National Cyber Leap Year Summit 2009

Co-Chairs' Report

September 16, 2009

Table of Contents

Summit Co-Chairs and Report Authors:

Hardware-Enabled Trust

- **Professor Fred Chong**, Director, Computer Engineering Program, Director, Greenscale Center for Energy-Efficient Computing, Professor, Department of Computer Science, UC Santa Barbara
- **Professor Ruby B. Lee**, Forrest G. Hamrick Professor in Engineering, Professor of Electrical Engineering and Computer Science, Director of Princeton Architecture Lab for Multimedia and Security, Princeton University
- **Dr. Claire Vishik**, Security, Trust & Privacy Policy & Technology Manager, Intel

Cyber Economics

- **Professor Alessandro Acquisti**, Associate Professor of Information Technology and Public Policy, Heinz College, Carnegie Mellon University
- **Dr. William Horne**, Research Manager, Systems Security Lab, HP Labs
- **Dr. Charles Palmer**, Senior Technical Advisor, Institute for Information Infrastructure Protection (I3P), Dartmouth and CTO for Security & Privacy, IBM Research

Moving Target Defense

- **Professor Anup K. Ghosh**, Chief Scientist & Research Professor, Center for Secure Information Systems, George Mason University
- **Dr. Dimitrios Pendarakis**, Research Staff Member & Manager, Secure Systems Group, IBM T.J. Watson Research Center
- **Professor William H. Sanders**, Donald Biggar Willett Professor of Engineering, Director Coordinated Science Laboratory and Information Trust Institute, University of Illinois

Digital Provenance

- **Mr. Eric Fleischman**, Technical Fellow, Boeing
- **Mr. Hugo Teufel III**, Director, Advisory Services, PricewaterhouseCoopers
- **Professor Gene Tsudik**, University of California, Irvine

Nature-Inspired Cyber Health

- **Professor Dipankar Dasgupta**, Professor, Department of Computer Science, University of Memphis, Director, Center for Information Assurance, Director, Intelligent Security Systems Research Laboratory
- **Dr. Steven Hofmeyr**, Research Engineer, Lawrence Berkeley National Laboratory
- **Professor Leor Weinberger**, Assistant Professor of Chemistry and Biochemistry, UC San Diego

1. Introduction

> "America's economic prosperity in the 21st century will depend on cybersecurity."
>
> President Obama, May 29, 2009

The Nation's economic progress and social well-being now depend as heavily on cyberspace assets as on interest rates, roads, and power plants, yet our digital infrastructure and its foundations are still far from providing the guarantees that can justify our reliance on them. The inadequacy of today's cyberspace mechanisms to support the core values underpinning our way of life has become a national problem. To respond to the President's call to secure our nation's cyber infrastructure, the White House Office of Science and Technology Policy (OSTP) and the agencies of the Federal Networking and Information Technology Research and Development (NITRD) Program have developed the Leap-Ahead Initiative. NITRD agencies include AHRQ, DARPA, DOE, EPA, NARA, NASA, NIH, NIST, NOAA, NSA, NSF, OSD, and the DOD research labs.)

As part of this initiative, the Government in October 2008 launched a National Cyber Leap Year to address the vulnerabilities of the digital infrastructure. That effort has proceeded on the premise that, while some progress on cybersecurity will be made by finding better solutions for today's problems, some of those problems may prove to be too difficult. The Leap Year has pursued a complementary approach: a search for ways to avoid having to solve the intractable problems. We call this approach changing the game, as in "if you are playing a game you cannot win, change the game!"

During the Leap Year, via a Request for Information (RFI) process coordinated by the NITRD Program, the technical community had an opportunity to submit ideas for changing the cyber game, for example, by:

- Morphing the board: changing the defensive terrain (permanently or adaptively) to make it harder for the attacker to maneuver and achieve his goals, or
- Changing the rules: laying the foundation for cyber civilization by changing norms to favor our society's values, or
- Raising the stakes: making the game less advantageous to the attacker by raising risk, lowering value, etc.

The 238 RFI responses that were submitted were synthesized by the NITRD Senior Steering Group for Cybersecurity R&D and five new games were identified. These new games have been chosen both because the change shifts our focus to new problems, and because there appear to be technologies and/or business cases on the horizon that would promote a change:

- Basing trust decisions on verified assertions (Digital Provenance)
- Attacks only work once if at all (Moving-target Defense)
- Knowing when we have been had (Hardware-enabled Trust)
- Move from forensics to real-time diagnosis (Nature-inspired Cyber Health)
- Crime does not pay (Cyber Economics)

As the culmination of the National Cyber Leap Year, the NITRD Program, with guidance from OSTP and the Office of the Assistant Secretary for Defense Networks and Information Integration, held a National Cyber Leap Year Summit during August 17-19, 2009, in Arlington, Virginia. Summit participants examined the forces of progress and inertia and recommended the most productive ways to induce the new games to materialize over the next decade. Two reports have been created as the result of the Summit:

1. **National Cyber Leap Year Summit 2009 Co-Chairs Report**: Written by the Summit Co-Chairs, this report presents the vision, the path, and next-step activities in the five game-changing directions as articulated by the Co-Chairs, based on the Summit discussions and Co-Chairs' expertise.
2. **National Cyber Leap Year Summit 2009 Participants' Ideas Report**: This report documents ideas that were introduced by participants and discussed and developed during the Summit. These ideas are presented to the community for inspiration and follow-on activities.

Taming this new frontier will require the contributions of many. The Summit, as the National Cyber Leap Year itself, should be seen as a tool for the community to use to build the shared way forward. The Summit reports clarify destinations with specific instantiations of the game changes and make the path plainly visible through practical action plans. For those who wish to begin immediately on next-step activities, the Summit community should be a great source of traveling companions.

The Summit's outcomes are provided as input to the Administration's cybersecurity R&D agenda and as strategies for public-private actions to secure the Nation's digital future.

More information about the National Cyber Leap Year and how to get involved can be obtained at: http://www.nitrd.gov.

The Summit was managed by QinetiQ North America at the request of the NITRD Program, Office of the Assistant Secretary of Defense Networks and Information Integration, and the White House Office of Science and Technology Policy. Ideas and recommendations expressed in this report are solely those of the report authors.

2. Hardware-Enabled Trust

New Game: *Knowing when we've been had.*

This section explores **Hardware-Enabled Trust** as a path to this new game.

2.1 Introduction

Hardware can be the final sanctuary and foundation of trust in the computing environment, based on the technologies that can be developed in the area of hardware-enabled trust and security. With cyber threats steadily increasing in sophistication, hardware can provide a game-changing foundation upon which to build tomorrow's cyber infrastructure. But today's hardware still provides limited support for security and capabilities that do exist are often not fully utilized by software. The hardware of the future also must exhibit greater resilience to function effectively under attack.

Vision

Within ten years, based on game-changing research:

- We will build a computer that will not execute malware, just as the human body can harbor certain viruses without ill-effect.
- We will build hardware that is itself more trustworthy.
- We will be able to determine, by technical means, whether to trust a device, a software package or a network based on dynamically acquired trust information rooted in hardware and user-defined security policies.
- We will build a computer that functions even under attack, through built-in resiliency that guarantees critical services in the face of compromise.

For hardware-enabled trust, the game is ripe for change. Continually increasing transistor densities have ushered in an era of multi-core microprocessors and increasing computing power in embedded processors. This has dramatically reduced the cost of hardware support for trust and security. At the same time, the increasingly pervasive nature of networked computing devices in our digital society has made trust in these devices and the environment where they operate more critical than ever. This pivotal balance between cost and benefit suggests that a tipping point is at hand, security will be increasingly demanded throughout the ecosystem. Industry must be prepared to provide security based on new developments in research, some of which are outlined in this report.

The strategies we propose are game-changing because they do not assume a perfect world. Rather, we assume all software is vulnerable to attacks, and even hardware/physical attacks are likely with the billions of diverse mobile computing and communications devices that can be lost or stolen. However, our computing devices, not only the datacenters, networks, or cloud servers, will be designed from hardware on up, to provide intrinsic security. We believe the role of hardware in establishing a safer computing environment will grow. This is a game change from predominantly software security solutions; hardware security solutions will be harder to break, increasing the "work factor" for attackers and therefore serving as an additional deterrent. Hardware computer and device architecture will provide fundamental features to enable us to build more trustworthy software and systems, and hardware itself will be built to be more trustworthy. For example, new technologies will ensure that hardware will not inadvertently leak secrets or execute malware (even if penetrated by malware), and it will execute security-critical tasks even if partially compromised. With enhanced research activity in this area, basic security mechanisms will be seamlessly provided without impacting the performance, energy consumption, cost, and usability of commodity computers and on-line services. Moreover, intrinsically secure computing devices will be able to share provable trust information, confirming their trustworthiness.

We propose three very promising game-changing technology strategies:

- End-to-End Trust
- Enabling Hardware to Thwart Attacks
- Hardware-enabled Resilience

Cutting across these technology strategies is the need for broader capabilities and collaboration to facilitate development and adoption of these ideas.

- Academia and industry collaboration is essential for executing a plan to improve hardware security and hardware-enabled trust. Therefore, it is imperative that appropriate models of cooperation are developed immediately to incentivize the participants to engage in research, development, and testing of technologies and approaches to achieve these goals. These models need to permit the immediate start of the collaboration and ensure long-term commitment of all parties. They may include direct pilot program funding of industry-academia teams, funding of industry-academia research consortia, and tax incentives for companies that work closely with academia to prototype promising research as a first step towards deployment. Approaches to organizing international initiatives in this area also need to be defined quickly.
- Sustained research funding for the next 10 years in Hardware-enabled Trust technologies should be clearly announced, to draw researchers to establish careers and initiate ambitious research projects in this important area.
- An independent testing and rating institute should be established to inform device users of the comparative security positioning of their devices, in order to help increase consumer awareness of the security and trust features in hardware.
- Adoption of technologies developed by academia and industry research efforts needs to be encouraged. Consequently, it is advisable to establish re-usable testbeds to ensure that new technologies developed by research programs are ready for deployment.
- In order to achieve end-to-end trust, interoperability across trust domains is required. Consequently, the development of standards to enable trust infrastructure and sharing of trust properties, including security properties in hardware, is essential to achieve the stated goals.
- We need to establish public-domain data repositories of the expected behavior of software applications and hardware devices, for jump-starting research in this area.
- We propose establishing annual conferences for Hardware-enabled Trust sponsored by government organizations as well as professional organizations such as IEEE, ACM or IACR. The goal of these conferences is to foster collaboration among industry, academia, and government in the design of coherent architectures that incorporate

promising research technologies, to promote publication of prototype implementations resulting from industry-academia teams, and to facilitate their widespread deployment.

We address game-changing technologies in sections 2.2 through 2.6, and cross-cutting ideas in section 2.7 and 2.8.

2.2 Idea: End-to-End Trust

What does the change look like?

We define "end-to-end trust" as the ability to secure trust in a distributed heterogeneous environment. Technologies discussed in this section are hardware-based or hardware-enabled. End-to-end trust is a collection of technologies, behaviors, implementations, and infrastructure approaches that, when used consistently, can enable a predictable level of trust in the ecosystem. The key to building end-to-end trust lies in identifying "trust properties" that can be used to determine the state of health of the system. For a device, these properties can include evidence of the authenticity of the hardware, proof that the devices and its software have not been compromised, and other general or domain specific "trust" information. In addition, we need protocols to exchange this evidence of "trust" information, approaches to dynamic measurement of the health of the system, and a way to compose and evaluate the resulting "trust messages" in order to make a determination of the trustworthiness of a device.

Currently, devices and networks contain some information about their trustworthiness that helps them operate in their silo. Examples of such information include the fact that a mobile phone ID has not been blacklisted (e.g., because the device was stolen) as a condition of connecting to a mobile network; evidence of up-to-date security patching when connecting to a corporate network, or comparisons of configuration measurements obtained at run time with those stored in a TPM. This information, however, is not sufficient to evaluate the trust state of a device or network, and is expressed and transmitted in terms and via protocols that are not interoperable. For example, when exchanging security-critical messages between a smart phone and a PC, each device currently has limited ability to determine the trustworthy state of the other and to communicate it.

If we could ensure that all communicating devices are trustworthy and operate in a trustworthy environment (that is, if we could establish end-to-end trust), the computing environment would be safer because each entity participating in a transaction could either vouch for its "trustworthy" state, or refuse to participate in a transaction without obtaining evidence of remediation. In order to achieve this, we need a foundation of interoperability across devices, systems and networks. Our vision is of a future where each device, from a sensor to a PC or server, and each network can be trusted based on a set of hardware-enabled "trust attributes" that could be exchanged over common protocols using appropriate trusted infrastructure.

2.2.1 Description

In recent years, some attention has been dedicated to the study of trust in hardware, defined as adherence to expected behaviour, in components and systems. As the behavior and composition of computing and electronic systems became more complex, manufacturing of the system components was globally deployed, and connectivity of these systems became nearly universal; the definition of the expected behaviors grew more complicated. As a result, some authors now use trust establishment and the relationship between trustworthiness and trust as the basis of the definition of trust, e.g.,

Trust: the degree to which the user or a component depends on the trustworthiness of another component. Trust and trustworthiness are assumed to be measured on the same scale. Ideally, we trust a system or a device because it is trustworthy; a trustworthy system is trustable. Currently, the limited ability to establish and communicate trust that exists in generic devices and platforms is not sufficient to provide adequate levels of assurance across the computing environment. When determining a system's level of trust, some questions can be formulated reflecting the main areas of concern:

- Is the system secure (does it have the expected security properties and configuration)?
- Is it in good standing (did it sustain attacks or unauthorized modifications)?
- Is it trustworthy (can it be trusted for the types of tasks it is expected to perform)? Even if "broken" in the future, can past operations be trusted?
- Is the system genuine (does it comprise only authentic components that are correctly implemented and are those parts assembled in a genuine way for both hardware and software)?
- What was its path from manufacturing to deployment?
- Was the design of its elements compliant with the best industry and technology practices?

If technology were available that could dynamically answer most of these questions in an automated fashion for all components of the ecosystem, the state of security assurance would be more on par with the dynamic nature of today's computing environment.

Consequently, we think that game-changing innovation could be introduced by focusing on the following activities:

- A canonical set of security and trust properties supplemented by domain-specific information. These properties will attest to the trusted state of a device or a system, will be available dynamically to help discover trustworthy resources, and will be rooted in hardware.
- Protocols to communicate this information
- Infrastructure to verify and transmit this information
- Process to compose elements of this information into evidence or a "trust message" and evaluate it
- Approach to support dynamic measurement of relevant parameters to ensure trust information is refreshed as appropriate.
- Interoperability to support this functionality across various silos.
- Privacy-protecting, un-forgeable, and trustable device identify to support reliable attribution and manage connections to and from the device.

Additional technologies that help support end-to-end trust:

- System "DNA" that vouches for system authenticity from the bottom up. The "DNA" itself may not be disclosable for privacy reasons, but it needs to be able to attest to the system's success in passing authenticity tests.
- Authoritative whitelist repository of the signatures for components and software signatures

Ideally, devices will be able to exchange trust messages prior to accepting connections or messages.

On a non-technical level, we must create incentives to define and deploy trust technologies (as well as to understand why the first generation of these technologies was not commonly deployed). We also need to ensure that trust features do not require significant tradeoffs in performance and usability.

