Program Description

Our nation relies increasingly on its computers and networks in commerce and for national defense. But these systems will not provide the needed service when subjected to attacks, which are occurring at an ever increasing frequency and with more severe effects[L01]. Worms and viruses are a common occurrence and although most of the worms launched so far have been less malicious than theoretically possible, the damage in terms of lost service and human time to recover the systems has totaled in the billions of dollars. Much more devastating worms have been predicted. Furthermore, attacks that have resulted in millions of credit card numbers being stolen have been observed. It is likely that what we have observed so far is just the beginning of a long and costly engagement with attackers.

The "defense" community has been active in the development of tools to cope with these attacks. But based on the success of new attacks, it seems the defense is at best one step behind the attack community; we say "at best" because even old attacks succeed on many sites. Given the ineffective experience in coping with attacks, it is essential that new techniques be developed. Most of the existing techniques are ad hoc, being designed to handle specific classes of attacks. UC Davis proposes the Center for Computer Security Research (CCSR) to develop a new science of computer security that will (1) have broad applicability across the space of systems that currently exist or are likely to be developed over the next 10 years, (2) have broad applicability over the current class of attacks, but also to attacks yet to be observed or suspected, (3) permit the evaluation of defense systems both analytically and experimentally. For cryptography, such a science has been under development, with significant results, for several decades. But in an effective defense system, cryptography is only part of the solution. Other issues relating to good system design to prevent attacks but also to the "coping" with attacks, as inevitably will occur with legacy systems, are also extremely important.

To achieve the vision of CCSR, UC Davis has assembled an experienced team with skills that cover the space of computer security, including research academic institutions (Stanford, Colorado, Georgia Institute of Technology, Idaho), teaching institutions (Sacramento State, San Jose State), research institutes (SRI International), and commercial organizations (Symantec and NetSquared). Together with partners who will not draw funds from the CCSR (Cisco, Hewlett Packard, Intel, NetScreen, Network Associates, CloudShield), the team is committed to developing a *science of computer security*, which can be applied to the development and evaluation of secure systems against the *threats of the future*.

1. Research Thrusts

To motivate the research, we have identified four thrusts which are suggestive of application areas for the research: (1) The design of resilient computers and networks which can stand up to a wide variety of "outside" attacks, such as worms, distributed denial of service attacks (DDOS), and attacks that attempt theft of data; the emphasis is on unknown attacks; (2) Coping with the insider threat, which costs many organizations more than the outsider; (3) Providing significant automation towards the security management of large systems; (4) Improving the privacy of sensitive data.

As we discuss below, these thrust area, representative of the difficult security problems facing the security community, are addressed through a common body of techniques and tools that the CCSR will develop. Furthermore, a set of interoperative models is proposed to support methods to reason about the security of systems through analytic methods and systematic experimental evaluation.

Brief descriptions of the four research thrusts follow, providing the objectives of each thrust and the general approach towards a scientific methodology each suggests.

1.1. Towards resilient networks and systems

The resilient networking thrust of our proposed Center will investigate all aspects of research related to the design, deployment, analysis, and operation of network architectures and their protocols which lead to their secure and robust operation. The networking thrust will tackle a broad spectrum of research problems, ranging from wireline to wireless networks, from network design to operation, from network security to network resilience and robustness, and from cryptography to intrusion detection to protocol vulnerability analysis, from novel network architectures to retrofitting the existing network infrastructure to improve its security and robustness properties.

Wireline and wireless networks. Network security research is a young discipline, and most of the attention in this field has been devoted to wireline networks. However, wireless networks, including cellular networks, ad hoc networks, sensor networks, IEEE 802.11 networks, etc., are proliferating at a rapid rate. The wireless communication channel presents some very serious network security and vulnerability problems. Oftentimes, due to performance requirements, wireless network configurations are unencrypted and very easy to tap into by an eavesdropper. Mobile ad-hoc wireless networking protocols allow for widespread mobile connectivity, but conventional protection mechanisms involving centralized policy, monitoring and enforcement schemes are useless in the face of rapidly changing network topologies. Untethered, self-contained wireless clients introduce new and novel forms of denial-of-service attack that aim to consume the limited battery power of mobile devices. Hence, we propose to conduct network security research that will address not only the general (wireline) networks, but wireless networks as well. Some of our preliminary research work targeted to the wireless channel can be found in [JNFWCH02].

