processing Systems (NIPS), Vancouver, B.C., Canada
- American Bar Association and the Association for Computing and Machinery Computing and the Law Conference in San Francisco, CA
- American Medical Informatics Association, Biomedical & Health Informatics (AMIA 2007), Chicago, IL
- Annual Workshop on Flow Analysis, Flocon 2008, Savannah, GA
- Center for Information Technology Research in the Interest of Society (CITRIS), London, England
- Collaborated on a joint DARPA proposal in with Lockheed Martin's Advanced Technology Lab
- DIMACS Workshop on Information Security Economics, Hanover, NH
- EEE Global Telecommunications Conference, Washington, DC
- eFraudNetwork Conference, Boston, MA
- Electronic Crimes Task Force, DHS, United State Secrete Service, San Jose, CA
- Grace Hopper Celebration of Women in Computing, Orlando, FL
- Hybrid Control, HYCON, L’Aquila, Italy
- Identity Theft Technology Council (ITTC), DHS–SRI International, San Mateo, CA
- IEEE Global Communications Conference (IEEE GLOBECOM'07), Washington, DC
- IEEE International Computer Software and Applications Conference, Beijing, China
- IEEE International Conference on Computer Vision, Rio de Janeiro, Brazil
- IEEE International Conference on Image Processing in San Antonio, TX
- IEEE International Conference on Industrial Informatics (INDIN 2007), Vienna, Austria
- IEEE International Workshop on Software Patterns: Addressing Challenges, Beijing, China
- IEEE Real-Time and Embedded Technology and Applications Symposium 2008, St. Louis, MO
- IEEE Symposium on Security and Privacy, Oakland, CA
- Intellectual Property Scholars Conference (IPSC), DePaul University School of Law, Chicago, IL
- International Association of Privacy Professionals (IAPP) Privacy Academy 2007, San Francisco, CA
- International Association of Privacy Professionals (IAPP) Privacy Summit 2008, Washington, DC
- International Conference on Pervasive Computing Technologies for Healthcare, Tampere, Finland
- International Workshop on Model-Based Trustworthy Health Information Systems (MOTHIS), Nashville, TN
- Medical Device Plug-and-Play Interoperability (MD PnP'07), Boston, MA
- NIPS workshop on "Machine Learning in Adversarial Environments, Whistler, B.C., Canada
- NSF Data Confidentiality Workshop, Arlington, VA

- Power Systems Conference: Advanced Metering, Protection and Control, Communication and Distributed Resources, Clemson, SC
- Privacy Enhancing Technologies Symposium (PETS 2008), Leuven, Belgium
- Richard Tapia Celebrating Diversity in Computing Conference, Orlando, FL
- Symposium at CISCO Inc, Milpitas, CA
- Symposium on Usable Privacy and Security (SOUPS), Pittsburgh, PA
- Technology audit of San Francisco, CA city surveillance camera systems in conjunction with the Samuelson Law Clinic at UC Berkeley
- Telecom Policy Research Conference, George Mason Law School, Arlington, VA
- Telecommunications Policy Research Conference (TPRC), The National Center for Technology & Law, George Mason University School of Law, Arlington, VA
- The Joint Conference on High Confidence Medical Devices, Software, and Systems (HCMDSS'07)
- TRUST workshop, Curriculum Development in Security and Information Assurance, San Jose, CA
- TRUST workshop, Information Assurance Capacity Building Program, San Jose, CA
- UC Berkeley Center for Law and Technology/TRUST June '07 Privacy Workshop, Berkeley, CA
- Usability, Psychology, and Security 2008 (UPSEC), San Francisco, CA
- USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET), SF, CA
- W2SP 2007: Web 2.0 Security and Privacy 2007, Oakland, CA
- Women's Institute in Summer Enrichment (WISE), Berkeley, CA
- Workshop on the Economics of Information Security, CMU, Pittsburgh, PA

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 attended conferences, supported workshop, personal development courses, and social and professional memberships. Professional development activities support student development of cross-domain and multi-domain knowledge, professional development, student success, and retention–all of which benefit TRUST and the student learning experience and impact provided by the Center.

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	Information Assurance Capacity Building Program (IACBP)
Led by	Sigurd Meldal and Mark Stamp (San Jose State University)
Intended Audience	Participants of the 2006 and 2005 IACBP at CMU
Approx Number of Attendees (if appl.)	16

Information Assurance Capacity Symposium is outreach to Hispanic Serving Institution (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. Participants first attend a one-month summer school at Carnegie Mellon followed by a symposium at San Jose State the next summer. The symposium date was June 14-15, 2007 and all participants in the 2005 and 2006 summer schools were invited.

The symposium (1) showcased the participants' achievements after the summer school, (2) further updated their expertise, and (3) reinforced connections with industry. Program materials generated by this program were cataloged on the TAO portal.

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.)	35

On April 18, 2008 TRUST organized the first 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.

The CDSIA 2008 had 35 participants registered, from 16 universities (14 of the 23 universities of the CSU), half of these universities are HSI institutions, and the remainders are all Associate members of the Hispanic Association of Colleges and Universities (HACU).

Four TRUST partner institutions (San Jose State (host), Stanford, UC Berkeley and Vanderbilt) participated in CDSIA 2008.

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.)	10

SECuR-IT is a ten-week residential program with paid internship co-located at Stanford and San Jose State. The program dates were June 3 – August 10, 2007.

SECuR-IT is a Graduate Student Academic Immersion with Internship Program. 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, 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
- Ten week, paid 40-hour per week internship.

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, and Policy and Legal Compliance. Participating TRUST industry partners supporting this program were: Ebay, Sun Microsystems, Yahoo, and Rapport.

This is a 40-hour per week obligation to internship, research, and learning activities. Students who participate in SECUR-IT should view this program as a full-time summer experience and are required to participate in San Jose State residential cohort, attend courses, and be employed as an intern by a participating SECuR-IT industry partner. Internships are paid employment and student housing at San Jose State University was provided.

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

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

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 International Identity Theft Technology Council (ITTC)
Led by	John Mitchell (Stanford) Liaison to ITTC
Intended Audience	Academics and Industry Professionals
Approx Number of Attendees (if appl.)	120 for the February 13, 2007 meeting

The DHS-SRI International Identity Theft Technology Council (ITTC) is a working forum where experts and leaders from the government, private, financial, IT, venture capitalist, and academia and science sectors come together to address the problem of identity theft and related criminal activity on the Internet.

These quarterly workshops, look to identify proactive IT security solutions and assist in the acceleration of its development and deployment into the market place. Seasoned IT security, 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 International Identity Theft Technology Council.

Activity Name	Trustworthy Interfaces for Passwords and Personal Information (TIPPI)
Led by	Dan Boneh (Stanford)
Intended Audience	Academics and Industry Professionals
Approx Number of Attendees (if appl.)	60

Despite tremendous advances in computer technology in general and information security in particular, users still typically provide personal information and credentials such as passwords the same way they did 30 years ago: through a text interface that they assume they can trust. Today, that trust assumption clearly can no longer be relied on.

Many security protocols have been proposed to protect credentials and personal information, but few are used in practice. A major reason is that the protocols have not been implemented in a way that ensures that they are actually used. For instance, a rogue Web site can still just ask the user for her password, regardless of how sophisticated a protocol the correct site employs.

The purpose of the workshop is to facilitate an effective solution to these problems by bringing together the designers of the cryptographic protocols with the implementers of the user interfaces. The third TIPPI workshop was held on June 22, 2007 at Stanford University. TIPPI brings together academic researchers and industry personnel in a forum for sharing ideas. The TRUST Center benefits from the workshop outputs in the forms of research papers and presentations and modules for the TAO Portal.

Activity Name	IT Security Entrepreneurs' Forum (ITSEF)
Led by	John Mitchell (Stanford)
Intended Audience	Academics and Industry Professionals
Approx Number of Attendees (if appl.)	210 for the March 11, 2008 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 is brought forth during a critical time as attacks and emerging threats continue to increase in sophistication and frequency against our nation's IT and Telecommunication critical infrastructures. The time is now to accelerate the search and implementation of "best of class" solutions that are being overlooked into our country's operating.

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 Workshops and Symposiums.

<u>Learning Materials and Technology</u>
Year-3 efforts included a major redesign and reconfiguration of the TAO portal, including meta data technology and information architecture. Year-3 effort added an additional 25 Learning Modules from all six TRUST research trust in addition to Education and Outreach to the TAO portal.

Workshop and Symposiums
Trustworthy Interfaces for Passwords and Personal Information (TIPPI) had an attendance of 60 academic and industry professionals. TIPPI program included topics on trustworthy systems. The 2007 workshop materials have been linked to the TAO portal.

The Information Assurance Capacity Building Program (IACBP) at Carnegie Mellon will again have participation by TRUST faculty and the Information Assurance Capacity Symposium (IACS) at San Jose State University as outreach and follow-up to the 2005– 2007 Information Assurance Capacity Building Program cohort. The summer 2008 IACBP program will generate learning modules that will be placed into the TAO portal.

The Education Community Development (ECD) continues to grow. The ECD is community of educators that utilize and contribute to Trusted system topics had three separate meeting during Year-3: the Spring 2007 meeting was held at San Jose State University with 20 participants; the Summer 2007 meeting was held at CSU Northridge with 20 participants and the April 2008 meeting was held at San Jose State University having 30 participants.

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, Computer Alliance for Hispanic Serving Institutions, Richard Tapia Celebration in Diversity in Computing, Grace Hopper Celebration of Women in Computing, Society of Hispanic Professional Engineers, National Science Digital Library Annual Meeting, TechLeaders: Anita Borg Institute, BEARS: Berkeley EECS Annual Research Symposium, San Francisco Electronic Crime Task Force Meetings, and the Department of Homeland Security SRI International Identity Theft Technology Council.

The Summer undergraduate research experiences SUPERB-IT at UC Berkeley will continue for summer 2008 and the Vanderbilt University undergraduate research experiences program called SIPHER will support TRUST students in the Summer 2009. Both SUPERB-IT and SIPHER are student research activities that support the centers' research thrust and goals. Both SUPERB-IT and SIPHER have supported the Center's goal of increasing the number of underrepresented minority groups and women that are conducting research in Trusted systems research.

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 Trust Academy Online will continue to develop. Course modules and learning objects will be developed as educational deliverables of each TRUST research trust. As the review process continues, refinement will be made to the module design and the portal.

TRUST Summer Programs will continue at UC Berkeley, Vanderbilt, and Cornell. TIPPI is expected to have a fourth workshop in 2008.

TRUST visibility and influence in Education Community Development is continuing to grow as TRUST participation in educational conferences, workshops, panel discussions, and Silicon Valley Industry Group activities take hold.

The Information Assurance Capacity Symposium at San Jose Sate has NSF funding through 2008.

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 10 graduate students in the summer of 2007 to 20 graduate students for the summer of 2008. Summer 2009, the SECuR-IT internship model will be expanded to TRUST partner campuses Cornell University (Financial Infrastructures) and Vanderbilt University (Health Infrastructures).

There are two education initiatives that are new and under development:

- **BUILD-IT: Bridges to Underrepresented Institutions for Long Term Development in Information Technology**. BUILD-IT is an outreach workshop targeted at the computer science, information technology and computer-security faculty of Historical Black Colleges and Universities (HBCUs) and Hispanic Severing Institutions (HSIs). The directive of the conference is to reach out to and educate the invited faculty to the research activities and opportunities presented by TRUST faculty and partner institutions. Faculty participating in this conference will be selected for their potential to influence and direct HBCU and HSI graduate student candidates to TRUST partner campuses.

- **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: 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.

  The STARS program lead is Dr. William McHenry, Project Director of the Science and Diversity Center and Executive Director of the Mississippi e-Center at Jackson State University. TRUST Executive Director of Education, Dr. Kristen Gates is active with the STARS STC partners planning group. First-year STARS funding will support two TRUST summer students.

Other Education and Outreach opportunities:

- TRUST education and outreach initiatives are actively engaged in leveraging TRUST multidisciplinary program with other NSF sponsored research as well as developing long-term community partnerships.

- The TRUST initiative, *BUILD-IT: Bridges to Underrepresented Institutions for Long Term Development in Information Technology* will support TRUST's goal to increase the number of women and underserved researchers—both student and faculty across the centers' campus partners.

- TRUST has collaborated with University of Southern California's Information Sciences Institute (ISI) on the NSF proposal called *Collaborative: CT-L: Beyond Testbeds - Catalyzing Transformative Research and Education through Cyber security Collaboratories*. The CT-L project will leverage and expand TRUST the Summer SUPERB-IT program at UC Berkeley and create new undergraduate research opportunities at San Jose State University.

- TRUST has submitted an NSF Integrative Graduate Education and Research Traineeship (IGERT) proposal called *Interdisciplinary Graduate Student Traineeship in Cyber Security and Trustworthy Systems*. Students selected for this program will participate in a new and innovative multidisciplinary educational program, having five components specific to the education and training of the next generation of cyber security professionals: 1) a thematic learning cohort; 2) creation of an academic program specialization in Cyber Security and Trustworthy Systems; 3) curriculum development 4) internships; and 5) professional development activities.