Finally, in terms of system development, it would be attractive to build systems from the bottom up with constraints that can enforce safer behavior at all levels. Object-oriented programming can provide some guidance for approaches in this area. Hierarchical trust models that are currently used in most systems have multiple dependencies (software needs to trust other software and operating systems). It would be interesting and productive to consider replacing these models with a new generation of trust models rooted in hardware.

2.2.2 Inertia

Why haven't we done this before?

The previous stage of the development of the connected heterogeneous distributed environment was dedicated to increasing reach (including computing power) and connectivity, and ensuring a sufficient level of interoperability to support connectivity on a larger scale. Elements of trust that were developed were domain-oriented (e.g., adapted for mobile phones operating on the same types of networks). At the time these elements were designed, interaction among different classes of devices was not yet envisioned as very important.

Additionally, large-scale trusted environments represent a new area. The first generation of technologies in this field were designed without much consideration of issues of deployment and flexibility, as evidenced, e.g., by limited user and ecosystem adoption of Trusted Computing Group (TCG) technologies.

2.2.3 Progress

Why is it technically and environmentally feasible now? What would mitigate our doubts?

Today, the security concerns caused by the diversity of devices and networks in computing environments are coming to the fore. As the computational capabilities of devices other than PCs increase, so do the security threats and malware adapted for these devices. The community is more ready than ever before to address the issue of trust in a comprehensive fashion. There is a clear need to extend end-to-end trust to electronic processes. In a process where diverse devices, e.g., PCs and smart phones, are equally engaged, trust needs to be extended to all entities taking part in the process.

Moreover, the existence of the first generation of trusted computing technologies highlights both the promise and serious shortcomings of the current approaches. We can improve the next generation of technologies based on the lessons learned up to now, resulting in more successful adoption. Drawbacks of the current generation of technology as well as the role of the ecosystem and infrastructure have also been highlighted and the technical community has benefited from these lessons.

Finally, each domain (from mobile telephony to computer platforms and TCP/IP networks) has developed some level of understanding of what trust means in that environment and limited mechanisms to support a certain level of trust (e.g., the ability to blacklist rogue devices in mobile networks). This knowledge can be used to build a viable approach with some level of commonality across domains, as well as domain-specific features.

2.2.4 Action Plan

What is a reasonable path to pursue? What will accelerate the change?

Short term:

- Establish an operational pilot implementing these concepts, in the short term based on the Trusted Computing Group (TCG) technologies.
- Develop a national trust infrastructure testbed that deploys and instruments TCG-based technologies and associated system architectures (Network Admission Control [NAC] and/or Trusted Network Connection [TNC], trusted virtualization) in a distributed testbed for academic, industry, and government research collaboration. Initially, the testbed can help assess deployability of currently available trusted computing technologies, including Trusted Platform Module Management (TPM)-based device identity, NAC, hypervisor-based attestation, and domain isolation, and include all currently available platforms, e.g., PCs, servers, routers, mobile phones. More importantly, the testbed needs to be designed so that it can be re-used for multiple projects and adapted for the next generation of the technology, to ensure early assessment of new ideas. The estimated cost is two million dollars over three years.
- Establish a study group to assess (beyond currently available market research) the reasons why TCG concepts have not gained more traction and whether it is possible to incentivize the stakeholders to enable operating systems, infrastructure, viable applications, and a sustainable model of deployment or if new approaches are more viable. This group will also generate recommendations for the new generation of technologies. This group needs to include an active participation of verticals, such as health, financial, Supervisory Control and Data Acquisition (SCADA), in order to assess the views of the users (as opposed to designers and developers) of trusted computing technologies. The estimated cost is two hundred thousand dollars.
- Establish a representative task force that will assess steps necessary to develop end-to-end trust technologies, starting with the common and domain-specific set of trust properties. The task force, international in scope and comprising industry, academia, and government representatives, must propose concrete recommendations and solutions. It needs to include a forum of verticals to establish a general approach to standards in inter-trustability.

Long term:

- Establish an independent body (Non-government organization co-funded by industry and government) to test and score trustworthiness of devices and networks; this could be performed by a security testing institute, described later in this report. This organization could also provide tools for self-testing where appropriate.
- Ensure continuous operation of a trusted infrastructure testbed where new technologies in this area could be tested in order to detect issues at early stages. The testbed (described at its initial stage in the section above) will be operated by a consortium consisting of academia, industry, and government. See above for cost.

- Develop standards for inter-trustability (canonical set of trust properties, domain-defined parameters, dynamic measurements, attestation protocols, and infrastructure requirements). The developers of these standards should collaborate with the testbed group.
  - Identify and develop standards for device identification – Include a way to describe a system or device top-to-bottom starting at integrated-circuit (IC) levels
- Define additional infrastructure services needed to carry out end-to-end trust, such as:
  - Compromised or revoked key replacement services
  - Repository of white-listed hardware and software entities and their signatures or other representative artifacts
  - Deployment service to build support for these features in the ecosystem, such as various operating systems.
- Sustain the level of research activity in the area of end-to-end trust by supporting research aimed at defining trust technologies for the new generation of devices and networks. The research needs to include innovative trust models, novel approaches to encryption, ways to capture and communicate trust "messages," and privacy-enhanced trust representations. The estimated cost is $10 million over three years.

Enabling Hardware to Thwart Attacks

Today, we do not know whether there is malware in our computers. We also cannot prevent inadvertent information leakage from our correctly-executing hardware. Tomorrow, our vision is to design computers that will *not* execute malware. Our personal computing devices will enable us to control the protection of our private information stored in on-line storage systems, and will not leak information through side-channel attacks.

We propose three major directions for hardware to counter attacks:

1. Trustworthy hardware that will not leak information.
2. Hardware that will not execute malware.
3. Hardware-assisted secure storage and self-protecting data.

2.3 Idea: Trustworthy hardware that will not leak information

Today, attackers can obtain secret or sensitive information from our computers by side-channel attacks without breaking any rules or security policies, but just by observing hardware behavior. This undermines strong cryptographic protections and strong software isolation provided by Virtual Machine (VM) technology.

Tomorrow, we want computers to have *leak-free hardware*, where trustworthy hardware components and systems do not leak information. Only very slow or inaccurate side-channel attacks would be possible - hence, significantly increasing the work factor for the attacker and changing the game.

2.3.1 Description

Hardware features, introduced into microprocessors and embedded processors to increase performance or decrease energy consumption, can be used by an adversary to leak information through side channels and covert channels. Since the microprocessor's clock is typically the fastest clock in a system, this can result is much faster information leakage than previously seen with covert or side channels based on operating-system, disk, or input-output devices. Inside-channel attacks, the attacker does not break any rules in the security policies implemented by the system, but merely observes the power usage or execution time of the victim computer. In recent software cache-based side-channel attacks, the attacker was able to deduce the entire secret encryption or signing key by just observing his own cache access times or the execution times taken by a software encryption program. Current software solutions are ad-hoc – each software implementation of an algorithm using secrets must be rewritten and carefully retargeted for each hardware implementation. The goal of this research direction is to design hardware that is itself more trustworthy in that it does not inadvertently leak secrets.

Hardware processor features like simultaneous multithreading (SMT) or hyper-threading, speculation, and branch prediction can also be used to leak secrets in covert-channel or side-channel attacks. Computer architects must be incentivized to design security-aware computer architecture that does not compromise security for the sake of performance, or vice versa.

Can we design hardware components, like cache subsystems, that are inherently leak-free? Alternatively, can we develop Computer-Aided Design (CAD) methodology and tools that can automatically convert circuits so that they thwart side-channel measurements of power usage or timing? For example, a post-design phase may take any circuit block, inject random values at the entry to a circuit block and remove them at the output to “mask” the circuit from side-channel measurements. Even if this masking is not perfect, it will significantly increase (by orders of magnitude) the attackers’ work factor to mount a successful side-channel attack.

Since many of these leak-free hardware techniques employ randomization to thwart attackers and increase their work-factor, there is synergy between this research thrust and that of Moving Target Defense.

2.3.2 Inertia

It is hard to enumerate all the potential side channels in a computer. This will be compounded by new technologies and paradigms such as multi-core chips, virtual machines, and cloud computing. In general, there is a tension between optimizing designs for performance or for security - although there are now a few concrete counter-examples that this need not be so. When attackers had much easier attack paths, they did not bother with side-channel attacks. Users are unaware of the existence of these side-channel attacks and the serious damage that can be done if critical data, like secret encryption keys, are leaked - voiding the confidentiality and integrity previously provided by strong encryption.

2.3.3 Progress

Why technically and/or environmentally is it feasible now?

Technically, we now have concrete examples that it is possible to defeat side-channel attacks and improve performance at the same time. For example, by rethinking the well-studied field of cache architectures (perhaps the most critical component for hardware performance), researchers were able to thwart attackers as well as improve cache performance, by using dynamic randomization of memory to cache mappings coupled with innovative circuit and micro-architectural optimizations. Other approaches have even used gate-level techniques to make simple microcontrollers leak free. Only very slow, hard to achieve, side-channel attacks may remain, increasing the difficulty of a successful side-channel attack by orders of magnitude.

Also, masking techniques for protecting an arbitrary hardware circuit by the use of tools that can transform a given net-list (hardware implementation block) to a “masked net-list” have been designed, and research is advancing in this area.

Environmentally, side-channel attacks in the commercial world have been very damaging, and thus there is increased interest in thwarting these attacks in both the private and public sectors.

2.3.4 Action Plan

- Provide tax incentives and other policy incentives for hardware companies to work with academia to design hardware that treats security as a first-class design goal, together with performance, power consumption and cost. Hardware must be designed to withstand a suite of attacks.
- Create a suite of security tests that a designer can use to test his design against known and potential side-channel attacks. Note that these tests may have to be written so they can be re-targetable for different processors and implementations. This same suite of tests can also be used by an independent organization to characterize or qualify the security of a given computer. This can be part of the National Information Safety Board (NISB), described later.
- Establish a competition to design new cache architectures, new memory architectures, or new processor architectures by collaborative industry-academic teams. Implement the best ones in open cycles the government may have in trusted fabrication lines. In alternate years, focus the competition on breaking the design. Hence, the competition will alternate between a design competition and a break-the-design competition.

2.3.5 Jump-Start

- In 90 days, we envision an RFP for academic-industry teams to build prototype hardware subsystems (e.g., caches) that defeat side channel attacks while improving performance. Fund a pilot program where a few industry-academia teams implement credible prototypes of selected design proposals that seem to be widely deployable. Silicon-based chip prototypes should be completed in two years, while FPGA-based prototypes should be completed in one year. The implemented designs are then made available to the public for extensive testing and attacking. The cost of this pilot program over 3 years is estimated to be $10M.

2.4 Idea: Hardware that will not execute Malware

Today, hardware blindly executes any software, including malware. Tomorrow, the game change we propose is that even if the computer is penetrated by malware, this malware will not be executed. This assumes an imperfect world where malware can exist in a computer but do no damage, just as viruses may exist in a healthy human body but not cause illness. Hardware will be designed to continuously measure and monitor normal behavior, and thus thwart the execution of many types of malware. Hardware will be designed to instinctively protect overuse of its resources, or other actions that damage the health and welfare of the system. In this regard, there is synergy between this research thrust and Nature-inspired Cyber Health.

2.4.1 Description

Program measurements often indicate a relatively narrow range of normal behavior. Models of normal behavior can be developed, often called process characterization in embedded control systems. Strong deviations from this normal behavior can be prevented by continuous hardware monitoring. Such deviant behavior can be quarantined and further checked, e.g., for legitimate but irregular or infrequent behavior, or legitimate surges of activity, before being allowed to execute. If this works, hardware will refuse to execute malware. Hardware can also do runtime checking of code integrity, not just load or launch time code integrity checking.

Hardware can also protect the measurements themselves as well as the process, so that it cannot be subverted by an attacker, including by a compromised operating system. The methodology must include sanitization features for malware already present and allow measuring of different parameters, state or dynamic paths. In addition to the normal behavior of software, hardware can also monitor the normal behavior of hardware. To detect hardware Trojans, state machines can be developed to characterize normal hardware behavior.

Hardware can also be designed to mimic instinctive behavior that protects against predatory actions, such as depletion of its resources or damage to its critical functions.

2.4.2 Inertia

False positives can be disruptive at worst and annoying at best. False negatives are also possible with undetected attacks because the system's behavior still looks normal. Continuous monitoring may take up significant bandwidth for storing the data. Not enough compute power, bandwidth or storage was available for such continuous monitoring in the past.

The research into instinctive computing is in its early stages.

2.4.3 Progress

There has been quite a bit of research in anomaly detection and characterizing normal behavior for networks and for software, but less has been done for characterizing hardware behavior. Deployment of existing research is in a primitive state. Some companies have products that provide robust characterization of certain embedded systems, e.g., process control systems. Such protection provides defenses against attacks that have not been seen before and thus are not in attack-signature databases.

Hardware collection can automatically collect the data and is non-by passable. Multi-core chips and cheaper hardware allow this monitoring and collection to be done without impacting performance. Hardware can protect the measurements and measurement procedure – an advantage over pure software monitoring.

By adding such measurement technology to SCADA systems, we can significantly improve the protection of security-critical process control systems of critical infrastructures such as the power grid.

2.4.4 Action Plan

- Support a 5-year research plan into software and hardware monitoring and normal behavior characterization of systems, from embedded systems to multicore systems to distributed systems. Identify the parameters and the methodology to measure varying system characteristics, and the normal behavior of a user or group of users. Research the methodologies for hardware to collect the measurements and to efficiently stream data for off-line analysis. This will also improve collection of audit trails.

- Together with the thrust for Nature-inspired Cyber Health, put out an RFP for new computer architectures that embody instinctive self-protection mechanisms.

2.4.5 Jump-Start

- In 90 days, issue an RFP for academic-industry teams to develop prototype systems which apply software behavior measurement and monitoring technology to selected low-entropy systems. For example, this could be a SCADA system or a web server. The implemented designs are then made available to the public for extensive testing and attacking. The design should be tested to see that it can handle legitimate peak loads, or irregularly scheduled jobs or activity. Identify what should be and what can be measured. The cost of this jump-start program is estimated to be $6M over 3 years.
- Make available the sanitized data characterizing different programs and systems that already exists in industrial and government labs, to jump-start the research in normal behavior characterization and methodologies, and attack patterns, for real applications.

2.5 Idea: Hardware-assisted secure storage and self-protecting data

Today, users are concerned about storing their secret or sensitive information in on-line storage (e.g., in Cloud storage). Privacy concerns prevent consumers from storing sensitive data, and confidentiality concerns prevent companies from storing proprietary information. Our vision is to develop user-controlled secure storage technologies that prevent adversaries from being able to view or modify such data, which could be stored in essentially un-trusted storage and transmitted over public networks. Of course, the Cloud storage provider should also make every attempt to provide secure and reliable on-line storage. In the longer term, we propose new architectures for self-protecting data, which can essentially be stored anywhere.

2.5.1 Description

One idea is to build fundamental architecture for secure key management and cryptographic processing into commodity client devices. Thus, each client device will have a minimal set of hardware-rooted, non-by passable security mechanisms that allows secure storage, retrieval or regeneration of a user's keys. The user's computing device can then automatically and seamlessly encrypt and hash all its sensitive information before transmitting it over public networks to on-line storage servers. This leverages decades of work in cryptography and security protocols to provide confidentiality and integrity of protected information, with new hardware-based key management techniques. Because the user may have many keys and key chains – including symmetric keys, private keys, public keys, and their certificates – the key management architecture should allow the keys themselves to be stored securely in on-line storage, i.e., not subject to device storage restrictions. This idea uses hardware-rooted trust in client devices to provide user control of Cloud storage, where the user may not completely trust the storage itself.