Network design and operation. It is important to not only design and build secure networks (the off-line problem), but also to develop sound and (on-line) algorithms for their secure operation. Our investigation will encompass both design and operation algorithms. Oftentimes, security breaches can be traced back to faulty assumptions about one's security policy. Automated network policy discovery can highlight instances in which the actual policy in force on a live network is different than the policy that administrators have assumed. An automated policy discovery and recommendation system would allow for the prevention of many attacks from the outset. Today's most common network security model relies upon security control at the Internet boundary; typically each boundary is protected by a firewall with a static filtering policy based upon IP headers. A more robust model would include adding additional protection devices deep in the network configured to cooperatively enforce security policy at all internal points in the network as well. Dynamic response to automated, distributed attack would be possible even if the secure network perimeter were bypassed.

Network security and network robustness. Traditional network security research considers secure design and operation of networks. However, our Center will also investigate the science and technology towards architecting fault-tolerant networks. The objective is to design and operate robust networks that are resilient to network faults (e.g., equipment failure) and that can quickly and efficiently recover from disasters (natural or otherwise) by rerouting traffic to other secure and operational parts of the network. Rather than having to eliminate all forms of attack to be successful, these systems would continue to operate even after an unexpected security breach has occurred. To address this goal, we shall expand on the research literature on fault-tolerant networking, including some of our contributions [RM99]. Incorporating redundant network services built upon heterogeneous collections of software and hardware, independently produced yet providing similar functionality would allow service to continue invulnerable to trivial attack replay. Variations of successful attacks launched against a heterogeneous set of servers would allow for automated attack learning and automated signature generalization that could be dynamically enforced by boundary devices [JRCDLMR02].

Cryptography to intrusion detection to protocol vulnerability analysis. Several members of our research team have made sustained contributions over the past decade on research problems on cryptography, intrusion detection, and protocol vulnerability analysis [MHL94,BD96,S96,KRL97, KBRLT01]. Our example contributions include the first network security monitor, the first distributed intrusion detection system, pioneering work in specification based intrusion detection, IPsec and IP thumbprinting for traceback. Today these techniques are included in a variety of commercial security products, which are used to protect large institutions as well as individual home and small office users. We shall continue to investigate this wide range of research problems in network security to improve our knowledge and understanding of the field. In particular, building upon our work in specification based intrusion detection, models of policy, monitoring and protection that are based upon correct operations of a system, rather than on the ad-hoc knowledge of experts will be a major focus. This approach allows formal methods and reasoning to be applied to real systems that secure computers and networks.

1.2. The Insider Threat

Numerous studies [LS00], some very recent, report that insiders pose a continuing and, in many cases, an increasing threat to organizations. Damage done by malicious insiders is far more severe than damage from outside attackers due to an insider's privileged access and sophisticated knowledge of the system. This threat continues to dominate the threat from outside despite the continued trend to increasing

networking of systems. The proposed center will investigated prevention methods and detection techniques that mitigate insider attack, a continuing and largely unsolved problem. Novel new insider misuse prevention and detection research would be based upon the following approaches.

Subtle insider misuse detection through analysis of access to fine-grained structures. Previous attempts to solve the insider problem through the monitoring and analysis of user accesses to file objects have failed. We believe that insider misuse is generally associated with unauthorized access to data objects which are at a finer granularity than files and that the misuse is generally associated with inappropriate access through specific applications. One preliminary approach models objects in question as semi-structured data, and provides a level of access pattern granularity where misuse, even if subtle, can be noted.

Built upon formal models. Another complementary approach is based on a solid formal model for data organization and querying. Through this model, access control policies can be specified by an expert, learned through policy discovery, or a combination of both. The policy could be expressed in terms of access paths to sub-documents, where paths contained in a policy define the documents that are accessible and how they are accessed. Policy statements are quite general, and reflect the users performing the access, the roles of the users, conditions (which can be very general, as they are expressed in first order logic), and temporal issues (such as the interval over which accesses are made). Policies that are relevant to document access control could be specified in a formal language, permitting one to reason about properties such as soundness, completeness and consistency. Policies could also be compared in a meaningful way, e.g., to determine the distance between policies. The latter permits policies to be ordered, e.g., according to priority.

Built upon well-known security principals. Although a data model and policy specifications are new and novel, the approach to policy enforcement and misuse detection must be based on sound and well-tested security principles: least privilege, a data access path model based on the concept of trusted paths, role-based access control, anomaly-based intrusion detection, and specification-based intrusion detection where specifications are policy statements, scoring misuse based on departure from historical profile, and response to misuse.