- TRUST will submit a *Broadening Participation in Computing (BPC)* proposal to the NSF in May 2008. This program proposal will propose the creation of an alliance for Cyber Security and Trustworthy Systems education, consisting of a consortium of four-year colleges, two-year community colleges, and community stakeholders. The purpose of the alliance will be to increase the number of women and underrepresented groups that earn graduate degrees in cyber security and trustworthy systems as well as the recruiting, retaining and advancing women and underrepresented groups into the professoriate.

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. To facilitate technology transfer from the research community to the industrial community a number of the investigators on this proposal, led by Sastry and Sztipanovits, have created the Embedded

Systems Consortium for Hybrid and Embedded Research (ESCHER), a non-profit organization that provides a repository for the tools and algorithms developed by researchers and establishes case-studies for design. TRUST will utilize ESCHER as a repository for developed tools and reference solutions. 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, General Motors, 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) Short courses held at major ACM and IEEE conferences as well as 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, and development of significant TRUST spin-offs (e.g., the AF-TRUST-GNC center for the U.S. Air Force), the exploratory work on a center for research on trustworthy electronic health records, and the 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

Goals	Objectives	Evaluation Criteria	Frequency
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.

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)

Body Sensor Technology Transfer		
Led by	Cornell University	
Organizations Involved		
	Name	Address
1	Cornell University	Ithaca, NY 14850
2	Qualcomm	San Diego, CA 92121

Prof. Wicker's group at Cornell has been in contact with Don Jones of Qualcomm to negotiate a collaboration between the medical sensor network group at Cornell and Qualcomm who is developing an ultra low power body area network technology.

Financial Services Industry Research and Development		
Led by	Cornell University	
Organizations Involved		
	Name	Address
1	Cornell University	Ithaca, NY 14850

As reported at the recent TRUST review meeting, we have started a vigorous 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. Birman has been invited to speak at the FSTC annual meeting in Napa on this topic, in June 08.

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

Knowledge transition takes other forms as well. The TRUST research team is prominent in 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 cybersecurity 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.

The third Trustworthy Interfaces for Passwords and Personal Information (TIPPI) workshop was held on June 22, 2007 at Stanford University with about 60 participants. In the first three years of this workshop, researchers have shared many different ideas about how to improve the situation with user interfaces for authentication, and industry efforts are starting to move along toward implementing some of them. We look forward to further reports both from the research community and developers on new ideas as well as progress in the field. The fourth TIPPI Workshop is planned for June 2008.

Visitor Monitoring		
Led by	Cornell University	
Organizations Involved		
	Name	Address
1	Cornell University	Ithaca, NY 14850
2	Johnson Museum	Ithaca, NY 14850

The Cornell team is applying results from TRUST to developing software and components on our existing testbed for visitor monitoring in the Johnson Art Museum on Cornell's campus. These include link encryption, power saving/management, and other components, which will be also applied to the medical monitoring network.

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 EMR		
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.

Book Series		
Led by	University of California, Berkeley	
Organizations Involved		
	Name	Address
1	University of California, Berkeley	Berkeley, CA 94720

We have entered into a new book series with the scientific publisher Springer to more widely disseminate our research results.

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.

Architectural Modeling and Policy Languages		
Led by	Vanderbilt University	
Organizations Involved		
	Name	Address
1	Stanford University	Stanford, CA 94305
2	Vanderbilt University (ISIS)	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.

Domain Analysis		
Led by	Vanderbilt University (ISIS)	
Organizations Involved		
	Name	Address
1	Vanderbilt University Medical Center	Nashville, TN 37235
2	Vanderbilt University (ISIS)	Nashville, TN 37235

In order to define a precise architectural model for EMR systems in general, and the MyHealth at Vanderbilt Patient Portal in particular, TRUST researchers have been organizing a series of meetings with VUMC personnel including Patient Portal designers, developers, and other associated personnel. The objective of these meetings for the TRUST researchers was to understand this domain deeply, so that the modeling language being developed, as well as the actual models, constitutes a high quality abstraction layer. Conversely, VUMC personnel gained insight into Model Integrated Technology with special emphasis on the benefits it can provide in developing EMR systems.

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.

Security Co-Design Toolbox		
Led by	Gabor Karsai, Vanderbilt	
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 basis 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.

Sensor-Based Remote Health Care System Deployment		
Led by	Vanderbilt University	
Organizations Involved		
	Name	Address
1	Vanderbilt Home Care Services, Inc.	Nashville, TN 37232
2	Vanderbilt University	Nashville, TN 37235
3	Cornell University	Ithaca, NY 14850
4	University of California, Berkeley	Berkeley, CA 94720

Industry and medical center collaboration has been established in the area of medical sensing system, with the most important development being an agreement with Vanderbilt Home Care Services, Inc. to test TRUST technology in a realistic medical environment. Additionally, researchers from TRUST have worked with the care givers at Vanderbilt Home Care Services on understanding in-home patient care scenarios. TRUST researchers have accompanied the nurses to visit the patient homes and the assisted living facilities to get first-hand experience in terms of an appropriate target group who will benefit most from our patient monitoring system, the medical data that are critical for their health, and the sensor devices that are feasible for deployment. All participants gained deeper understanding of the wide variety of issues that are raised by remote patient monitoring.

Vulnerability Analysis		
Led by	Vanderbilt University	
Organizations Involved		
	Name	Address
1	Vanderbilt University Medical Center	Nashville, TN 37235
2	Vanderbilt University (ISIS)	Nashville, TN 37235

TRUST and MyHealth researchers and developers have formed a study group on understanding scenarios in Patient Portal use cases that can have potentially negative consequences. A large group of people have been participating in these ongoing meetings including Dr. Jim Jirjis, the project manager of the Patient Portal, the Chief Security Officer of Vanderbilt, Gay Smith from the Vanderbilt privacy Office, lead developers of the Patient Portal and the Vanderbilt internal EMR system, as well as representatives from the legal office, the medical library, patient billing, etc. All participants gained deeper understanding of the wide variety of issues that are raised by publishing medical data on the web. Several issues were uncovered that would otherwise may have remained hidden.

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 cybersecurity 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 following activities:

- **Summer Experience, Colloquium and Research in Information Technology (SECuR-IT)** – SECuR-IT is a 10-week residential program for graduate students with paid internship co-located at Stanford University and San Jose State University. This is a 40-hour per week obligation to internship, research, and learning activities. Students who participate in SECUR-IT participate in San Jose State University residential cohort, attend courses, and are employed as an intern by a participating SECuR-IT industry partner. TRUST has arranged for participants to access student housing at San Jose State University and is coordinating with industry partners to place students in pain internships Participating technology companies include TRUST industry partners Intel, Yahoo, Sun Microsystems, Symantec, Visa International, and Xilinx.

- **Silicon Valley Industry Computer Security Curriculum Group** – The Industry-Backed Computer Security Curriculum is a document created by the Industry/Academic work group. The group meets on a monthly basis with representatives from Silicon Valley industries, Stanford University, UC Berkeley, San Jose State University, and TRUST.

The group's charge is to develop a plan for an industry-backed computer security curriculum and collect contributions from academic and industrial contributors. This curriculum is designed to address computer security as well as additional topics such as risk management, legal issues, and regulatory compliance—all of which are considered to be essential knowledge areas for computer security professionals.

Additionally, the Center plans on expanding the collaborative research being conducted in support of the Air Force Team for Research in Ubiquitous Secure Technology for GIG/NCES (AF-TRUST-GNC) and the International Collaboration for Advancing Security Technology (iCAST). For AF-TRUST-GNC, TRUST researchers are providing expertise and conducting research on Air Force trusted computing needs. 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	Number of Knowledge	Annual

Objective	Metric	Frequency
Knowledge Transfer Collaborators	Transfer collaborators	
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
- ESCHER Research Institute
- Hewlett Packard
- IBM
- Intel
- Microsoft

- Pirelli
- Qualcomm
- Sun
- Symantec
- 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. For example, several technology companies in the Silicon Valley area are allocating internship slots to graduate students for the TRUST Summer Experience, Colloquium and Research in Information Technology (SECuR-IT) program coordinated by Stanford University, San Jose State University, and the University of California, 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 cybersecurity 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, Taiwan, and the United Kingdom.

6 DIVERSITY

6.1 Goals and Objectives

No changes are anticipated. Below is the centers current activity.

The overall TRUST goal is to have no weak links left in the education of our society about the technical, compositional, privacy, economic and legal aspects of trusted information systems. To this end, we will begin locally but spread our outreach as far as we can along as many diverse axes as we can.

To meet this objective, the center has delivered the following programs:

- Capacity Building Program for Faculty from Historically Black and Hispanic Serving Institutions: Information Assurance Capacity Building Program at San Jose Sate University

- Summer Research Experience for underrepresented minority groups and women: SUPERB-IT at UC Berkeley and SECuR-IT at Stanford and San Jose State University

- Women Research Programs: supporting underrepresented minority groups and women in Information Technology: Women's Institute in Summer Enrichment (WISE) at UC Berkeley.

- Community Outreach at all TRUST campus

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	Every 3 Years

Goals	Objectives	Evaluation Criteria	Frequency
		alumni	
Curriculum Development	NSA certified program in IA modules	Accreditation, Modules transferred to other campuses	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
SIPHER-TRUST	Research opportunities for minority grad students at non-partner institutions	Exit surveys, Tracking surveys of alumni, Repeat visits	Every 3 Years
SUPERB-TRUST	Research opportunities for minority undergrad 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.

Additional promotion and recruitment has been performed at conference and workshop attendance. During 2007-2008, Dr. Kristen Gates, Executive Director of Education and Outreach attended the following conferences, workshops and meetins: DHS-SRI International Identity Theft Technology Council, Engineering Education NSF Awardees Conference, Grace Hopper Celebration of Women in Computing, National Science Digital Library Annual Meeting, Richard Tapia Celebration in Diversity in Computing, San Francisco Electronic Crime Task Force Meetings, Society of Hispanic Professional Engineers, CSUS Science Educational Equity (SEE) Program, TechLeaders: Leading Across Culture, Anita Borg Institute, TechLeaders: Power and Influence, Anita Borg Institute.

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

We have the concrete goal of having 30% being women and 10% being under-represented minorities among all the participants in TRUST—faculty, students, and Center staff. In addition, we will direct our outreach activities, starting locally at each campus and then as our curriculum and research gets more integrated we will also broaden the scope to TRUST-wide activities. The center will also make special attempts at outreach to Native American populations and disabled Americans.

The sections below describe some of the Center's activities which are contributing to the development of US 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 Internship for HBCU Faculty in Trusted Systems and Cyber Security
TRUST at the UC Berkeley campus, hosted Dr. Md Adbus Salam from Southern University. Dr. Salam's research and summer activity included the development of course curriculum and participation in several research workshop and conferences.

Summer Research in Information Assurance for HBCU/HSI Faculty
As a National Security Agency-designated Center of Academic Excellence (CAE) in Information Assurance Education, Carnegie Mellon has developed and offers during the summer an intensive, month-long, in-residence summer program to help develop Information Assurance education and research capacity at colleges and universities designated as Minority Serving Institutions – specifically, Historically Black Colleges and Universities (HBCUs) and Hispanic Serving Institutions (HSIs). The first two offerings of this program have been a resounding success. Carnegie Mellon has forged strong ties with a number of minority serving institutions and has significantly increased their ability to address Information Assurance in their computer science and information systems curricula. TRUST Center partner San Jose State University participated with Carnegie Mellon and hosted the 2007 Information Assurance Capacity Building Program (IACBP) follow-up workshops.

Curriculum Development for Minority Serving Institutions
San Jose State created a new course titled, CMPE 025: *The Digital World and Society*.

> *CMPE 025: The Digital World and Society -- 3 Units*
> The secure, effective, and ethical use of information technology. The effect of such technology on people and institutions. Technology-related challenges to society and policy. Frameworks for the analysis of information technology with respect to its cultural, historical, environmental, and spatial contexts.

Summer Undergraduate Research Opportunities: SUPERB-IT
The REU program at UC Berkeley, Summer Undergraduate Program in Engineering Research at Berkeley–Information Technology (SUPERB-IT) offers a group of talented undergraduate engineering students the opportunity to gain research experience. The program's objective is to increase diversity in the graduate school pipeline by affirming students' motivation for graduate study and strengthening their qualifications. SUPERB-IT participants spend eight weeks at UC Berkeley during the summer working on exciting ongoing research projects in information technology with EECS faculty mentors and graduate students. Students who participate in this research apprenticeship explore options for graduate study, gain exposure to a large research-

oriented department, and are motivated to pursue graduate study. TRUST is dedicated to developing a research experience for undergraduates from institutions serving under-represented groups during an eight-week summer term. SUPERB-IT 2007 had two participants. SUPERB-IT 2008 has six confirmed participants.

Women's Institute in Summer Enrichment

WISE is a one-week residential summer program on the UC Berkeley 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. This was the second offering of this program-- summer of 2007 had 24 participants with twelve speakers. The Institute emphasizes the inclusion of women and underrepresented graduate students, post-docs and junior faculty.

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: 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.