The Cloud storage provider should also provide the user convincing assurances of secure and reliable on-line storage. Storage area networks (SAN) and network attached storage (NAS) are no longer disk drives hiding behind server systems, but rather full-fledged network nodes themselves. They need to selectively share data with multiple clients with different security needs. Users must be allowed to specify access control policies. Data leakage attacks through covert channels must be controlled. The storage device and controller architecture should be revamped to include these new security needs.

Another idea is to implement a new architecture that provides self-protecting data. For example, this may be based on secure objects. The object defines the set of allowed operations for different users and programs. The allowed operations may be based on the requester presenting certain capabilities or tokens. Here, we envision hardware support for data encapsulated as self-protecting objects that control access to and operations on the data (through encryption or other means). Today, we must be concerned with securing storage and restricting the flow of data. Self-protecting objects would free us from these concerns and make it impossible for an attacker to exploit purloined objects without appropriate credentials. This game change could be a key enabler in ensuring the integrity and confidentiality of data for future systems ranging from a national health database to secure cloud services for small businesses.

The key to this game change is the shift from building systems that protect data flow and access, to data objects that inherently protect themselves. We envision cryptographic techniques that allow objects to be encrypted such that access to operations on those objects will depend upon the identity or key of a component. This game change has a strong relationship with the Digital Provenance area and depends upon similar issues of identity establishment and key distribution. Data objects also facilitate maintenance of digital provenance. These could leverage the data-tagging mechanisms used in previous architectures.

2.5.2 Inertia

Building secure key management and cryptographic hardware in commodity client devices is likely to increase the cost. Commodity devices are very cost-sensitive and even minor additions are difficult to include.

Building flexible security policy enforcement in storage device controllers was not done in the past because disk designers assumed that this would be done by front-end server systems.

New architectures like self-protecting objects require the entire software infrastructure to be built and applications to be migrated or recompiled. Data encapsulation requires increased storage capacity and data bandwidth at all levels of the system. Both increases have traditionally been viewed as impractical. Performance and power overheads for encryption and decryption have also traditionally limited commodity application.

2.5.3 Progress

To “do it right” rather than applying successive band-aids, the time may have come to jump-start a new architecture that includes security as one of its primary goals, together with performance and power. This new architecture may start with self-protecting data, rather than with computation.

Substantial increases in storage capacity and data bandwidth make encapsulation plausibly practical today and into the future. Advances in identity-based encryption, group-based cryptography, and metadata representation show promise and perhaps could be combined to achieve a practical cryptographic solution to achieve self-protecting objects.

2.5.4 Action Plan

Jump-Start: Put forth an RFP for a $10 million program to build prototypes to implement fundamental and flexible key-management mechanisms in commodity devices, e.g., a notebook PC, or an iPhone or cell phone. This may involve extensions to a microprocessor chip and the development of a chip prototype. Another pilot team may use TPM technology. Connect this client device prototype to cloud storage services and subject it to public testing and attacks.

Jump-Start: In 90 days, produce a Request for Proposals (RFP) for a $2 million program to build prototype systems that emulate self-protecting objects through trusted components and protocols that control every object operation. Complete these prototypes in one year to facilitate experimentation and innovation.

Long-Term: Sustained investment of $20 million over five years is needed to develop new architectures based on self-protecting objects, cryptographic techniques, and integrated hardware mechanisms to efficiently implement this game-changing technology.

2.6 Idea: Hardware-Enabled Resilience

Today, a compromised system does not guarantee the integrity or availability of critical services. Tomorrow, we envision resilient computer hardware that can guarantee the execution of critical services even while compromised. This will significantly increase the work factor of attackers by protecting critical services from corruption or denial of service. Today, if we find out that we’ve been had, there is no easy way to get back to a pristine state. Tomorrow, the hardware (together with trusted software) will restore this pristine state.

We envision future systems that provide these guarantees by leveraging techniques traditionally applied to achieve fault tolerance and apply these techniques to protect critical services from attack. This game change could be a key enabler for a future Internet immune from malware disruption.

2.6.1 Description

Specifically, we envision future systems that incorporate the following techniques in hardware: redundancy, diversity, check-pointing and recovery, and self-repair and evolution.

One instantiation of such future systems would be a multi-core processor incorporating some or all of these techniques. Multiple cores can be used to redundantly execute critical code, and majority-voting mechanisms can be used to inhibit compromised cores. The cores can be designed differently to provide hardware diversity and prevent simultaneous compromise of multiple cores. Hardware support for efficient check-pointing and recovery can be used to supplement redundancy by tolerating compromises. When an attack is detected, hardware can potentially roll back to an uncompromised state and replay computation with attack vectors removed.

Finally, in the spirit of moving-target defense, aspects of the design could be reconfigurable at the gate or architectural level. This reconfigurability can be exploited to repair malicious or vulnerable hardware, as well as evolve structures to become more resistant to attacks as vulnerabilities become known.

2.6.2 Inertia

In the past, these security techniques have been regarded as too expensive in terms of silicon resources, performance, and power. Resistance from the microprocessor industry to implement this game change could derail the change. Cost constraints from the embedded processor area could derail change in the majority of market segments. General lack of demand for advanced security features in many markets leads to lack of incentives to develop and implement such features.

2.6.3 Progress

Silicon cost of implementing these mechanisms has decreased dramatically. At the same time, the safety of computation has become increasingly important as our devices are used to process critical business transactions, control essential utilities, and manage our daily lives.

Microprocessor vendors are looking for ways to add value to their transistor-rich products. Hardware-supported security and resilience, if given proper incentives and optimized for low cost (area, performance, and power), could add such value.

2.6.4 Action Plan

Jump-Start: In 90 days, we envision a $1 million RFP for academic-industry teams to build prototype systems using Field Programmable Gate Array (FPGA) technology and discrete components. We expect such prototypes to be completed in one year to provide platforms for experimentation and innovation. We further recognize a pressing need to establish benchmarks and metrics in the first year to test and quantify developing systems' resilience to attack.

To enable industry adoption of this game change, realistic prototypes and industry involvement are essential. A major research program with substantial, sustained investment is needed to develop credible silicon-implementing resilience techniques. We envision a five-year, $50 million research program to support academic-industry teams to build multi-core chip prototypes incorporating resilience mechanisms.

An orthogonal action to provide incentive for this game change involves creating an independent, government test organization to evaluate and score product systems (see below).

Crosscutting Ideas

2.7 Idea: National Information Safety Board (NISB)

Establish an independent testing organization to test and rate the safety and security of computer systems. This organization will be an NGO co-funded by industry and government. We could call it the NISB, analogous to the National Transportation Safety Board (NTSB). The NISB rating would give manufacturers an incentive to improve the security of their hardware and system designs. Consumers would also be made more aware of the importance of security and the security differences among systems.

2.7.1 Description

The NISB would, with expert help from its members (government and industry, possibly academia), develop metrics for evaluating system safety and security. These metrics might be focused on vertical domains such as cell phones, financial applications, and Internet services. Different system configurations would be tested such as best practices and usermisconfigurations, analogous to NTSB crash tests conducted with and without seat belts. Scoring should be relatively coarse-grained since the metrics and tests cannot be comprehensive. The NISB will also define test strategy, including self-certification and other methods, in order to increase the reach and responsiveness. The NISB could also investigate critical security breaches and issue findings as to causes and mitigations, as well as potentially issue safety recalls and require reporting of incidents. The NISB will issue recommendations for research strategy in the area of security metrics; it may also fund or otherwise support such research.

2.7.2 Inertia

Security has not traditionally been regarded as a high priority in the evaluation of computer systems. Metrics for computer system safety and trustworthiness are difficult to define in a consistent and meaningful way. A process for applying these metrics that is objective and adapted to today’s fast product cycles is difficult to devise. Given short production cycles and lifetimes of most ICT products, it has been difficult to develop and sustain an operation that could cover a large subset of tens of thousands of systems and devices.

2.7.3 Progress

The National Cyber Leap Year (NCLY) Summit, with the support of the current Administration, hopes to make security a critical focus in system evaluation. Efforts such as the Common Vulnerabilities and Exposures (CVE) database, provide a starting point for vulnerabilities and test development. The size and diversity of the IT market requires a cooperation of all the stakeholders – industry, academia, government – and new models for defining metrics and designing and administering tests, including via self-testing.

2.7.4 Action Plan

Jump-Start: In 90 days, establish a charter for the NISB.

In one year, recruit core participants and develop an initial set of metrics and tests for computer systems in different vertical markets.

The NISB will require sustained and increasing support to maintain test coverage of an ever-changing set of systems and to investigate serious security incidents. We expect the NISB will require at least as much funding as the NTSB, which has an annual budget of $79 M (2008). This funding can be shared by government and industry. The NISB needs to be very nimble, in response to the constantly changing market, be able to recruit new expertise quickly, and develop creative approaches to testing. An NGO is the best form for such an organization.

2.8 Idea: Establish an Influence Cell

Today, too many leap-ahead ideas die on the vine because they are not implemented by vendors or because those who procure IT place more importance on factors other than security. Tomorrow, an Influence Cell will help attract the right players and influence methods to greatly increase the likelihood of change.

2.8.1 Description

The Influence Cell is a small office of about five people supported by consultants who are national influence experts. They mine the results of this Summit and other sources in search of game-changing opportunities. As they identify opportunities, they leverage the expertise of key leaders such as the Summit attendees and a cadre of senior people who participate in advisory boards and science boards. Those people work with the experts to select the most promising actionable opportunities. The influence experts conduct analyses to judge why past efforts have failed and craft strategies to guide action and measurement of progress.

2.8.2 Inertia

There have been successes in IT change, often implicitly using influence and change management. However, sometimes the successes have involved massive investment (e.g., high-performance supercomputing initiative) and sometimes they have simply been ad-hoc efforts that succeeded without explicit influence strategies. Most government efforts to shape the commercial IT marketplace (or even government IT practices) have fallen well short of their objectives.

2.8.3 Progress:

Over the past few years, much progress has been made in turning influence from an art into a science, e.g., see Influencer: The Power to Change Anything[1]. Furthermore, environmentally, the climate for improved IT security is far stronger now than ever before, e.g., never before have the White House, Congress, the Office of Management and Budget (OMB), and Federal agencies placed security so high on their agenda – the Comprehensive National Cybersecurity Initiative (CNCI) is utterly without precedent and offers a potent window of opportunity. Vendors have security higher on their agenda than ever before; however, commercial pressures have too often prevented security capabilities from being built, bought, or used.

2.8.4 Action Plan

- In 30 days, charter a team with a $500,000 budget to plan an Influence Summit and establish an Influence Cell
- In 60 days, hold an Influence Summit of 20 select people to critique and develop three proposed influence pilots
- In 90 days, establish an Influence Cell with a $3 million annual budget
- In a year, perform an independent assessment of the office; the assessment recommends whether the office should continue and, if so, how it might be improved.

2.9 Conclusion

Cyber threats are causing a steep increase in financial and economic losses to a degree that has raised cybersecurity to one of the top national priorities. With this increase in awareness and motivation comes a corresponding increase in opportunity. Decreasing hardware costs change the business equation. Industry sees the need for change. The time is right to begin to use hardware to enforce trust and support security at a much broader level than today.

[1] Patterson, K., Grenny, J., Maxfield, D., and McMillan, R. (2008). New York, McGraw-Hill, VitalSamarts, LLC

Developing strategies to take advantage of security and trust capabilities of hardware is extremely important. Coordinated action is needed now to create a foundation for a safer computing environment that is crucial for the national economy.

This report identifies a family of jump-start activities that, with modest investment, can catalyze the actions needed to leap ahead to develop much stronger hardware foundations. Pilots and prototypes can guide the way for longer term technological investments while developing the partnerships and organizations to leverage new incentives, provide critical community standards and services, and reshape cultures. The path to a secure future begins with these next steps.

3. Cyber Economics

New Game: *Crime doesn’t pay*

This section explores **Cyber Economics** as a path to this new game.

3.1 Introduction

The economics of cybersecurity reflects the recognition that information security problems are, fundamentally, issues of misaligned incentives and misallocated resources - and therefore economic problems that require economic, more than merely technical, game changing solutions. Accordingly, the Cyber-Economics group at the 2009 National Cyber Leap Year Summit identified four economic strategies through which research and policy efforts may spur game changes in cybersecurity:

1. MITIGATING INCOMPLETE INFORMATION: Mitigate incomplete and asymmetric information barriers that hamper efficient security decision-making at the individual and organizational levels.
2. INCENTIVES AND LIABILITIES: Leverage incentives and impose or redistribute liabilities to promote secure behavior and decision making among stakeholders.
3. REDUCING ATTACKERS’ PROFITABILITY: Promote legal, technical, and social changes that reduce attackers’ revenues or increase their costs, thus lowering the overall profitability (and attractiveness) of cybercrime.
4. MARKET ENFORCEABILITY: Ensure that proposed changes are enforceable with market mechanisms.

In addition to the above four directions, the Cyber-Economics group observed that the purpose of information and communication technologies is not to provide perfect security, but to enable society to accomplish other objectives. This implies that we should not focus on absolute but on relative concepts of security and reliability in an unavoidably, necessarily insecure world – a fundamental issue of costs and benefits: How good does the security of various systems have to be? Or, in other words: What level of technical insecurity do we decide to accept, and live with? (This approach is further discussed in the National Cyber Leap Year Summit 2009 Participants’ Ideas Report under the “Swimming with the Sharks” section.)

This document offers a number of research and policy recommendations based on a subset of the ideas discussed at the Summit that relate to the four strategies identified above. Further insight on the discussions of the costs, benefits, feasibility or potential unintended consequences of various game-changing ideas and recommendations can be found in the National Cyber Leap Year Summit 2009 Participants' Ideas Report.[2]

3.2 Mitigating Incomplete Information

Meaningful economic analysis of security problems cannot be achieved without representative data, rigorous metrics, and theoretical models to analyze those data. Unfortunately, security relevant data today is either scattered or, in fact, largely unavailable. Organizations have few incentives to publicly share information about attacks and infection rates. Consumers are both unmotivated and unable to report data to anybody who would be in a position to react. Furthermore, for the data that we do collect, we often lack a clear understanding of how to process that information for security signals -- or even if we are collecting the right information. Our limited understanding of how to construct security metrics that are economically meaningful also contributes to this lack of clarity around what data to collect and how. These deficiencies create a significant barrier to efficient policy making, corporate investment, and individual behavior: Do we invest enough, too little, or too much in security? Are the costs of a certain technology worth the risk it mitigates? How can we make more efficient security investments?

Several ideas proposed and debated by the Cyber Economics group focused around these problems. The consensus among participants was that better data and metrics would enable a variety of economic analyses that could vastly improve the state of security. We could more easily evaluate the cost-benefit tradeoffs of various approaches to security, have a clearer understanding of risk, and enable better security investment decision making. We could understand the behavior of attackers better to assist in investigations and forensic analysis of security incidents.

3.2.1 Recommendations

Based on the discussion at the Summit, we recommend a combination of short-term and medium terms research and policy initiatives. Such initiatives are only briefly highlighted here. As noted above, the interested reader should find additional details and possible jump-start plans in the National Cyber Leap Year Summit 2009 Participants' Ideas Report.