1.3. Towards Automated System Administration

Managing the security of interconnected entities requires an understanding of the goals of security, of how the entities work, and of how their interconnection works.

The definition of security is critical. The difficulty arises when entities have different definitions of security. Consider a joint project that a commercial firm is doing with a university. Parts of the project are proprietary data, and the company will keep those confidential. Other parts are academic research, which the university wishes to make available to all. The entities collaborating on the project have different security goals. Reconciling these goals requires integrating the security policies of the two entities in a way acceptable to both.

Security models describe the security-relevant workings of the entities. These models capture the desires security policies to which the entities conform. The reconciliation of disparate models, as discussed above, is a component of security management. But we also will examine the accuracy of the models: how well does the model capture the *reality* of the system's protection goals, and how can we augment the model to bring its level of abstraction closer to the reality of the system? For example, consider a system that implements the Clark-Wilson model [CW87], requiring that only certain users can execute certain transaction procedures. When mapping the model into the system, the mapping abstracts details of the system. In security, the details may cause problems as well. For example, a buffer overflow would be abstracted out under most mappings, but would present a vulnerability that an attacker could use to violate the security of the system. We plan to use assurance techniques, including layering mappings (so there are multiple layers of abstraction, and mappings between adjacent layers) to validate that the models accurately reflect the system structure.

A similar approach will help us understand how the interconnection of the various entities affects conformance to a set of security goals. The interconnections cause the composition of policies and models, and it must be shown that the composition of two entities deemed secure is secure (because the composition of secure systems may not be secure [M87]). Again, our approach will reconcile issues arising at multiple layers of abstraction, to provide convincing evidence of conformance to security goals.

We note that other management functions, such as patching, upgrading systems, and making resources available (or unavailable) due to changes in policy are important parts of security management.

These also involve changes at several layers of abstraction, from model (because they affect the availability of resources or entities) to implementation (because they introduce new programs or alter existing ones).

1.4. Privacy Enhancement

Privacy is an issue, although increasingly important, has not received the attention afforded to security. CCSR will develop new techniques to determine when data is to be released and the degree of trust to be placed in data. Besides an issue in its own right, for example what data in a medical informatics system is releasable and to whom, it is an issue in the other thrusts of the CCSR. Sharing of data among enterprises about incidents is essential in handling an incident. However, enterprises might want to release only the data needed by other enterprises to assist it in tracking down an attacker. The effect of the attack, e.g., on specific IP addresses, might not be necessary so sanitization of the data is possible. Also, the trust [LWM01,CDM01] to be accorded to data might depend on the security state of the data source. Thus this thrust will influence the creations of models that bear on policy, trust, operation status and sanitization.

CCSR will also study cryptographic techniques for designing efficient privacy-preserving protocols across a wide variety of applications, building on the work of Chaum [CE86] and others[BF99]. In a distributed key generation protocol, two or more parties cooperate to produce a new public key, such that each party holds a share of the private key. This allows the parties to jointly decrypt or sign messages thereafter, while the private key has never and will never be in any one place at any time. In an electronic voting protocol, many parties input secret votes, and the output is a tally of the votes efficiently computed while maintaining the privacy of the individual votes. Two main design paradigms have been successful for electronic voting: "mix-based" schemes that compute directly on decrypted but sanitized inputs; and "homomorphic" schemes that compute indirectly on specially encrypted inputs. General techniques are known for assuring privacy to all participants in essentially all protocols, but these generic methods are usually not efficient.

2. Approach: Towards Interoperative Models in Support of Systematic Design and Evaluation for Computer Security

The above thrusts and other research in the synthesis and evaluation of secure systems need a unified approach. The CCSR will focus its research around a collection of interoperative models to be developed and demonstrated during the research. Through the models it will be possible to carry out the reasoning and experimental evaluation needed to provide real assurance about secure systems. Moreover, the models can serve as specifications for modules that comprise a secure system, bringing principles of software engineering to bear on the design of secure systems. Focusing on model creation and analysis, the elements of our approach are as follows:

Models of vulnerabilities: Most security problems are due to vulnerabilities [B99], which can reside in designs, protocols, implementation, or configuration tables, all of which we will consider. The research will identify classes of vulnerabilities and determine their possible manifestations in various kinds of system objects. For example, many security-critical programs suffer from such vulnerabilities as buffer overflow, race conditions, or unexpected action due to improper input. From the models of vulnerabilities, it will be possible to analyze an object to determine if it contains vulnerabilities and how the vulnerabilities are exploitable.

