

National Cyber Leap Year Summit 2009 Participants' Ideas Report

**Exploring Paths to New Cyber Security
Paradigms**

September 16, 2009

Table of Contents

Introduction	12
1 Moving Target Defense	15
1.1 Idea - Mutable Networks: Frequently Randomized Changing of Network Addresses and Responses	15
1.1.1 Description.....	16
1.1.2 Inertia.....	16
1.1.3 Progress	17
1.1.4 Action Plan	17
1.1.5 Jump-Start Plan.....	17
1.2 Idea - Diversity in Software	19
1.2.1 Description.....	19
1.2.2 Inertia.....	19
1.2.3 Progress	19
1.2.4 Action Plan	20
1.2.5 Jump-Start Plan.....	20
1.3 Idea - Robust Random Authentication	20
1.3.1 Description.....	20
1.3.2 Inertia.....	21
1.3.3 Progress	22
1.3.4 Action Plan	22
1.3.5 Jump-Start Plan.....	22
1.4 Idea - Resilient Cryptographic Systems	23
1.4.1 Description.....	23
1.4.2 Inertia.....	24
1.4.3 Progress	25
1.4.4 Action Plan	25
1.4.5 Jump-Start Plan.....	25
1.5 Idea - Connectivity Diversity.....	26
1.5.1 Description.....	26
1.5.2 Inertia.....	27
1.5.3 Progress	27

1.5.4	Action Plan	28
1.5.5	Jump-Start Plan	29
1.6	Idea - Decoys	29
1.6.1	Description	30
1.6.2	Inertia	30
1.6.3	Progress	30
1.6.4	Action Plan	31
1.6.5	Jump-Start Plan	31
1.7	Idea - Configuration-Space Randomization for Infrastructure	32
1.7.1	Description	32
1.7.2	Inertia	33
1.7.3	Progress	33
1.7.4	Action Plan	33
1.7.5	Jump-Start Plan	33
1.8	Idea - Distributed Data Shell Game	34
1.8.1	Description	35
1.8.2	Inertia	35
1.8.3	Progress	35
1.8.4	Action Plan	35
1.8.5	Jump-Start Plan	36
1.9	Idea - Security on Demand	36
1.9.1	Description	37
1.9.2	Inertia	38
1.9.3	Progress	38
1.9.4	Action Plan	39
1.9.5	Jump-Start Plan	39
1.10	Idea - Terrorist Organization Model	39
1.10.1	Description	40
1.10.2	Inertia	40
1.10.3	Progress	40
1.10.4	Action Plan	40
1.10.5	Jump-Start Plan	40
1.11	Idea - Smart Motion Adaptation Management	41

1.11.1	Description	41
1.11.2	Inertia	41
1.11.3	Progress.....	42
1.11.4	Action Plan	42
1.11.5	Jump-Start Plan	42
2	Cyber Economics	44
2.1	Idea - Data & Metrics for Cybersecurity Analysis	44
2.1.1	Description.....	44
2.1.2	Inertia.....	45
2.1.3	Progress	45
2.1.4	Action Plan	45
2.1.5	Jump-Start Plan.....	47
2.2	Idea – Vendor Incentives and Accountability	47
2.2.1	Description.....	47
2.2.2	Inertia.....	48
2.2.3	Progress	48
2.2.4	Action Plan	49
2.2.5	Jump-Start Plan.....	50
2.3	Idea – Cyber “NTSB”	50
2.3.1	Description.....	50
2.3.2	Inertia.....	50
2.3.3	Progress	50
2.3.4	Action Plan	51
2.4	Idea – Cyber “Interpol”	51
2.4.1	Description.....	51
2.4.2	Inertia.....	51
2.4.3	Progress	52
2.4.4	Action Plan	52
2.5	Idea – Cyber Insurance.....	53
2.5.1	Description.....	53
2.5.2	Inertia.....	53
2.5.3	Progress	53
2.5.4	Action Plan	54

2.6	Idea - Empowering ISPs, Registrars, and Registries	54
2.6.1	Description.....	54
2.6.2	Inertia.....	54
2.6.3	Progress	55
2.6.4	Action Plan	55
2.7	Idea - Property Rights of Personal Information	56
2.7.1	Description.....	56
2.7.2	Inertia.....	56
2.7.3	Progress	57
2.7.4	Action Plan	57
2.8	Idea – Infrastructure Diversity	58
2.8.1	Description.....	58
2.8.2	Inertia.....	58
2.8.3	Progress	59
2.8.4	Action Plan	59
2.9	Idea - Multiple Networks	60
2.9.1	Description.....	60
2.9.2	Inertia.....	60
2.9.3	Progress	61
2.9.4	Action Plan	62
2.9.5	Jump-Start Plan.....	62
2.10	Idea - 911 Cyber	63
2.10.1	Description	63
2.10.2	Inertia	63
2.10.3	Progress.....	63
2.10.4	Action Plan	64
2.11	Idea - Swimming with the Sharks.....	64
2.11.1	Description	64
2.11.2	Inertia	64
2.11.3	Progress.....	64
2.11.4	Action Plan	65
2.12	Idea – Minimize and Target Authentication	65
2.12.1	Description	65

2.12.2	Inertia	65
2.12.3	Progress.....	65
2.12.4	Action Plan	66
2.13	Other Ideas	66
3	Digital Provenance	68
3.1	Idea - Stable Network Identity	68
3.1.1	Description.....	68
3.1.2	Inertia.....	68
3.1.3	Progress	68
3.1.4	Action Plan	69
3.1.5	Jump-Start Plan.....	69
3.2	Idea – Data Provenance Security	69
3.2.1	Description.....	69
3.2.2	Inertia.....	69
3.2.3	Progress	69
3.2.4	Action Plan	69
3.2.5	Jump-Start Plan.....	69
3.3	Idea - Data Provenance Definition and Management	70
3.3.1	Description.....	70
3.3.2	Inertia.....	70
3.3.3	Progress	70
3.3.4	Action Plan	70
3.3.5	Jump-Start Plan.....	70
3.4	Idea - Reputation Engine	70
3.4.1	Description.....	70
3.4.2	Inertia.....	71
3.4.3	Progress	71
3.4.4	Action Plan	71
3.4.5	Jump-Start Plan.....	71
3.5	Idea - Trustworthy Systems	71
3.5.1	Description.....	71
3.5.2	Inertia.....	71
3.5.3	Progress	71

3.5.4	Action Plan	71
3.5.5	Jump-Start Plan.....	72
3.6	Idea - Government Role.....	72
3.6.1	Description.....	72
3.6.2	Inertia.....	72
3.6.3	Progress	72
3.6.4	Action Plan	72
3.6.5	Jump-Start Plan.....	72
3.7	Idea - Trusted Path (TP)	72
3.7.1	Description.....	72
3.7.2	Inertia.....	73
3.7.3	Progress	73
3.7.4	Action Plan	73
3.7.5	Jump-Start Plan.....	73
3.8	Idea - Global Identity-Based Cryptography	73
3.8.1	Description.....	73
3.8.2	Inertia.....	73
3.8.3	Progress	73
3.8.4	Action Plan	73
3.8.5	Jump-Start Plan.....	74
4	Nature-Inspired Cyber Health	75
4.1	Idea - Distributed Defense	75
4.1.1	Description.....	75
4.1.2	Inertia.....	76
4.1.3	Progress	76
4.1.4	Action Plan	76
4.1.5	Jump-Start Plan.....	77
4.2	Idea - Centers for Cyber Disease Control (CCDC) and Prevention	77
4.2.1	Description.....	77
4.2.2	Inertia.....	78
4.2.3	Progress	79
4.2.4	Action Plan	79
4.2.5	Jump-Start Plan.....	80

4.3	Idea - Using Attack Vectors.....	80
4.3.1	Description.....	80
4.3.2	Inertia.....	82
4.3.3	Progress	82
4.3.4	Action Plan	82
4.3.5	Jump-Start Plan.....	82
4.4	Idea - Missing-Self Paradigm	83
4.4.1	Description.....	83
4.4.2	Inertia.....	84
4.4.3	Progress	84
4.4.4	Action Plan	84
4.4.5	Jump-Start Plan.....	85
5	Hardware-Enabled Trust	86
5.1	Idea - End to End (e2e) Trust.....	87
5.1.1	Description.....	87
5.1.2	Inertia.....	87
5.1.3	Progress	88
5.1.4	Action Plan	88
5.1.5	Jump-Start Plan.....	88
5.2	Enable Hardware to Counter Attacks.....	89
5.2.1	Description.....	89
5.2.2	Inertia.....	89
5.2.3	Progress	89
5.2.4	Action Plan	90
5.3	Sub-Idea - Enable Hardware to Counter Attacks—Hardware that does not leak, hardware defenses for information-leakage attacks.....	90
5.3.1	Description.....	90
5.3.2	Inertia.....	90
5.3.3	Progress	90
5.3.4	Action Plan	90
5.3.5	Jump-Start Plan.....	91
5.4	Sub-Idea - Enable Hardware to Counter Attacks—Continuous hardware monitoring of normal behavior.....	91

5.4.1	Description.....	91
5.4.2	Inertia.....	91
5.4.3	Progress	91
5.4.4	Action Plan	91
5.5	Idea - Trustworthy Storage and Data	92
5.5.1	Description.....	92
5.5.2	Inertia.....	92
5.5.3	Progress	93
5.5.4	Action Plan	93
5.5.5	Jump-Start Plan.....	93
5.5.6	Comments	94
5.6	Idea - Resilience	94
5.6.1	Description.....	94
5.6.2	Inertia.....	94
5.6.3	Progress	95
5.6.4	Action Plan	95
5.7	History of Idea Development.....	95
5.7.1	Leap-ahead, Long Term Goals – 10 year	95
5.7.2	Initial Ideas	96
5.7.3	Focus Areas	99
5.7.4	Game Changing Ideas.....	99
6	Additional Ideas	100
6.1	Idea - Virtualisable Network Architecture	100
6.1.1	Description.....	100
6.1.2	Inertia.....	100
6.1.3	Progress	100
6.1.4	Action Plan	100
6.2	Idea - Global Electronic Identity Management System	101
6.2.1	Description.....	101
6.2.2	Inertia.....	101
6.2.3	Progress	102
6.2.4	Action Plan	102
6.3	Idea - Global Post-Quantum Secure Cryptography Based on Identity	103

6.3.1	Description.....	103
6.3.2	Inertia.....	103
6.3.3	Progress	104
6.3.4	Jumpstart Activities.....	104
6.4	Idea - Evaluating the Effectiveness of Data Depersonalization Techniques and It's Impact on the Community.....	105
6.4.1	Description.....	105
6.4.2	Inertia.....	105
6.4.3	Progress	105
6.4.4	Action Plan	105
6.4.5	Jumpstart Activities.....	105
6.5	Idea - Measuring the Impacts of Unauthorized Information Disclosure.....	106
6.5.1	Description.....	106
6.5.2	Inertia.....	106
6.5.3	Progress	106
6.5.4	Action Plan	106
6.5.5	Jumpstart Activities.....	106
6.6	Idea - Semiconductor Intellectual Property Protection	106
6.6.1	Description.....	106
6.6.2	Inertia.....	107
6.6.3	Progress	107
6.6.4	Action Plan	107
6.6.5	Jumpstart Activities.....	107
6.7	Idea - Dynamic Distributed Key Infrastructures (DDKI).....	107
6.7.1	Description.....	108
6.7.2	Inertia.....	110
6.7.3	Progress	110
6.7.4	Action Plan	112
6.7.5	Jumpstart Plan.....	112
6.8	Idea - Removing Barriers to Entry for Crypto Products into Federal Use.....	113
6.8.1	Description.....	113
6.8.2	Inertia.....	113
6.8.3	Progress	113

6.8.4	Action Plan	113
6.8.5	Jumpstart Plan.....	114
6.9	Idea – Real-Time Internet “MRI” (Orthogonal View).....	114
6.9.1	Description.....	114
6.9.2	Inertia.....	114
6.9.3	Progress	115
6.9.4	Action Plan	115
6.9.5	Jumpstart Plan.....	115
APPENDIX A:	Acronyms	116

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 cyber security 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 Cyber Security 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 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 cyber security 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 Summit participants.