1. **Interdisciplinary Workshop on Economic-relevant Cybersecurity Data and Metrics**. Lack of reliable and representative data about cybersecurity incidents, investments, and trade-offs hampers decision making at the government, corporate, and individual levels. Furthermore, lack of agreement on standardized and lack of reliable and representative data about cybersecurity incidents economic relevant metrics impairs our ability to analyze the data. To address this problem, we recommend holding an interdisciplinary

[2] We would also like to refer the interested reader to other reports that, in recent years, have covered similar issues: "National Cyber Security Research and Development Challenges," I3P, 2009; "Toward a Safer and More Secure Cyberspace," CSTB and DEPS, 2007; "Report to the President on Cyber Security: A Crisis of Prioritization," PITAC, 2005; "Ensuring (and Insuring?) Critical Information Infrastructure Protection", Rueschlikon Conference on Information Policy, 2005; and "Four Grand Challenges in Trustworthy Computing," Second Conference on Grand Research Challenges in Computer Science and Engineering, 2003.

workshop to address how we can choose, collect, standardize, and share data on incidents, attacks, infection rates, and security related losses. The workshop would bring together initiatives already discussed in related, but scattered, efforts (such as those by CERT, various ISPs, [central] banks, as well as specific efforts by Securitymetrics.org, the Open Web Application Security Project, the Center for Internet Security, MetriSec, and so forth), extending them to focus on the economic significance and purpose of those metrics. The workshop would address both theoretical and practical challenges. Workshop attendees should also include government representation from the Office of Science and Technology Policy (OSTP), the Council of Economic Advisor, and the National Bureau of Economic Research, in addition to academic and industry representation.

2. **Incentivizing Information Sharing.**
   We recommend focusing policy efforts towards the creation of a public repository of data, by incentivizing (or, in fact, mandating) public and private sector organizations to share and disclose data on incidents, attacks, infection rates, and – where possible – security related losses. By making such information publicly available, policy-making, research, and the industry as a whole would benefit. While details about this idea are available in the National Cyber Leap Year Summit 2009 Participants' Ideas Report, we note here that one of the primary challenges in fostering this change lies in finding proper mechanisms to incentivize firms to contribute such information, striking the right balance between protecting firms' confidentiality (through anonymization and aggregation) and guaranteeing sufficient social utility of detailed micro data.

3. **Cyber-NTSB.**
   We suggest the genesis of an entity similar to the National Transportation Safety Board (NTSB) - an independent Federal agency charged by Congress with investigating every civil aviation accident in the United States or significant accidents in the other modes of transportation, and with issuing safety recommendations aimed at preventing future accidents. A similar organization in the cybersecurity area would be charged with investigating major security breaches and incidents and issuing public recommendations aimed at preventing similar attacks. US businesses may be obligated to cooperate with investigations. This organization would help solve the information gap that currently hinders efficient security decision making.

4. **Funding for Cyber Economics and Understanding Security Behavior.**
   We recommend that more funding become available to address some of the basic issues in cyber economics, such as metrics, incentives, liability, and enforcement. Furthermore, we note that the mere provision of additional information for market players may not change their incentives to act on security threats and therefore their behavior: economic incentives fall flat without understanding of human factors, because of behavioral and cognitive biases that affect individual consumers and corporate agents. Hence, we also recommend promoting interdisciplinary research on cybersecurity behavior, spanning psychology, human-computer interaction, security usability, human factor, behavioral economics, and behavioral decision research. Because of the need for fundamental answers to some of these questions, funding in the above areas should be directed primarily towards academic, nonprofit, and industrial research organizations through programs with interdisciplinary requirements and interdisciplinary review processes.

However, in addition to traditional research funding models, funding agencies could experiment with non-traditional approaches, such as requiring that government-funded cybersecurity research projects or procurements include a cyber economics component or collaboration, inter-agency sponsorships, or Small Business Innovation Research (SBIR) and Small Business Technology Transfer (STTR) programs used to draw in small businesses to help address these research challenges.

Numerous other ideas were discussed at the Summit. They include: requesting that vendors of security systems disclose an impact analysis of the benefits and costs of their products to various stakeholders, similar to an Environmental Impact Analysis (including forecasted users efforts); and developing a risk adjusted return on investment methodology. Details about and critiques of these ideas are also available in the National Cyber Leap Year Summit 2009 Participants' Ideas Report.

3.3 Incentives and Liabilities

Generating incentives to jump-start a market (for instance, a cyber-insurance market for vendors or for business users of security products) or imposing costs to affect behavior (for instance, in the form of liabilities or accountability requirements) are traditional economic tools. They work by influencing stakeholders' profit functions, therefore affecting their behavior so that they conform to some desired objectives. Understanding and influencing security stakeholders' incentive structures is, therefore, a key to getting stakeholders to behave in a way that will improve overall security; for example, economists argue that the burden of preventing distributed denial of service attacks should fall on the operators of the networks from which the attacks originate - because they are arguably the entities facing the lowest costs to remedy the situation. However, what form of product or service accountability in the realm of cybersecurity may be actually beneficial has been a hotly debated topic in the literature and at the Summit; complicating factors include the inter-reliance between different information systems and components, the unknown liability path for open source products, the need not to stifle innovation, and the risk of imposing substantial compliance costs to vendors without actually improving overall security.

3.3.1 Recommendations

Based on the discussion at the Summit, we recommend a combination of short-term and medium terms research and policy initiatives:

1. **Workshop on Accountability in Cybersecurity.**
   We recommend convening a multidisciplinary workshop, perhaps even an annual series, on the technologies and policies to support accountability in cybersecurity. Established secure development practices standards for hardware and software would be evaluated and best practices identified. Furthermore, additional research on the issue and consequences of vendor liability and accountability would be encouraged.

2. **Vendor Accountability.**
   Also inspired by the results of the aforementioned workshop, we envision producers of software and hardware collaborating with consumers and industry groups on sets of baseline security and privacy development practices and guaranteed security capabilities. Initially, this effort will focus on key industries - such as healthcare and cyber-physical systems (where some efforts are already underway) - and incentives for following these

practices would likely take the form of consumer (industrial, trade association, government, etc.) procurement guidelines. (There are already efforts among some of the DHS working groups on how acquisition teams can support this effort.) In absence of industry compliance, however, regulators may have to consider stronger forms of accountability (such as vendor liability), guided by the results of research regarding complex trade-offs associated with imposing software liabilities. For instance, one idea discussed at the summit was that vendors may be held responsible for adhering to recognized software engineering standards (such as not allowing buffer overflows), and they would not be allowed to disclaim warranties of merchantability or fitness for the very purposes described in their products' operating manuals.

3. **Empowering ISPs, Registrars, and Registries.**
   We suggest considering policies to empower ISPs, Registrars, and Registries to take action to stop cyber-criminal activities – while also considering their accountability for failures to address clearly abusive behaviors. A pilot government program could be enacted to reward ISPs who discover, repair, and clean computers infected with crimeware – in order to realign the incentives of ISPs to keep their networks from harboring crimeware. The results of the pilot could be used to evaluate whether such interventions might be helpful on a large scale.

4. **Fostering the Cyberinsurance Market.**
   We recommend developing appropriate policies to incentivize the growth of a healthy cyberinsurance market. The goal of a cyberinsurance market is not just to spread security risk but also promote the adoption of security best practices and efficient levels of investment in cybersecurity - in both the business consumer and in the vendor communities (insured parties have an incentive to lower insurance cost by increasing their prevention investments). An additional social value of cyberinsurance would provide reliable risk pricing for internal decision-making. Gathering better data and metrics on incidents, attacks, and infection rates may make the growth of a cyberinsurance market more likely. However, one of the primary challenges is to define harmonized guidelines and train vendors and consumers to implement and audit adherence to those guidelines. Another challenge lies in dealing with cumulated risk (through diversity and/or appropriate financial instruments). The goal would be to adapt insurance models to the specific characteristics of cyber-risk.

Other ideas discussed at the Summit included establishing property rights for personal information. Details about recommendations and ideas are discussed in the National Cyber Leap Year Summit 2009 Participants' Ideas Report.

3.4 Reducing Attackers Profitability

From an economic perspective, defenders of information systems can use two (complementary) strategies: increase the costs incurred by the attacker when he is trying to breach a system, or reduce the benefit the attacker expects to receive by mounting a successful attack. These goals can be achieved through technology and policy/policing - and economic theory can help clarify where different approaches will be most effective, differentiating across strategies by highlighting their various costs and benefits.

Several ideas along these lines were discussed at the Summit, including: having a more diverse technology infrastructure, including multiple, isolated, purpose-built virtual networks;

considering ways to disrupt the attackers with decoys; and moving away from identity-based authentication to behavior-based authentication. Details about these ideas are discussed in the National Cyber Leap Year Summit 2009 Participants' Ideas Report, and further interdisciplinary and cross-sector research in these areas is recommended.

3.5 Market Enforceability

Incentives, liabilities, and changes in the attackers' profit function must be "real" - not just hypothetical - to actually affect market players. This requires efforts aimed towards the actual enforcement of the game-change - including legal, regulatory, and institutional mechanisms. Ideas in this area discussed during the summit included an international enforcement agency to address current international law and treaties and the creation of a centralized organization where stakeholders who traditionally have no effective means to report and obtain recourse for security incidents. Details about their benefits, costs, and feasibility are discussed in the National Cyber Leap Year Summit 2009 Participants' Ideas Report. Again, further interdisciplinary and cross-sector research in these areas is recommended.

4. Moving Target Defense

New Game: *Attacks only work once if at all.*

This section explores **Moving Target Defense** as a path to this new game.

4.1 Motivation

4.1.1 The Need

In the current game, our systems are built to operate in a relatively static configuration. For example, addresses, names, software stacks, networks, and various configuration parameters remain relatively static over relatively long periods of time. This static approach is a legacy of information technology system design for simplicity and elegance in a time when malicious exploitation of system vulnerabilities was not a concern.

In order to be effective, adversaries must know a particular vulnerability of a system. The longer the vulnerability of a system exists, the more likely it is to be discovered and then exploited. Many system vulnerabilities are published by researchers and software vendors in order for system owners to patch those vulnerabilities. A system that remains unpatched is vulnerable to exploitation. Vulnerabilities that are not publicly disclosed are called zero-day vulnerabilities, and are known to a limited set of people. Zero-day vulnerabilities present a large risk to system owners because without knowledge of the vulnerability, they have no way to patch it.

It is now clear that static systems present a substantial advantage to attackers. Attackers can observe the operation of key IT systems over long periods of time and plan attacks at their leisure, having mapped out an inventory of assets, vulnerabilities, and exploits. Additionally, attackers can anticipate likely responses and deploy attacks that escalate in sophistication as defenders deploy better defenses. Attackers can afford to invest significant resources in developing attacks since they can often be used repeatedly from one system to another.

Current approaches to addressing this problem are to remove bugs from software at the source, patch software as rapidly and uniformly as possible, and identify malicious attacks against software. The first approach of perfect software development does not scale to complete protection because the complexity of software precludes perfection. The second approach of patch distribution is now standard practice in large enterprises and has proven difficult to keep ahead of the threat. It also does not provide protection against zero-day attacks. The last approach is predicated on having a signature or definition of the malicious attack in order to find it and potentially block it. However, the speed and agility of adversaries as well as simple polymorphic mechanisms that continuously change the signatures of attacks renders signature-based approaches largely ineffective.

The magnitude of this problem suggests that we need a radically new approach, or “game change," for IT system defense. To visualize the elements of the new game, observe that for attackers to exploit a system today, they must learn a vulnerability and hope that it is present long enough to exploit. For defenders to defeat attacks today, they must develop a signature of mal-ware or attacks and hope it is static long enough to block that attack. Malware writers developed mechanisms to rapidly change malware in order to defeat detection mechanisms. We, as defenders, should learn from this approach, and build systems that rapidly change, never allowing the exploitation of a particular vulnerability to impair the ability of a system to perform its mission/function, or if exploited once, not allowed to be exploitable again. If done correctly, this “moving target" defense can present a formidable obstacle to attackers since they depend on knowing a system's vulnerabilities a priori.

Therefore, a game-changing approach to building self-defending systems can and must be developed. Protecting systems (thus avoiding exposed vulnerabilities) to the greatest extent possible should still be the first goal. However, recognizing that absolute perfection in software or hard-ware is untenable, we propose an alternate strategy that continuously shifts the attack surface of the system.

We call this game-changing approach "Moving Target Defense." An important benefit of moving target defense is to decrease the known attack surface area of our systems to adversaries while simultaneously shifting it; a key challenge of moving target defense is to ensure that our systems remain dependable to their users and maintainable by their owners. By making the attack surface of software appear chaotic to adversaries, we force them to significantly increase the work effort to exploit vulnerabilities for every desired target. For instance, by the time an adversary discovers a vulnerability in a service, the service will have changed its attack surface area so that an-other exploit against that vulnerability will be ineffective.

4.1.2 The Vision

The Moving Target approach will enable an "end state" in which systems can actively evade attacks, becoming substantially more secure even if they have vulnerabilities. It will result in systems in which fewer attacks will be successful. Those that are successful will be less likely to negatively impact a system's mission/function and less likely to be successful again, while other systems will automatically reconfigure themselves to be resilient to the same attack vector.

Moving target strategies employ architectures where one or more system attributes are automatically changed in a way that make the system attack surface area appear unpredictable to attackers. These strategies are beneficial at both the level of individual, high-value systems as well as large, national scale systems that may employ them collectively in a coordinated manner. They first make it much harder for attackers to identify vulnerabilities in targets and second, prevent them from repeating the attack on the same system or other similar systems.

4.1.3 Why Now?

The traditional approach to security, predicated on avoiding all vulnerabilities, has run out of steam, as evidenced by the dramatically growing number of successful attacks and the increase in severity of attacks that occur. All indications, including the growth in reported vulnerabilities, the growth in the number of software patches that are issued, and the number of virus definitions that are released, point to the end of pierce-and-patch as a successful defense mechanism.

Fortunately, at the same time, technological innovations over the last few years are about to make moving target defense strategies feasible and practically deployable. These include low-cost high-processing capacity, virtualization and workload migration on commodity systems, widespread and redundant network connectivity, address space layout randomization, just-in-time compilers, online software distribution, Software as a Service, and advanced analytics. These key technologies provide the computing power, functional redundancy, and diversity needed to enable the development of moving target defense mechanisms. These ingredients, together with development of the moving target approaches as illustrated in the next sections, will enable the end state of systems that achieve security by evading attacks.

4.2 The Moving Target Approach

The moving target approach has two components. In particular, any moving target approach needs a “space to move,” and a “mechanism to move.” The space to move can be achieved by building redundancy into the functionality of a system, and diversity can be applied when building the redundant functional components in order to make it such that a successful attack that is mounted on one part of the system is not successful on another part. But providing the space to move is not enough; the system must change its configuration so that an attack that works once (on that particular configuration) will not work a second time.

The movement may be achieved in different ways that can be classified according to different dimensions. First, it can be state-independent or state-dependent, where the state can depend on the system itself and the environment in which the system operates. The environment captures the behavior of the attacker, external factors that affect system performance, and accidental faults that may occur. Second, the movement may be time-triggered or event-triggered. In the time-triggered approach, system movement occurs on a schedule at predefined times, and the decision of whether it occurs at the particular point in time may or may not be state-dependent. The event-triggered approach is always state-dependent, and is triggered by a change in system, attacker, fault, or performance-related state. Note that it may be possible to build moving-target systems that are both time- and event-triggered, where both types of mechanisms are used to trigger movement of a system within its possible behavior space.