Models of threats: Ideally, once one's vulnerabilities are known, steps are taken to eliminate them. Oftentimes however, strict prevention isn't possible (e.g. no patch exists, management resources are limited or performance requirements can't be maintained). In these cases, models of threats allow administrators to map between vulnerabilities and particularly dangerous attack avenues. If no prevention measures exist, security monitors can be configured to detect an attack's earliest stages and emergency measures employed. Our team has considerable experience with models of attack and their specification [LC02,TL00].

Models of policy: When a system is said to be secure, it is secure with respect to a policy. There are numerous classical security policies [DDL01], but many enterprises are concerned with policies specific to their activities. A model is needed in which site-specific policies can be specified. Such a model will reflect the resources of concern, how they are to be used (safety property) and when they must be available (liveness property). Such policies will be referenced by other models, for example concerned with detection of misuse and responses to an incident that are policy-aware. The detection of insider

misuse and the determination of appropriate responses will also depend on a policy, which can be discovered by a combination of expert specification and data mining. Policies are also needed to identify the amount of sanitization needed when sensitive information is shared, perhaps among enterprises experiencing an attack.

Design paradigms: While significant work has been done in intrusion detection algorithm development, little work has been done in IDS architecture research. Sophisticated attackers use distributed and stealthy activity to mask malicious behavior. Distributed IDS systems are required to monitor for this behavior [NRL03]. These systems must be architected to be scaleable, efficient and robust to single point failures.

Detection of threats: Our team has considerable experience in detecting attacks. Signature based methods detect known attacks with low false positives. Statistical anomaly detection can detect unknown attacks but with significant false positives. Specification based approaches fall in between but require expert knowledge to construct. Models of how different approaches vary in their ability to detect specific threats will lead to correlation methods that can combine alarms from different IDS sensors, maximizing the strengths of each.

Response and recovery from threats: Models of vulnerabilities, threats and models of a sensor's ability to detect exploits naturally lends itself to automating the response to potential exploits or even attacks in progress. Patch distribution systems could be configured in response to discovered vulnerabilities. Temporary restrictive firewall rules could be enacted in the face of detected known threats. Planning methods are used to automate complex responses, where the system is moved to a secure goal state (based upon models of vulnerabilities) using sequences of response transitions (with effects based upon a response model) in the face of continuing attack (based upon a threat model) [BMRL03]. Game theoretical models can be incorporated to anticipate the likely steps of an attacker and select corresponding responses necessary to leap ahead.

Reasoning about the security effectiveness of systems: One of the initial applications of formal methods was to verify the security of a system, usually what is called its "design" where the property was usually one of the classical security policies. Since our approach involves many models, the verification we will focus on involves the consistency and completeness of models. Since the models are abstractions of the system and mostly expressed in logic, the verification should be feasible [PSHCD98,PDD97,WLS97] – as compared with code verification. We will consider the verification of models that heretofore have not been considered verifiable. For example, from a model of specification based intrusion detection, it should be possible to verify that a particular intrusion detection system detects all attacks in a class or, even, all attacks that cause a policy to be violated [SAKZL03]. Similarly, it should be possible to verify that a response system will never result in resources reflected in a policy to become unavailable.

Experimental evaluation: Formal reasoning and verification can provide strong guarantees about the soundness of a particular security model. Unless these models are translated into real systems that protect computing assets, however, this work may be perceived to have little real value. A major focus of the CCSR will be the development of experimental evaluation procedures and the maintenance of a facility to perform the testing. Wired and wireless testbeds, configured with applications, client hosts, servers and network infrastructure devices will be constructed to ensure that security methods developed are not only theoretically sound but translate into prototype implementations that can be tested in a live environment. To aid in preliminary testing, sanitized static datasets containing both normal background and malicious behavior will be constructed and maintained

3. Education

Our team involves 7 universities: UC Davis, Stanford University, University of Colorado, Georgia Institute of Technology, University of Idaho, California State University San Jose, and California State University Sacramento. Many surveys of U.S. Government agencies and information technology industry report a significant shortage of engineers with skills in most aspects of computer security. This proposal addresses this shortage across the spectrum of needs.