Summit Framework

The Summit utilized the Six Thinking Hats (see Edward de Bono's *Six Thinking Hats*) process and the Summit goals and deliverables to structure the working sessions. The Summit's goal was to clarify the vision by describing specific instantiations of the game changes, and to make the vision tangible by building practical action plans. To create maximum momentum, the participants were challenged to identify activities they can begin immediately. These are a smaller subset of the action plans. By considering forces of both progress and inertia, participants attempted to determine the most likely way forward.

The structure to capture each idea and associated questions below illustrate this thought process:

Idea: What does this change look like?

Description: Further explanation of the idea.

Inertia: Why have we not done this before? What would derail the change?

Progress: Why technically is this feasible now? Why environmentally is this feasible now? What would mitigate our doubts?

Action Plan: What are reasonable paths to this change? What would accelerate this change?

Jump-start Plan: Pieces of the action plan that can be started now.

1 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.

What is the new game?

In the current game, attackers win by taking advantage of the relatively static nature of our systems. For example, permanent, well known addresses, names, port numbers, etc. represent clearly identifiable parameters that turn vital servers and services into an easy target. Adversaries can plan at their leisure, relatively safe in the assumption that our key IT assets will look the same for a long time. They can map out our likely responses and stockpile a set of exploits that escalates in sophistication as we deploy better defenses. They can afford to invest significant resources in their attacks because they expect to persist in our systems for a long time. In the new game we win by increasing the randomness or decreasing the predictability of our systems. By making the cyber terrain appear chaotic to the adversary, we force him to do reconnaissance and launch exploits anew for every desired penetration; the attacker enjoys no amortization of development costs. The new game, in this context, consists of considering very dynamic rather than static network architectures. In other words, the new game is about real-time distributed monitoring, control and diagnosis of very dynamic and flexible cyber environments.

1.1 Idea - Mutable Networks: Frequently Randomized Changing of Network Addresses and Responses

- Create Virtual Machines (VMs) that are rotated and exposed to the attacker only for a limited time
- Applicable for short transactions
- Restart with different operating system
- Concerns
 - Virtualization performance
 - Total cost of ownership
 - Fixed patterns of management
 - Difficult to do root cause analysis because the Intrusion Detection System (IDS)/Intrusion Protection System (IPS) does not work
- Paths to This Change
 - Round robin address movement
 - Frequency-hopping analogies
 - Approaches that are unpredictable or not necessarily random to attackers
 - Redundancy, recovery, fast switching
 - Deployment on new architectures, e.g., the smart grid
 - Tunnels for hidden services

- Building on Content Delivery Network (CDN)
- Deployment on an overlay
- Derailers
 - Lack of demonstrated scalability
 - Lack of Internet Protocol Version 6 (IPv6) adoption because uses large address space
 - Architectural invariants, if any
 - Usability impact on systems

1.1.1 Description

A prerequisite for building successful cyber defense systems is to investigate effective countermeasures to scanning and reconnaissance attacks that allow for discovering network resources end-addresses and system fingerprint. Scanning and reconnaissance attacks are precursory steps to launching devastating attacks such as system penetration or denial of service. The objective is to provide the ability to dynamically change the external host interfaces such as names, IP addresses, and port numbers. Also, the external response behavior should be randomly changed to counter scanning worms, and reconnaissance and fingerprinting attacks. These changes are accomplished by continuously outdated the collected system information within a short time window, and deceiving attackers to fake targets for further analysis.

In this proposed approach, networked systems (i.e., end-hosts) will be assigned different addresses frequently based on random functions such as hash tables. One approach is to select interfaces using the randomized round robin technique. The change has to be done:

- On a high frequency basis to outperform automated scanner and worm propagation
- Quickly to minimize service disruption and delays
- Unpredictable to ensure that future IP addresses and keys are undiscoverable and irreversible (i.e., high entropy distribution)
- Operationally safe to preserve system requirements and service dependencies.

Redundancy can be added to this scheme using Virtual Machines to support recovery and diversity to the attack profile surface.

We have two mechanisms to randomize external system responses:

- First, as a short-term approach, session control responses such as Transmission Control Protocol (TCP) 3-way handshake, in network applications, will be intercepted and modified to give a false fingerprint identification in order to deceive and analyze the reconnaissance adversaries. However, in the long-term, it will be advantageous to have camouflaging capabilities integrated in the system session control.
- Second, firewalls will also deceive scanners by generating positive responses for all denied packets. The combination of these two techniques will give an effective motion target approach for countermeasure reconnaissance attacks.

1.1.2 Inertia

- Requires instantaneous update of network routing tables and security policies
- Scalability: How can such activities be done in a timely fashion for large networks?
- Lack of theoretical foundations to model and analyze network configurations

- Lack of efficient distributed configuration management that can orchestrate such dynamic changes without causing inconsistency and access or availability problems
- Lack of efficient network proxies and indirection technologies
- Lack of adoption of IPv6 to maximize IP addresses hopping
- Lack of efficient and scalable address translator Network Address Translation (NAT)
- Not capable of supporting multiple interfaces in MAC and IP level
- Lack of techniques to manage session and network perturbation as a result of dynamic changes such as service interruption due to mis-synchronization, and mis-configuration
- Requires maintenance of service dependency and system invariant
- Impact and overhead on operational system functionality, reliability and performance

1.1.3 Progress

- Availability of efficient and widely accepted virtualization configuration
- Ability of high-speed networks with rapid update capabilities
- Multi-switching hardware
- Recent improvement in computation including desktop, module checker, Boolean Satisfiability (SAT) solvers
- Better understanding of attacker tactics

1.1.4 Action Plan

- Leverage hashing technology, develop a function to generate network interface randomly considering the time, a shared secret key and service dependency
- Modify network protocols to support multiple simultaneous interfaces at the end hosts during the transient changing period
- Implement a distributed controller to coordinate the dynamic allocation and distribution of network address
- Implement rapid hot-swapping for router and host configuration changes
- Use OS and/or Kernel/Chip level direct reconfigurable address and translation tables
- Use software level retranslation for port connection
- Integrate this technique in Domain Name System (DNS) and Dynamic Host Configuration Protocol (DHCP) to support dynamic address-hopping technique

1.1.5 Jump-Start Plan

1.1.5.1 Technical Plan

- Use a simplified approach to implement the basic components of the system including pseudo random function, and centralized management controller
- Leverage open source OS such as Linux to make the necessary changes in the protocol stack to make IP tolerant to address-switching transient delay
- Using diversity of VMs to simulate different system responses (fake finger printing) and create a false identity
- Build proxies for address translations and redirection

- Use open source virtual router implementation to demonstrate configuration hot-swapping.

1.1.5.2 Experimentation Plan

- Identify testbed demonstration opportunities and demonstrate relevant capabilities using research networks (e.g.: Defense Research Engineering Network (DREN), DETER etc). DETER is a testbed for network security projects.
- The following use case studies should be implemented:
 - Use these test beds to implement the basic components of dynamic address motion and evaluate the effectiveness of this approach against random scanning and divide-and-conquer worms. The objective is to demonstrate the effectiveness of this approach to significantly slow down worm propagation by increasing uncertainty in scanning phase. Also solicit real worm traces from companies like Symantec and Cooperative Association for Internet Data Analysis (CAIDA) repositories.
 - Test the fingerprinting and firewall deceiving techniques against automated scanning tools like Network Mapper (Nmap) and Nessus, a network scanner tool, as well as using real scan traces from Semantic.

1.1.5.3 Team Collaboration and Bootstrapping

- Approach and engage potential collaborators from configuration management, network device vendors, ISPs and security operations and management industries through a series of talks and panel discussion during an invited 1-day workshop. We identified the following potential collaborators based on their relevance to the projects:
 - RedHat for using Linux in our short-term case study
 - Telcordia for automatic synthesis and verification of network configurations
 - Cisco for the network virtualization and hot-swap configuration capabilities
 - VMWare for integrating fingerprinting deception mechanism in the virtual machines
 - Symantec for test and evaluation using real scanning traces. We will also deploy this on a real operational network with collaboration with AT&T.
- Engage government agencies such as the National Security Agency (NSA) and Army Research Office (ARO) / Army Research Laboratory (ARL) to evaluate the potential of this idea on mission critical networks

1.1.5.4 Case Study

- Use an identified testbed (e.g., DREN or DETER) to evaluate the effectiveness of this approach against random scanning and divide-and-conquer worms. The objective is to demonstrate the effectiveness of this approach to significantly slow down worm propagation by increasing uncertainty in scanning phase. Solicit real worm traces from companies like Symantec and CAIDA repositories.
- Test the fingerprinting and firewall deceiving techniques against automated scanning tools like Nmap and Nessus tools as well as using real scan traces from Semantic.

1.2 Idea - Diversity in Software

1.2.1 Description

Currently, we live in a software monoculture - most computers run essentially the same software. This makes it easy for an attacker because the same attack vector is likely to succeed on most computers. If we make every computer run a subtly different version of the same software, a different attack vector is needed for different computers. From the perspective of the end-user, all the different versions behave in exactly the same manner, but they implement their functionality in subtly different ways.

As a result, any specific attack will succeed only on a small fraction of systems and will no longer sweep through the whole internet. An attacker would require a large number of different attacks and would need to target the specific software versions that are susceptible to each specific attack, which radically increases the cost to the attacker. The effective penalty to the attacker is the inability to amortize knowledge over a series of attacks - each attempt is distinct from any previous attempt or attempts. If multiple versions of the same software are run in parallel on a single computer, attacks could be detected in real-time when the behaviors of the versions diverge as the result of an attack that is successful on only some of the versions, but not on others.

1.2.2 Inertia

Until now, software was predominantly shipped "in boxes on a CD". Mass production of the CDs made it impractical to give every user a different version. But we are rapidly transitioning to software distribution over the network, where this is no longer a concern.

There is a cost associated with creating diversity. Until now, people have been oblivious to the risks and have not embraced the idea of paying for security. The tradeoff between security vs. performance is only now becoming better understood by a wider audience.

Until now, we have focused on creating the "best" version, e.g., in compiler optimizations. Only one of the versions can be the "best". So if we give a different version to every user, by definition, not everyone can have the "best" version. So there is a performance cost associated with this solution. There is an additional intrinsic cost of diversity - configuration management, centralized administration, etc. might become more onerous.