Together, options regarding how to create the space to move, and what mechanism is used to move it, define a particular moving target approach. To illustrate the possible universe of moving target approaches, consider that the space to move can be created by building redundancy into communication, computation hardware, software, sensing, and control mechanisms. The trick to building this redundancy is to find approaches that are space- and time-efficient, in the sense that they do not add unacceptable cost to a particular system implementation and operation, and do not de-grade its performance in an unacceptable way. It’s also important to find approaches that can be designed and implemented efficiently so they do not unacceptably add to development cost or time. This necessitates that mechanisms to create movement space be automated, in the sense that they are instantiated in compilers, operating systems, or applications.

Note that (identical) redundancy is not enough, because if the attacker finds a vulnerability in one copy, that same vulnerability could likely be exploited in all copies. Thus redundancy must be coupled with diversity, ensuring that it is unlikely that a vulnerability in one redundant part of a system could affect another redundant part. Artificially creating diversity in software and hardware is an active research area, and recent research results (e.g., John Knight, Genesis framework) show much promise.

There are also many possible ways to cause a system to adapt or "create movement" within its possible spaces of behavior. For example, a very simple approach would be to build a system that causes system configurations to cycle in an unpredictable way to an attacker (see the Configuration-Space Randomization for Infrastructure idea), through different redundant functional components (see the Diversity in Software idea), a communication that uses diverse communication paths (see the Connectivity Diversity idea). A more sophisticated approach would make use of state information (regarding the system, environment, and attacker) to adapt itself using the various moving target mechanisms in an intelligent manner (see the Smart Motion Adaptation/Management idea).

While significant “leap ahead” research and development are still needed to achieve such moving target functionality, success would be achieving systems that provide desired functionality at expected service levels even while under continuous cyber attack.

4.3 The Path towards the Change

As presented in the Motivation section, the time is right to create a moving target active defense capability. Many of the required base capabilities exist, including the ability to construct redundant functional components (e.g., well-provisioned networks, alternative and multiple processors, and alternative OS and other software components) and to create diversity. What is needed to achieve a moving target defense is 1) a process to systematically build systems that automatically incorporate these alternate functional components that ensures diversity among component functions, and 2) demonstrably sound methods to move through the functionality space, actively con-figuring the system in an "optimal" sense with respect to the chosen security metrics, attributes, and system performance parameters.

Section 4.4 presents examples of ideas that could contribute to this vision, either by generating a large movement space, providing a mechanism to move within the space (thwarting a potential attacker), or both. For more examples of ideas that explore this movement space, please see the Moving Target Defense section of the National Cyber Leap Year Summit 2009 Participants' Ideas Report.

4.4 Examples of Moving Target Mechanisms

A few samples of moving target defense mechanisms are described below, in no particular order, that we believe hold significant promise for game changing defense.

4.4.1 Data Chunking and Decentralization (Distributed Data Shell Game)

As servers and databases contend with ever-growing amounts of data, data protection becomes an increasingly difficult problem. Not only does more data create a larger attack surface and a higher payoff for successful attackers, but it also causes devastating business losses if data availability is compromised, whether through system failure, network failure, or forces of nature. Traditional methods of increasing data availability, however, also have associated costs: require the purchase of additional servers, entail higher management complexity, and demand more costly backup systems, etc.

Cloud storage has the ability to allay some of these concerns, but most existing data cannot easily be moved into third-party cloud storage systems due to system integration problems, security and privacy concerns, and increased complexity in reasoning about availability.

We propose to promote wider adoption of secure chunking and redundant decentralization for data at rest, and for research applying the same principles to live systems such as relational databases.

As a specific example, we wish to highlight the Tahoe grid file system (http://allmydata.org/), a cross-platform open-source software solution which demonstrates both secure chunking and redundant data decentralization. Instead of storing all files on a single server and then enhancing availability through (e.g., hot swap) server replicas, Tahoe promotes an explicitly secure, fault-tolerant model: stored files are broken into pieces, encrypted, and the pieces are redundantly stored across arbitrarily many servers. Redundancy is achieved through erasure coding and can be varied to provide the desired level of availability guarantees. For instance, in a grid of ten storage servers, every file might be stored on five servers, of which any three need to be available to reconstruct the whole file. In addition, despite the chunking, each file has multiple coherent "views": depending on which view he was provided, one user might only be able to read a file, while another might be allowed to write to it as well.

Wider deployment of this type of file storage system would have an immediate impact on the quality of modern data protection. Built-in fault tolerance lowers server costs by allowing any machine with an excess of unused disk space to join the storage grid; because files are encrypted prior to storage, the individual storage grid nodes need not be trusted. Most important, by spreading data across a number of (potentially heterogeneous) machines and coupling the process with strong encryption, data storage as a whole is transformed into a moving target. An attacker can no longer compromise a single storage server and obtain all the data; indeed, without possession of the encryption keys, even a complete compromise of all the grid nodes provides the attacker with no usable data.

Systems like Tahoe are making these methods immediately usable for securely and availably storing files at rest; we propose that the methods be further reviewed, written up, and strongly evangelized as best practices in both government and industry. In parallel, we propose that research be encouraged into the efficient application of the same methods to more complex data organization systems, chiefly relational databases.

4.4.2 Decoys

The core idea of Decoys is to dynamically deploy a large number of fake targets that appear to attackers indistinguishable from the real targets. Decoys serve two main purposes: first, they serve as sensors that can detect new attack activity, automatically analyze attack patterns and learn, predict and prevent attacks at the early stages, before attackers reach the real targets. Second, a large number of dynamically (re-)deployed decoys will significantly increase the attacker's work factor by sowing confusion and uncertainty, while raising his profile.

Since decoys should not be normally accessed, they immediately pinpoint ongoing attacker activity, which can range from mapping out networks to launching exploits or denial of service at-tacks. Attackers end up at decoys because they are reachable over shortcuts or because they may bypass common access-control patterns. Detecting new or newly initiated attacks, together with slowing down the attacker, the defense has more time to evaluate a response or to discover unknown attacks (zero-day vulnerabilities or new ways of evading firewalls, anti-virus, or access controls).

Decoys can take different forms to effectively protect various security targets. They can emulate systems, virtual machines, applications, data, networks, or sensors and actuators. Decoys increase the attack surface while decreasing the probability of a successful attack on the real target and hence reduce the attack Return on Investment (ROI). To significantly slow down and frustrate the attackers, the ratio of decoy to real targets must be very high – for example, on the order 10000:1. This in essence vastly increases the attack surface that an attacker needs to cover before eventually zooming in on the real target (see also Honeypots and Honeynets). There are several ways to “slow down” attackers at decoys; approaches range from simple shallow multisystem emulations listening ranges of unused network addresses, to full fake run-time environments with fixed IP and real business application configurations (traps, jails) that are more difficult to distinguish from real targets, even for attackers taking control of the decoy.

Several recent technological developments make the dynamic, large-scale deployment of effective and low-cost decoys feasible. In particular, widespread adoption of virtualization technologies addresses several important scalability challenges. First, cloning of VMs becomes as easy as "forking" a process - copy write memory and storage might allow instant cloning even of fully deployed VMs at run-time. Second, resource optimization based on hardware- or OS-level virtualization enables prioritization of real targets to limit the overhead of decoys and reduce cost. Such optimization should be carefully planned to avoid revealing opportunities for attackers to distinguish decoys from real targets (e.g., response time or other side-channels). Furthermore, default configurations of NATed, and encrypted communication channels with appropriate access controls prevent attackers from easily identifying decoys by observing network traffic. Finally, advanced analytical capabilities to correlate large traffic streams in real-time, enable learning by observing attacks on random decoys to protect the real target.

Some moving target defense concepts related to decoys are already employed in Internet honeypots and honeynets that predominantly serve as a resource for academic and commercial security research communities interested in learning about new attacks and offense tools. Currently, honeypots are not cost-effective to set up or maintain for individual organizations whose only interest is protecting their own networks. The research organizations that do deploy them find that pooling and sharing honeypot data present complex technical and legal concerns. Furthermore, honeypots currently are typically used as data collection points: they are not designed to pursue or actively combat attackers.

We observe that a number of Internet Service Providers (ISPs) and other organizations have large amounts of allocated but unused IP addresses; these addresses constitute the so-called dark IP space. Because the borders of the dark IP space are not generally known and are constantly changing, they could be a good target movement space; in some sense, the problem thus far has been proposing an interesting target to make use of the movement.

We think it would be beneficial to bring together an international coalition of ISPs willing to use their dark space to jointly combat the spread of computer malware. Participating ISPs would take two measures: first, deploy automated monitoring across some or all of their dark space, effectively turning every dark IP address into a simplistic sensor-only honeypot; and second, share a trusted backchannel that allows exactly one kind of automated message to be sent from any participating ISP to another – "your user at this IP address sent packets to an address in my dark space, and is therefore likely compromised."

We explicitly do not suggest a particular policy for dealing with such messages; it's not appropriate, for instance, to always immediately disconnect a user receiving a likely-compromised notice from another ISP. Individual ISPs can best determine how to react to such notices; some might choose to notify the IT departments of affected commercial customers about a potential systems breach, or choose to temporarily disconnect repeat residential offenders and notify them that their machines are compromised.

Modern computer worms most frequently spread by randomly choosing IP addresses to attack. Because the current dark IP space owned by ISPs is immense and in flux, the end result is that it's likely that a coalition of cooperating ISPs could use their combined dark IP space as a moving target for worm detection and mitigation, and realize far greater success than similar individual efforts that have been heretofore attempted.

4.4.3 Robust Cryptographic Authentication

Phishing attacks are already a severe problem in the modern threat landscape and continue to advance in scope and sophistication. Rudimentary, low-yield, and unfocused phishing attacks are giving way to attacks aimed at specific high-value targets, such as executives and other individuals with access to privileged information. Modern phishers have also demonstrated the capability to make use of stolen information to more effectively target their attacks; attackers who com-promised TD Ameritrade in 2007 and extracted 6.3 million customer e-mail addresses – but not the corresponding usernames and passwords – immediately launched a phishing campaign against those addresses to obtain the users' credentials.

Setting aside the question of user training, phishing is successful at a technical level because current authentication methods largely employ static credentials. The Web is certainly the most prominent example of the problem: user authentication on the Web is overwhelmingly performed through the entry of a username and a password into an HTML form sent to the server for validation. On Web sites that employ it, TLS protects the user's credentials from eavesdropping while in flight to the server, but it doesn't protect them from the server itself. In other words, it is critical to understand that regardless of whether encryption is used at the transport layer, the server ultimately receives the credentials in raw form, exactly as entered by the user.

Consequently, if the user can be convinced to visit a phishing site and enter his credentials, the phishers will immediately be able to use the credentials to log into the real system they were impersonating. This is an example where a static system configuration, in this case attributes used to establish trust between two communicating systems, presents a static target to attackers. Ro-bust cryptographic authentication would change the game by employing cryptographic methods which enable secure authentication without transmitting the raw credentials for validation.

A specific example is Secure Remote Password (SRP) (http://srp.stanford.edu) protocol, which is a royalty-free, patent-unencumbered form of zero-knowledge, password-authenticated key exchange. In practice, instead of transmitting a password, SRP constructs a special proof that the user indeed knows the password, and sends that proof to the server for validation. This proof has several highly desirable properties: it incorporates random numbers and is hence different each time it is constructed, it reveals zero bits of knowledge about the password to both passive eavesdroppers and active man-in-the-middle adversaries, and most important, it is useless to a server that doesn't already know the user's password. In other words, if an SRP proof is transmitted to a phishing site, the phishers do not learn the user's true password and are unable to reuse the proof itself to authenticate to the real site.

We do not mean to suggest that it would be an overnight endeavor to change password authentication on the Web to a proof-based protocol like SRP. We are, however, pointing out that an SRP-enabled Web makes phishing pointless; by making static credentials obsolete, SRP trans-forms Web authentication to a moving target and phishers are left with entirely useless information.

It is also noteworthy to highlight the problem regarding the general adoption of cryptography. Apart from the enormous success with TLS on the Web, modern cryptography has had remarkably little luck with widespread adoption, despite providing a plethora of increasingly sophisticated benefits to its users. Several decades of experience suggest that the chief culprit inhibiting its adoption is the problem of key management: understanding how to provision, use, revoke and expire cryptographic keys has proven to be insurmountably complex for all but the most technically advanced organizations.

In light of this, we wish to stress the importance of re-evaluating alternative trust models in order to find ways to remove obstacles for adopting cryptography. In particular, we propose that a significant effort be devoted to evaluating identity-based cryptography (IBE) as a viable key-management model for enterprises and other controlled, hierarchical institutions, in lieu of the individual web-of-trust model which has proven all but unusable at scale. While not intended to be a universal solution, it is our opinion that techniques like IBE hold the potential to immediately dismantle numerous roadblocks to widespread adoption of cryptography in industry and government, and in particular to curtail the widespread use of robust cryptographic authentication for systems and services (like e-mail) where current security guarantees clearly fall short. Robust cryptographic authentication acts as a defense amplifier for almost all of its deployed services, turning them intrinsically into moving targets. (Of course, an amplifier is not a silver bullet: if a system is vulnerable to authentication bypass, no amount of cryptography will make it more secure. For systems that are reasonably secure today, however, robust cryptographic authentication provides strong and clear benefits.)

In conjunction with IBE and SRP, we believe there is value in exploring the use of diversity in end user authentication for both human users, smart devices (sensors and actuators), and connections between different software components. Diversity can apply to the manner in which authentication is performed (what credentials it is based on), the time at which an authentication request is triggered, and the communication channel used to exchange the authentication information. As an example, authentication may be performed using a combination of multiple bio-metrics (e.g., face, voice, keystroke), multiple tokens (e.g., PC/phone signature, key fobs) and over multiple communication channels (e.g., Web, e-mail, voice, text). Furthermore, authentication may be requested not only at a defined log-on instance, but possibly during random times during a session. Employing the use of diversity in the means, timing, and avenue for authentication represents a robust moving target mechanism that can protect against attackers that success-fully compromise one component of a system or obtain one piece of information (for example, a password or a key fob).

4.4.4 Smart Motion Adaptation/Management

The various moving target defense mechanisms introduce the movement space within which the configuration of defended systems can dynamically vary. At a "micro" or very fast time scale, these mechanisms are expected to operate in a decentralized and autonomous manner, without requiring global coordination. However, at the "macro" or slower time scale, the overall moving target defense system capabilities require a smart management and motion adaptation capability that continuously collects and analyzes vast amounts of sensed information and computes optimal moving target strategies based on sound mathematical models, the value of different targets, and the evolving threats. These strategies drive the automatic actuation of defended systems to preserve their confidentiality, availability and integrity.

Smart moving defense management nodes collect various sources of distributed sensed data such as attack (e.g., intrusion, anomaly), performance, and availability, and evaluate optimal response strategies based on the attacker’s actions, the cost of responses, uncertainty in the observed data, and the risk associated with different systems. Computation algorithms for the response strategies will be based on a variety of modeling techniques, including but not limited to game theory, machine learning, statistical analysis, control theory, and cognitive reasoning that manages dynamic system behavior.

4.4.5 Action Plan

In this report, we identified an area of innovation – achieving moving target defense – that is ripe for innovation as a key game changer in cybersecurity. A number of moving target defense ideas were discussed, some of which are highlighted in this report, others of which are documented in the accompanying Moving Target Defense Participants’ Ideas Report. In this section, we propose an action plan for moving beyond the idea stage to enable a true game change. Our action plan was derived from prior successes and failures in technology transition. The history of science and technology shows most ideas never reach widespread use for myriad reasons, but one consistent theme of transition failures is the lack of a robust value chain of players to effect game change.