3.1. Graduate Education

The training of graduate students is essential to fill effectively the many open security slots at U.S. universities and the open slots for researchers in other organizations. The team collectively will be employing approximately 45 graduate students on a continuous basis. We hope that least 35 of them will graduate with PhDs during the proposed project. The broad and ambitious scope of the project should ensure that these students have across-the-board-skills in security in addition to skills in particular aspects of the security problem. Besides obtaining a strong research background, the project's emphasis on improving security management will provide the graduate students with good practical skills – an essential mix in today's environment where much of the research is driven by the need to confront security problems brought on by real incidents. In addition to university training, our research and commercial partners will employ our students (and those at other universities) as interns, providing them with additional exposure to real security administration problems. We expect the research associated with the proposed project will provide material for additional graduate classes, included classes in network security, informed responses to incidents, misuse detection and response to insider attacks, large scale security, security evaluation of systems. The team will make available course material and experimental equipment.

3.2. Team Members Collaboration on Undergraduate Education and Curricula Development

The team of universities also intends to employ undergraduates in the research activities. All of the participating universities have long standing programs that involve undergraduates in research. The security administration infrastructure to be developed at UC Davis, will be a live classroom for students to work on security problems of the next generation. Students can try out new defense techniques as new incidents inevitably occur. It is well known that "hands-on" experience is often the best way to "turn on" young students. Hence, as appropriate, we would make the infrastructure available to local high schools. Through the infrastructure, we will convince the young people that security is a challenging and rewarding activity, and one that offers good opportunities for lifetime employment.

The infrastructure will provide a resource to the University of Idaho, an EPSCoR school. The collaborative partnership with the other campuses participants and access to this testbed will increase the research capabilities of Idaho researchers, enhancing their research competitiveness. Access to the testbed and a collaboration with the campuses will enhance the educational opportunities for Idaho's Scholarship-for-Service students.

There is a serious need for undergraduate classes in computer security. The participating universities all offer popular classes in this area, but this is unique among U.S. universities. UC Davis offers an undergraduate security class that regularly draws 200 students per year. To expand on the growing importance of security in the undergraduate curriculum UC Davis is planning on offering an emphasis in security which would include several additional classes besides the basic class. Network security, cryptography, and security administration would be likely offerings. The faculty of the participating universities are committed to the development of such new classes, including the creation of class notes, textbooks (to complement the popular and recently published text by Prof. Matt Bishop [B03]), and laboratory equipment – all expected products of the proposed project. Besides classes dedicated to computer security, we intend to use security examples in other classes. For example, examples on mitigation of attacks using planning and game theory would be provided for a core Artificial Intelligence class, on recovery from incidents would be part of a class on Fault-Tolerance, on reasoning about security models would be part of a class on formal methods, on coping with wide-spread attacks would be part of a class on networks.

We indent to migrate the material for undergraduate security curriculum beyond the setting of research universities, in part through our partners California State University Sacramento (CSUS), whose PI is Prof. Cui Zhang and California State University San Jose (SJSU) CSUS and SJSU are committed to creating new classes in computer security to complement their very popular graduate classes. Another California State University Campus, Hayward, has expressed an interest in participating in the project which the team will consider.

The teaming universities plan to co-ordinate laboratory exercises among their classes. For example, one class may design and implement a web server, and another may analyze the source code and the configuration of the web server to determine whether the server meets the security-related requirements. As another example, one university's class may configure a network to provide certain services, and a given level of security. A second university's class would then analyze the configuration and the software using the flaw hypothesis methodology. Or, the second university might try to "break in" while the first one tried

to detect the attacks. After the exercise, the two classes would switch roles. This would provide real-life training in defending systems and analyzing attacks. Forensic studies of systems successfully compromised would enhance the value of these exercises.

3.3. Certificate Program in Computer Security and Outreach

Beyond degree programs, the team will develop a Certificate program in security. There are numerous such programs available through the U.S. emphasizing security administration, but we believe this project will demonstrate a new approach to the problem through models and tools that emphasize prevention and rapid and policy-aware response to incidents. Placing more science into security administration should be exciting material for a certificate class.

To make the results of our work available to other college students in the geographical areas surrounding the participating universities, we will organize annual workshops where people from academia (as well as industry and government can participate). Our goal is to raise the awareness of computer security issues in as many students as possible.