Security has in the past focused on "predictability" and testing. The idea of running completely different code on each individual computer requires a radical shift in thinking and culture and certification and accreditation, because, by definition, one can no longer test all of the versions, but one is required to trust the compiler.

Understanding the complexities of software and hardware dependencies among linked/embedded applications is not well preserved.

1.2.3 Progress

Distribution of a different program version to each and every customer becomes feasible when software is downloaded via the network rather than installed from a CD. We have just arrived at the point when many programs are now routinely installed only from the internet. For example, more than 400 million people have downloaded the Firefox browser.

Computers now have such high performance that paying a small performance overhead such as 5%-10%, for the extra security brought about by diversity, may be worth the cost in many contexts.

Compilers have advanced very significantly, so that automated generation of variants is now a reliable and predictable process. Even dynamic compilation is now routinely employed with very high reliability. For example, Apple has transitioned millions of users from the PowerPC to the Intel architecture using a fully automated just-in-time compiler without any reported incidents. The reliability of these compilers is stunning, considering that they have been able to automatically translate programs of the size of the Microsoft Office suite fully unattended, without any testing of the resulting output, and on-the fly.

Multi-core processors offering high degrees of parallelism (80 cores already announced by Intel) make it feasible to run several slightly different versions of just one program in parallel.

1.2.4 Action Plan

- Develop compilers which, instead of choosing the best path, preserve all legal alternative paths
- Develop a software distribution engine that queues up different variants of a software program so that the first requester gets the first version, the second requester the second, etc. The system would continuously generate new versions to queue up at the same rate as requests are coming in. For small programs, versions could be generated on-the fly at the time of the request, but for larger programs (e.g., Firefox or the Apache server), such versions would be generated ahead of time and queued up for delivery.
- Develop n-version systems that execute multiple versions of the same software in parallel for added resilience against attacks
- Develop randomization techniques that further increase the variability to an attacker without changing the functionality for the end-user
- Develop inventory management database to track how versions are distributed and provisioned. In many cases, no inventory management may be necessary at all. For example, we don't really care which version of Firefox any given user has.
- Tackle the hardest problem Commercial Off-the-Shelf (COTS) or layered/embedded multiple COTS

1.2.5 Jump-Start Plan

Pick an existing open-source project (Firefox, Apache) with documented past vulnerabilities. Modify the compiler used in its build process to generate many functionally equivalent versions simultaneously. Run old software versions with known vulnerabilities through the diversity mechanism and measure which proportion of attacks no longer succeed on the diversified code base.

1.3 Idea - Robust Random Authentication

1.3.1 Description

Tests to authenticate someone vary dynamically (at different points).

1.3.1.1 Concerns

- Usability, user acceptance
- Finite number of mechanisms
- Difficulty in delegating
- Take a small number

1.3.1.2 Mitigation

- Deploy ubiquitous Public Key Infrastructure (PKI). There are examples where this has been deployed
- Could use a fingerprint stored in Trusted Platform Module (TPM). This eliminates passwords and other weak forms of authentication.
- Provide diversity in end-user authentication for both human users, smart devices (sensors), and application software connections in manner, timing and channel. Apply a combination of multiple biometrics (e.g. face, voice, keystroke), multiple tokens (e.g. PC/phone signature, multi key fobs), and over multiple channels (e.g. web, email, voice, text) to authenticate not only at a defined log-on, but possibly during the session for validation. For the applications layer, use analogous continuous authentication (e.g. a low detectable, frequent challenge/response protocol possibly via keystroke, facial).

1.3.1.3 Benefits

- Raises the bar for any attacker attempting to steal a user's credentials, authorizations, or impersonate a user's identity by requiring the attacker to steal, counterfeit or spoof stronger credentials (not just user password and out-of-wallet information). Also the attacker must time this, not only at log on time, but over the entire user session at unpredictable times and over multiple channels.
- Increases privacy by reducing the spread of Personally Identifiable Information (PII) across multiple websites, as the user can be authenticated by a federated authentication; make possession of PII insufficient to gain control over a user's accounts or to be able to impersonate the user over the Internet because stronger credentials, such as biometrics, are required in addition to knowledge of PII, to be authenticated.

1.3.2 Inertia

- User acceptance and historical precedence
- Early immaturity (performance and cost) of biometrics
- Early cost and inconvenience of tokens (necklace problem - by necklace problem we mean that the early implementation of this approach required each website to provide their own token/credential, such as a One Time Password (OTP) token, so the user needed a growing number of tokens/credentials - one per website)
- When the Internet first got commercialized, there was not sufficient commerce to attract organized crime and it was not a sufficiently big problem to require more than ID and password over Secure Sockets Layer (SSL)
- Need for mutual authentication and ability to address man-in-the-middle, man-in-the-browser attacks
- Vulnerability in the initial registration/credentialing process

- Scalability to work with 10s of millions of users over 10s of thousands of sites

1.3.3 Progress

- Moore's law (decreased cost, increased capability) provides the necessary computational power for authentication devices at more affordable costs
- Advances in biometrics - improvements in performance at lower cost
- Advances in tokens and growing ubiquity of the smart phone making multiple channels, biometrics, device fingerprinting, geo-location all practical now
- Changing attitudes as cyber crime has dramatically risen. User acceptance and demand for stronger authentication is growing, as well as greater acceptance of white-listing, along with coincident improvements in browser design – greater isolation between browser sessions
- Growing willingness for key identity providers such as government and financial services to cooperate in initial user identification

1.3.4 Action Plan

- Work with the smart phone companies and carriers to incorporate FI-issued credentials and required access methods
- Utilize the Federal Federated Identity Management Bridge authentication as a foundation to grow upon, as well as other popular Identification schemes (e.g. CardSpace, Open ID)
- Prototype and validate in a test bed using a smart phone, with browser either on PC, or on the phone itself, with strong financial service user registration, credential issuing and verification
- Demonstrate that the prototype satisfies user acceptance, privacy, security and liability concerns, and works in the face of defined threat and red team attacks

1.3.5 Jump-Start Plan

- Build upon current smart phone designs and Wireless Fidelity (WIFI) authentication infrastructure services
- Pick a few compelling high assurance applications (e.g. from Government, Finance, and Healthcare) with friendly users (e.g. customer employees) to pilot

1.3.5.1 Use Case

As part of this effort we would include a number of examples and test cases that can serve as explicit illustrations of how the pilot can be expanded and used by a larger audience. One test case could be to have three or more financial institutions, at least one non-financial company and at least one government agency cooperate to use interoperate medium Federal Institute of Processing Standards (FIPS)/National Institute of Standards and Technology (NIST, Level 3) assurance credentials for login to multiple online sites. The scenario might include a member of another critical industry requiring high identity assurance, such as the healthcare industry. The scenario could also illustrate how authentication could be applied to smart devices such as power grid sensors.

1.4 Idea - Resilient Cryptographic Systems

Most cryptographic techniques, protocols and implementations today are brittle and vulnerable to catastrophic collapse of security due to a single point failure. This is in part because remote penetration, social engineering, insiders, supply chain modifications, and the age-old practice of bribery continue to provide successful means to bypass cryptography. Better cryptography, longer key lengths, algorithm composition, etc., do absolutely nothing to remediate these bypass vulnerabilities. The goal is develop a new generation of cryptographic systems that are resilient to multiple compromises. Although new cryptography can incorporate multiple hard mathematical problems, attention to the broader range of attack surfaces is necessary to staunch current hemorrhaging.

1.4.1 Description

Cryptographic systems can collapse due to failures in multiple dimensions, or attack surfaces, often beyond the crypto-analytic components. By making these dimensions impervious to single failures, attackers will face increased work factors. Below are listed dimensions of fragility together with approaches to improve resiliency.

1.4.1.1 *Randomizer Failure*

- Compensate with multiple random sources
- Utilize external sources of randomness
- Devise more resilient protocols to manage low entropy randomness

1.4.1.2 *Incorrect Implementations (Supply chain)*

- Develop independent implementations and compare their outputs
- Improve third party certification and accreditation
- Incorporate real time test vectors to check cryptographic operations actively

1.4.1.3 *Secret Key Compromise*

- Use techniques for split keys and distributing them to non-intersecting security domains
- Develop techniques for key agility
- Employ third party assistance in crypto computations (example. composite private keys)
- Deploy tamper resistant containers

1.4.1.4 *Side Channels and Covert Channels*

- Develop useful models of information leakage and cryptographic computational methods resistant to such leakage
- Devise techniques for reducing timing synchrony (consistent timings)
- Deploy techniques for power leveling
- Implement techniques for obfuscating hardware cache behavior
- Use encoded computation to maintain secrecy even in the presence of side channels leakage
- Improve virtual machine separation at hardware and software level, to reduce threat of cross-VM key ex-filtration

- Identify and construct minimal secure components from which larger secure computations can be built up

1.4.1.5 Software Bugs

- Write crypto code in safe abstraction-oriented programming languages designed for verifiability
- Require verified compilers
- Verify crypto code

1.4.1.6 Hardware Failure

- Use active checking to assure correct numeric calculations
- Design for minimizing catastrophic effects of faults, e.g., prevent "fault attacks", where a single bit flip causes a full key leak, as some current algorithms
- Use late binding logic, e.g., Field Programmable Gate Array (FPGAs), for crypto operations
- Perform computations redundantly on separate processing units with strategically different supply chains

1.4.1.7 Depot and Distribution Vulnerabilities

- Develop crypto systems using certified supply chains
- Institute certified tracing and handling for crypto systems
- Devise deployment mechanisms that enable rapid, or even dynamic, update of crypto algorithms or protocols

1.4.1.8 Weak Standards

- Engage broader communities in design of standards (pre competition)
- Use NIST competitions to "red team" algorithms

1.4.1.9 Loss of Physical Security

- Deploy anti-tampering techniques
- Use volatile storage for keys
- Develop techniques to reconstitute trust reactively in response to breach or proactively to assure system loyalty

1.4.1.10 Novel Attacks

- Exploit mathematical leverage beyond factoring
- Develop algorithms that resist quantum attacks

1.4.2 Inertia

- System security has been the weakest link
- The community is entrenched in private key trust model
- Government resistance to widespread distribution of more robust cryptographic systems
- Widespread deployment of current PKI models makes upgrading slow

- Misplaced belief that strategies such as algorithm composition, diversity, and frequent updating will provide more security when, in fact, they primarily introduce unneeded complexity, signatures, expense, updates and licenses (multiple vendors)

1.4.3 Progress

- Vibrant academic cryptography community
- New crypto models (e.g., elliptic curve cryptography, identity-based encryption, homomorphic encryption, leak-resistant crypto)
- New authentication schemes (e.g., multi-factor authentication, identity-based authentication, mutual authentication)
- Recent progress in verified compilers and verification of software and hardware
- Specialized programming languages for crypto (e.g., Cryptol)
- Trusted Platform Module (TPM) and Trusted Computing (TC) effort
- Greater integrated circuit capacity
- Weak system security renders more conventional crypto ineffective and creates a need
- New computational platforms (mobile, cloud) and convergence pose new challenges for crypto
- Considerable experience with deployed cryptographic systems

1.4.4 Action Plan

- Fund research to develop more resilient cryptography and an advanced implementation tool chain
- Fund research to develop wide-area collaboration systems to support design, development, implementation and management of cryptographic systems
- Establish a program for teaching crypto to advanced high school students, including a summer math camp
- Develop interoperable standards for resilient cryptographic systems across the vulnerability dimensions
- Weave resilient crypto into the fabric of system and network architectures (synergistic protection)
- Adopt new standards for government use to prime commercial build out
- Mandate use of more robust cryptography in areas requiring higher levels of assurance in the context of markets stratified by levels of information assurance necessary for safety and security

1.4.5 Jump-Start Plan

- Hold workshops on:
 - Resilient cryptography to mobilize the technical community
 - Verified adaptive programming languages for crypto
 - Hardware architectures to support resilient crypto
 - Application needs for early adopting sectors

- Announce a challenge competition for resilient crypto to engage a broad community in the development of new paradigms for resilient cryptographic systems
- Jump start research via new funding on advanced programming languages designed for crypto code
- Fund initial studies and research seedlings to explore the feasibility of resilient cryptographic algorithms, protocols, and software implementation tools in the context of critical sectors

1.4.5.1 Use Cases

- Implement stateless clients for financial transactions that leverage personal mobile hardware tokens. Use a thin client and flush all state after every transaction. Persistence occurs at server and the personal token hardware. Move the security onto personal hardware where it can be defended using resiliency techniques.
- Other areas include critical infrastructure, Supervisory Control And Data Acquisition (SCADA) systems, Voice Over Internet Protocol (VoIP) systems and electronic voting

1.5 Idea - Connectivity Diversity

Introduce duplicative, rotating network connectivity, redundancy in throughput, larger number of network traffic paths.