Technology ideas often fall into the so-called “valley of death” that exists between research and operational deployment because different communities in the information technology ecosystem do not communicate. In an effort to bridge the gap, the NCLY Summit brought together a number of different communities – government, defense contractors, major software and hardware companies, start-ups, and academia for a two-and-a-half-day brainstorming, collaboration session to enable each of the players to become vested in its outcome.

Our action plan calls for the creation of a value chain of partners for ideas. We expect this summit to create an online marketplace of ideas to form value chains that will propel the ideas to market.

For example, a university professor may look at the "diverse software" idea and write a grant proposal to a government agency to develop a just-in-time compiler. The professor may team with a system integrator to get an early prototype built and evaluated. A government agency with an operational mission may fund an advanced technology demonstration (ATD) to deploy the JIT compiler to measure the effect of diversifying software on a government enterprise services application. Given measurable success, the university may spin out a start-up that is funded by venture capital or angel investment. Alternatively, the university may license the intellectual property to a major software vendor that builds and releases a new version of a standard compiler with this innovation built-in. Independent software vendors (ISVs) will be able to use the JIT compiler as part of their software development and distribution model to enable users or purchasers to benefit from this form of moving target defense.

As the example illustrates, for an idea to successfully transition from research to widespread use, a value chain must emerge, with each player incentivized to promote the technology. This requires communication among academia, industry, government, venture investors, and end users, to enable a formidable change in the cybersecurity environment.

4.5 Recommendations

We make the following recommendations for an action plan:

Fund U.S. Government moving target defense research at $30M/year – Both basic research and advanced research projects should be funded by U.S. Government Agencies to jump-start research, development, testing, and evaluation of potential moving target defense. (Editorial note: U.S. Government Agencies with research portfolios in cybersecurity include, but are not limited to, DARPA, DHS, DOD Service research organization, IARPA, NIST, NSA, NSF. For additional information please see www.nitrd.gov.)

Construct a public moving target research testbed to enable experimentation and evaluation of moving target defense technologies with operational and, preferably, mission-critical soft-ware. Public research testbeds are a vehicle for research experimentation that enable both government funded and non-government-funded researchers to evaluate their technical approaches. The DARPA intrusion detection evaluation experiments and corpus of data from the last decade is an example of a successful public testbed that leveraged a small amount of government funding to enable a large community of researchers to evaluate different intrusion detection algorithms. Similar efforts have evolved in other communities such as speech recognition.

Encourage private-public partnerships to research, develop, and pilot technologies. In order to substantially increase the odds of a moving target defense technology to reach widespread adoption, it is imperative for researchers to involve end users. In the case of moving target defense, researchers need to work with private IT enterprises to ensure that proposed solutions meet operational needs. Likewise, private enterprises can significantly impact research and transition by conducting in-house pilots of the technology.

Create online collaboration and electronic communities – One key success of the summit was the disparate groups it brought together. In order to keep the momentum going, these communities need to continue to engage through online collaboration and electronic communities. We are pleased to report this is already taking place using wikis, Web pages, mailing lists, discussion groups, and LinkedIn groups.

Encourage private investment in moving target technologies – Most contracted research will ultimately fail to transition technologies to market because contracts end before technology is market-ready. Private investment communities including venture capital, angel groups, and state economic development authorities need to be engaged. Fortunately, these groups are actively engaged in looking for the next great idea and teams. It is proposed to exploit commonly used forums, such as conferences, to promote the use of online communities. For example, the DeVenCi (Defense Venture Catalyst Initiative) introduces venture capitalists to emerging government-funded technologies. The MAVA (Mid-Atlantic Venture Association) and other industry groups hold similar forums.

Encourage open source software efforts – Open source software is a well-established method for gaining widespread adoption of software. The summit has promoted many of the attributes of open source software – collaboration by a number of different groups, online publishing and evaluation, no proprietary intellectual property, and continuing collaboration. In security, open source software has made significant impacts. For example, in the early part of this decade, the original FireWall ToolKit was open source software that enabled other government-funded technologies in open-source information assurance to gain widespread adoption (e.g., SE-Linux ex-tension)

5. Digital Provenance

New Game: *Basing trust decisions on verified assertions.*

This section explores **Digital Provenance** as a path to this new game.

5.1 Introduction

This report on digital provenance game-changers was drafted in response to a requirement of the National Cyber Leap Year (NCLY) Summit. The Summit and this report are consistent with the President's Cyberspace Policy Review, which provides for the development of "a framework for research and development strategies that focus on game-changing technologies that have the potential to enhance the security, reliability, resilience, and trustworthiness of digital infrastructure." (Policy Review, p. 38. Emphasis added.) [3] The Cyberspace Policy Review calls for building a "cybersecurity-based identity management vision and strategy that addresses privacy and civil liberties interests, leveraging privacy-enhancing technologies for the Nation." (Id.) [4]

5.2 Definitions

Digital Provenance (DP) is a set of technologies, incentives, and policies which, in combination, provide an appropriate level of attribution to users of -- and/or resources accessible via -- the Internet, allowing for trust decisions to be based on verified identity assertions.

Identity is a unique reference to a distinct (possibly composite) entity. It is a recursive concept based on the context; any attribute of an entity may be considered an identity. **Provenance** of an object is the set of identities, labels, and events associated with the object.

5.3 Vision

We envision an end state in which DP enables identification, authentication, and reputation for entities and objects with appropriate granularity at many layers of the protocol hierarchy. For example, networked entities will be capable of authenticating the origin(s) and integrity of communications traffic. Also, DP will enable users to identify and authenticate the origins of data objects. This mitigates spoofing, phishing, denial of service (DoS), and impersonation attacks.

5.4 Privacy and Transparency

Today’s Internet offers a certain measure of ad hoc anonymity. Introduction of DP can result in loss of anonymity, i.e., the increase in trust would be undercut by the decrease in privacy. Government surveillance of individuals, whether for law enforcement or intelligence purposes, would become easier and more comprehensive. Industry would be able to better market products

[3] See also, Policy Review, pp. 32-33.
[4] See also, Policy Review, pp. 33-34.

because of the tremendous insight into the most personal aspects of individuals' lives. At the same time, individuals and their actions would be significantly more exposed than they are today.

A shift from anonymity to complete transparency and, presumably, complete accountability would be a game-changer for cyberspace but not necessarily one that Americans would accept. Scalability beyond the United States would be an issue, with European data protection laws likely conflicting with such an approach.

To protect the individual from the transparency resulting from DP, strong information governance (privacy) constraints must be established. Further, it may be advisable to place retention periods on information associated with individual provenance, in order to limit individual exposure. Such constraints and restrictions may be less important if DP is limited to business-to-business or intra-/inter-government transactions. Also, different levels of DP granularity (e.g., organization, job function, age) can be used to mitigate exposure of the individual and obtain certain measures of privacy.

5.5 Action Plans

5.5.1 Stable Network Identity

Identity and location are currently combined in many systems or domains (e.g., telephony, Internet). This semantic overloading complicates security, particularly in authenticating origin. For example, an IP address currently performs two distinct functions: 1) it indicates the topological location of its network interface within the larger network system context, and 2) it provides a host identity function. As a result, the Internet architecture relies on IP addresses to identify hosts but when a remote network protocol stack receives a packet, it currently cannot reliably prove that the sender of that packet is located at the source IP address identified in the packet header. Consequently, IP networks are vulnerable to a variety of impersonation attacks at the network layer and above, including network penetration, denial-of-service, phishing, spam, and routing reset attacks.

The Internet Engineering Task Force[5] (IETF) has recognized since the early 1990s that the semantic overloading of IP addresses is a fundamental weakness of the IP protocol family. This "IP Identity Problem" undermines Internet security (particularly in mobile environments) by allowing IP identifiers to be easily spoofed. It also impacts application- and session-layer coherence, enabling session hijacking. Separating the identifier function from the locator function within routing (network topology) is an increasingly popular approach. Unfortunately, the IETF has historically chosen to make local ad hoc changes to a few specific protocols (e.g., Mobile IPv6).

[5] http://www.ietf.org/

The Host Identity Protocol[6] (HIP) has been proposed as a strategic mechanism to separate end-point identifier and locator functions based on the current TCP/IP architecture. HIP uses public keys to serve as endpoint identifiers within the IP security[7] (IPsec) framework. Deploying HIP represents an opportunity to correct a fundamental weakness of the IP protocol family thereby solving multiple problems with a single solution. HIP provides applications with a unique cryptographic identity that is closely coupled with IPsec. It provides a stable identity that existing authentication and authorization systems can leverage. It provides session protection services to thwart session hijacking and other risks. HIP also provides a needed framework for reducing the impact of IP mobility and multihoming upon applications. It can be deployed incrementally without disrupting a host's existing internet stack or requiring applications or Internet routers to change. It can be implemented in proxies for legacy or embedded hosts that cannot upgrade. HIP is poised to move to the standards track in the IETF.

The Open Group has defined the Secure Mobile Architecture[8] (SMA) that uses HIP as its foundation for mobile security. At least one corporation has deployed SMA within its factories in order to secure Supervisory Control and Data Acquisition (SCADA) machine control systems. Like that of other machine controllers, the security of SCADA devices presumes physical security (isolation) that becomes threatened when these devices become networked. SMA uses HIP proxies to provide a "virtual enclave" capability that is compatible with existing SCADA infrastructure. The virtual enclave securely restores the isolation of past systems while still allowing communications between physically separate clusters of existing SCADA components over modern shared network infrastructures.

In order to make any change to core Internet protocols: 1) use cases must be clarified, 2) deployment paths must be carefully calculated, 3) operational concerns must be addressed, and 4) software must be exceptionally robust. We recommend that OSTP or NITRD sponsor these steps supporting a transition to HIP. We recommend that OSTP or NIRTD organize future workshops to jumpstart large-scale architectural experimentation and/or trial deployment of HIP. We also recommend that a government security organization (e.g., NSA) evaluate HIP and SMA security to verify its robustness and efficacy.

While this discussion has been IP-centric and has focused on HIP, the broader issue of trustworthy entity identification needs to be addressed. This includes entities at various levels of the protocol stack, applications, users, and data at rest. We recommend that OSTP or NITRD also consider the broader identification issue.

5.5.2 Data Provenance Management[9]

[6] http://www.ietf.org/html.charters/hip-charter.html

[7] See RFC 4301

[8] http://www.opengroup.org/products/publications/catalog/e041.htm

[9] The action item "Data Provenance Management" is a merger of the ideas "Data Provenance Definition and Management" and "Data Provenance Security." See the National Cyber Leap Year Summit 2009 Participants Ideas Report.

One of the fundamental challenges of (and prerequisites for) DP is the effective definition and management of DP meta-data (DPMD) (e.g., identity and action labels, [cryptographic] tags, etc.) attached to objects. DPMD reflects the chain of custody of, and its effects on, the object.[10]

Several technological factors incentivize the timeliness of DPMD definition and management. First, a number of domain-specific markup languages (e.g., XML) have been developed and widely deployed. This should facilitate the design of a DP-centric markup language. Second, there are numerous techniques for cryptographically binding an object and its origin (i.e., via various flavors/breeds of digital signatures). Also, the growing availability, affordability, and popularity of time and location services (e.g., NTP, GPS) facilitate corresponding origin/DP attributes.

The definition of DPMD must be scalable and extensible; this is likely to be a major challenge in its own right. Moreover, since DPMD includes (possibly fine-grained) identifying information, its exposure prompts privacy concerns that must be addressed from the outset.[11]

Initial exploratory efforts should be undertaken in the research community and fast-tracked into relevant standardization bodies (e.g., IRTF/IETF and W3C). Standardization efforts must involve the relevant industry segments (e.g., major network appliance and web browser vendors). In particular, to facilitate user acceptance, web browsers would need to incorporate DPMD support as early as possible. This, in turn, will trigger the need for understanding and designing usable security techniques for the presentation and user-driven management of DP meta-data. In the same vein, OS vendors must be engaged to include DPMD in file systems. In parallel, pre-existing relevant government and industry standards and software (e.g., government meta-data working group) must be revised to meet DP needs.

Once DPMD standardization efforts get off the ground, privacy and security of the DPMD itself needs to be addressed. This includes, but is not limited to, authorization and access control of entities (e.g., users, processes, hosts and groups thereof) with respect to DPMD of objects. Note: DPMD definition and management does not apply to objects themselves but to the associated meta-data.

Scalability is one of the crucial factors and challenges in designing an effective privacy and security architecture for DPMD. Other challenges are likely to stem from diverse international rules, laws, and regulations as well as from limited (or no) experience with similar technology. At the same time, certain research advances may facilitate bootstrapping the technical work: (1) attribute-based cryptography, (2) key-private encryption, and (3) group/ring signatures. Even more importantly, increasing recognition by the government and industry of the need for DP will play an important role in fostering the necessary research, standardization, and development efforts.

[10] Where the term “object” applies to messages, packets, software, database records, files, devices, etc.

[11] Technologists should consider privacy protections as they design, build, and deploy systems. DHS Privacy Office, Privacy Technology Implementation Guide (August, 2007) http://www.dhs.gov/xlibrary/assets/privacy/privacy_guide_ptig.pdf.

Initial efforts should be strategically focused on relatively narrow and manageable aspects. One natural first step is to design and architect secure DPMD support for immutable objects (e.g., software or music/video content). These efforts could be subsequently expanded into the domain of append-only objects exemplified by log files. In parallel, the research community should work on developing a general model for securing DPMD.

Aforementioned efforts need to be coordinated and integrated with other work on trusted systems (especially the Hardware-Enabled Trust group from this summit). In the context of secure network-layer DP, hardware acceleration techniques need to be developed for routers. Finally, government and stakeholders should work to establish both policy and legal frameworks for resolving DP conflicts.

5.5.3 Other DP Action Plans

Digital provenance is a broad domain that has many dimensions. Additional dimensions not noted in this report can be found in the appendix. By no means is this an exhaustive list.

Additional dimensions include measuring the credibility of principals or entities by tracking popularity, participant responses, and feedback scoring and associating those responses with an identity. These perceptions can therefore form part of a provenance framework. Also, DP can establish a foundation to assure the integrity and origin of data objects that have passed through multiple hands. As well, it can provide non-spoofable interfaces (trusted paths) between users and trustworthy systems that provide assurance that a specific user is the origin of resulting interactions with the trustworthy system.

5.6 Conclusion

At its core, DP must address at least two of three things to succeed as a game-changer. It must either create greater validity in assertions of identity or greater trust in the system, and it must do so in a way that protects privacy. Stable network identity and data provenance management address both the need for greater validity and the need for greater trust. By thoughtfully considering the privacy implications of these game-changers during development, and careful application of information governance principles upon testing and deployment, we believe that stable network identity and data provenance management can be successful in changing the game for cybersecurity in the United States.

6. Nature-Inspired Cyber Health

New Game: *Moving from forensics to real-time diagnosis.*

This section explores **Nature-Inspired Cyber Health** (renamed from Health-Inspired Network Defense) as a path to this new game.

6.1 Introduction

Our working group focused on Nature-inspired approaches to cybersecurity because we believe that one of the best ways to generate novel ideas is to look to natural systems for inspiration—these systems evolved to face specific threats and have undergone millions of year of evolutionary selection to select the best fit. There are many natural systems that are far more complex than our cyber-systems but are none-the-less extremely robust, resilient, and effective. One notable example is the biological immune system that many organisms use to defend against invaders. Systems such function remarkably well in distributed, complex and ever-changing environments, even when subject to a continuous barrage of attacks. They exhibit a wealth of interesting mechanisms that could be the inspiration for many new methods for securing cyber-systems.