The STARS program lead is Dr. William McHenry, Project Director of the Science and Diversity Center and Executive Director of the Mississippi e-Center at Jackson State University. TRUST Executive Director of Education, Dr. Kristen Gates is active with the STARS STC partners planning group. First-year STARS funding will support two TRUST summer students.

TRUST Speakers Series at UC Berkeley

The TRUST Speakers Series began in fall 2007. The program will be a weekly event on the University of California, Berkeley campus. The fall 2007 series hosted thirteen speakers with a total attendance of 715 participants. The spring 2008 series will host fourteen speakers with a projected attendance of 640.

Fall 2007 Schedule	**Spring 2008 Schedule**
Object Capabilities for Security David Wagner, UC Berkeley	*Compressed Sensing Meets Machine Learning -- Classification of Mixture Subspace Models via Sparse Representation* Allen Yang, UC Berkeley
A High Assurance Least Privilege Separation Kernel and its Application Cynthia E. Irvine, Naval Postgraduate School	*Davis Social Links: P2P, Online Social Network, and Autonomous Community* Felix Wu, UC Davis
Can Systems and Networks Really Be Trustworthy? Peter Neuman, SRI	*BitBlaze: a Binary-centric Approach to Computer Security* Dawn Song, UC Berkeley
Technologies for Massively Scalable VPNs David McGrew, Cisco	*The anatomy of a Privacy Breach* Rebecca Herold
Authentication Without Identification Anna Lysyanskaya, Brown University	*Managing Multiple Perspectives on Trust* Clifford Neuman, US, Information Sciences Institute
Need Credit? No Identity? No Problem!	*Security with Privacy: Respectful Cameras and*

Fall 2007 Schedule	Spring 2008 Schedule
Chris Hoofnagle, UC Berkeley	*Actuator Networks* Ken Goldberg (IEOR and EECS and iSchool) and Jeremy Schiff (EECS), UC Berkeley
Distributed Wireless Sensors on the Human Body Ruzena Bajcsy, UC Berkeley	*Broadcast Encryption and Traitor Tracing for Content Protection* Hongxia Jin, IBM
Experiences With Countering Internet Attacks Vern Paxson, UC Berkeley / International Computer Science Institute / Lawrence Berkeley National Laboratory	*One Laptop per Child: Bringing to the children of the world an innovative and secure educational tool* Andriani Ferti, One Laptop Per Child Foundation
POTSHARDS: Secure Long Term Archival Storage Without Encryption Ethan Miller, University of California, Santa Cruz	*Improving the Robustness of Private Information Retrieval* Ian Goldberg, University of Waterloo
Privacy Tools for the End User Jessica Staddon, PARC	*Handling New Adversaries in Wireless Ad-hoc Networks* Virgil Gligor, Carnegie Mellon University
Building Reliable Voting Machine Software Ka-Ping Yee, University of California, Berkeley	Roy Maxion, Carnegie Mellon University
Quantifying Strengths and Risk Assessments of Software Protections George Cybenko, Dartmouth College	*Predicate Encryption: A New Paradigm for Public-Key Encryption* Jonathan Katz, University of Maryland
Two Techniques for Programming by Sketching Rastislav Bodik, UC Berkeley	Steve Gribble, University of Washington

6.5 Diversity Activity Impact

The goal of TRUST diversity activities is to concretely impact the number of women and personnel from under-represented groups and address issues of diversity in technical fields. Ultimately, we would like to see TRUST diversity activities positively change findings such as the following from the National Research Council of the National Academy of Sciences study *To Recruit and Advance: Women Students and Faculty in Science and Engineering*:

> “Although women have made great strides in becoming full members of the science and engineering (S&E) enterprise, they are still underrepresented among graduate students and postdoctorates and among faculty in science and engineering programs.” (NRC, 2006:1)*

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

The Women’s Institute in Summer Enrichment: WISE supports the development and advancement of women academics and researchers in the field of Information Technology and Trusted Systems.

SUPERB-IT and SIPHER: Both programs have the objective of increasing the number of students in underrepresented minority populations and women applying to graduate research programs and hopefully conducting graduate level research at a TRUST institution.

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. The series, in addition to having on campus presentations, will archive presentations on the TRUST portal. The speaker's series is learning exchange for professionals and academics in the security profession.

6.6 Diversity Metrics/Indicators

The tables below provide detail on the gender, race, and US citizenship breakdown of TRUST participants in WISE, SECuR-IT, SUPERB-IT, IACBP and CDSIA programs during the June 1, 2007 to May 31, 2008 reporting period.

WISE 2007

Constituency	**Gender**		**Race**					**US Citizen**		**Total**
	M	F	White	African American	Asian	Hispanic	Other	Y	N	
Faculty	2	9	3	1	3	1	2	10	1	11
	18%	*82%*	*27%*	*9%*	*27%*	*9%*	*18%*	*43%*	*4%*	*48%*
Graduate Students	0	10	4	2	6	0	0	10	0	10
	0%	*100%*	*40%*	*20%*	*60%*	*0%*	*0%*	*43%*	*0%*	*43%*
Research Scientists	0	0	0	0	0	0	0	0	0	0
	0%	*0%*	*0%*	*0%*	*0%*	*0%*	*0%*	*0%*	*0%*	*0%*
Post Doctorates	0	2	2	0	0	0	0	1	1	2
	0%	*100%*	*100%*	*0%*	*0%*	*0%*	*0%*	*4%*	*4%*	*9%*
	2	**21**	**9**	**3**	**9**	**1**	**2**	**21**	**2**	**23**
TOTALS	*9%*	*91%*	*39%*	*13%*	*39%*	*4%*	*9%*	*91%*	*9%*	*100%*

SECuR-IT 2007

Constituency	**Gender**		**Race**					**US Citizen**		**Total**
	M	F	White	African American	Asian	Hispanic	Other	Y	N	
Graduate Students	2	3	1	0	4	0	0	3	2	5
PhD	*20%*	*30%*	*10%*	*0%*	*40%*	*0%*	*0%*	*30%*	*20%*	*50%*
Graduate Students	4	1	2	0	2	1	0	5	0	5
MS	*40%*	*10%*	*20%*	*0%*	*20%*	*10%*	*0%*	*50%*	*0%*	*50%*
	6	**4**	**3**	**0**	**2**	**1**	**0**	**8**	**2**	**10**
TOTALS	*60%*	*40%*	*30%*	*0%*	*60%*	*10%*	*0%*	*80%*	*20%*	*100%*

SUPERB-IT 2007

Constituency	**Gender**		**Race**					**US Citizen**		**Total**
	M	F	White	African American	Asian	Hispanic	Other	Y	N	
Undergraduates	0	2	0	1	1	0	0	1	1	2
	0%	*100%*	*0%*	*50%*	*50%*	*0%*	*0%*	*50%*	*50%*	*100%*
	0	**2**	**0**	**1**	**1**	**0**	**0**	**1**	**1**	**2**
TOTALS	*0%*	*100%*	*0%*	*50%*	*50%*	*0%*	*0%*	*50%*	*50%*	*100%*

IACBP 2007

Constituency	Gender		Race					US Citizen		Total
	M	F	White	African American	Asian	Hispanic	Other	Y	N	
Faculty	10	6	5	4	6	1	0	16	0	16
	62%	*38%*	*31%*	*25%*	*38%*	*6%*	*0%*	*100%*	*0%*	*100%*
TOTALS	**10**	**6**	**5**	**4**	**6**	**1**	**0**	**16**	**0**	**16**
	62%	*38%*	*31%*	*25%*	*38%*	*6%*	*0%*	*100%*	*0%*	*100%*

CDSIA 2008

Constituency	Gender		Race					US Citizen		Total
	M	F	White	African American	Asian	Hispanic	Other	Y	N	
Faculty	26	9	12	4	14	5	0	35	0	35
	74%	*26%*	*34%*	*11%*	*40%*	*14%*	*0%*	*100%*	*0%*	*100%*
TOTALS	**26**	**9**	**12**	**4**	**14**	**5**	**0**	**35**	**0**	**35**
	74%	*26%*	*34%*	*11%*	*40%*	*14%*	*0%*	*100%*	*0%*	*100%*

Visiting Faculty from Minority Serving Institution (MSI)

Constituency	Gender		Race					US Citizen		Total
	M	F	White	African American	Asian	Hispanic	Other	Y	N	
Faculty	1	0	0	0	1	0	0	1	0	1
	100%	*0%*	*0%*	*0%*	*100%*	*0%*	*0%*	*100%*	*0%*	*100%*
TOTALS	**1**	**0**	**0**	**0**	**0**	**0**	**0**	**1**	**0**	**1**
	100%	*0%*	*0%*	*0%*	*100%*	*0%*	*0%*	*100%*	*0%*	*100%*

6.7 Next Reporting Period Diversity Plans

The recruitment of women and underrepresented minorities is a collaborative and ongoing process. The TRUST recruitment strategy for enhancing diversity is based on recommendations developed by the National Research Council as part of the study *To Recruit and Advance: Women Students and Faculty in Science and Engineering* (NRC, 2006: 47)* and includes the following recommendations:

- Advise and mentor prospective and current women and underrepresented minority undergraduate, graduate students and postdocs.

- Networking with faculty at community colleges and other four-year institutions to broaden the search for prospective recruits.

- Invite women and underrepresented minority students to participate in research opportunities.

- Participate in bridge programs, lectures and seminars.

- Broaden admission criteria and cast a wider net in recruiting students.

Listed below are new and continuing efforts that we have made towards this goal:

- We our continuing commitments to support underrepresented undergraduate summer students at all our sites (SUPERB-IT, SIPHER, San Jose State University, and Smith College).

- We are actively participating in National conferences and workshops for underrepresented faculty and students. Dr. Kristen Gates, Executive Director of Education attended the following conferences, workshops and meetings during the 2007-2008 reporting period: DHS-SRI International Identity Theft Technology Council, Engineering Education NSF Awardees Conference, Grace Hopper Celebration of Women in Computing, National Science Digital Library Annual Meeting, Richard Tapia Celebration in Diversity in Computing, San Francisco Electronic Crime Task Force Meetings, Society of Hispanic Professional Engineers, CSUS Science Educational Equity (SEE) Program TechLeaders: Leading Across Culture, Anita Borg Institute, TechLeaders: Power and Influence, Anita Borg Institute.

- Dr. Ruzena Bajcsy together with Dr. Richard Tapia of Rice University, Dr. Roscoe Giles of Boston University, and Dr. Cynthia Lanius of Drexel University have created the Empower Leadership: Computing Scholars of Tomorrow Alliance (EL Alliance). This program will engage underrepresented minority students in computing disciplines at majority institutions in a nationwide network. The network, composed of dozens of leading universities, professional societies, laboratories, research centers, and corporations, will involve students in research opportunities, professional development, mentoring programs, and support to keep the students excited and motivated as they pursue computing careers.

- The WISE one-week summer institute at UC Berkeley had 23 registered participants out of which 21 were women (graduate students and junior faculty). The WISE 2008 program is scheduled for June 8-13 at Cornell University in Ithaca, NY.

- In the summer of 2007, UC Berkeley hosted Dr. Md Adbus Salam from Southern University. This visit was sponsored by NSF Quality Education for Minorities (QEM) Program. Dr. Salam’s research and summer activity included the development of course curricula and participation in several research workshop and conferences sponsored by TRUST.

- Smith College is an active participant in TRUST activities. Dr. Judith Cardell of Smith College participated in TRUST research trust, Secure Sensor Networks. WISE Summer 2008 will be hosted at Cornell University and Dr. Cardell will be a program organizer and speaker at the event.

- As a follow-up to the Carnegie Mellon University Capacity Building Program for Faculty from Historically Black and Hispanic Serving Institutions, the Information Assurance

Capacity Building Program workshop (IACBP) was hosted at San Jose State on June 14-15, 2007.

- Dr. Ruzena Bajcsy of UC Berkeley with Ms. Maryanne McCormick of UC Berkeley created a new Computer Science course called Trustworthy Systems: the societal/ethical impact of information technology applications. This course provides an interdisciplinary introduction and overview of the societal and ethical implications of trustworthy systems in information technology in society and is open to undergraduate majors.

- We are engaged in continuous efforts of fundraising that should increase and extend our outreach efforts. TRUST has applied for an Integrative Graduate Education and Research Traineeship (IGERT) grant from NSF. The proposal, *IGERT: Interdisciplinary Graduate Student Traineeship in Cyber Security and Trustworthy Systems*, is lead by TRUST Principal Investigator Dr. Shankar Sastry. This project would support interdisciplinary teams of students studying both the technical and non-technical aspects (e.g., law, policy, usability, privacy, security, economics) of trustworthy systems and cyber security.

*National Research Council (NRC). 2006. To recruit and advance women students and faculty in US Science and engineering/Committee on the Guide to Recruiting and Advancing Women Scientists and Engineers in Academia, Committee on Women in Science and Engineering, Policy and Global Affairs, National Research Council of the National Academies.

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 the University of California, Berkeley. Additional Center leadership and management is provided by the Chief Scientist, Prof. Fred Schneider from Cornell University; the Executive Director, Larry Rohrbough, from the University of California, Berkeley; the Education Director, Dr. Kristen Gates from the University of California, Berkeley; the Outreach Director, Prof. William Robinson from Vanderbilt University, the Program Manager, Gladys Khoury from the University of California, Berkeley, and the Program Coordinator, Sally Alcala, from the University of California, 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, Program Manager, and university Principal Investigators.