3.4. Minority Student Participation

The UC Davis Department of Computer Science in the College of Engineering participates in several programs designed to broaden opportunities in engineering, mathematics and computer science for underrepresented groups. Among these are:

The UC Davis MESA Engineering Program (MEP), http://mesa.engineering.ucdavis.edu/, is an academic program that supports educationally disadvantaged students in attainment of four-year degrees in engineering or computer science. MEP's years of success have increased opportunities for thousands of students--especially those from groups historically underrepresented in math-based fields. MEP is a rigorous academic program that uses various components to support students including an academically-based peer community to provide mutual student support and motivation.

The UC Davis Women in Engineering (WIE), http://wie.ucdavis.edu/programs/index.htm , encourages women to pursue advanced degrees in engineering. It provides a supportive environment and opportunities for women and provides students with activities that reinforce and enhance their interest in engineering. WIE has focused on changing the academic environment for women in engineering by researching barriers to their success and developing programs that reduce or eliminate these barriers. Our programs for students focus on participation in their engineering education from pre-college through the graduate level.

The Minority Undergraduate Research in the Physical and Mathematical Sciences Program (MURPPS), http://maxwell.ucdavis.edu/~murpps/ , is a UC Davis undergraduate mentoring program designed to increase the number of disadvantaged students who pursue graduate studies in the physical and mathematical sciences by offering students the chance to work with professors on research projects relevant to their major. The goal of MURPPS is to help create a diverse post-graduate population in the Physical and Mathematical Sciences.

Women's Engineering Link (WEL) is a graduate-to-undergraduate mentoring program designed to address factors that cause undergraduate attrition and low rates of perseverance to graduate school. WEL's goals - first, to act as a retention tool for undergraduate students and, second, to enhance the pipeline into graduate school - are addressed through the mentoring of a mid-level undergraduate by a graduate student in (usually) the same major. The undergraduate student earns credit for working on various research projects for the graduate student for approximately three to five hours per week.

Through these programs, Karl Levitt and Matt Bishop, two of this project’s co-PIs, have been paired with students who have participated in computer security projects under their tutelage. In addition, several of the UC Davis graduate students who would participate in the proposed project co-supervise minority undergraduate students, introducing them to computer security research. The funding of this project will encourage continued and greater participation in these programs.

3.5. Participation of Grades 7-12

Currently security is not a topic covered in middle or high school. CCSR believes that computer security can be introduced into this curricula as the students will find the technical material current and

stimulating and, also, through computer security the students will learn about the ethical issues associated with computing and the Internet. We would not include in the curricula specific instructions on worm and virus creation, but stress the need for defense measures. We imagine that the students would apply the scientific methods of CCSR to secure their schools' systems as well as their own computers through personal firewalls.

4. Management Plan

4.1. The Team

The large scope and technical diversity of CCSR require that we implement a strong management plan to achieve success. UC Davis has assembled a team with the diverse research skills and experience in solving practical security problems required to create scientific methodologies that will significantly improve the security of large systems facing future security threats. Our management structure is intended to coordinate the activities of the participating campuses and other organizations (research not-for-profit and commercial) supporting UC Davis:

- UC Davis: PI – Prof. Karl Levitt (CCSR Director); co-PIs: Professors Biswanath Mukherjee, Matt Bishop, Felix Wu, Michael Gertz: plus other faculty (Raju Pandey, Matt Franklin), Premkumar Devanbu, Kwan Liu Ma, Zhendong Su), senior research staff (Drs. Jeff Rowe, Poornima Balasubramanyam, Tye Stallard) and students: Oversee the research and establish research direction, technical and development lead for the creation of the prototypes to demonstrate the scientific principles, integrate the research results and tools developed by the other organizations.
- SRI International: PI – Dr. Patrick Lincoln (CCSR co-Director), Sr. Staff: Phil Porras, Drew Dean, Grit Denker, Jonathan Millen, Vitaly Shmatikov, Steven Cheung: provide significant expertise in the areas of formal methods, secure system design, privacy enhancement. Based on the Emerald System, will assist in the creation of the cooperative architecture of the resilient network research thrust; significant technical assistance in other research tasks
- University of Idaho (EPSCoR partner): PIs – Professors James Alves-Foss, Deborah Frincke: Technical lead on formal analysis of models, significant participation in other research tasks
- Stanford University: PI: Prof. John Mitchel; Sr. Staff: Profs: Dan Boneh, David L. Dill, Mendel Rosenblum: participate in the development of models, on formal methods for evaluation of secure systems, on static analysis for exposing security-related vulnerabilities, on cryptographic approaches to privacy enhancement.
- University of Colorado: PI – Prof. Dennis Heimbinger; participate in the development of new architecture principles needed to realize resilient systems.
- Georgia Institute of Technology: PI – Prof. Wenke Lee; participate in the development of scientific principles underlying the detection and response to attacks; develop a scientifically-based methodology for testing the security of systems.
- California State University Sacramento (non-PhD granting partner): PI – Professor Cui Zhang: Creation of new undergraduate classes in computer security that utilize the principles to be developed in CCSR.
- California State University San Jose (non-PhD granting partner): PI – Professor Melody Moh: Creation of new undergraduate classes in computer security that utilize the principles to be developed in CCSR.
- Symantec Corp.: PIs – Carey Nachenberg, Juanita Koilpilla; Sr. Staff – Bruce McCorkendale, Gregg Hunter, Brian Nernacki: Provide significant assistance on mitigation of programmed attacks, current security administration practice, real data from ISPs, current approach to worm and virus mitigation, commercial monitoring and filtering devices for early prototypes.
- NetSquared: PI – Todd Heberlein: Provide technical assistance on real incidents, on monitoring techniques