The compelling similarities between the problems we face in cybersecurity and those faced by biological systems have sparked investigative research to analyze how biological immunology concepts can be applied to cybersecurity. Immuno-computing or Artificial Immune Systems (AIS) emerged in the 1990s as a new computational intelligence field. For example, as long ago as 1996, an attempt was made to define the equivalent of the biological "self" for a computer system.[12] This led to a novel approach to anomaly and intrusion detection, which has spawned a new paradigm in cybersecurity research.[13] Ongoing research into the analogy between cybersecurity and immunology continues to result in useful ideas. The work in the noted references[14] laid the foundations for a possible architecture and the general requirements for an immunity based intrusion detection and response system. We discuss the application of biological immunology to cybersecurity in section 6.2.1.

[12] S. Forrest, S. A. Hofmeyr, A. Somayaji, and T. A. Longstaff. A sense of self for Unix processes. In Proc. of 1996 IEEE Symposium on Computer Security and Privacy, 1996

[13] D. Dasgupta. Book-Immunological Computation, CRC press, September 2008

[14] P. D'haeseleer, S. Forrest, and P. Helman. An immunological approach to change detection: algorithms, analysis, and implications. In Proc. of the 1996 IEEE Symposium on Computer Security and Privacy, IEEE Computer Society Press, Los Alamitos, CA, pp. 110-119, 1996; S. Forrest, S. A. Hofmeyr, A. Somayaji, and T. A. Longstaff. A sense of self for Unix processes. In Proc. of 1996 IEEE Symposium on Computer Security and Privacy, 1996; S. Forrest, S. Hofmeyr, and A. Somayaji. Computer Immunology. In Communications of the ACM, Vol. 40, No. 10, pp. 88-96, 1997; A. Somayaji, S. Hofmeyr, and S. Forrest. Principles of a Computer Immune System. 1997 New Security Paradigms Workshop pp. 75-82 (1998); S. A. Hofmeyr and S. Forrest. Immunity by Design: An Artificial Immune System. In Proc. of 1999 GECCO Conference, 1999; D. Dasgupta. Book-Immunological Computation, CRC press, September 2008; P Matzinger. The Danger Model in Its Historical Context, Scandinavian Journal of Immunology, 54: 4-9, 2001

Although most of the currently active research into cybersecurity that is inspired by nature focuses on the immune system, there are many other natural systems that could serve as inspiration for cybersecurity. In this report, we present two novel concepts, both inspired by a study of natural systems. The first concept (section 6.3) involves developing a national information-sharing and warning system for cybersecurity, using the Centers for Disease Control (CDC) as a model. The second concept (section 6.4) is a controversial one that involves using attack vectors to secure vulnerable computers. In particular, we suggest new approaches to mitigate the effect of cyber-worms by emulating biological concepts concerning 'phage therapy[15]' that uses viruses to attack bacterial pathogens, 'oncolytic viral therapy'[16] that uses viruses to attack cancerous tumors, and interfering particle therapy[17] that uses sub-viruses to attack pathogenic viruses.

The ideas detailed in this report emerged during the course of a three (3) day summit that involved a productive cross-section of industry, academia and government. In addition to the two (2) novel concepts of a Cyber-CDC and using attack vectors, the participants took a fresh look at previous approaches such as defining a sense of self for computer systems. Notes covering all the discussions and outcomes can be found in the National Cyber Leap Year Summit 2009 Participants' Ideas Report.

One of the most important recommendations to emerge from this working group is that cross-disciplinary research has the potential to truly change the game for cybersecurity. We believe it is vital to promote such research through the establishment of communities, research programs, and even multi-disciplinary institutes focused on cybersecurity.

6.2 Immuno-inspired Defense Research

We have only just begun to mine the wealth of possibilities provided by the correspondence between biological immunity and cybersecurity. There is much to learn and much to potentially gain. We continue to believe that the Biological Immune System (BIS) is one of the best existing examples of an effective defense mechanism for a complex system, and an ongoing study should continue to yield new insights that could be game-changing for cybersecurity.

6.2.1 Biological Immunity

The BIS is a robust defense system that has evolved in vertebrates to protect them from invading pathogens. To accomplish its tasks, the BIS uses sophisticated detection and response mechanisms and follows differential response pathways, i.e., depending on the type of pathogen, the way it enters the body and the damage it causes, the immune system uses various mechanisms for detection, recognition and subsequent destruction of the invader or

[15] A Sulakvelidze, Z. Alavidze and J. G. Morris. Bacteriophage Therapy Antimicrobial Agents And Chemotherapy, Mar. 2001, p. 649–659 Vol. 45, No. 3

[16] D. M. Nettelbeck and D.T. Curiel. Tumor-Busting Viruses: Zapping Cancer Cells with Viruses" Scientific American Magazine, October 2003

[17] LS Weinberger, DV Schaffer, AP Arkin, Theoretical design of a gene therapy to prevent AIDS but not human immunodeficiency virus type 1 infection. Journal of Virology, 77: 10028-10036, 2003

neutralization of its effects. In medicine, historically, the term immunity refers to the condition in which an organism can resist disease, more specifically infectious disease. However, a broader definition of immunity is a reaction to foreign (or dangerous) substances. Cells and molecules responsible for immunity constitute the BIS, and the collective coordinated response of such cells and molecules in the presence of pathogens, is known as the immune response.

The BIS can be envisioned as a multilayer protection system, where each layer provides different types of defense mechanisms. For example, skin and mucus membranes provide the first level of defense by blocking/filtering out many bacteria, fungi, etc. There are three (3) main layers which include the anatomic barrier, innate immunity (nonspecific) and adaptive (specific) immunity. Innate (non-specific) immunity and adaptive (specific) immunity are inter-linked and influence each other.

Once adaptive immunity recognizes the presence of an invader, it triggers two types of responses: humoral immunity and cell-mediated (cellular) immunity, which act in a sequential fashion. Innate immunity is usually directed against an invading pathogen; however, if the pathogen evades the innate defenses, the body launches an adaptive and specific response against it.

Signaling is essential for activating and coordinating biological defenses. Signaling also allows a cell to transfer information about its internal state to its environment, where it can be recognized by cells in the Immune system. Furthermore, signaling results in changes to the cell, allowing it to appropriately respond to a stimulus.

6.2.2 Digital Immunity

From an information-processing perspective, there are several immunological principles that make the system very appealing, which include distributed processing, pathogenic pattern recognition, multi-layered protection, decentralized control, and diversity and signaling.

Understanding the immune mechanisms on the abstract level could result in the development of novel approaches to solve problems of cybersecurity: early and dependable detection and recognition of information attacks, rational utilization of the network resources for minimization of the damage and fast recovery, and development of successful ways to prevent further attacks.

The elements of the BIS act in concert, mediated by a variety of communication mechanisms, for example, signal diffusion and dialogue. This signaling plays a major role in sharing and transmitting information during an immune response and is essential to the effectiveness of the immune defenses. By contrast, in cybersecurity many available tools (e.g., firewalls, file integrity checkers, virus scanners, intrusion detection systems, anti-malware software, etc.) operate independently and neither exchange data nor have consistent security policies. Each of them may have been developed by a different vendor, perhaps even competitors in the industry who do not follow the same standards. Since there is no consistent data exchange between these tools, many attacks remain unnoticed due to lack of correlation. Moreover, security administrator intervention is usually required to analyze the acquired data and make decisions about what actions may need to be taken to prevent a compromise, or to recover from one. Human intervention is slow and limited in scope, and is one of the major bottlenecks in building survivable autonomic systems. Integrating the many disparate tools using both feed forward and feedback signaling mechanisms in a cyber defense system should greatly enhance our understanding of real attacks and their sources (connecting dots!), and enable us to build faster, more responsive security systems.

Another promising development is trying to understand how we could use the Danger Model concept[18] to refine the accuracy of cybersecurity response, because not all abnormal events (non-self) represent attacks – only a small percentage of such events are of real concern. Simple observations can be used to trigger a chain of defensive actions, but the challenge is clearly to define what constitutes suitable danger signals[19].

Action Plan

Increased research efforts are needed to explore immunological principles to automatically detect situational changes, determine imminent danger and mitigate cyber attacks, for example:

- Thwart malicious attacks through signaling, implementation of diversity and immunogenic detection as hardware-software solutions. Rapidly regenerate (self-healing) survivable capabilities in mission critical systems after an sophisticated attack
- Evolve immunity to attacks through evolutionary computing to create new deceptions (gaming strategies) as new threats emerge. Self-learning while monitoring insider activity and develop profiles for appropriate and legitimate behavior (modeling).
- Signaling and Message-passing(SM): Integrating the many disparate security tools using both feed forward and feedback signaling mechanisms in a cyber defense system should help to ensure tolerance and identify attacks while minimizing false alarms (i.e. improve judgments between dangerous attacks and benign anomalies).
- Decentralized Control (DC): The immune system uses distributed control mechanisms for learning, memory and associative retrieval to solve recognition and classification tasks. There is no single organ that controls the immune response; rather it handles the antigenic challenges through collaborative interaction. A similar strategy (distributed control mechanisms for monitor and response) needs to be pursued as a game changing strategy in cyber defense in order to avoid a single point of failure and to enable robust decision making.
- Missing Self Paradigm: The missing self hypothesis from immunology literature may shed new light to secure host systems, in particular, to validate, authenticate and permit codes, data and scripts to execute in a machine. Different techniques are used to preserve integrity at the process, system and communication levels. For example, Trusted Platform Module (TPM), Intel Trusted Execution Technology and Windows Vista Kernel Mode security ensure system level integrity and security; whereas, Digital Signature, Code Signing, Watermarking, Integrity Checker, Magic Cookies, etc. address file

[18] P Matzinger. The Danger Model in Its Historical Context, Scandinavian Journal of Immunology, 54: 4-9, 2001

[19] T. S. Guzella, T. A. Mota-Santos and W. M. Caminhas. A novel immune inspired approach to fault detection. Lecture Notes in Computer Science, Proceedings of ICARIS 2007; J. O. Kephart. A biologically inspired immune system for computers. In R. A. Brooks and P. Maes, (Eds.), Artificial Life IV. Proceedings of the 4th International Workshop on the Synthesis and Simulation of Living Systems, MIT Press, Cambridge, MA, pp. 130–139, 1994

integrity of data and executables in transit. This missing self concept will be investigated as the first line of defense as an additional layer integrity checker universal for all systems (reference National Cyber Leap Year Summit 2009 Participants' Ideas Report, Missing Self Paradigm).

We believe that there needs to be an emphasis on the ongoing exploration of this area, especially cross-disciplinary research bringing together computer scientists, biologists and immunologists. New insights and game-changing ideas often come from the intersection of radically different research fields. We believe that steps should be taken to strongly encourage such cross-disciplinary research. For example, research funding is all too often focused on a single discipline – we would like to see funding availability that makes it easy to collaborate across disciplines, or even requires such collaboration. One possibility is the establishment of a research institute intended to bring together scientists from many different disciplines to attack the problem of cybersecurity. This is a model that has been successfully tried with the study of Complex Systems (for example, the Santa Fe Institute).

Jump-Start Plan

Specifically, we propose the following Jump-Start Plan:

1. Conduct small workshop to convene a group of immunologists, technologists and government officials to develop strategic plan for cybersecurity coverage.
2. Fund the first phase of research to determine the feasibility of immunological approaches as a game changing strategy in cyber defense. The new game is about real-time distributed monitoring and protecting very dynamic and flexible cyber environments.

6.3 Cyber CDC

Novel ideas for cybersecurity emerge not only from biomimicry, but also from a study of how society reacts to health threats. The government plays a key role in protecting society from various health threats, and equally should do so for cyber-health. One powerful idea is to have a government organization that is the cybersecurity equivalent of the CDC. The CDC plays a critical role in protecting the health of nation's citizens through the gathering and dissemination of information, epidemiological studies, education and, most importantly, through prevention and control of disease, especially communicable diseases.

The stated mission of the CDC is: "to collaborate to create the expertise, information, and tools that people and communities need to protect their health – through health promotion, prevention of disease, injury and disability, and preparedness for new health threats."

We propose a Cyber-CDC (CCDC) which has an analogous mission to enable an enhanced level of cyber-health. We are particularly interested in how the CDC addresses the issues of communicable diseases, since these are the most similar in societal impact to cyber-threats. Communicable diseases are an issue for more than the infected individual because of the possibility of transmission – similarly, compromised computers can be used as launching points for cyber-attacks, and consequently affect the health of the whole of cyberspace.

There are already organizations that gather information on cyber-threats and help to coordinate response strategies. Most of these are private, either commercial, for example, Symantec's Security Focus, or non-profit, such as CERT/CC. US-CERT is very similar to the idea of the CCDC which is, a government organization that is tasked with "providing response support and defense against cyber attacks for the Federal Civil Executive Branch (.gov) and information sharing and collaboration with state and local government, industry and international partners."[20]

The major difference between US-CERT and our proposed CCDC is that US-CERT focuses on protecting government entities and has no power over the civil populace, unlike the CDC, which can act to prevent disease threats, through approaches such as quarantine and vaccine distribution. We argue that an effective new approach to cyber-health must include a government organization with power equivalent to that wielded by the CDC.

6.3.1 Key Functions

There are several key functions that can be most obviously drawn from the analogy but there could be more that need consideration, and the precise details of each function need to be carefully deliberated. Our initial proposal includes the following list:

- Data collection: Gather local data on cyber threats and cybersecurity outbreaks, analogous to the information collected by medical doctors' offices. There are legal reporting requirements for most communicable diseases – we propose that something similar is needed for cyber-threats. There are already many organizations that collect data about cyber-threats – these should be leveraged as much as possible. Further work will also be needed to adequately anonymize data so that private organizations and individuals are protected, but the data are still useful for cyber-threat analysis. Legislation to protect privacy – similar to that provided by HIPAA – will also have to be enacted.
- Data dissemination: Provide data about the spread and danger of threats. These data can help communities and organizations plan in response to threats.
- Cyber-threat analysis: Diagnose and investigate cyber-threats in the community. Where possible, the CCDC should try to verify outbreaks of new cyber-threats and understand the extent and impact of these outbreaks. This will probably require extensive partnership with the private sector, to leverage existing expertise and react as fast as possible.
- Intervention analysis and recommendations: Provide a cost/benefit analysis of interventions and make recommendations, especially to less sophisticated users. This is a difficult task that will require much input from academia, private industry and other organizations. We suggest that the CCDC be a repository and clearing house of such information, rather than a prime developer of it. It should also vet all information to ensure that any recommendations are grounded in fact, and not driven by the business objectives of vendors. US-CERT already does much of this.
- Education: Educate people about cyber-health issues and link individuals to needed cyber-health services. This education should be made interactive, with a website where individuals could test the health of their computers, and subsequently be presented with

[20] US Cert government organization, http://www.us-cert.gov/aboutus.html

strategies for remediation if it was determined that the computer was vulnerable. US-CERT already provides websites such as[21] the Stay Safe Online website that provide useful information and that can form the basis for this function. Organizations like InfraGuard,[22] the National White Collar Crime Center and others can collaborate to disseminate online safety information to the populace to help combat Internet fraud.