Since the last reporting period, Dr. William Robinson of Vanderbilt University joined the Center Executive Board as Outreach Director replacing Prof. Ruzena Bajcsy from the University of California, Berkeley. Also, Prof. Adrian Perrig replaced Prof. Mike Reiter as Principal Investigator at Carnegie Mellon University. Prof. Reiter took a faculty appointment at the

University of North Carolina, Chapel Hill but has agreed to remain active in Center research, education, and leadership activities for this year.

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.

As an example, members of the Executive Board worked extensively this past year to address two significant management and leadership changes—identifying a new Center Outreach Director and new Principal Investigator for TRUST partner Carnegie Mellon University.

Dr. William H. Robinson, an Assistant Professor of Electrical Engineering and Computer Engineering at Vanderbilt University, is the TRUST Outreach Director. Dr. received his B.S. in electrical engineering from the Florida Agricultural and Mechanical University in 1996, his M.S. in electrical engineering from the Georgia Institute of Technology (Georgia Tech) in 1998, and his Ph.D. in electrical and computer engineering from Georgia Tech in 2003. Not only is Prof. Robinson actively involved in research and teaching in the areas of system reliability and security, he is engaged in outreach programs such as the National Society of Black Engineers (NSBE) and is the coordinator for the Alfred P. Sloan Foundation Minority Ph.D. Program in the Department of Electrical Engineering and Computer Science at Vanderbilt University.

Prof. Adrian Perrig, and Associate Professor at Carnegie Mellon University, is the Principal Investigator of TRUST partner Carnegie Mellon University. Prof. Perrig has appointments in the departments of Electrical and Computer Engineering, Engineering and Public Policy, and Computer Science and is also the technical director of Carnegie Mellon University's Cybersecurity Laboratory (CyLab). Prof. Perrig earned his Ph.D. in Computer Science from Carnegie Mellon University, and spent three years during his Ph.D. degree at University of California at Berkeley. He received his B.Sc. degree in Computer Engineering from the Swiss Federal Institute of Technology in Lausanne (EPFL). Prof. Perrig's research interests revolve around building secure systems and include Internet security, security for sensor networks and mobile applications, and trusted computing.

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

The last TRUST External Advisory Board meeting took place October 11-12, 2007 in Ithaca, NY.

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

The last TRUST Industrial Advisory Board meeting took place May 23, 2007 in Berkeley, CA.

7.7 Center Strategic Plan Changes

Changes to the TRUST Strategic Plan will be indicated within that document.

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. These publications are grouped by Peer Reviewed Publications, Books and Book Chapters, and Non-Peer Reviewed Publications.

8.1.1 Peer Reviewed Publication

- Peter Boonstoppel, Cristian Cadar, Dawson Engler RWset: Attacking Path Explosion in Constraint-Based Test Generation ETAPS Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2008) Budapest, Hungary, March-April 2008

- Dawson Engler and Daniel Dunbar Under-constrained execution: making automatic code destruction easy and scalable, to appear: International Symposium on Software Testing and Analysis (ISSTA), 2007

- C. Jackson, A. Barth, A. Bortz, W. Shao, and D. Boneh Protecting Browsers from DNS Rebinding Attacks In proceedings of the 14'th ACM conference on Computer and Communications Security (CCS), 2007 http://crypto.stanford.edu/~dabo/abstracts/dnsrebind.html

- Kumar, Tal Garfinkel, D. Boneh, and T. Winograd Reducing Shoulder-surfing by Using Gaze-based Password Entry. In proceedings of the 2007 Symposium On Usable Privacy and Security (SOUPS) http://crypto.stanford.edu/~dabo/abstracts/eyepassword.html

- Boneh, C. Gentry, and M. Hamburg Space-Efficient Identity Based Encryption Without Pairings.
- In proceedings of FOCS 2007 http://crypto.stanford.edu/~dabo/abstracts/bgh.html

- Michael Martin and Monica S. Lam. Automatic generation of XSS and SQL injection attacks with goal-directed model checking. To appear in the Proceedings of the 17th Usenix Security Symposium, 2008.

- C. Unkel and M. S. Lam Automatic Inference of Stationary Fields: a Generalization of Java's Final Fields. In Proceedings of the 35th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, (San Francisco, CA, 10-12 January 2008). ACM, 2008.

- M. S. Lam, M.C. Martin, V. B. Livshits, and J. Whaley Securing Web Applications Using Static and Dynamic Information Flow Tracking, In ACM Sigplan 2008 Workshop on Partial Evaluation and Program Manipulation, (Keynote address), January 2008.

- Michael Martin, V. Benjamin Livshits, and Monica S. Lam Finding Application Errors and Security Flaws Using PQL: a Program Query Language. In Proceedings of the Conference on Object-Oriented Programming, Systems, Languages and Applications (OOPSLA '05), October 2005.

- V. Benjamin Livshits and Monica S. Lam Finding Security Vulnerabilities in Java Applications Using Static Analysis In Proceedings of the 14th USENIX Security Symposium, August 2005.

- Monica S. Lam, John Whaley, V. Benjamin Livshits, Michael C. Martin, Dzintars Avots, Michael Carbin and Christopher Unkel. Context-Sensitive Program Analysis as Database Queries In Proceedings of the 24th SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, June, 2005. (Invited Tutorial).

- S. Bugrara and A. Aiken Verifying the Safety of User Pointer Dereferences. http://theory.stanford.edu/%7Eaiken/publications/papers/oakland08.pdf Proceedings of the IEEE Symposium on Security and Privacy, to appear, May 2008.

- Korolova, S. Nabar, and Y. Xu Link Privacy in Social Networks. Proceedings of the 21st International Conference on Data Engineering (ICDE), 2008. http://theory.stanford.edu/~rajeev/papers.html

- E. Stinson and J.C. Mitchell, Characterizing Bots' Remote Control Behavior, 4th GI Int'l Conf. on Detection of Intrusions \& Malware, and Vulnerability Assessment (DIMVA), Lucerne, Switzerland, July, 2007.

- Jackson, D. Boneh, and J.C Mitchell, Transaction Generators: Rootkits for the Web, 2nd USENIX Workshop on Hot Topics in Security (HotSec '07).

- A.Barth, J.C. Mitchell, A. Datta and S. Sundaram, Privacy and Utility in Business Processes, 20th IEEE Computer Security Foundations Symposium (CSF 20), Venice, July, 2007.

- Roy, A. Datta, A. Derek and J.C. Mitchell. Inductive Trace Properties Imply Computational Security, 7th International Workshop on Issues in the Theory of Security (WITS'07), Braga, Portugal, March, 2007. (Competitive submission workshop)

- Roy, A. Datta, A. Derek, J.C.~Mitchell, J.-P. Seifert, Secrecy Analysis in Protocol Composition Logic, 11th Annual Asian Computing Science Conference (ASIAN'06), Tokyo, December, 2006.

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

- Cynthia Kuo, Ahren Studer, Adrian Perrig. "Mind Your Manners: Socially Appropriate Wireless Key Establishment for Groups". ACM Conference on Wireless Network Security (WiSec), ACM, April, 2008.

- Jonathan M. McCune, Bryan Parno, Adrian Perrig, Mike Reiter, Arvind Seshadri. "How Low Can You Go? Recommendations for Hardware-Supported Minimal TCB Code Execution". ACM Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), ACM, March, 2008.

- Jason Franklin, Vern Paxson, Stefan Savage, Adrian Perrig. "An Inquiry into the Nature and Causes of the Wealth of Internet Miscreants". ACM Conference on Computer and Communications Security (CCS), ACM, November, 2007.

- Arvind Seshadri, Mark Luk, Ning Qu, Adrian Perrig. "SecVisor: A Tiny Hypervisor to Provide Lifetime Kernel Code Integrity for Commodity OSes". ACM Symposium on Operating Systems Principles (SOSP), ACM, October, 2007.

- Cynthia Kuo, Adrian Perrig, Jesse Walker. "Low-cost Manufacturing, Usability, and Security: An Analysis of Bluetooth Simple Pairing and Wi-Fi Protected Setup". Usable Security (USEC), February, 2007.

- Michael Merideth. "Tradeoffs in Byzantine-Fault-Tolerant State-Machine-Replication Protocol Design". Technical report, Institute for Software Research, Carnegie Mellon University, CMU-ISR-08-110, March, 2008.

- David Brumley, Pongsin Poosankam, Dawn Song, Jiang Zheng. "Automatic Patch-Based Exploit Generation is Possible: Techniques and Implications". 2008 IEEE Symposium on Security and Privacy, April, 2008; To appear at the 2008 IEEE Symposium on Security and Privacy, Oakland, CA.

- Kenneth Debelak, Larry Howard, Yuan Xue, Christina Lee, Janos Sztipanovits. "Introducing security in a chemical engineering design course using adaptive online learning". International Conference on Engineering Education, July, 2006.

- Adrian Lauf, William H. Robinson. "Distributed Intrusion Detection System for Resource-Constrained Devices in Ad Hoc Networks". *Elsevier Journal of Ad Hoc Newtorks*, 2008.

- Aniruddha Gokhale, Joe Hoffert, Douglas Schmidt. "A QoS policy configuration modeling language for publish/subscribe middleware platforms". DEBS '07: Proceedings of the 2007 inaugural international conference on Distributed event-based systems, 140--145, June, 2007.

- Janos Laszlo Mathe, Jan Werner, Yonghwan Lee, Bradley Malin, Akos Ledeczi. "Model-based design of clinical information systems." Unpublished article, 2008; Under submission to Methods of Information in Medicine.

- Janos Laszlo Mathe, Sean Duncavage, Jan Werner, Akos Ledeczi, Bradley Malin, Janos Sztipanovits. "Towards the security and privacy analysis of patient portals". *ACM SIGBED Review*, 4(2):5, 2007.

- Michael Merideth, Michael Reiter. "Write markers for probabilistic quorum systems". Technical report, Computer Science Department, Carnegie Mellon University, CMU-CS-07-165, November, 2007.

- Michael Merideth, Michael Reiter. "Probabilistic opaque quorum systems". Distributed Computing: 21st International Symposium, DISC 2007, 403-419, September, 2007.

- Shanshan Jiang, Yanchuan Cao, Sameer Iyengar, Philip Kuryloski, Roozbeh Jafari, Yuan Xue, Ruzena Bajcsy, Stephen Wicker. "CareNet: An Integrated Wireless Sensor Networking Environment for Remote Healthcare". BodyNets, 2008.

- Shanshan Jiang, Yuan Xue, Douglas Schmidt. "Minimum Disruption Service Composition and Recovery over Mobile Ad Hoc Networks". MOBIQUITOUS, 2007.

- Amin Aminzadeh Gohari, Venkatachalam Anantharam. "New Bounds on the Information-Theoretic Key Agreement of Multiple Terminals". To appear in ``Proceedings of the IEEE International Symposium on Information Theory", Toronto, Canada, 2008.

- Adam Barth, Collin Jackson, John C. Mitchell. "Securing Frame Communication in Browsers". Talk or presentation, 3, April, 2008.

- Patrice Godefroid, Michael Levin, David A Molnar. "Automated Whitebox Fuzz Testing". Talk or presentation, 3, April, 2008.

- Daniel Dunbar, Cristian Cadar, Peter Pawlowski, Dawson Engler. "Effective Testing via Symbolic Execution and Input Recombination". Talk or presentation, 3, April, 2008.

- Jonathan M. McCune, Bryan Parno, Adrian Perrig, Michael Reiter, Hiroshi Isozaki. "Flicker: An Execution Infrastructure for TCB Minimization". Talk or presentation, 3, April, 2008.

- Michael Merideth, Michael Reiter. "Write Markers for Probabilistic Quorum Systems". Talk or presentation, 3, April, 2008.

- Ken Birman, Mahesh Balakrishnan, Tudor Marian, Hakim Weatherspoon. "Maelstrom: An Enterprise Continuity Protocol for Financial Data Centers". Talk or presentation, 3, April, 2008.

- Mikhail Lisovich, Stephen Wicker. "Power Consumption Monitoring - An Emerging Threat to Privacy". Talk or presentation, 3, April, 2008.

- Panagiotis Papadimitriou, Hector Garcia-Molina. "Detecting Data Leakage". Talk or presentation, 3, April, 2008.

- Deirdre Mulligan. "Security Breach Notification Laws: A "Race-to-the-Top"?". Talk or presentation, 3, April, 2008.

- Janos Laszlo Mathe, Jan Werner, Yonghwan Lee, Bradley Malin, Akos Ledeczi, John C. Mitchell, Janos Sztipanovits. "Experimental Platform for Model-Integrated Clinical Information Systems". Talk or presentation, 3, April, 2008.

- Annarita Giani, Gabor Karsai, Tanya Roosta, Aakash Shah, Bruno Sinopoli, Jon Wiley. "A Testbed for Secure and Robust SCADA Systems". Talk or presentation, 3, April, 2008.