- Concerns
 - Performance, traffic engineering, limited physical diversity
 - Requires communication between multiple parties
 - Routing/complex communication
 - Limited physical diversity
 - Peer-to-peer communication risk
 - Keeping it simple would make it easier to penetrate
- Mitigation
 - Frequency hopping is an example
 - Commercial products that changes port numbers, IP addresses (e.g., Network Address Translation (NAT))
 - Ubiquitous connectivity
 - Enhancements to IP routing protocols
- Useful help from other groups
 - Cyber-economics group can help by developing economic/business models for assured services that satisfy both network providers and mission-critical users
 - We need an economic model for Service Level Agreements (SLAs) with provider having incentives to meet SLAs; it is a real "pain" when they don't

1.5.1 Description

Connectivity diversity (or path diversity) refers to the ability to provide multiple physical and virtual paths between information sources and users. It includes physical path, transmission

media, logical path, provider (carrier), and technology diversity. Also included is the capability to create unpredictable and dynamic paths using intelligent Sense-and-Respond mechanisms that minimize the opportunity for single-points of failure. This makes Denial of Service (DoS) attacks and Man-in-the-Middle (MiM) attacks more difficult to achieve because the path that data packets travel through the network changes in unpredictable ways. End systems do not need to know the algorithm for the path changes; only the network equipment including edge routers needs to know this. Although the technology exists for path diversity and re-routing, the Game Change is to change paths "unpredictably" (from an attacker's perspective) with Sense-and-Respond intelligence.

The business case / benefits for connectivity diversity (in addition to the cyber-security benefits) includes the use of path diversity as a mechanism to support disaster recovery / continuity of operations Disaster Recovery (DR) / Continuity of Operations Plan (COOP).

1.5.2 Inertia

Why have we not done this before? What would derail the change?

- Concerns about end-to-end performance from a user perspective. This includes network performance/overhead to dynamically change the paths without disrupting ongoing data flows/connections.
- Complexity of creating and managing multiple diverse paths between endpoints
- Network providers provide reliable service using lowest-possible cost physical media, not diverse or redundant path
- Network planning and traffic engineering becomes complex
- Multi-vendor solutions create operational support expenses issues as well as more cost up front

1.5.3 Progress

Why technically is this feasible now?

- Network providers now provide foundational technologies (Multi-protocol Label Switching (MPLS), anycast/multicast, IPv6)
- Management and monitoring tools are becoming more sophisticated and autonomous, allowing control at a segment-by-segment level
- Cloud and Service Oriented Architecture (SOA) technologies combine with architecting at the "Services" level of abstraction (vice the technology level), allowing "Services" to be created and accessed independent of the underlying technology
- Dynamic Domain Name Service (DDNS) is available
- Connectivity is becoming ubiquitous, with multiple paths available between endpoints (fiber, copper, wireless point-to-point, cellular, 802.11 (WiFi) and 802.16 (WiMax), satellite, Broadband over Powerline)
- Self-healing network technologies are available

Why environmentally is this feasible now?

- Many enterprises are already providing limited connectivity diversity for DR/COOP

- Many network providers are competing in the same market, creating redundant paths between endpoints
- Provider networks are designed with redundant and diverse paths embedded internally
- Customers are willing to pay for assured services - commercial business models exist e.g. Quality of Service (QoS)

What would mitigate our doubts?

- Availability of bandwidth enables over-provisioning to mitigate performance problems
- Failover techniques such as SONET Rapid Path Restoration (RPR) have shown that switchovers can happen instantaneously with near-zero performance impact
- Planning tools that allow prediction of path performance before alternate path selection can be created using current/near-term technology
- Management tools can select from pre-defined alternate paths can be created to minimize traffic engineering and management complexity
- Network providers are already using vendor-diversity to avoid sole-source issues and provide different cost/benefit tradeoffs at the different network layers
- Mission-critical users are less cost-sensitive when buying assured services - different business cases exist

1.5.4 Action Plan

What are reasonable paths to this change?

- Pre-planned disaster recovery scenarios taking advantage of resilient connectivity already exist in some places; these can be leveraged as examples of what's already being done
- Large scale demonstrations can be created on test networks (DREN, Global Environment for Network Innovations (GENI), very high-speed Backbone Network Service (vBNS+), Planet Lab, Emulab/DETER, etc.) in support of cyber-exercises. These demonstrations should be done in conjunction with other cyber infrastructure workshops, cyber war-gaming exercises, etc.
- Incremental network planning steps can be made less complex using "brute-force" techniques – over-provisioning, QoS and dedicated Virtual Local Area Network (VLAN).
- An "overlay" approach can be used, starting with a small number of diverse paths and overlaying additional path/segment diversity to build in greater and greater levels of robustness
- Management tools that can orchestrate the required level of dynamicity may need to be developed and rigorously tested - vendors would have a critical role here

What would accelerate the change?

- Availability of more sophisticated routing protocols that embed significant connectivity diversity and control within the network layer equipment (analogous to Hot Standby Router Protocol (HSRP))
- Providing significant incentives to network providers for implementing increased levels of diversity (or, conversely, providing significant disincentives when lack of diversity leads to reliability, availability or performance issues (strong SLAs))

- Evolving network overlays such as Smart Grid control or Healthcare Information interchange could be designed with the necessary sensors for dynamic path diversity "built-in"

1.5.5 Jump-Start Plan

Pieces of the action plan that can be started now:

- The academic and open-source software community should prototype a solution using sense-and-respond intelligence for a quick proof of concept using open-source routing software (Zebra or Quagga)
- A consortium of government (possibly including NATO nations), industry and academia should identify test bed demonstration opportunities and demonstrate relevant capabilities using research networks (e.g., DREN, vBNS+, DETER, etc.)

An example use-case is to have a network with multiple physical and logical paths available using current routing and recovery techniques, engage NSA or other skilled red team to perform a Distributed Denial of Service (DDoS) attack targeted at denying service at a target host; then enabling connectivity diversity and performing the same DDoS attack - access to the host should remain available using other network paths and media. (This use case / test case should prove the hypothesis of defeating DDoS attacks.)

Start longer-term research efforts by building collaborative teams such as:

- Engage network providers (e.g., Verizon, AT&T, etc.) to determine their current/planned future state and their approaches for responding to security events, to create a synergistic vision and collection of Best Practices related to path diversity
- Engage Management Systems vendors (e.g., CA, HP, IBM, etc.) about extending capability of management platforms to provide connectivity diversity control using Sense-and-Respond methods
- Engage network equipment vendors (e.g., Cisco, Juniper, etc.) for discussions of embedding capability within network equipment
- Engage Internet Engineering Task Force (IETF) to develop standards for diverse connectivity routing platforms

1.6 Idea - Decoys

Most applications, systems and networks are not perfectly secure. Hence, it is a matter of time until they can be compromised in a targeted attack. The core idea of decoys is to distinguish attackers from authorized users and additionally provide a large number of decoys (fake targets) to attackers while only providing the real targets to authorized users. As a consequence, attackers will be slowed down (probably confused or discouraged) by interacting with fake targets and defense will be able to easier distinguish authorized from unauthorized activities, i.e., detect new attack activity. Ideally, this mechanism will be invisible to the authorized user.

- Value - Defense can detect new attack activity, automatically analyze new attacks, and learn predict and prevent attacks based on early attack stages before the attacker reaches the real target. The result is containment of risk from imperfect networks, systems, and applications by deflecting and mitigating attacks as they develop.

- Concerns
 - Legal barriers
 - Management (ability to detect real system in an emergency)
 - Scalability
 - Cost
- Mitigation
 - Virtualization: ability to create multiple decoys, easily
 - Attempt to change legal framework

1.6.1 Description

Decoys provide several advantages to defenses in cyberspace. First, they can decisively delay and confuse attackers by presenting them with fake targets. Second, since decoys are not usually accessed, any such access points to ongoing attacker activity, which can range from mapping out networks to launching exploits or denial of service attacks. Detecting new or newly initiated attacks, together with slowing down the attacker, the defense wins valuable time to prepare a response or to study attacker's behavior to discover unknown ways of attackers (unknown 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 fake systems, virtual machines, applications, data, or networks.

Attackers end up at decoys because the decoys are reachable over shortcuts or they may bypass common access control patterns. The 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 real: decoy targets must be very low, for example on the order 1:10000. This, in essence, creates a large additional attack surface that an attacker needs to cover before eventually zooming in on the real target (c.f., Honey pots and Honey nets). There are several ways to 'slow down' attackers at decoys; they reach from simply shallow multi-system emulations listening on 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.

1.6.2 Inertia

Why have we not done this before? What would derail the change?

- Manageability of creating, destroying, migrating decoys and tracking decoys
- Cost or lack of scalability of building decoys and maintaining them in the 'image' of evolving targets. This requires extremely fast and low-overhead cloning of systems.
- Legal: If users end up at decoys instead of real services there could be legal consequences, especially for critical services (e.g., controller applications, data base applications, financial transaction servers, emergency services based on VoIP).

1.6.3 Progress

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

- Virtualization answers several important scalability questions:
 - Cloning of VMs becomes as easy as "forking" a processes (copy on write memory and storage might allow instant cloning even of fully deployed VMs at run-time)
 - Default configurations of NAT-ed, and encrypted communication channels with appropriate access controls prevents attackers from easily sort out decoys by observing network traffic
 - Optimization based on hardware or OS level virtualization enables to prioritize real targets to limit the overhead of decoys. Such optimization might offer opportunities for attackers to distinguish Decoys from real targets (e.g., response time or other side-channels).
- Advanced analytical capabilities to correlate large traffic streams in real-time enable real-time learning by observing attacks on random decoys to protect the real target

1.6.4 Action Plan

What are reasonable paths to this change? What would accelerate this change?