Professor Karl Levitt, PI, will have overall technical responsibility for the project. For the past 17 years, he has been a Professor of Computer Science at UC Davis and founded its Computer Security Laboratory which now includes 10 faculty with significant interest in computer security, 4 senior staff members, 30 current graduate students (20 pursuing the PhD degree), and 2100 sq. feet of laboratory space. Prior to joining UC Davis, for 20 years Prof. Levitt was a member of the SRI International

Computer Science Laboratory, the final 4 as laboratory director. Thus he has 37 years of experience in research (computer security, formal methods, fault-tolerance, software engineering, computer architecture, programming languages), 20 of which have also involved significant responsibility in management and team building. At UC Davis and SRI he has been the PI for numerous projects, many of which were funded at a level exceeding $1M. Professor Levitt will report directly to Dean E. Lavernia of the UC Davis College of Engineering.

Dr. Patrick Lincoln will be co-director. He is Director of SRI International's Computer Security Laboratory and has done extensive research in formal methods applied to the verification of systems with respect to security and fault-tolerance properties. In addition, he has carried out research on the design of next generation secure systems.

Ms Patty Graves will have administrative responsibility, the duties including scheduling meetings, coordinating interaction among the team members and other participants, maintaining the project web site and issuing newsletters, coordinating the annual Security Industry Day.

UC Davis' co-PIs (Professor Biswanath Mukerhee, Matt Bishop, Felix Wu, Michael Gertz), Sr. faculty (Raju Pandey, Matt Franklin), Premkumar Devanbu, Kwan Liu Ma, Zhendong Su) have extensive experience in computer security and related research topics and have been PIs for numerous projects. Most of the co-PIs and other faculty have been active participants in UC Davis' Computer Security Laboratory for up to 15 years. The PIs from the other team members are all experienced in computer security and related fields, and have worked with UC Davis on numerous funded projects.

The team will also involve the following companies, all of whom are well-known for their security products and services: Hewlett Packard, NetScreen, Network Associates, Cisco, CloudShield, Honeywell, Intel, Promia, Captus, Silicon Defense, NetScreen, Sandia Corp., Lawrence Livermore National Laboratory. These companies will provide technical advice, products that could enhance the development of initial prototypes of defense systems, and software that can be the basis for initial security administration systems.

4.2. Team Management Structure

As indicated, Prof. Karl Levitt will be the PI for the proposed project, overseeing the research, prototype development and management activities. Dr. Patrick Lincoln will be co-Director, overseeing the research on model development and reasoning. Other faculty and senior staff will be responsible for the core research areas, prototype development, educational activities, and outreach. Levitt and Lincoln will be involved in setting the objectives and overall approach for each of the technical areas.