- Allen Yang, Sameer Iyengar, Shanshan Jiang, Philip Kuryloski, Yanchuan Cao, Roozbeh Jafari, Yuan Xue, Ruzena Bajcsy, Stephen Wicker, S. Shankar Sastry. "Deploying Distributed Real-time Healthcare Applications on Wireless Body Sensor Networks". Talk or presentation, 3, April, 2008.

- Tanya Roosta, Sameer Pai, Phoebus Chen, S. Shankar Sastry, Stephen Wicker. "The Inherent Security of Routing Protocols in Ad-Hoc and Sensor Networks". Talk or presentation, 3, April, 2008.

- David Wagner. "Keynote Speech–California Top-To-Bottom Review of Voting Systems". Talk or presentation, 3, April, 2008.

- Adrian Lauf, Richard A. Peters, William H. Robinson. "A Distributed Intrusion Detection System for Resource-Constrained Devices in Ad Hoc Networks". Talk or presentation, 3, April, 2008.

- Sergio Bermudez, Stephen Wicker. "Taking Advantage of Data Correlation to Control the Topology of Wireless Sensor Networks". Talk or presentation, 3, April, 2008.

- Hui Qu, Stephen Wicker. "Co-designed anchor-free localization and location-based routing algorithm for rapidly-deployed wireless sensor networks". *Information Fusion*, 2008.

- Hui Qu, Stephen Wicker. "A Combined Localization and Geographic Routing Algorithm for Rapidly-Deployed Wireless Sensor Networks". *International Journal of Distributed Sensor Networks*, 4(1):44-63, 2008.

- Hazer Inaltekin, Mung Chiang, H. Vincent Poor, Stephen B. Wicker. "On the Asymptotic Behavior of Selfish Transmitters Sharing a Common Wireless Communication Channel". IEEE ISIT 2008, 2008.

- Coalton Bennett, Judith Cardell, Stephen Wicker. "Residential Demand Response Wireless Sensor Network". Fourth Annual Carnegie Mellon Conference on the Electricity Industry, Carnegie Mellon University Department of Electrical Engineering, 5, March, 2008.

- Sergio Bermudez, Stephen Wicker. "Taking Advantage of Data Correlation to Control the Topology of Wireless Sensor Networks". International Conference on Telecommunications, 2008.

- Alvaro Cardenas, Saurabh Amin, S. Shankar Sastry. "Secure Control: Towards Survivable Cyber-Physical Systems". First International Workshop on Cyber-Physical Systems (WCPS2008), IEEE, June, 2008.

- Sameer Pai, Tanya Roosta, Stephen Wicker, S. Shankar Sastry. "Game Theoretic Modeling of Trust in Networks of Bayesian-Learning Sensors". Unpublished article, 2008.

- Sameer Pai. "Using Social Network Theory Towards Development Of Wireless Ad Hoc Network Trust". Talk or presentation, 11, October, 2007.

- Mikhail Lisovich, Sergio Bermudez, Stephen Wicker. "Reconfiguration in Heterogeneous Mobile Wireless Sensor Networks". ISWPC 2008, May, 2008.

- Hazer Inaltekin, Stephen B. Wicker. "A One-shot Random Access Game for Wireless Networks". Symposium on Information Theory in Wirelesscom, 2005.

- Hazer Inaltekin. "The Analysis of a Game Theoretic MAC Protocol for Wireless Networks". IEEE Secon 2006, 2006.

- Hazer Inaltekin, Tom Wexler, Stephen B. Wicker. "A Duopoly Pricing Game for Wireless IP Services". IEEE Secon 2007, 2007.

- Mikhail Lisovich, Stephen Wicker. "Power Consumption Monitoring - an Emerging Threat to Privacy". Unpublished article, 2008.

- Hazer Inaltekin, Stephen B. Wicker. "Random Access Games: Selfish Nodes with Incomplete Information". IEEE Milcom 2007, 2007.

- Annarita Giani, Gabor Karsai, Tanya Roosta, Aakash Shah, Bruno Sinopoli, Jon Wiley. A Testbed for Secure and Robust SCADA Systems, 14th IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS) Saint Louis 2008

8.1.2 Journal Articles

- Amin Aminzadeh Gohari, Venkatachalam Anantharam. "Information-Theoretic Key Agreement of Multiple Terminals - Part II: Channel Model". *Submitted to ``IEEE Transactions on Information Theory"*, 2008.

8.1.3 Books and Book Chapters

- None to report.

8.1.4 Non-peer Reviewed Publications

- None to report.

8.2 Conference Presentations

The following is a list of conference presentations made by TRUST Center personnel during this reporting period.

- Annarita Giani, Gabor Karsai, Tanya Roosta, Aakash Shah, Bruno Sinopoli, Jon Wiley. A Testbed for Secure and Robust SCADA Systems, Talk or presentation, 3, April, 2008.

- Adam Barth, Collin Jackson, John C. Mitchell. Securing Frame Communication in Browsers, Talk or presentation, 3, April, 2008.

- Aaron Burstein. *Conducting Cybersecurity Research Legally and Ethically*, Talk or presentation, 15, April, 2008; Accepted for presentation to the first USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET '08) on April 15, 2008, San Francisco. Publication forthcoming.

- David Wagner. *Keynote Speech–California Top-To-Bottom Review of Voting Systems*, Talk or presentation, 3, April, 2008.

- David Wagner. *California Top-To-Bottom Review of Voting Systems*, Talk or presentation, 3, April, 2008.

- Sergio Bermudez, Stephen Wicker. *Taking Advantage of Data Correlation to Control the Topology of Wireless Sensor Networks*, Talk or presentation, 3, April, 2008.

- Adrian Lauf, Richard A. Peters, William H. Robinson. *A Distributed Intrusion Detection System for Resource-Constrained Devices in Ad Hoc Networks*, Talk or presentation, 3, April, 2008.

- Tanya Roosta, Sameer Pai, Phoebus Chen, S. Shankar Sastry, Stephen Wicker. *The Inherent Security of Routing Protocols in Ad-Hoc and Sensor Networks*, Talk or presentation, 3, April, 2008.

- Allen Yang, Sameer Iyengar, Shanshan Jiang, Philip Kuryloski, Yanchuan Cao, Roozbeh Jafari, Yuan Xue, Ruzena Bajcsy, Stephen Wicker, S. Shankar Sastry. *Deploying Distributed Real-time Healthcare Applications on Wireless Body Sensor Networks*, Talk or presentation, 3, April, 2008.

- Janos Laszlo Mathe, Jan Werner, Yonghwan Lee, Bradley Malin, Akos Ledeczi, John C. Mitchell, Janos Sztipanovits. *Experimental Platform for Model-Integrated Clinical Information Systems*, Talk or presentation, 3, April, 2008.

- Jonathan M. McCune, Bryan Parno, Adrian Perrig, Michael Reiter, Hiroshi Isozaki. *Flicker: An Execution Infrastructure for TCB Minimization*, Talk or presentation, 3, April, 2008.

- Deirdre Mulligan. *Security Breach Notification Laws: A "Race-to-the-Top"?*, Talk or presentation, 3, April, 2008.

- Panagiotis Papadimitriou, Hector Garcia-Molina. *Detecting Data Leakage*, Talk or presentation, 3, April, 2008.

- Mikhail Lisovich, Stephen Wicker. *Power Consumption Monitoring - An Emerging Threat to Privacy*, Talk or presentation, 3, April, 2008.

- Ken Birman, Mahesh Balakrishnan, Tudor Marian, Hakim Weatherspoon. *Maelstrom: An Enterprise Continuity Protocol for Financial Data Centers*, Talk or presentation, 3, April, 2008.

- Michael Merideth, Michael Reiter. *Write Markers for Probabilistic Quorum Systems*, Talk or presentation, 3, April, 2008.

- Daniel Dunbar, Cristian Cadar, Peter Pawlowski, Dawson Engler. *Effective Testing via Symbolic Execution and Input Recombination*, Talk or presentation, 3, April, 2008.
- Patrice Godefroid, Michael Levin, David A Molnar. *Automated Whitebox Fuzz Testing*, Talk or presentation, 3, April, 2008.
- Prateek Saxena, Dawn Song. *BitBlaze - Binary Analysis for COTS Protection and Malicious Code Defense*, Talk or presentation, 21, February, 2008.
- Anupam Datta. *Privacy and Utility in Business Processes*, Talk or presentation, 10, October, 2007.
- Akos Ledeczi. *Model-Based Design Environment for Clinical Information Systems*, Talk or presentation, 10, October, 2007.
- Ken Birman. *Quicksilver Scalable Multicast*, Talk or presentation, 10, October, 2007.
- Yee Jiun Song. *The Building Blocks of Consensus*, Talk or presentation, 10, October, 2007.
- Michael Reiter. *Probabilistic Opaque Quorum Systems*, Talk or presentation, 10, October, 2007.
- Aaron Burstein. *Network Security and the Need to Consider Provider Coordination in Network Access Policy*, Talk or presentation, 10, October, 2007.
- Bob Mungamuru. *Competition and Fraud in Online Advertising Markets*, Talk or presentation, 10, October, 2007.
- Deirdre Mulligan, Stephen Wicker. *Can Selective Sensing Protect Democratic Principles and Enhance Policing?*, Talk or presentation, 10, October, 2007.
- Mark Luk. *Don’t Sweat Your Privacy: Using Humidity to Detect Human Presence*, Talk or presentation, 10, October, 2007.
- Tanya Roosta. *A Model-based Intrusion Detection System for Wireless Process Control Systems*, Talk or presentation, 10, October, 2007.
- Weider D. Yu. *ARSL: A Language for Authorization Rule Specification in Software Security*, Talk or presentation, 11, October, 2007.
- Sameer Pai. *Using Social Network Theory Towards Development Of Wireless Ad Hoc Network Trust*, Talk or presentation, 11, October, 2007.
- Elizabeth Stinson. *Characterizing the Remote Control Behavior of Bots*, Talk or presentation, 11, October, 2007.
- Arnab Roy. *Inductive Proofs of Computational Secrecy*, Talk or presentation, 11, October, 2007.
- Adrian Perrig. *Portcullis: Protecting Connection Setup from Denial-of-Capability Attacks*, Talk or presentation, 11, October, 2007.

- Ruzena Bajcsy. *Monitoring Elderly*, Talk or presentation, 11, October, 2007.
- Michael Martin. *Automatically Generating Attacks on Webapps with Model Checking*, Talk or presentation, 11, October, 2007.
- Sameer Pai. *Using Social Network Theory Towards Development of Wireless Ad Hoc Network Trust*, Talk or presentation, 11, October, 2007.
- Cody Hartwig. *Towards Automatic Discovery of Deviations in Binary Implementations*, Talk or presentation, 11, October, 2007.
- Aaron Burstein. *Toward a Culture of Cybersecurity Research*, Talk or presentation, 10, August, 2007.
- John McHugh. *Monitoring your network for fun and prophet[sic]*, Talk or presentation, 3, May, 2007.
- Kevin Fu. *Vulnerabilities in First Generation RFID-enabled credit cards*, Talk or presentation, 22, May, 2007.
- Edward A. Lee. *Is Truly Real-Time Computing Becoming Unachievable?*, Talk or presentation, 3, April, 2007.
- Trust seminar presentation. *Brian Chess*, Talk or presentation, 12, April, 2007.
- Brian Chess. *Selling Security to Software Developers*, Talk or presentation, 12, April, 2007.
- Edward A. Lee. *Is Truly Real-Time Computing Becoming Unachievable?*, Talk or presentation, 3, April, 2007.
- Matt Bishop. *Elections and Computers: A Match Made in ... Someplace?*, Talk or presentation, 19, April, 2007.
- John Mitchell. *Developing an Industry Supported Computer Security Curriculum*, Talk or presentation, February, 2007.
- Ruzena Bajcsy, Kristen Gates. *TRUST Summer Study Programs*, Talk or presentation, 19, March, 2007.
- Larry Howard. *Disseminating Learning Materials:TRUST Academy Online (TAO)*, Talk or presentation, 19, March, 2007.
- Kristen Gates. *TRUST Education and Outreach*, Talk or presentation, 19, March, 2007.
- Anthony D. Joseph, Vern Paxson, Robbert van Renesse. *Network Defense Research*, Talk or presentation, 19, March, 2007.
- John Mitchell, Doug Tygar. *Online ID Theft, Phishing, and Malware*, Talk or presentation, 19, March, 2007; Presented at the TRUST March 2007 NSF Site Visit/All Hands Meeting, Berkeley, CA.
- Mike Reiter. *Trustworthy Systems*, Talk or presentation, 19, March, 2007.