- Develop real-time 'multi'-cloning of VMs or applications at minimal cost and in various depths (OS/Application simulation --> full cloning)
- Develop OS/Apps that automatically create shadow decoys for data and executable files to confuse attackers (data) or increase cost of planting Trojans. Could be seen as a form of file-system randomization.

1.6.5 Jump-Start Plan

Pieces of the action plan that can be started now:

- Create decoys or "shadow" services for systems or VMs on demand for high value targets. Leverage existing honey pot technology, such as Honey nets and Black-hole sensor systems (e.g., see Internet Motion Sensor). Configure the decoys according to the perceived threat if required (e.g., make sure the attacked service or OS is emulated or simulated sufficiently to not raise suspicion).
- Analyze distributed attacks detected at sensors to layout the best positions for in-line network Intrusion Prevention Systems (IPS). Then, enable decoys to create detectors or simply signatures on-the-fly. Finally, configure IPS at those strategic network positions and provision them with those newly created signatures or detectors. Virtualized environments offer sufficient capabilities to instantiate network IPS, e.g., on open source industry standard such as Xen, using Domain0 network interception, or VMware using the VMSafe introspection APIs. Real-time stream analytics can analyze decoy sensor data even in case of broad attacks on-the-fly and correlate it with network layout information to determine strategic intersection points for the IPS.
- Test signature and detector creation in a small setting, then run large scale tests to validate and optimize the positioning of IPS for different network topologies, e.g., use private virtualized testbeds.
- Later steps would include moving from the black-hole/honey pot approach that traps random attacks to a close-target approach that can protect individual systems (identified by IP address) or applications (IP address + protocol + port number). This requires (a) sophisticated real-time analytics that safely differentiate between attackers and authorized

'clients', and (b) a balancer that forwards requests from authorized clients to the real target and requests from potential attackers to decoy copies of the target.

1.6.5.1 Use Cases

- First layer of defense, slowing down attackers and offering a pre-warning system
- Contain risk (raise cost of attackers) of unnoticed compromise of high-value targets through zero-day exploits or other vulnerabilities by external attackers
- Safely study and analyze new attacks in cyber space to create models for attack prediction, prevention, mitigation, and response

1.7 Idea - Configuration-Space Randomization for Infrastructure

1.7.1 Description

Configuration is the glue that logically integrates components to support end-to-end services. It defines the logical structures and relationships at and across multiple protocol layers. Acquiring this information is critical for attack planning, e.g., for identifying high-value targets, the paths to reach them, the intermediate components to compromise, and customizing attacks to each target. We propose to make this information much harder for an adversary to acquire by randomizing it, but, doing so in such a way, that end-to-end services continue to be available. This is analogous to address-space randomization for software that makes it much harder to plan buffer overflow attacks and frequency hopping that makes it difficult to plan jamming attacks on communication links.

Notes:

- A medium-scale infrastructure can contain 100,000 configuration variables defined in the configuration files of its components. Thus, there is a very large space of possible configurations. Rapidly “moving” between different points in this space can make it very hard for an adversary to guess the correct configuration, and rapidly invalidate his “map” of the configuration.
- The idea can be used to protect infrastructure at any layer: physical, MAC, network, virtual private networking, messaging, peer-to-peer and application. Examples of configurations that one can change are addressing, security policies (firewall rules), virtual networking architecture, routing protocol architecture, and virtual server architecture.
- The idea is orthogonal to diversity because one can change configuration without diversity and still confuse an adversary
- The idea is intended not only to resist but also survive intrusions and contain their damage
- NOTE: A capability to find a new configuration satisfying end-to-end requirements is a useful one for other approaches to moving-target defense. For example, if a new virtual machine replaces an existing one, its needs to be configured to support all services that depend on it. In general, its configuration is not identical to that of the virtual machine it just replaced.

1.7.2 Inertia

- Infrastructure design, computing configurations consistent with end-to-end requirements, and debugging configurations to enforce these have been very hard problems. Today, these are manually solved. Dynamic reconfiguration has, therefore, been inconceivable.
- Governance, especially in a collaboration environment is hard. If there is no centralized configuration authority, then reconfiguration that is consistent with intended policies of all collaborators requires agreement of all of these.
- Scalability, cost and operational impact and corporate acceptability have to be proved.

1.7.3 Progress

- Modern model-checkers and SAT-based constraint solvers allow one to efficiently compute configurations satisfying end-to-end requirements. These can solve millions of constraint in millions of variables in seconds.
- Modern fault-tolerance protocols (including routing protocols for networks) allow millisecond-scale reconfiguration. Of course, these must be correctly configured or recovery is precluded in spite of availability of redundant resources.
- Virtualization has become widely available, accepted and efficient
- Resources have become much cheaper allowing us to create diversity and redundancy
- There are well defined interfaces to infrastructure components for their control and configuration

1.7.4 Action Plan

- Understand business case for idea in consultation with administrators that operate real infrastructure. An example of this would be Defense Information Systems Agency's (DISA) or the National Security Agency's (NSA) collaboration infrastructure that use host and network virtualization.
- Develop faster methods of translating end-to-end requirement/specification into configurations
- Develop faster safe reconfiguration methods, i.e., for changing configuration without disrupting mission-critical services or introducing security breaches
- Develop distributed reconfiguration methods
- Develop cooperative reconfiguration methods to allow implementation of idea across administrative boundaries
- Quantitatively evaluate effectiveness of idea with mid-term and final "exams". "Exams" will be administered by red teams.

1.7.5 Jump-Start Plan

- Realize the IETF spirit of rough consensus and running code
- Team with administrators of real collaboration infrastructure e.g., from DISA and NSA. These use both host and network virtualization.
- Team with red-team experts at these organizations
- Identify the security and functionality requirements that administrators most care about

- Create a test bed with e.g., routers and virtual machines, and set up these requirements. This test bed can be set up in e.g., DETER, or in partnership with a company with large laboratory capabilities.
- Define and implement configuration randomization plan
- Quantitatively evaluate increase in red-team's difficulty in successfully violating security or functionality requirements. Also, assess performance impact.

1.7.5.1 Use-cases/Scenarios

- A worm may try to locate the address of a server offering a particular service. But it may need to compromise other machines before it can attack the server. Before, the adversary has had a chance to compromise other machines, our system would have randomly moved the service to another machine, so the attack would be rendered ineffective.
- Host-to-host traffic is randomly made to flow through tunnels and firewall policy is changed to permit only tunnel traffic. Then, an adversary's packets are blocked.
- The layering of IPSec tunnel architecture over the IP network is randomly changed. If an adversary had planned on sniffing at a component where IPSec traffic is decrypted, that plan would be invalidated.

1.8 Idea - Distributed Data Shell Game

Break data into pieces and move it around. The results will ensure all aspects of CIA: Confidentiality, Integrity and Availability. The process obscures data thereby assuring confidentiality. Any violations of a piece of data's integrity will result in failure to recombine. Availability is enhanced by distributing the risk across locations and allowing recovery when a location is lost. The addition of cryptography to the system will further increase confidentiality and privacy.

- Break data of interest to the attacked in multiple pieces, spread them to different – redundancy scattering – fragments have to be operated on. Use different keys.
- Bit torrent
- IP issues – originally driven by the need to compress data
- Low hanging fruit

Concerns

- Larger bandwidth costs
- Law enforcement issues: how do you recover data
- How to write applications (legacy)
- Cultural problems
- Cost

Mitigation

- Improving data de-duplication and redundancy
- Low cost storage
- Already proven (bit torrent, cloud computing)

- Data vanishing
- Easy APIs

1.8.1 Description

- Break up data into pieces and distribute those pieces to different locations, which could be logical or physical. Individually the pieces reveal little to an adversary. They can only be combined at the time of proper authentication.
- To add another hurdle to the attacker, the locations of the pieces change periodically. The rate of this change will be based on the level of risk. For example, the rate of location switching increases as the number of incidents increases or as the value of the data increases.
- Cryptographic techniques can be added at the time of the data separation or at later stages in the process.
- Design into the system an audit trail that shows what has accessed and combined the data.

1.8.2 Inertia

- Cost of storage
- Infrastructure-centric data model
- Cost of bandwidth
- Performance hits on the database
- Increase in network latency
- Culture of people seeking local control over data

1.8.3 Progress

- Lower cost of storage
- People are getting used to storing their data remotely both at an individual and corporate level
- More suppliers of bandwidth for data movement
- Distributed data bases are becoming more accepted
- Network management is driving up network efficiency

1.8.4 Action Plan

- Demonstrate the new capability to national leaders in a major test range. Use NSA's red team to attempt to identify the moving data. Identify the additional work effort needed by the attackers to reach the data.
- Market the idea as a business continuity capability that allows a business to recover operations when one location is lost. Other locations will have other pieces of the data and can recalculate and re-assemble the data. This distributes the risk of a failure at any one location, and highlights its benefit for information availability.
- Promote the value of the system for being able to detect the integrity of the data. You can't reassemble the data, if any of the pieces has been compromised.
- Emphasize to early adopters its value for reducing concerns with data destruction and archiving because the data at any one location is of no value -- one can leave it behind

1.8.5 Jump-Start Plan

- Develop a limited demonstration of a few elements of the solution leveraging currently available technology such as the Tahoe File System
- Go to industry standards group and show them what was accomplished
- Make the information available to the consumer and vendor community with the goal of creating a consumer demand

1.8.5.1 Use Case

- Human resources (HR) and financial data are two of the most critical assets of any company. Both types of data, which are both competition-sensitive and personally private, need to be accessed frequently by authorized users. The confidentiality and continuous availability of this data must be assured for business operations. Currently this data is centrally stored.
- Users at the corporation or its partners gain access to the data base, and often copy the data into their local space. This exposes more data than necessary to users, and fosters uncontrolled distribution of copies.
- By distributing this data into dispersed locations, its confidentiality is assured. Yet, by allowing authorized users at either the corporate site or partner sites to access a recombination of individual data records assures its access to those who need to use it. There are certain times when this data's sensitivity is more critical and its loss presents even greater risk than normal; for example, just prior to running an earnings report. At this time, the locations of the data are changed, i.e. the data becomes a moving target.

1.9 Idea - Security on Demand

Change the current mindset from security needs to keep bad guys out to assuming that we are essentially in a fundamentally insecure environment. Therefore, if you need security (trustworthiness), you need to do things differently. The "things you would do differently" will present a computing void to the adversary (i.e., if he breaks in he will not find the address book, which will reside on the detached stick; if a zombie is installed, most of the time, he will not have a fully functional network to propagate-in general, he will have access to useless information, resources etc, or things that will become useless within a short period of time). You will dynamically constitute a "trustworthy cocoon" – on demand, to run the application that needs higher security. The cocoon will include the application as well as the infrastructure you need to use that application, and the trustworthiness will be verifiable. At the same time, the cocoon will take a different shape (variant) each time, and each cocoon will be short lived, and exposed to public networks for a short duration.

Note this is not a silver bullet for all problems – this technique will work better for applications that do not need long duration sessions.