- Leadership: provide leadership in organizing effective public and private sector strategies to address community cyber-health problems. Most corporate entities and government organizations have policies about cybersecurity, but there are few commonalities and too little focus on coherent strategies that could benefit all organizations. The government should take the lead through an organization such as the CCDC.
- Coordination of preventative actions: The CCDC should have the power to not only suggest response strategies, but coordinate them and ensure they are carried out, for example, the equivalent of quarantining and vaccination strategies. Many other active strategies are required, for example, aggressive cyber patrolling to police fraud online.

We believe that the time is right for the establishment of a new government body with a much more extensive remit to ensure the cyber-health of the whole civil society. There is increasing public awareness of the impact of cybersecurity issues and the growing reliance on cyberspace for the economy means that it is imperative that we as a nation improve the state of our cyber-health. We suffer daily losses from the damage caused by cyber-health issues, and we suffer potentially vast losses from forgone opportunities – cyber-threats discourage online business and slow the pace of innovation and economic growth. Furthermore, our technology has advanced dramatically, with both increased processing power and greater connectivity, which should enable us to gather and analyze data in ways not possible, even recently.

6.3.2 Action Plan

Our immediate recommendation is to establish a community of interested parties to further investigate the idea of the CCDC and explore the means to make the vision a reality. Such a community should include academia, industry and government, with experts from a broad cross-section of disciplines. Because of the scope of this idea, we will need cybersecurity experts, legislators, lawyers and many others. In particular, we need the input of officials from public health organizations such as the CDC, who can provide firsthand experience of the issues involved with establishing and running such an organization.

We recommend a series of workshops bringing together interested parties to address a variety of issues, including:

- Determining what data should be required, how the data can best be anonymized, and how it can be accumulated and understanding the privacy issues

[21] Stay Safe Online website, http://staysafeonline.org
[22] InfraGurard Organization, http://www.infragard.net/

- Determine the legislation that would be needed to report requirements and what those requirements should look like. An important comparable area involves the regulations that pertain to the standardization of human pandemic vaccines. Specifically, the World Health Organization (WHO) has developed standardized protocols for data collection and reporting in real-time to the international community, in a pandemic emergency.[23] These protocols could be useful to study and analyze those protocols in order to evaluate and update the network security protocols.
- Determining the precise role and scope of the CCDC
- Applying the lessons learned from public health organizations, such as the CDC and the WHO
- Identifying potential partnerships and soliciting input from industry and government about the effect and practicality of reporting requirement
- Determining if there are any suitable business models for the involvement of private organizations
- Investigating interventions than could be carried out on a national level, similar to quarantining or vaccinations

The goal of the series of workshops will be to develop a plan for the establishment of the CCDC that is supported by all parties and can be enacted promptly by the government.

6.4 Attack-vector approaches

In this section we propose a controversial offensive (rather than defensive) approach to cybersecurity. The idea is to deliberately (and without requiring consent) infiltrate vulnerable computers using the same vectors that attackers use, and once inside, to secure those computers. Importantly, the 'offensive' approaches presented in this section are not mutually exclusive to the cyber-CDC or artificial immune system 'defensive' approaches explained above. In fact, the offensive strategies below are completely compatible with the approaches above. An example of an offensive approach that uses attack vectors in the biomedical field is the strategy behind the Oral Polio Vaccine (OPV) which has been chosen for the world-wide eradication campaign against Polio. OPV was chosen, in-part, due to its ability to 'shed' and thus spread through a population 'passively' immunizing individuals who were not directly immunized (i.e., the passively immunized individuals were not required to consent to the intervention).

Given the expanding threat environment, we believe that in highly specialized circumstances - where there are tangible cyber threats to our nation's economy and security that cannot be dealt with via other methods - offensive strategies that use attack vectors may be warranted. We

[23] Regulatory Preparedness for Human Pandemic Influenza Vaccines, World Health Organization, Regulatory Preparedness for Human Pandemic Influenza Vaccines, Genova, 2007

believe that a rational approach is to fund the study of attack vector approaches[24] so that potential technical hurdles may be overcome if the development and deployment of attack-vectors is required.

6.4.1 Motivation

The motivation for proposing attack-vector approaches (like OPV) stems from the knowledge that hordes of vulnerable computers exist on the internet. These vulnerable computers are a huge problem because they are subverted to form vast botnets that have enormous power to do harm in cyberspace. Although the technology to secure many of these computers exists, it is virtually impossible to widely deploy by traditional means – they remain insecure for a variety of reasons, including ignorance or apathy of the user/administrator. Clearly, malevolent attackers (hackers, black hats, etc.) have vectors into these vulnerable computers – we propose using those same vectors to better secure vulnerable systems.

Biological organisms also frequently encounter attackers, such as viruses, and in some cases have evolved into active counter attacking viruses to mitigate the effects of the pathogenic virus (e.g., defective interfering particles). Based on this biological precedent (and the OPV example), it seems prudent to explore the use of attack vectors for securing cyber-systems. These offensive approaches could be used in conjunction with current defensive approaches to enhance computer security, and the BIS-like defensive approaches described above. Below, we briefly review strategies for using biological attack vectors and propose analogous cyber attack-vector technologies that could be explored or have been proposed and should be explored further.

6.4.2 Biological Underpinnings (purely defensive systems are inadequate)

Even the most sophisticated biological defensive systems - the adaptive immune systems of mammals evolved over millions of years of selection - are frequently invaded and over run by parasites (e.g., microorganisms or viruses) and often fail to protect the organism from eventual death. Periodically, the failure of a defensive or immune system to protect the organism leads to large scale epidemics (e.g., the 1918 influenza pandemic) and even species extinctions. The inability of the immune system to defend against all possible attacks is a major driving force behind biomedical research and the pharmaceutical industry's goal to create better vaccines and therapies for infectious diseases. While there are clearly examples of parasites and hosts developing a symbiotic relationship, in almost all cases a parasite's evolution into a benign symbiote occurs only after multiple generations (e.g., millions of years) of 'predator-prey' dynamics where the parasite continually decimates the host population until a symbiotic relationship evolves. If effective and cyber immune systems are produced, we are likely to see similar 'predator-prey' dynamics in the cyber world and we must consider whether large portions of a network should be sacrificed to mass 'extinctions' while waiting for the cyber defense systems to 'evolve' to mitigate the parasite. Attack-vector approaches have the potential to

[24] D. J. Albanese, M. J. Wiacek, C. M. Salter, J. A. Six. The Case for Using Layered Defenses to Stop Worms, Network Architecture and Applications Division of the Systems and Network Attack Center (SNAC), NSA, Report # C43-002R-2004, July 2004.

effectively hasten the ‘parasite evolution’ process to become a more benign vector. Here, we propose that cybersecurity consider a pharmaceutical-like offensive approach to attacking the invading pathogens by using ‘attack vectors’.

We base our proposal on a number of biomedical therapies have been proposed and developed over the past century that are overtly offensive (rather than defensive) in nature. These active/offensive therapies are designed to attack the invading pathogen rather than merely enhance the host’s defenses against the pathogen. We will very briefly review two of the offensive biomedical attack approaches that could act as inspiration for cyber attack-vector approaches: (1) ‘phage therapy’, a historical competitor to antibiotics that instead uses viruses that infect only bacteria;[25] (2) oncolytic ‘virotherapy’, an anticancer approach using engineered viruses that replicate only in tumors and thereby destroy the tumors;[26] and (3) conditionally replicating viruses (a.k.a. ‘interfering particles’), an approach developed using non-pathogenic sub-viruses, that can only replicate in the presence of full-length parent viruses, and diminish the effects of the pathogenic parent virus by resource competition.[27] These interfering particles are found to naturally arise for some viruses in nature.

Importantly, phage therapy, oncolytic virotherapy, and interfering particle therapy all share a common theme that differentiates them from more standard pharmaceuticals or even cybersecurity approaches: they semi-autonomously replicate or spread through the infected population. The ability to semi-autonomously spread through the infected population is the major ‘game-changing’ departure that attack-vector approaches utilize over conventional cybersecurity approaches.

Nature has evolved similar offensive tactics against many viruses, including evolving beneficial sub-viruses that attack viral pathogens (e.g., defective interfering particles that compete with viruses).[28] We propose to develop similar attack-vector approaches for specific cyber parasites that are considered dangerous enough to potentially cause large scale network failures.

Essentially, the two examples of attack-vector approaches proposed below use the same vectors as malevolent entities use but in this case, to do good (e.g., patch). Importantly, the patching is done without the user’s consent and for the greater good of the network as a whole. Again, this patching, without the user’s consent, for the ‘greater good’ argument, is rooted in the OPV approach for worldwide eradication of Polio.

[25] A Sulakvelidze, Z. Alavidze and J. G. Morris. Bacteriophage Therapy Antimicrobial Agents And Chemotherapy, Mar. 2001, p. 649–659 Vol. 45, No. 3

[26] D. M. Nettelbeck and D.T. Curiel. Tumor-Busting Viruses: Zapping Cancer Cells with Viruses" Scientific American Magazine, October 2003

[27] LS Weinberger, DV Schaffer, AP Arkin, Theoretical design of a gene therapy to prevent AIDS but not human immunodeficiency virus type 1 infection. Journal of Virology, 77: 10028-10036, 2003

[28] J. Li and P. Knickerbockera. Functional Similarities between Computer Worms and Biological Pathogens. computers & security 26, 338–347, 2007

Description

Initially, three potential attack-vector approaches were discussed. The first idea, designing and using Good Worms (aka ‘gworms’)[29] was deemed too problematic and resulted in few benefits. (gworms are actually an old idea which had a known historical implementation in 2003; the Welchia gworm detected and terminated the Blaster worm by patching the system and rebooting; Welchia did more harm than good due to the increased network traffic it generated).

The remaining two ideas were considered more feasible and potentially game-changing:

1. ‘Piggybacking’: ride the worm (analogous to defective interfering particle approaches in biology)
2. Drive-By Downloads

Briefly, the piggyback idea is to replace the worm payload with a ‘rider’. The rider prevents host damage but still allows network spread. Analogous to the biological case,[30] the rider goes where the worm goes, at the same rate as the worm spreads (although this latter hypothesis clearly requires further careful research). The idea would be to load the riders onto select honeypots and wait for worms to exploit these honeypots. Once exploited, the honeypots would transmit the rider along with the worm. The piggyback has numerous benefits over the gworms: (1) The rider is dormant until activated by the worm (i.e. only act when harm is happening); (2) it is much easier for the piggyback to match the spread rate of the worm; (3) the rider contains no exploit or transmission code that can be copied by black hats; (4) there are fewer ethical and legal roadbloacks as compared to gworms (although the ethical and legal roadblocks may still be significant); and (5) there is a possibility for designing piggyback worms that could spread with worms even when the specific vulnerability being exploited is not a known priori.

There are significant challenges to designing and implementing a piggyback. First, there are nontrivial technical challenges to payload replacement of a worm with a rider. Second, methods must be developed to constrain damage caused by the worm. Third, if payload replacement is achieved and damage can be constrained, the code must still be optimized to react very quickly against fast-moving worms. Fourth, if controlling the spread rate of the worm is desirable (again research is required to test this), implementation technologies will need to be developed. Finally, although the ethical and legal considerations for the piggyback are less than the gworms, there are still significant legal & ethical challenges to implementing the piggyback approach.

The second attack-vector idea which we believe to be novel and potentially game-changing is the benign ’drive-by downloads.’ The kernel of this idea is to have ‘good’ (i.e. benevolent) webservers exploit the same vulnerabilities that malicious webservers exploit when clients visit, but instead of installing malware on a client’s machine these benevolent servers install

[29] F. Castañeda, E. C. Sezer, and J. Xu. Worm vs. Worm: Preliminary Study of an Active Counter Attack Mechanism, Proceedings of the 2004 ACM workshop on Rapid malcode, October 29-29, 2004, Washington DC, USA

[30] LS Weinberger, DV Schaffer, AP Arkin, Theoretical design of a gene therapy to prevent AIDS but not human immunodeficiency virus type 1 infection. Journal of Virology, 77: 10028-10036, 2003

‘whiteware’, which could patch vulnerabilities, clean off malware, etc. The important departure from previous approaches is again that whiteware is installed without a client’s consent using the client’s inherent vulnerabilities. Drive-by downloads have a number of pros and cons. The pros are that drive-by downloads could patch vulnerabilities that cannot be fixed by gworm or piggyback approaches, and they would address a common vector used to spread botnets. Importantly, the drive-by download approach is NOT a viral-like or worm-like approach in that it doesn’t generate harmful spreading via increased network traffic. The cons of drive-by downloads are that penetration and auto-patching could be harmful, and they could be useless if a system is already compromised. Clearly, drive-by downloads would pose serious ethical and legal problems, but these problems may be mitigated by the non-viral, non-worm status of the drive-by download method.

Why haven’t we used attack vectors before?

As previously mentioned, a gworm has already been deployed on at least one occasion, however in a vigilante fashion and with less than an optimal outcome (Welchia vs. Blaster worm, 2003). However, to the best of our knowledge, there has not been a deployment of a piggyback or a drive-by download approach. Clearly, serious ethical and legal issues surround these approaches and public perception, liability, legality, and the specter of side-effects and lack of efficacy are major barriers to deployment of these potential attack-vector technologies. As mentioned above, there are also major technical challenges (e.g., payload replacement). However, today we enjoy far more advanced technical tools than only a few years ago (e.g., virtual machines, increased computing power). In terms of perception, today there is an increased awareness that cybersecurity issues are a very real threat to our nation’s economy and security and we face the expanding problem of botnets and malware that may radically change society’s cost-benefit calculus. It is possible that these considerations may make the use of attack vectors more palatable in highly specialized circumstances.

6.4.3 Action Plan

Increase research focused on the area of attack-vector technology. Specifically, we believe increased research efforts in three areas would be most prudent:

1. Technical feasibility of payload replacement
2. Epidemiological models & simulations of attack vector scenarios
3. Investigations into the non-technical aspects (e.g., legality) of attack vectors usage

6.4.4 Jump-Start Plan

Specifically, we propose a jump-start plan with two components:

1. Conduct small workshop on how to use attack vectors, bringing together technologists, lawyers, and government officials
2. Fund early-stage research to determine the feasibility of attack vectors approaches (e.g., payload replacement)

APPENDIX A: References

Digital Provenance

- Center for Democracy and Technology, *Privacy and the White House Cyberspace Policy Review*: http://www.cdt.org/security/20090619_cybersec_actions.pdf
- Center for Strategic and International Studies, *Securing Cyberspace for the 44th Presidency*: http://csis.org/files/media/csis/pubs/081208_securingcyberspace_44.pdf
- Intelligence and National Security Alliance, *Critical Issues for Cyber Assurance Policy Reform*: http://www.insaonline.org/assets/files/INSA_CyberAssurance_Assessment.pdf
- Internet Security Alliance, *The Cybersecurity Social Contract Policy Recommendations for the Obama Administration and 111th Congress*:
- http://www.whitehouse.gov/files/documents/cyber/ISA%20-%20The%20Cyber%20Security%20Social%20Contract.pdf
- White House, *Cyberspace Policy Review: Assuring a Trusted and Resilient Information and Communications Infrastructure*:
- http://www.whitehouse.gov/assets/documents/Cyberspace_Policy_Review_final.pdf
- David Brin, *The Transparent Society*, (Addison-Wesley, 1998)
- Ann Cavoukian, Ph.D., *"Privacy by Design"* (IPC, Ontario)
- Daniel J. Solove, *The Digital Person: Technology and Privacy in the Information Age* (New York University Press, 2004)