- Janos Sztipanovits. *Electronic Medical Record (EMR) Project*, Talk or presentation, 19, March, 2007.
- Stephen Wicker, Deirdre Mulligan, Judy Cardell. *Sensor Networks and Embedded Systems*, Talk or presentation, 19, March, 2007.
- Stephen B. Wicker. *TRUST Center Activities*, Talk or presentation, 19, March, 2007.
- S. Shankar Sastry. *TRUST:Team for Research in Ubiquitous Secure Technologies Overview*, Talk or presentation, 19, March, 2007.
- Stephen Wicker. *Sensor Networks: Technology Transfer*, Talk or presentation, 19, March, 2007.
- Doug Tygar, John Mitchell. *ID Theft Technology Transfer*, Talk or presentation, 19, March, 2007.
- Mike Reiter. *Integrative Projects Ideas*, Talk or presentation, 21, March, 2007.
- Deirdre K. Mulligan. *Knowledge Transfer - Policy*, Talk or presentation, 19, March, 2007.
- Deirdre K. Mulligan. *Policy Outbrief*, Talk or presentation, 21, March, 2007.
- Chris Karlof. *End user security outbrief*, Talk or presentation, 21, March, 2007.
- Stephen Wicker. *Sensor Networks and Embedded Systems: Breakout Session Report*, Talk or presentation, 21, March, 2007.
- S. Shankar Sastry. *TRUST: Team for Research in Ubiquitous Secure Technologies: Home Work Assignment*, Talk or presentation, 21, March, 2007.
- Larry Rohrbough. *Knowledge Transfer*, Talk or presentation, 21, March, 2007.
- Kristen Gates. *TRUST Education and Outreach Year 3 Projects*, Talk or presentation, 19, March, 2007.
- Gabor Karsai. *TRUST Knowledge Transfer EMR Project*, Talk or presentation, 19, March, 2007.
- Amin Aminzadeh Gohari, Venkat Anatharam. *Unconditionally Secret Key Agreement using Public Discussion*, Talk or presentation, 15, February, 2007.
- Peter Fuhr. *Industrial Wireless Systems: Implications for Everyone*, Talk or presentation, 8, February, 2007.
- Alfonso Valdes. *Using Model Based Intrusion Detection for SCADA Networks*, Talk or presentation, 18, January, 2007.
- Rik Farrow. *Security is Broken*, Talk or presentation, 31, January, 2007.
- Deirdre Mulligan, Jack Lerner. *Taking the long view on the Fourth Amendment: Stored Records and the Sanctity of the Home*, Talk or presentation, December, 2007.

- Robbert van Renesse. *Cross-cutting Opportunities in Network Defenses*, Talk or presentation, 21, January, 2007.

- Sigurd Meldal. *SECuR-IT: A Summer School and Immersion Program*, Talk or presentation, 19, January, 2007.

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 30, 2007: Berkeley law attorney Director Deirdre Mulligan participated in the KPPC Patt Morrison radio show, which discusses current public affairs. The major issue put forth on the radio show was how Southern California Edison is seeking to install smart meters, the aim of which is to reduce consumer costs and saving energy. However, installing these "smart meters" poses privacy concerns as customers' activities within the home will be monitored.

- September 9, 2007: Berkeley law attorney Deirdre Mulligan and Clinical student David Snyder co-authored an op-ed published by the *San Francisco Chronicle* titled "Why Security and Liberty Fit Hand In Glove" which discussed the tension between increased security and civil liberties and why the two concepts need not be mutually exclusive. The op-ed discussed video surveillance as well as other types of surveillance as examples of security measures that encroach on individual liberties, noting that the need to "watch the watchers" and ensure such tools are not abused is a vital part of living in a democratic society.

- December 10-11, 2007: Berkeley Law attorney Chris Jay Hoofnagle made two presentations to a Federal Trade Commission workshop focused on the use of the Social Security number. At the workshop, Hoofnagle discussed synthetic identity theft, a practice where impostors create fictitious identities for financial gain.

- December 17-18, 2007: Berkeley Law attorney Deirdre Mulligan and social science researcher Jennifer King presented at the Department of Homeland Security's public workshop "CCTV: Developing Privacy Best Practices." The workshop was aimed to discuss and develop best practices that protect individuals' privacy for jurisdictions deploying video surveillance systems.

- December 20, 2007: Berkeley Law attorney Deirdre Mulligan participated as a panelist in the "Privacy and the Network of You" discussion hosted by Sun Microsystems. She discussed how privacy is not just a matter of regulation but that it is also a matter of creating markets that allow for technological innovations: "…privacy is something that is expected as part of the fabric in this environment."

- January 30, 2008: Berkeley Law social scientist Jennifer King recommended some tools for consumers to strengthen their privacy settings online in a *Wall Street Journal* article "It's Hard to Hide From Your 'Friends.'" The article discussed how several online services, such as Facebook and Google, have allowed anyone to access personal information about people they know.

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	She made preeminent contributions to her fields, and to the world	Elected to the American Academy of Arts & Sciences	Monday, April 28, 2008	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	Apostol, Alexander (Smith)	B.S.
2	Attaluri, Srilatha (San Jose State)	M.S.
3	Avula, Dharani (San Jose State)	M.S.
4	Bowers, Kevin (Carnegie Mellon)	Ph.D.
5	Cai, Fangli (San Jose State)	M.S.
6	Chan, Wing On (San Jose State)	M.S.
7	Cheong, Elaine (Berkeley)	Ph.D.
8	Chow, James (Stanford)	Ph.D.
9	Dao, Thang (San Jose State)	M.S.
10	Derek, Ante (Stanford)	Ph.D.
11	Dilys, Thomas (Stanford)	Ph.D.
12	Duncavage, Sean (Vanderbilt)	Ph.D.
13	Emerson, Matthew (Vanderbilt)	Ph.D.
14	Hosur, Prachi (San Jose State)	M.S.
15	Hosur, Vadiraj (San Jose State)	M.S.
16	Jothiram, Vijayalakshimi (San Jose State)	M.S.
17	Jyotula, Deepika (San Jose State)	M.S.
18	Kasivwanathan, Nagapriya (San Jose State)	M.S.
19	Kumani, Alpana (San Jose State)	M.S.
20	Li, Yaping (Berkeley)	Ph.D.
21	Livshits, Vladimir (Stanford)	Ph.D.
22	Mazareeb, Seif (San Jose State)	M.S.
23	McGhee, Scott (San Jose State)	M.S.
24	Mishra, Shree (San Jose State)	M.S.
25	Nargundkar, Shruti (San Jose State)	M.S.
26	Nayak, Ellora (San Jose State)	M.S.
27	Nguyen, Khoi (San Jose State)	M.S.
28	Oprea, Florin (Carnegie Mellon)	Ph.D.
29	Pandya, Vaibhav (San Jose State)	M.S.
30	Patil, Amita (San Jose State)	M.S.
31	Philip, Roney (San Jose State)	M.S.

	Student Name	Degree(s)
32	Sathyanarayana, Sreevathsa (San Jose State)	M.S.
33	Shanmugam, Basu Gopal (San Jose State)	M.S.
34	Shroff, Shenal (San Jose State)	M.S.
35	Trutoiu, Laura (Smith)	B.S.
36	Tsay, Elbert (San Jose State)	M.S.
37	Veerabhadraiah, GangaBhavani (San Jose State)	M.S.
38	Vellanki, Laxmisamyukta(San Jose State)	M.S.
39	Venkataramu, Ramya (San Jose State)	M.S.
40	Venkatesan, Ashwini (San Jose State)	M.S.
41	Wu, Taojun (Vanderbilt)	Ph.D.
42	Zhang, Ying (San Jose State)	M.S.

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 Knowledge Transfer	N
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	eBay	Company	San Jose, CA	Dave Cullinane	Education	Y
8	General Electrical Capital	Company	McKinney, TX	James Beeson	Education	N
9	Greater Bay Bank	Company	Palo Alto, CA	Jason Hoffman	Education	Y
10	Hewlett-Packard	Company	Palo Alto, CA	Rich McGeer	Research Knowledge Transfer	N
11	ING	Company		Robert Weaver	Education	N
12	Intel	Company	Santa Clara, CA	Anand Rajan	Research Knowledge Transfer	N
13	Jefferson Wells	Company	Brookfield, WI	Jeffrey Camiel	Education	N
14	Microsoft Research	Company	Redmond, WA	Mike Schroeder	Research	N
15	Oracle	Company	Redwood Shores, CA	Mary Ann Davidson	Knowledge Transfer	N
16	Pirelli Research Laboratory	Company	Berkeley, CA	Marco Sgroi	Research Knowledge Transfer	N
17	Rapport, Inc.	Company	Redwood City, CA	Andrew Singer	Education	Y
18	Silicon Valley Bank	Company	Santa Clara, CA	Andrew Neilson	Education	Y
19	Sun Microsystems	Company	Menlo Park, CA	Emil Sarpa	Research Education	Y
20	Symantec	Company	Santa Monica, CA	Ken Baylor	Research	N

					Knowledge Transfer	
21	Tata Consultancy Services	Company	Chennai, India	Sanjay Bahl	Education	N
22	United Technologies	Company	East Hartford, CT	Clas Jacobson	Research Knowledge Transfer	N
23	Visa International	Company	San Francisco, CA	George Sullivan	Research Knowledge Transfer	N
24	Yahoo Inc.	Company	Sunnyvale, CA	Mark Seiden	Education	Y
25	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

Carnegie Mellon University:
Anupam Datta – Anupam Datta is a Research Scientist at Carnegie Mellon University. He obtained MS and Ph.D. degrees from Stanford University and a BTech from IIT Kharagpur, all in Computer Science. Dr. Datta's research interests are in security, cryptography and privacy. He has authored over 20 papers in topics including security analysis of network protocols, theory of cryptography, languages for privacy policy specification and enforcement, and software system security. Dr. Datta is the General Chair for the 2008 IEEE Computer Security Foundations Symposium and has served on program committees for a number of security conferences including the 2007 IEEE Symposium on Security and Privacy.

Virgil Gligor – Professor of Electrical and Computer Engineering at Carnegie Mellon University. Virgil D. Gligor received his B.Sc., M.Sc., and Ph.D. degrees from the University of California at Berkeley and taught at the University of Maryland between 1976 and 2007. He was an Editorial Board member of the ACM Transactions on Information System Security and several IEEE Transactions (e.g., Dependable and Secure Computing, Computers, and Mobile Computing) and he is currently the Ediro In Chief of IEEE TDSC. Over the past three decades, his research interests ranged from access control mechanisms, penetration analysis, and denial-of-service protection to cryptographic protocols and applied cryptography. He was awarded the 2006 National Information Systems Security Award jointly given by NIST and NSA in the US for his contributions to security research.

Bruno Sinopoli – Bruno Sinopoli received his M.S. and Ph.D. in Electrical Engineering from the University of California at Berkeley, in 2003 and 2005 respectively. Previously he received the Dr. Eng. degree from the University of Padova. Dr. Sinopoli is assistant professor in the Department of Electrical and Computer Engineering at Carnegie Mellon University. His research interests include networked embedded control systems, distributed estimation and control, hybrid systems with applications to wireless sensor-actuator networks and system security. Dr. Sinopoli was awarded, jointly with Dr. Schenato, the 2006 Eli Jury Award for outstanding research achievement in the areas of systems, communications, control and signal processing at UC Berkeley.

Stanford University:
Monica Lam – Monica Lam is a Professor in the Computer Science Department at Stanford University since 1988. She received a B.Sc. from University of British Columbia in 1980 and a Ph.D. in Computer Science from Carnegie Mellon University in 1987. She has worked in the areas of compiler optimization, software analysis to improve security, simplifying computing with virtualization.

Her contributions compiler optimizations include software pipelining, data locality, and parallelization. The SUIF compiler infrastructure developed by her research group has been widely used by compiler researchers all around the world. She helped found Tensilica in 1998, which specializes automatic generation of configurable processor cores and compilers from a high-level description.

Her contributions in program analysis for security include tools for automatically detecting cross-site scripting and SQL injection bugs in Java/JSP web applications, which was based on a novel context-sensitive pointer alias analysis. Other contributions include the bddbddb (BDD-based Deductive DataBase) analysis system, the PQL program query language, the Diduce dynamic root-cause analyzer, the Clouseau C++ memory leak detector, and the Cred buffer overrun detector. She co-authored Compilers, Principles, Techniques, and Tools (2nd Edition), also known as the Dragon book, which was published in 2006.

In the area of simplifying computing, her Collective project developed the concept of a livePC: subscribers of the livePC will automatically run the latest of the published PC virtual images with each reboot. This approach allows computers to be managed scalably and securely. In 2005, the group started a company called moka5 to transfer the technology to industry.

Monica is an ACM Fellow. She received an NSF Young Investigator award in 1992, the ACM Most Influential Programming Language Design and Implementation Paper Award in 2001, an ACM SIGSOFT Distinguished Paper Award in 2002, and the ACM Programming Language Design and Implementation Best Paper Award in 2004. She was the author of two of the papers in "20 Years of PLDI--a Selection (1979-1999)", and one paper in the "25 Years of the International Symposia on Computer Architecture".

She chaired the ACM SIGPLAN Programming Languages Design and Implementation Conference in 2000, served on the Editorial Board of ACM Transactions on Computer Systems and numerous program committees for conferences on languages and compilers (PLDI, POPL), operating systems (SOSP), and computer architecture (ASPLOS, ISCA).