- Separate VM for each application that can be run on a USB device (a stick with enough CPU/memory to run Linux) – e.g., Spyros Rosetta
- Leverage emerging processor architecture like Intel Virtualization Technology (VT)/Active Management Technology (AMT) or Advanced Micro Devices (AMD) Pacifica to establish a trusted path from the USB device to the laptop/desktop

- Use the laptops capability to do IO Kernel-based Virtual Machine (KVM) + Network))
- Low hanging fruit

Concern

- These things can be attacked also
- Smart sticks could be poisoned – could be shipped with malware
- Inability to pass data between different domains
- Acceptability

Mitigation:

- Proven, devices exist
- CM is manageable

1.9.1 Description

(Concept of Operations (CONOPS)) What will it look like?

Imagine the future where traditional desktop/laptop computers have become the chassis on which key chain computing Secure Digital Input/Output (SDIO) devices with enough CPU and memory to run Linux and at least one VM can be plugged in – the laptop/desktops will only be used to provide the IO/peripheral functions to the key chain devices. You will have one dedicated device for each of your critical applications (e.g., email, banking, Google app client etc.) running a VM specialized to run that app – (e.g., all other services and ports disabled). A verified version can be preloaded to the device, but the VM can also have software to load variants of the app (leverage SW diversity) from a "trusted source" (see below for how to get to that source). It is also conceivable that the device will only have a very basic loader – and when you connect to the network you will be pushed a secure variant of the entire VM.

If you need to use banking, you plug in the "banking app" device. The device-chassis pair engages in attestation checking (TPM and other HW support in modern processor architectures). If the check succeeds, the device boots up. Then, from the device's memory, a functionally equivalent variant of that application could be loaded to run on the device. (Alternatively, as noted before, a variant can be downloaded after you have network access).

When the device boots up, you (the user) request a protected path (imagine establishing a VPN tunnel) to destination from your network provider. For this to work, like a telephone network, the chassis must have a dial tone – i.e., instead of always on broadband, the chassis is connected to the ISP with a very basic highly controlled channel. If your request for secure path is granted, you have a fatter pipe, but also with VPN-type protection. You can have better QoS if you pay more:

- Then you use your application to do your transaction, save data on the device (or copy if you need to save VM (actually data for VM if any) on that, hung up on the protected path and unplug.
- Analogous things could be done at the server side too. Imagine the enterprise procuring CPU/servers from the cloud, and establishing links between them on demand.

Benefits

- The application is online for a short duration (short exposure for adversary)

- You are not connecting to the chassis unless you verify its attestation.
- You run a different variant each time.
- You procure a secure link each time.
- Enterprise management and IP rights management become easier (when the application is pushed to the stick device).

1.9.2 Inertia

- Concerns/inertia
- Device technology was not mature (CPU/memory on stick)
- Virtualization technology was not there
- Bandwidth on Demand (BoD) was not there
- The concept has not been demonstrated/evaluated for scale/complexity

Derailers

- There is a bootstrap issue – easy to see the client side CONOPS. If we make the server/services moving, how do we connect the client and server in a trusted way? Man in the middle?
 - Mitigation approaches: Secure Directory/discovery services that becomes available with ISP dial tone, leverage Uniform Resource Name (URN), Digital Object Identifier (DOI) handles etc.
- Education/Acceptance – how to get vendors/users/service providers accept this?
 - Mitigation approaches: for end users, make it easy/transparent; for providers/vendors: show them that there is cost savings or additional revenue stream (new services, control spam, better protection against botnets etc.)
- What if the smart stick is shipped with bad code?
 - Mitigation: What if MS (or choose your favorite vendor) ships your favorite product on a media that you paid for? This is no different, and no worse.
- What if the chassis computer being attacked (corrupt, rootkitted, recruited by botnet)?
 - Mitigation: The proposed solution is no worse than what we currently have. BoD limits exposure/usefulness of these attacks. Processor architecture (and other mechanisms can be engineered – prior work exists) will facilitate isolation of all communication from keyboard to the stick

1.9.3 Progress

Feasible Technology

- VM, BoD, attestation techniques are here now
- Mechanisms to create SW diversity automatically and at a low cost and with different vulnerability mix has been demonstrated (Just-in-time (JIT)), link/load level transforms, compilers)
- Cloud computing, spread spectrum/"hopping" techniques are commercially available

Environmentally Feasible

On the environmental front: realization that we are under attack, and perfect security that will prevent that is a pipe dream.

1.9.4 Action Plan

- Need to serve a wide range of users (Grandma to mission critical)
- Need to engage different stakeholders: Government services (enterprise applications), big defense contractors (mission critical applications), academic/industrial research, network providers, hardware (processor and SDIO manufacturers)
- Assemble a dream team: one intellectual lead (who is in there not sell products, but get paid for the R&D); one service/SW vendor (to offer their software on Security on Demand (SoD) sticks or a defense contractor for transition to mission critical application; one network provider to offer BoD; one hardware vendor to offer new hardware platforms; one academic research institution to liaison with academic research/open source community
- For longer term, the dream team will develop SoD applications for the proposed Healthcare Information Network or the emerging Smart Grid

1.9.5 Jump-Start Plan

Do an advanced technology demonstration (ATD) pilot on a moderate scale: choose one application (a good attack target such as outlook and exchange), give 500 random volunteers the stick device loaded with SoD client and host a dynamically managed SoD exchange servers in multiple clouds. Use BoD among the Exchange servers, allow volunteers to request for protected service from the ISP. Engage a red team to attack the clients. This project is shovel ready (BBN Technologies and CSC inputs to the NITRD Conference Leap Year processes provide more detail, prior work from a SANS can be rolled in as well) and can be started in the next 60-90 days. The project will have a 9 month development phase (to work out the right scope and remaining engineering) followed by a 9 month field trial.

1.9.5.1 Use Case

- Need a sponsor to convene the team of various stakeholders including the application owner, hardware vendor, network provider and architect/integrators
- The Outlook-Exchange target application may not be a good example-- perhaps a specialized browser for doing financial transactions is a better one where the client state can be at various places (adds one dimension for varying the application)

1.10 Idea - Terrorist Organization Model

Use the decentralized nature of terrorist groups and cells as a reference model for a new information system. Terrorist groups are hard to penetrate, not susceptible to large losses if a subpart is compromised, and can work autonomously with a very small rule set. This model is a "game changing" idea in that it approaches computer and network science in a radically different manner.

- Study terrorist model and why it is hard to penetrate, how it is resilient, if one gets captured, all get captured
- Concerns

- Revolutionary change compared to the current hierarchical model
- Cultural resistance
- Lots of unknowns

Mitigating

- Coalition: sharing networks (concept worked on by NATO) – low hanging fruit
- Gaming industry – massive multiplayer online games
- Lessons learned from mobile ad-hoc networks (MANETs)
- Cultural acceptance from the new generation

1.10.1 Description

This is a fundamentally different approach to information systems as compared to today's hierarchical models.

- Rather than linear command/control relationships, tight lines of communication, and high dependence on the successful operations of other groups (processes) the terrorist model has very loose ties, autonomy of parts, and self organized leadership
- It also has other attributes that make it very resilient to penetration and disruption such as "tribal leadership" or "headless organizations"

1.10.2 Inertia

There would be significant cultural resistance to this approach, due to the many decades of development invested in the current architectures and reference models. Also, the idea of "terrorist groups" is offensive to many and might hamper good innovation and creativity. There are many unknowns and not much literature on the specifics of how these groups communicate and protect themselves.

1.10.3 Progress

Some applications use an early and crude application of this methodology such as Massive Multiplayer Online Games (MMOGs), disposable hardware devices, social networking sites, web 2.0, and the portion of our society known as "Generation Y". Service Oriented Architectures (SOAs) might also provide some insight into how this model might work due to the "loose coupling" of services offered by SOA.

1.10.4 Action Plan

Need to better understand how terrorist groups organize, how their information networking evolves, why they are hard to penetrate, where the resilience comes from, and how the capturing of one person or cell has little impact on the entire operation. These groups might follow the principles of complex and chaotic systems, which could in turn provide insights for a new reference model for information systems.

1.10.5 Jump-Start Plan

- Use existing sharing networks and systems such as that being developed by NATO, lessons from MMOGs, or even concepts from MANETs as a basis for developing an experimental framework or model
- Leverage the different cultural values of the Y Generation, and create a Facegroup page, Wiki, or other virtual meeting place where this idea can be discussed and fleshed out

- Obtain funding from Department of Homeland Security (DHS)/Science and Technology (S&T) for a pilot in this area, and establish a public/private consortium to develop proof-of-concept technical solutions

1.11 Idea - Smart Motion Adaptation Management

Redundancy and diversity in SW, infrastructure and resources create the space where defended systems can shape shift. Develop a sound model to manage the movement in that space such that it is unpredictable to the attacker. Use variety of modeling techniques including but not limited to game-theory, machine learning, statistical, control theory, cognitive reasoning and planning to develop the algorithms that manage the dynamic system behavior.

- Model based motion management, M^4
- Inertia:
 - Hasn't been enabled in terms of mechanisms
 - Scalability
 - False positives – problem common in these approaches
 - Practitioners have good ideas looking for a fit
 - Does the model fit reality?
- Progress:
 - Provability feature
 - Way to adapt
 - High speed processing
 - Bayesian decision trees
 - Advanced reasoning engines

1.11.1 Description

The "smart management" will use the various options and/or possibilities unleashed by other techniques. For example, how to place replicas, which address/port to use, which variant to use, how to configure the network (overlay/interconnection) etc. all dynamic adaptation decisions will be governed by this smart management mechanism.

Benefits

- System dynamically configures itself for optimal security-performance trade off
- Proactive (as opposed to reactive - limit exposure)
- Adaptation is based on sound theory – easier to establish the operating regions (bounds, control theoretic proofs that certain bad conditions will never arise)
- Performance improvement
- Financial impact and brand protection

1.11.2 Inertia

- The degrees of freedom to navigate and the space to manage was smaller or not there – it is now (or we can see how it can be) with the other techniques before

- Mathematical formalisms were not mature
- Processor speed/capacity to run the compute intensive adaptation management decision making algorithms
- The communication bandwidth needed to compute the decisions was not there (wide area, reliable, ubiquitous, high bandwidth)

Derailers

- Decision cycle time: need to move faster than the attacker
- Complexity of the algorithms: self explanatory
- Model fidelity: how do we know that the model fits reality
- Uncertainty/incompleteness of observations/measurements driving the decision model
- Attacks on the management may lead to delayed or plain incorrect decisions
- Acceptance (validation) of automated adaptation management can be tricky (how do I know it will do the right thing?)

1.11.3 Progress

Feasible Technology

- Proof of concept of various types of adaptation management capability (algorithms, models) and architecture (hierarchical, centralized, peer to peer) demonstrated
- Diversity/redundancy space to manage now available

Environmentally - The stakeholders are more receptive now – with the adaptation space growing large, smart management is inevitable.