- Overall objectives and approach for the CCSR: Karl Levitt, Patrick Lincoln, Matt Bishop, Biswanath Mukerhee, Todd Heberlein
- Architecture development: Felix Wu, Dennis Heimbinger
- Threat modeling: Matt Bishop
- Policy modeling: Phil Porras
- Model development: Michael Gertz, Jonathan Millen
- Model reasoning: Jim Alves-Foss, John Mitchell
- Monitoring technology: Phil Porras, Drew Dean
- Static analysis of programs: Zhendong Su, David Dill
- Control Theoretic Approach to Attack Mitigation: Poornima Balasubramanyam
- Automated response research: Felix Wu
- Automated recovery research: John Knight
- Machine learning and datamining: Wenke Lee
- New techniques to enhance security administration: Premkumar Devanbu, Juanita Koilpilla
- Worm mitigation research: Jeff Rowe, Phil Porras, Carey Nachenberg
- Insider misuse research: Matt Bishop, Michael Gertz
- Cryptographic research, e.g., for privacy enhancement: Dan Boneh, Matt Franklin
- Information Visualization: Felix Wu, Kwan-Liu Ma
- Prototype development management: Tye Stallard
- Curricula development: Matt Bishop, Cui Zhang, Melody Moh, Deborah Frincke,
- Education outreach: Matt Bishop, Poornima Balasubramanyam
- Workshop Development: Tye Stallard

At the start of the project, the focus will be on defining the goals of the overarching problems to ensure that the research efforts have clear objectives. A week-long workshop will be organized for this purpose. It is expected that each research area at each university and company will conduct weekly research meetings where relevant graduate students, PIs, and other research personnel will discuss latest research results and new research ideas. The technical leads for the research areas will work closely with one other, meeting periodically (possibly monthly) to ensure fruitful collaboration, and will meet often with the leads for the overarching problems to assure research progress to the creation of the demonstration prototypes.

We will form a Technical Advisory Board (TAB) consisting of industry leaders (CTOs, VP-Engineering, etc.) to guide our research activities and technical vision. The TAB members will close the loop to keep our research on track with respect to meeting real-world needs for security technologies at an affordable cost. We have close ties with many companies with security products or services, such as: Netscreen, Cisco, HP, Intel, Honeywell, Network Associates; some of these will be partners in the research – as subcontractors or technical partners. We will tap into these companies for corporate support of our research project as well as for informal and formal advice. Similarly, successful conduct of this project will facilitate interactions among these companies (through our project's forum), thereby leading to a win-win situation for all concerned. The project team will meet with the TAB semi-annually to present progress. At the beginning of the project, the TAB, having significant expertise in many aspects of security, especially security administration, will help establish objectives for the CCSR. We expect to take on the problem of automating security administration, a problem that most members of the TAB acknowledge to be important but difficult. We expect to benefit from their attempts to provide some automation to security administration.

5. Knowledge Transfer Plans

The increasing interest in computer security provides the CCSR with many opportunities for the transfer of knowledge and actual prototypes that demonstrate the value of the scientific principles in support of security to be created.

Education material: We have outline our plans regarding the development of security-related curricula for the various levels: middle school, high school, certificate granting institutions, community colleges, undergraduate, and graduate.

Scientific principles: We expect the CCSR to produce new methodologies for the creation and evaluation of secure systems. These principles will be exposed to the community through annual workshops to be organized by the CCSR.

Prototypes: An important product of the CCSR will be working prototypes (including testbeds) that can be used on an experimental basis. Initially we will install the prototypes, especially those in support of resilient systems and security management, at the university sites. All of the participating universities manage large networks that are continuously under attack and could benefit from prototypes that detect and respond to such attacks and, also, support the overall management of these systems. For example, UC Davis' network has 40,000 users and 100 subnetworks; we have been encouraged by the Campus' Information Technology Department to install security-enhancing systems. Being both part of the University of California, UC Davis has a working relationship with Lawrence Livermore National Laboratory, who is eager to install the prototypes. The insider misuse detection and response thrust will produce working prototypes to deal with this important problem, systems that various Government agencies (FBI, NSA, CIA) have indicated a strong interest in acquiring. The Northern California Universities associated with CCSR have worked for various California State Government Agencies (including DMV, Franchise Tax Board, CalPers), who have expressed interest in installing the prototypes to be created. SRI International has numerous contacts with U.S. Government Agencies (DISA, AF Rome Labs, NSA), who have previously used SRI's research prototypes and have a continuing interest in the kinds of systems the be created by the CCSR. Symantec, our industry partner, has been a pioneer in offering security-enhancing products and has long terms to offer advanced products of the kind that can benefit from the new principles and prototypes to be offered by the CCSR. Other industry partners, who will participate on a non-funded basis, have long term relationships with the CCSR team and will be provided with the opportunity to install and evaluate the prototypes. All of the prototypes will be made available to the community through the CCSR's website.