Rajeev Motwani – Rajeev Motwani is a Professor of Computer Science at Stanford University, where he also serves as the Director of Graduate Studies. He obtained his Ph.D. in Computer Science from Berkeley in 1988. His research has spanned a diverse set of areas in computer science, including databases, data mining, and data privacy, web search and information retrieval, robotics, computational drug design, and theoretical computer science. He has written two books -- Randomized Algorithms published by Cambridge University Press in 1995, and an undergraduate textbook published by Addison-Wesley in 2001. Motwani has received the Godel Prize, the Okawa Foundation Research Award, the Arthur Sloan Research Fellowship, the National Young Investigator Award from the National Science Foundation, the Distinguished Alumnus Award from IIT Kanpur, the Bergmann Memorial Award from the US-Israel Binational Science Foundation, and an IBM Faculty Award. He is a Fellow of the Institute of Combinatorics and serves on the editorial boards of SIAM Journal on Computing, Journal of Computer and System Sciences, ACM Transactions on Knowledge Discovery from Data and IEEE Transactions on Knowledge and Data Engineering. Motwani serves on various industry boards and advisory boards, including Adchemy, Agitar, CastTV, Coral8, DotEdu Ventures, Fatdoor, Flarion, Fraudwall, Google, Jaxtr, Jumpstartup Ventures, Mimosa Systems, Neopath Networks, Revenue Science, Snaptell, Stanford Student Enterprises Ventures, uGenie, and Xambala. He is a charter member of TIE (The IndUS Entrepreneurs) and on the board of BASES (Business Association of Stanford Engineering Students).

UC Berkeley:

Yale Braunstein – Yale M. Braunstein is a Professor at the School of Information at the University of California, Berkeley. He received a B.S. degree from Rensselaer Polytechnic Institute and a doctorate in economics from Stanford University. He is the author or co-author of over 50 articles in the fields of economics and information science and has served as a consultant for several corporations and government agencies in the U.S. and internationally.

As an economist, Yale focuses on competition in information products and services, in particular on how new generations of products and technologies alter the commercial landscape for incumbent players. His research areas include economies of scale and scope, pricing, market structure, and the economics of intellectual property rights. His work has been published in the major scholarly journals in economics, information science, and legal policy.

Yale has also developed financial, forecasting, tariff, and valuation models in areas that include cellular, fixed, and international telecommunications; cable, satellite, and IP television; and broadband. This work has been used by applicants for licenses, regulators, and policy makers in the U.S., Brazil, Canada, China, Ireland, Israel, Sweden, Ukraine, and the UK.

Yale has been a visiting scholar and guest lecturer in China and Germany and at the East-West Center in Hawaii. Working with faculty at the Center for Digital Technology and Management (CDTM) in Munich, he co-developed and co-taught the course "Realizing Digital Convergence" which was simultaneously offered in Berkeley and Munich with lectures delivered live over the web in both directions.

Dawn Song – Assistant Professor at University of California, Berkeley. She obtained her PhD in Computer Science from UC Berkeley (2002). Her research interest lies in security and privacy issues in computer systems and networks. She is the author of more than 60 research papers in areas ranging from software security, networking security, database security, distributed systems security, to applied cryptography. She is the recipient of various awards and grants including the NSF CAREER Award, the IBM Faculty Award, the George Tallman Ladd Research Award, the Sloan Award, and the Best Paper Award in USENIX Security Symposium.

Vanderbilt University:

Bradley Malin – Bradley Malin is an assistant professor of biomedical informatics at the Vanderbilt University Medical Center. His primary research focus is on data privacy and management issues in biomedical research and clinical management systems. He is the author of numerous scientific articles on data privacy, fraud detection, and surveillance within various technologies, including text databases, biomedical databases, and face recognition systems. His research on the re-identification and privacy protection of patient-specific genomic database records has received several awards from the American Medical Informatics Association and International Medical Informatics Association. Brad holds a bachelor’s in molecular biology, a master’s in public policy and management, a master’s in computer science ("data mining and knowledge discovery"), and a doctorate in computer science ("computation, organizations, and society") from Carnegie Mellon University.

Prior to joining Vanderbilt, he was a graduate researcher in the Data Privacy Laboratory at Carnegie Mellon University.

Appendix B: Center Organizational Chart

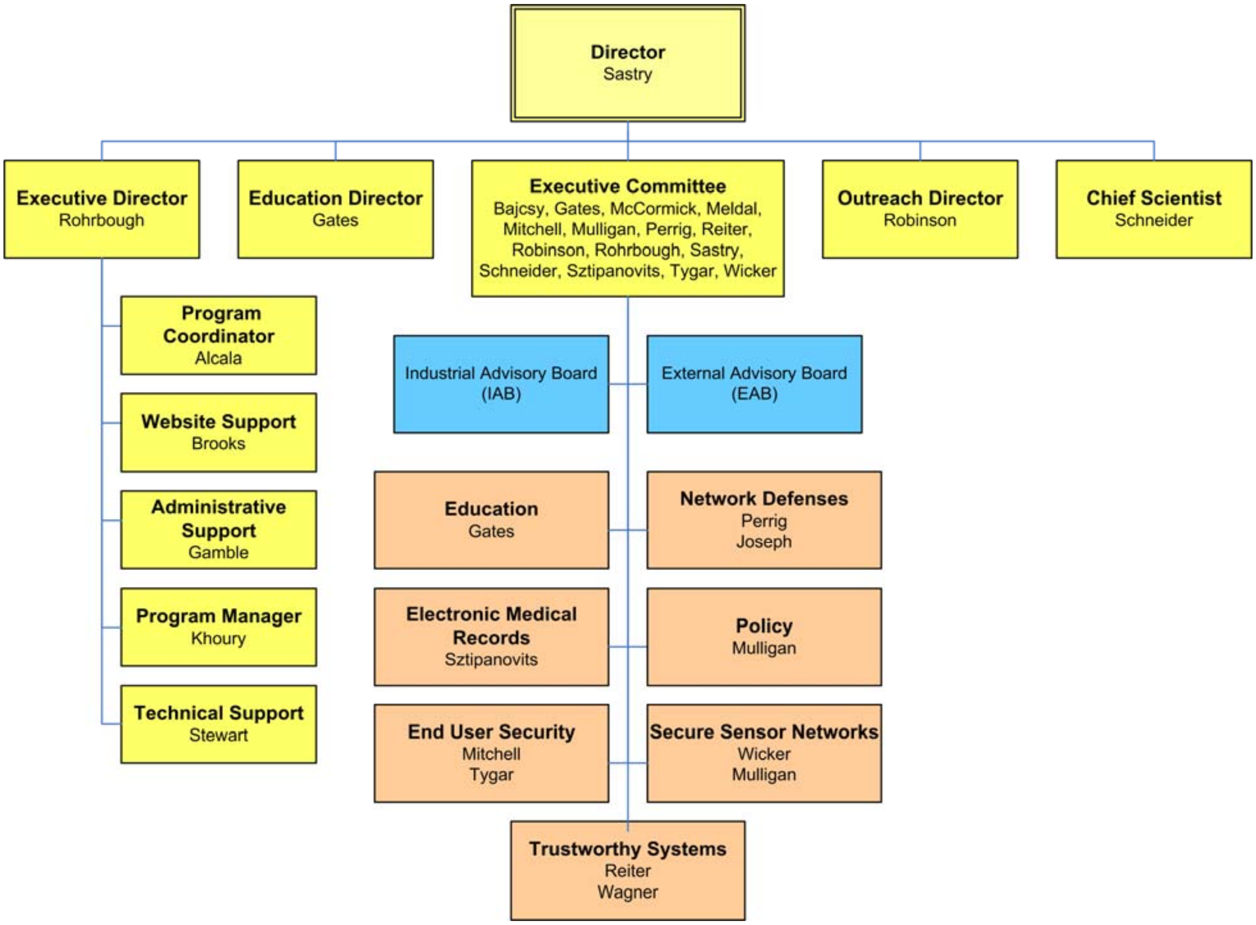

Appendix D: Media Publicity Materials

Flyers for three TRUST Education and Outreach programs conducted in the Summer 2007 are included on the following pages. Program flyers are for SECuR-IT, SUPERB-IT, and WISE.

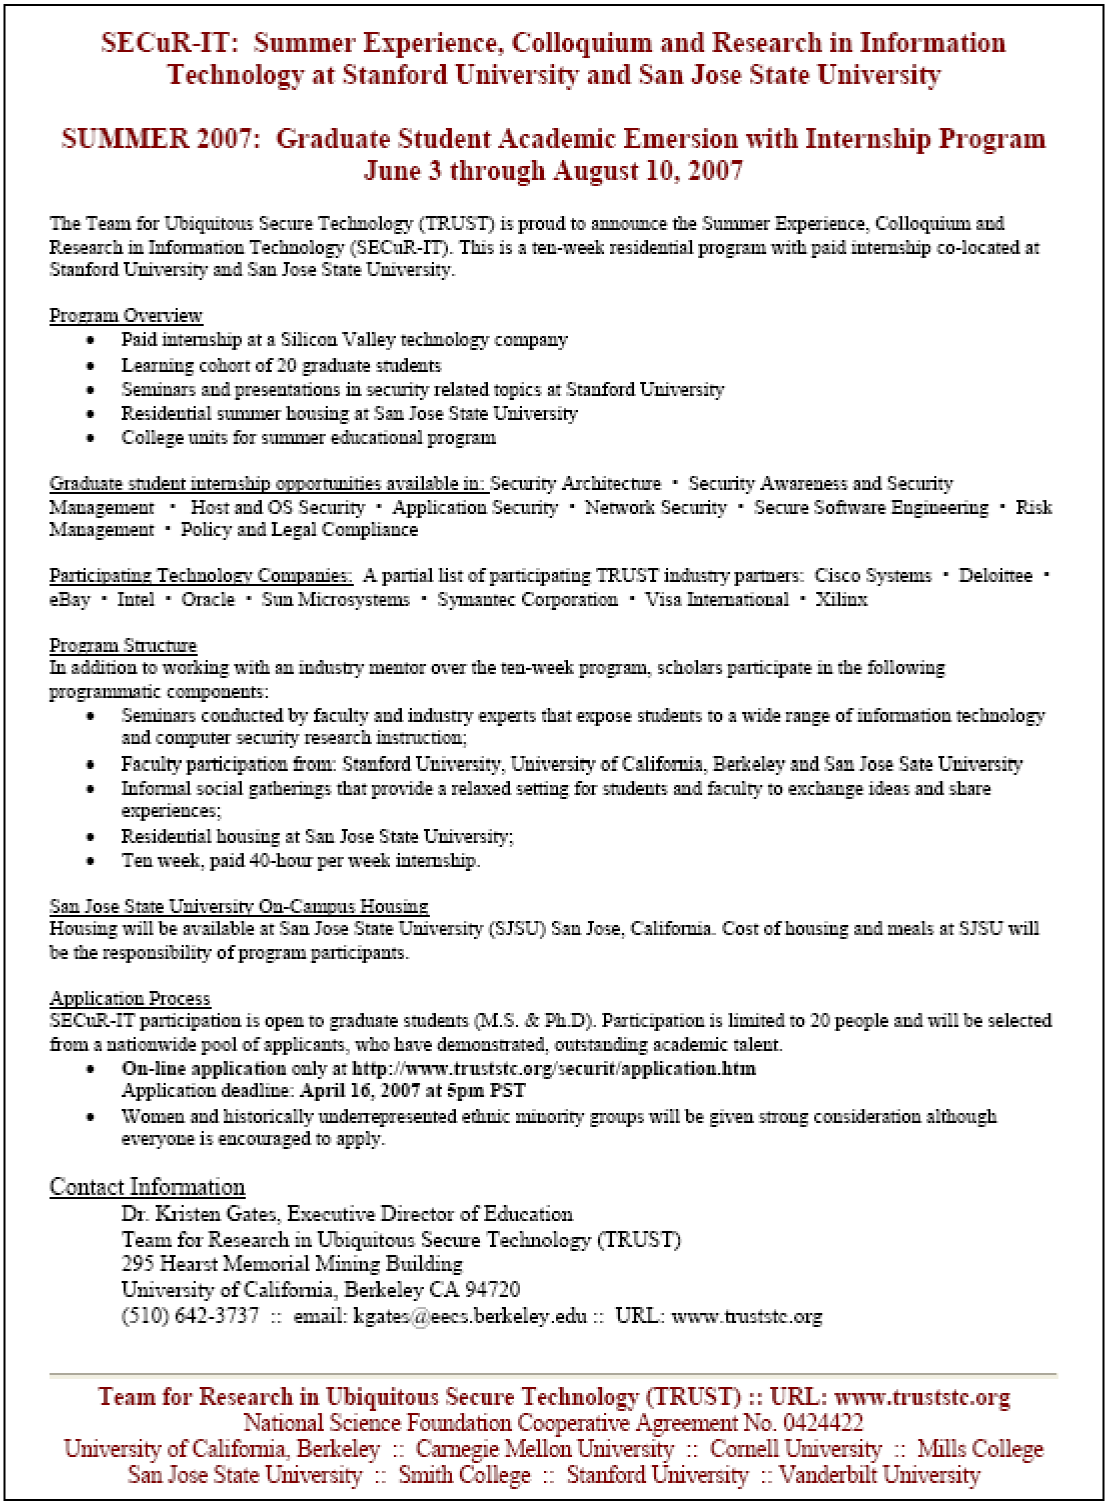

SECuR-IT: Summer Experience, Colloquium and Research in Information Technology at Stanford University and San Jose State University