1.11.4 Action Plan

- Identify a transition target (smart grid/Healthcare Information Network) – build the new entity such that it has smart dynamism built in
- Grid or HIN with smart management cannot be built in one step – attempt to reach interim milestones: First build a smart management mechanism that works in a passive mode (it gets all the data, does all the computation, produces results – but does not control the system – the results are for humans to validate the mechanism). As the second milestone, use the smart management mechanism as an expert assistant – it will offer suggestions to real operators/controllers, and perform some tasks automatically, but under operator's supervision – operator needs to check off first. The final milestone is to make the smart management system fully operational – the operators will still have an override switch.
- Assemble a team to work on this. A number of past Defense Advanced Research Projects Agency (DARPA) and National Science Foundation (NSF) funded projects developed and demonstrated building block capabilities that can be used.

1.11.5 Jump-Start Plan

- Developing a moderate scale smart management architecture can start within the next 60-90 days. Existing (e.g., DETER, Planet Lab) and planned (National Cyber Range) testbeds can be used to provide venue for testing. After the initial proof of concept, make this framework open such that "expansion technology" vendors can contribute their

technology and create their own experiment to see how the smart management mechanism can effectively manage it, what the issues (performance, new vulnerability) are so that new research can start to address them.

- Different increments with increased scale, increased scope (more dimensions to manage). Initial candidates of "expansion technologies" that can be integrated with the initial smart management architecture framework are "software diversity" and "infrastructure diversity".
- Validate each increment (test, red team)
- Dream team for the pilot: one team experienced in building adaptive and survivable system architecture, technology providers in the software and infrastructure diversity, a government and private sector stakeholder who could use the smart management capability and provide the use case/threat requirements etc., and a red team like IV&V.
- The first step is to identify a sponsor and put together the dream team

2 Cyber Economics

New Game: Crime doesn't pay

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

What is the new game?

Today cyber crime pays. So does cyber-espionage. Security and privacy failures are often due to perverse incentives. Understanding the incentive structure is a key to getting stakeholders to behave in a way that will improve overall security. Cyber crime and cyber-espionage are attractive because the cost to engage in them is very small compared to the return on investment. Attack development costs can be amortized over both time and space. The cyber resources upon which the illicit activities are built are cheap, even free, thanks to webmail and botnets. Risk also is low when other people's assets are used to launch attacks. These advantages, however, may be more fragile than they look, as they are sensitive to slight perturbations in the economy of cost and exposure. In the new game we even the odds and make cyber malefactors take more risk at a lower rate of return.

2.1 Idea - Data & Metrics for Cybersecurity Analysis

2.1.1 Description

Markets do not work efficiently under incomplete information. Such is the case of the market for cybersecurity. Notwithstanding recent progress in the economics of cybersecurity, we still lack empirical and theoretical tools - reliable and exhaustive data and rigorous metrics on cybersecurity incidents, attacks, and infection rates - to make the right decisions. This greatly limits the types of security economic analyses that can be performed at the policy, corporate, and individual levels. We cannot answer even simple questions such as: How secure am I? Am I spending too much or too little on security? Is the cost of technology X worth the risk it mitigates? Are the costs of a certain technology worth the risk it mitigates? How can we make more efficient security investments?

What does the change look like?

The game change consists of incentivizing (through government subsidies, best practices, or mandated through regulatory intervention) information sharing among private and public sector entities, in order to create a public repository of data on incidents, attacks, and infection rates, and, where possible, security related losses. This data would enable a variety of applications and more finely tuned policy making - such as more accurate cyber risk management or cyber insurance. Furthermore, information on the security status and policies of consumer-facing businesses should be made publicly available: better security information may also allow the individual and organizational buyer to make informed decisions when choosing applications, vendors and information systems. This could raise vendors' incentives for developing secure products and services.

2.1.2 Inertia

Why have we not done this before?

Firms have few incentives to disclose information about their rates of attack and infections. In fact, firms believe that disclosing this information may uncover weaknesses that could be further breached, or that the disclosure might adversely affect their brand. As a result, it is likely that such disclosure would have to be mandated through legislation. While existing laws (often at the state level) mandate disclosure of information regarding so-called “data breaches,” such initiatives do not address the broader issue of cyber-attacks and infections.

Furthermore, the field of cyber-security metrics is still developing. While more efforts are being aimed at developing rigorous metrics, as a community, we still debate what information should be collected, and which of that information has economic relevance, how that information could be used, or how to deal with uncertainty and inaccuracy associated with such information. For instance, we lack a taxonomy of incidents that is commonly embraced across technology, policy, legal and economic communities.

What would derail this change?

Firms may continue to be wary of sharing their information publicly because of reputational or legal motivations, or may have incentives to provide incomplete and inaccurate data. While data could be made anonymous to protect the confidentiality of the firm, the type and depth of information required for meaningful analyses may need to be so extensive as to make de-anonymization of ostensibly anonymous reports a practical threat. If firms do not have sufficient guarantees that the release of such data does not jeopardize their confidentiality and brand, then they may exert significant efforts to resist any legislation promoting such disclosures.

Furthermore, a challenge lies in the need to develop appropriate standard/interoperable formats to make sure that the quality of reports is consistent.

2.1.3 Progress

What technologies are emerging that makes this change look doable now?

We are collecting more and more information than in the past. Centralized event and log collection is becoming increasingly popular, and we could learn from growing experience with data collection in certain areas of incident monitoring, response, and analysis. Furthermore, there have been significant advancements in storage systems and data mining. Such technologies could be valuable in making this game change real.

What environmental (business, political) changes are pointing in this direction?

The existence of widespread privacy breach notification laws makes the idea of disclosure more palatable to industry. It has been noted in the literature that such legislation has improved the overall security of systems that manage private information, and that disclosing breach events may not be as damaging as once believed.

2.1.4 Action Plan

First, we would need to rigorously define the scope of the solution: What types of information should be collected? What are the mechanisms for obtaining the information? Who would collect and host the information? How do we assure that we are getting the right information? What measures can we take to improve the quality and availability of information we collect today?

What inhibitors are present that preclude the sharing or disclosure of information collection and can these be overcome without unintended consequences? What safe harbors can we provide to data collectors to prevent disclosure?

We could explore data collection in a limited context – for instance, at a university, or in the public sector - as a test bed for the approach. Then, we would identify a data schema for one class of incident (such as phishing, data breaches, or denial of service attacks). We would identify methods for collecting the data defined by that schema, identify a model and methodology for sharing that information across technology, policy, and legal entities, and obtain feedback from the disparate parties as to the utility of the data collected. We would revise the data schema accordingly, and then extrapolate from the study: does a model emerge? Do we need different schema for different incidents? Do we need different collection methods? What measures do we also need to incorporate to assure data origin integrity (or chain of custody)? We could then report on results and carry to a larger test bed or carry to industry as a recommended best practice or NIST requirement.

We could also start with a less ambitious step – such as a detailed survey conducted by a credible agency – from which to build the larger efforts described in this section.

Another approach, related to the banks' online fraud losses, would have the Financial Services ISAC (or other appropriate body) collect data on losses from banks due to online banking-related fraud and the number of customers affected. Losses should be broken down according to recovered and un-recovered losses. The FS-ISAC could then aggregate the figures across banks and publish the totals on a quarterly basis. The data could be further broken down by bank size, geographic region, etc. The goal would be to set into motion a repository of public information on the overall impact of online fraud, which accounts for the majority of direct consumer losses due to Internet insecurity.

An additional (and complementary) path would focus on the development and deployment of large-scale empirical and experimental research testbeds in Cyber Security Economics, modeled after comparable initiatives in the technical side of cyber security (e.g. DETER and the National Cyber Range).

What would accelerate this change?

We could leverage or expand existing legislation, or build on existing organizations that collect data about breaches, in order to gradually bring this idea into existence.

Safe harbor considerations that assure data collectors will not be penalized for disclosure could play an important role in enabling change. Clearly defining what constitutes appropriate aggregation and anonymization of otherwise sensitive/protected/incriminating information would also help.

What are the missing technical pieces?

We should agree on the proper articulation of the right set of security metrics – and prior to identifying metrics, examine the data to be collected, and determine what metrics can be defined from these data. Furthermore, we should investigate how actual market players (both at the corporate and individual levels) make use of, and act upon, security information – this implies could be achieved by promoting interdisciplinary research on cyber-security and privacy spanning psychology, HCI, human factor, behavioral economics, and behavioral decision research.

2.1.5 Jump-Start Plan

- Coordinate with DoD and others where work has already begun
- Plan and announce a conference to unite the various multidisciplinary research communities with the goal of defining a path forward
- Learn from current approaches (e.g., CERT)
- Plan and announce a subsequent NSF program focused on the many research and development challenges posed by this topic
- Organize an interdisciplinary workshop to address how to choose, collect, standardize, and share data on incidents, attacks, infection rates, and security related losses. The workshop would bring together and extend ideas and 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, OWASP, the Center for Internet Security, MetriSec, etc.), focusing on the economic significance and purpose of those metrics. The workshop would address both theoretical and practical challenges. Workshop attendees could include government representation from the Office of Science and Technology Policy (OSTP), the Council of Economic Advisor, the National Bureau of Economic Research, and academic and industry representation in an effort to define rigorous economic metrics, data standards, and data collection strategies for the field of cyber-security.
- Implement the National Computer Security Survey initiated (but canceled) by the Bureau of Justice Statistics at the DOJ and the National Cyber Security Division (NCSA) at the DHS, in order to assess the threat.
- Empower the Federal Communication Commission (FCC) to expand the Automated Reporting Management Information System to include information on system security of network service providers (furnishing information about the network security of different providers will enable consumers and business to choose secure providers. This information could also enable periodic updating of an authoritative threat assessment.)
- The Federal Trade Commission could evaluate claims regarding consumer goods and should be funded to provide more information in the virtual realm (for example, a 2006 Harvard University analysis concluded that TRUSTe seals, which appear to vouch for trustworthiness, were instead correlated with malicious computer code and exploitive privacy policies. Additionally, a 2008 Cambridge University study found that among e-commerce sites that were subverted, those not publicly identified were significantly more likely to be subverted again).

2.2 Idea – Vendor Incentives and Accountability

2.2.1 Description

Understanding and influencing stakeholders' payoff structure through incentives and accountability is one means to getting them to behave in ways that will improve overall security and increase social welfare. For example, economist Hal Varian has argued that the burden of preventing distributed denial of service attacks should fall on the operators of the networks from which the attacks originate. However, what form of vendor incentives and product or service

accountability may be beneficial in the realm of cybersecurity remains a hotly debated topic in the literature and among policy makers.

What does the change look like?

Producers of software and hardware openly collaborate with consumers on sets of baseline security and privacy development practices and functional capabilities. Incentives for following these practices initially take the form of best practices and procurement guidelines, or certification processes. The research community teams with the producers and policy-makers on the investigation of the benefits, trade-offs and potential unintended consequences associated with a regulatory framework of accountability for software products and cybersecurity solutions. Incentives and liability would initially focus on vendors policing themselves better (or coming up with reasonable metrics or standards to improve the quality of their products) - with the caveat that, if they do not do it, then government will have to step in (Common Criteria was a step in that direction, but it needs a strong vendor involvement). In other words, in the absence of industry compliance, regulators may have to consider stronger forms of accountability, such as vendor liability, guided by the results of research in this area.

2.2.2 Inertia

Why have we not done this before?

It is not clear that market forces alone can drive vendors to invest optimally in the security of their information products. For instance, if consumers do not understand or consider security features when purchasing cyber products and services, a competitive vendor will only face limited incentives to allocate more resources to improve the security and reliability of its products. However, publicity about discovered vulnerabilities and attacks seems to have proven to be effective means of changing vendor behavior, at least in some cases.