SUMMER 2007: Graduate Student Academic Emersion with Internship Program
June 3 through August 10, 2007

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 20 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

Participating Technology Companies: A partial list of participating TRUST industry partners: Cisco Systems • Deloittee • eBay • Intel • Oracle • Sun Microsystems • Symantec Corporation • Visa International • Xilinx

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 20 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/application.htm**
  Application deadline: **April 16, 2007 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

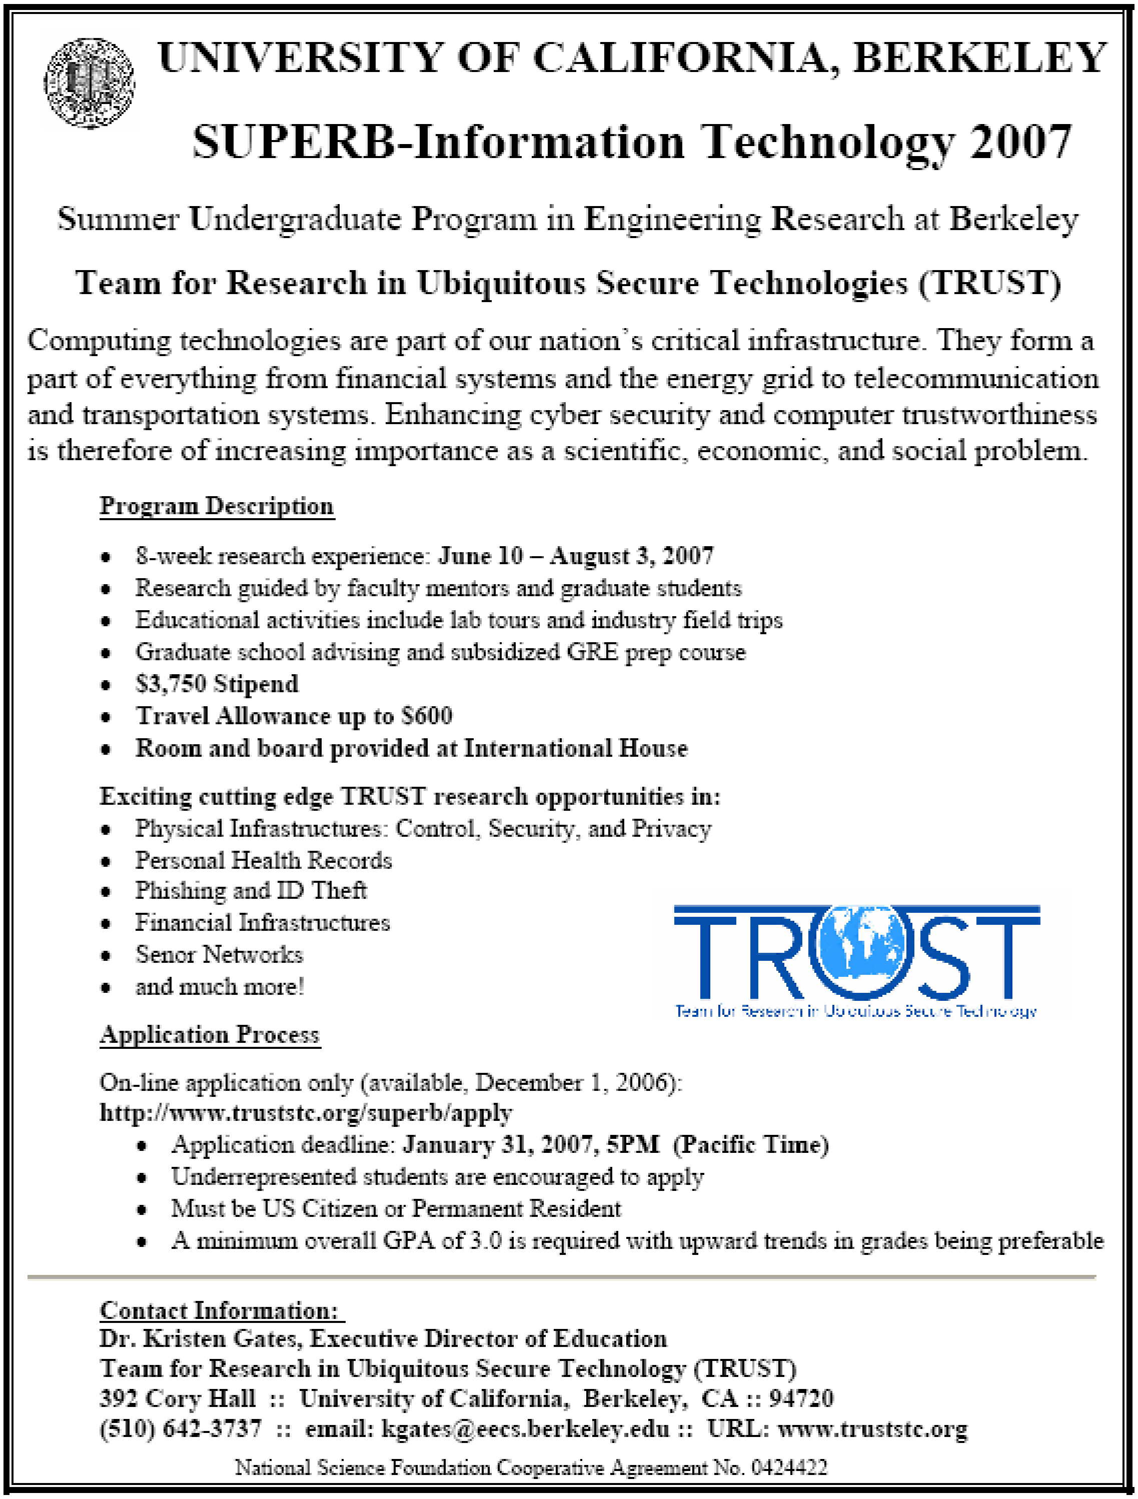

UNIVERSITY OF CALIFORNIA, BERKELEY

SUPERB-Information Technology 2007

Summer Undergraduate Program in Engineering Research at Berkeley

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 cyber security and computer trustworthiness is therefore of increasing importance as a scientific, economic, and social problem.

Program Description

- 8-week research experience: **June 10 – August 3, 2007**
- 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
- **$3,750 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, 2006):
http://www.truststc.org/superb/apply

- Application deadline: **January 31, 2007, 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

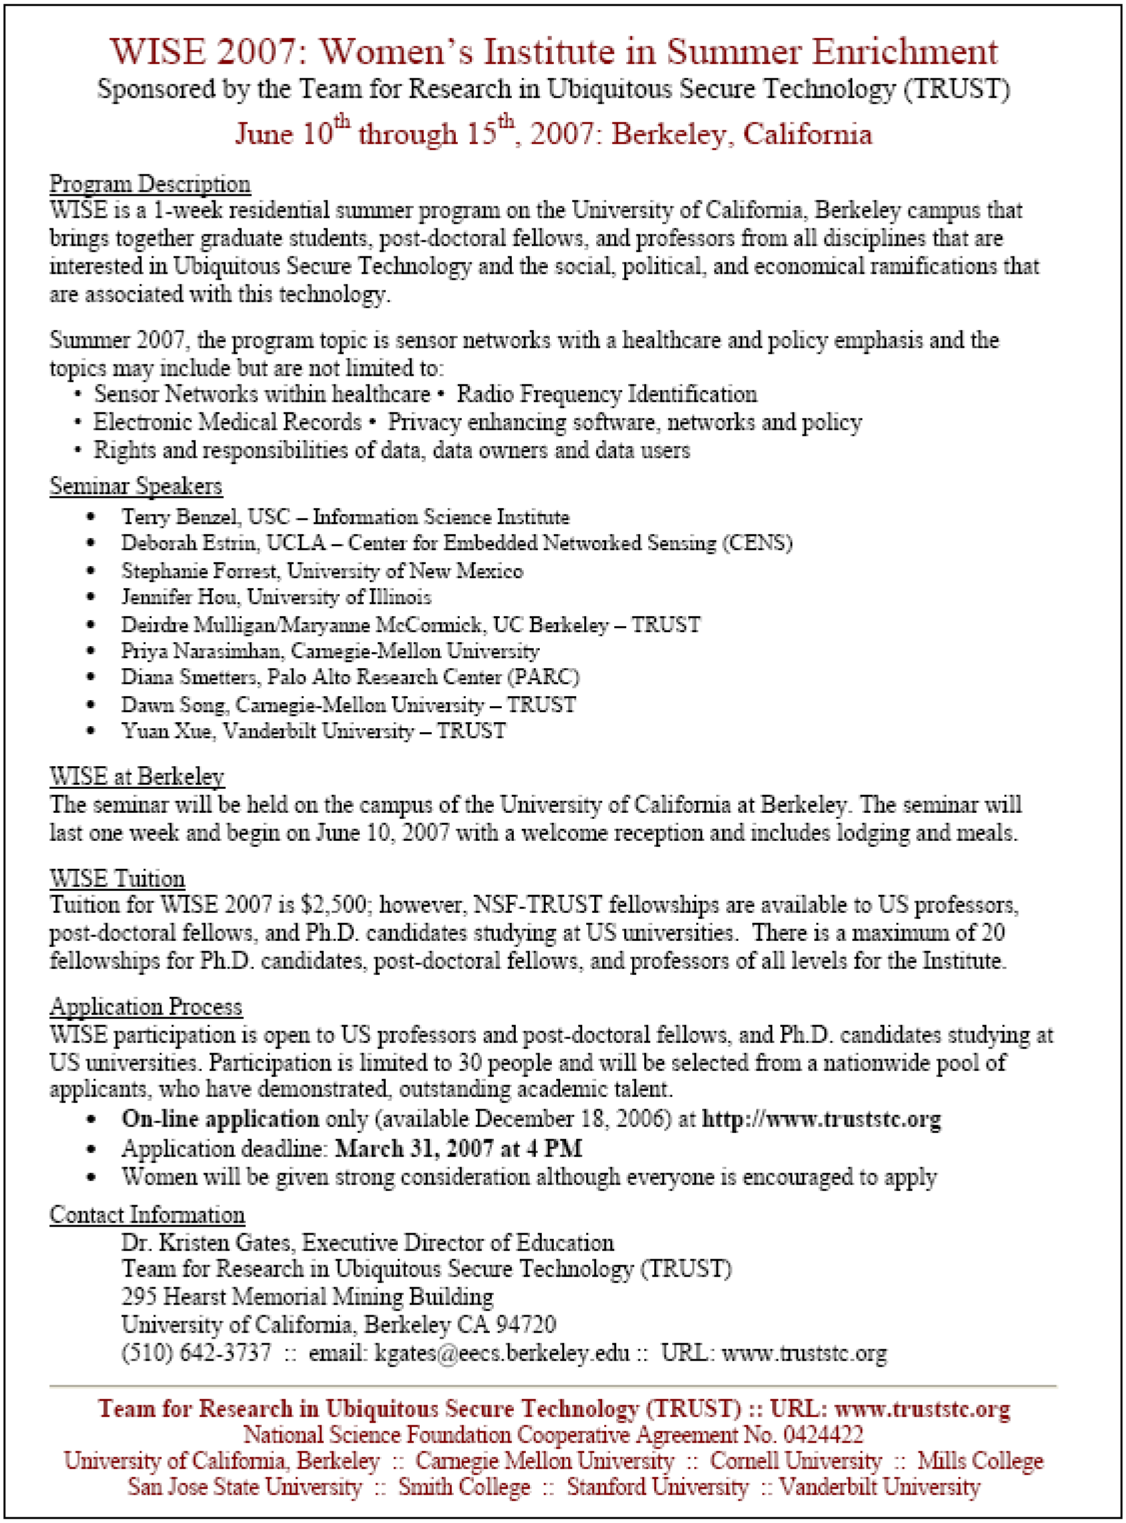

WISE 2007: Women's Institute in Summer Enrichment

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

June 10th through 15th, 2007: Berkeley, California

Program Description

WISE is a 1-week residential summer program on the University of California, Berkeley 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.

Summer 2007, the program topic is sensor networks with a healthcare and policy emphasis and the topics may include but are not limited to:

- Sensor Networks within healthcare • Radio Frequency Identification
- Electronic Medical Records • Privacy enhancing software, networks and policy
- Rights and responsibilities of data, data owners and data users

Seminar Speakers

- Terry Benzel, USC – Information Science Institute
- Deborah Estrin, UCLA – Center for Embedded Networked Sensing (CENS)
- Stephanie Forrest, University of New Mexico
- Jennifer Hou, University of Illinois
- Deirdre Mulligan/Maryanne McCormick, UC Berkeley – TRUST
- Priya Narasimhan, Carnegie-Mellon University
- Diana Smetters, Palo Alto Research Center (PARC)
- Dawn Song, Carnegie-Mellon University – TRUST
- Yuan Xue, Vanderbilt University – TRUST

WISE at Berkeley

The seminar will be held on the campus of the University of California at Berkeley. The seminar will last one week and begin on June 10, 2007 with a welcome reception and includes lodging and meals.

WISE Tuition

Tuition for WISE 2007 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 18, 2006) at **http://www.truststc.org**
- Application deadline: **March 31, 2007 at 4 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