On the other hand, the academic and policy debates have shown that regulatory interventions in this area may produce a number of unintended (and undesirable) consequences. Furthermore, the concept of product security is often a poorly specified goal, with few robust tools or processes to rigorously define it: the complexity of interaction of different software components from different vendors has made progress in development, testing, auditing, and forensics very slow. The open-source movement poses a similarly vexing problem – where would any accountability claims fall in an open source environment? As a result of these and other issues, efforts towards this goal are expensive, long-term projects and neither market forces nor regulatory bodies have supported them.

What would derail this change?

Vendors may resist initiatives that establish baseline security and privacy practices and capabilities, and would certainly oppose strong-handed initiatives aimed at establishing a liability regime. On the other hand, too much focus on vendor accountability may slow down innovation, by pushing vendors to re-allocate resources away from R&D, and forcing them to engage in lengthy debates on issues of public policy.

2.2.3 Progress

What technologies are emerging that make this change look doable now?

Advances in secure software engineering, software vulnerability detection, software analysis tools, software testing and assurance, and security incident forensics are key to this idea. Some

promising progress has been made, but this is clearly an area for further research and development investment.

What environmental (business, political) changes are pointing in this direction?

Increasing numbers of business-critical applications are raising the cost of a security incident. The rising threat of botnets has added to the pressure on system owners to secure their systems, at their own cost.

2.2.4 Action Plan

What are the reasonable paths towards bringing about that change?

Initially, the efforts should focus on safety-critical industries such as healthcare, cyber-physical systems, and critical infrastructure. The industry groups for those areas are well established and some are already working toward this goal.

It would also be important to recognize and learn from approaches and standards that are already available or under development, such as ISO SC27 and the Common Criteria Development Board (who has expressed interest in moving from document-focused standards to criteria that would actually focus on vendor techniques that lead to more secure software). Government efforts could build on such criteria and standards to achieve real-world results, in addition to considering additional research.

One suggestion resulting from the Summit discussion defined two paths to accountability: the supplier could either document their adherence to a very specific set of checkpoints during development (and be held responsible for adhering to recognized standards of good software engineering – such as not allowing buffer overflows), or they could accept the responsibility to demonstrate that their product development process meets or exceeds the same level of security. Other suggestions focused on phasing out “blanket” disclaimers and adopting product features that the vendor advertises or describes in the product manual as potential points of accountability: in other words, it should not be possible for technology vendors to disclaim warranties of merchantability or fitness for the very purposes described in their products' operating manuals.

Other models exist, but all of them will require a strong focus on the research and development of enabling technologies.

What would accelerate this change?

Limiting the scope of potential accountability (for instance, tying accountability to licensing terms, or advertised/documented functionality), establishing time-bounds for indemnification, or defining clear standards on obviously negligent practices (for instance, tied to well known classes of vulnerabilities) may help address vendors' concerns with calls for open-ended accountability.

Convening a multidisciplinary workshop, perhaps even an annual series, on the technologies and policies to support accountability in cybersecurity (TAPSAC) would provide a jump start and an ongoing drive for this research and development process. Establish and evaluate a secure development practices standards for hardware and identify best practices.

Select a subset of well-defined vulnerabilities, and require that this particular flow be subject to some liability.

2.2.5 Jump-Start Plan

- Plan and announce the first Technologies and Policies to Support Accountability in Cybersecurity (TAPSAC) conference
- Plan and announce a subsequent NSF program focused on the many research and development challenges posed by this topic

2.3 Idea – Cyber “NTSB”

2.3.1 Description

Currently, when a major breach or security incident happens, there is limited information about the root cause of the vulnerabilities which led to the incident. Often this information is gathered somewhere, but is held confidential. As a result, it is difficult for other organizations to learn from those mistakes and improve the quality of their own systems.

What does the change look like?

We envision the establishment of an entity similar to the National Transportation Safety Board (NTSB). The NTSB is 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. The NTSB is also charged with issuing safety recommendations aimed at preventing future accidents. US businesses would be obligated to cooperate with investigations. A similar organization in the field of cybersecurity would be charged with investigating major breaches and incidents, and issuing public recommendations aimed at preventing similar attacks. Such an organization may coordinate large-scale, representative public-private surveys (on the model of the CSI-FBI).

The ensuing opportunity for informed public discussion of cybersecurity risks and threats may help improve information security awareness among consumers and raise the value proposition of cyber-security.

2.3.2 Inertia

Why have we not done this before?

For all of the reasons described in “Data & Metrics for Cybersecurity Analysis” above, firms have no incentive to disclose this information.

What would derail this change?

As with “Data & Metrics for Cybersecurity Analysis” above, concerns about the confidentiality and the impact on the business of the victim of the attack may make firms strongly object to this approach.

2.3.3 Progress

What technologies are emerging that makes this change look doable now?

This is not a technology problem per se, but advancements in tools such as log collection would make forensic analysis easier.

What environmental (business, political) changes are pointing in this direction?

As the scope of breaches and security incidents becomes increasingly larger, there is a need to understand the root causes of such incidents.

2.3.4 Action Plan

What are the reasonable paths towards bringing about that change?

As with “Data & Metrics for Cybersecurity Analysis” above, we need to define the scope of the solution: What organization would be responsible for doing these investigations? How would this organization interact with existing law enforcement organizations? What is the scale and type of security breach that would warrant an investigation? What are the mechanisms for obtaining the information? How can the results of the investigation be shared?

We could explore this approach in a limited context -- for instance, in an industry where a security breach might affect safety -- as a trial to experiment with the approach.

What would accelerate this change?

In order to bring this idea incrementally into existence, as with “Data & Metrics for Cybersecurity Analysis” above, we could leverage or expand existing legislation and organizations that collect and disseminate information about breaches.

In addition, we might be able to leverage existing private organizations that do this kind of forensic investigation today. Furthermore, having a well defined set of best practices for forensics would be helpful for this idea.

What are the missing technical pieces?

None were identified.

2.4 Idea – Cyber “Interpol”

2.4.1 Description

Currently, when hackers use a trail of computers in many different countries it's hard to trace them because of jurisdictional issues. Getting permission takes so long that the trail is often cold.

What does the change look like?

The creation of an international body for the monitoring and reporting of cyber attacks and cyber security incidents, with powers to enforce international treaties in the area of cyber-crime.

Getting a multilateral treaty in place that would let authorized investigators from partner countries file a report with foreign authorities that they are investigating a crime and get access to or investigate the foreign computers in real-time could help to make international investigations more effective.

2.4.2 Inertia

Why have we not done this before?

Some existing organizations are already partially addressing this problem (for instance, Interpol itself is increasingly involved in preventing online crimes), but it was debated among the Summit participants how effectively existing organizations currently satisfy this purpose. The disruptions brought by cyber incidents have only recently reached a sufficient pain threshold so as to raise the problem to a political level. That has led to the need for discovery and education on the part

of international governments which has been a slow process. The issue has been further delayed by the potential association with cyber warfare.

What would derail this change?

Such a change would raise clear geopolitical concerns. These include jurisdictional conflicts, offensive cyber operations secrets, and cultural differences around what is deemed acceptable behavior. There are also technical challenges, primarily around attribution and privacy protection. If any legal action is to be taken, within any jurisdictional scope, there is a prerequisite of precisely knowing the alleged offending party(ies) and the victim(s). The process of creating draft international laws, discussing them among the potential signatories, negotiation of terms, and the eventual adoption of the laws via treaty is an extremely lengthy process. This change is extremely time-sensitive, both due to its necessity and the support provided by the current positive political views. It is also unclear to this group what other ongoing activities may exist in this area.

2.4.3 Progress

What technologies are emerging that makes this change look doable now?

The development of cyberforensics (and associated training of law enforcement in this area).

What environmental (business, political) changes are pointing in this direction?

There is growing international political will to support some kind of cyber rules of conduct. The attacks on Estonia and increasingly on various entities within the United States have further raised the issue's profile. Countries around the world are being driven towards defining a cyber-warfare doctrine as well as preparing (both defense and offense) for major cyber incidents. Furthermore, as notes above, Interpol itself is increasingly involved in preventing online crimes.

2.4.4 Action Plan

What are the reasonable paths towards bringing about that change?

Existing international law enforcement cooperation agreements for combating organized crime, financial services fraud, and others should be studied and either adjustments to those agreements or wholly new ones should be proposed. As preliminary steps, Summit participants suggested scholarship programs for students who combine criminal justice and computer security at the graduate or undergraduate level; "Yellow Ribbon" campaigns to encourage those with military experience to consider computer security; recommending that the new GI Bill offer extra incentives to students to combine these majors; and internship programs for students in computer security in Interpol with the Scholarship for Service or Center of Excellence populations.

What would accelerate this change?

Legal action implies identification, which in cybersecurity means attribution. This is a broad challenge (consider, for instance, privacy concerns as well as the technological challenges), even if attribution were limited to "locate the alleged offender to a nation or state". To address the privacy challenges posed by fine-grained (deep packet) network monitoring, the application of modern techniques in privacy-preserving traffic monitoring and analysis would be beneficial.

What are the missing technical pieces?

None identified.

2.5 Idea – Cyber Insurance

2.5.1 Description

A healthy cyber insurance market for both users and vendors of security products would emerge, promoting best practices and efficient levels of investment in cyber security. Insurance is one way to spread risk across multiple institutions and enforce sensible security standards (insured parties have an incentive to lower insurance cost by increasing their prevention investments). Furthermore, one of the social values of cyber insurance would lie in risk pricing for internal decision-making. This, in turn, may require non-traditional products such as risk pooling among trading partners.

2.5.2 Inertia

Why have we not done this before?

The lack of the cyber equivalent of actuarial data makes it very difficult to write cyber-insurance policies so that they efficiently cover the right kinds of eventualities. Furthermore, previous efforts to spur cyber-insurance markets have suffered due to the qualifying requirement of having pre-existing security investment. The limited number of such initiatives, so far, and their lack of diversity, results in a limited ability to correlate risk factors. Other concerns relate to the risk that cyber-insurance efforts may encourage mere compliance rather than improvements in security. Furthermore, known economic failures in insurance markets -- such as moral hazard (for instance, an entity taking chances knowing that they are “covered”) -- may reduce the probability that a cyber-insurance market would actually improve overall information security for firms and the nation as a whole.

What would derail this change?

Currently, the efforts in this area have been small, thereby making scalability an unknown. Since most vendors and service providers who produce things that would be objects of cyber insurance are international in scope, and because the activity that led to a cyber insurance claim may be from a foreign source, international law and jurisdictions will come into play.

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.

We also noted that some participants expressed the fear that advancements in cyberinsurance would make firms complacent with transferring all of their information security risks to a single policy, without improvements in actual security.

2.5.3 Progress

What technologies are emerging that makes this change look doable now?

Advances in data mining, anonymization techniques, cyber forensics, and security best practices.

What environmental (business, political) changes are pointing in this direction?

Increases in consumer security consciousness, and consumer frustration in attempting to improve security, have further accented the complexity of securing our systems. This increased

complexity makes it harder for individual entities to optimally choose the appropriate level of investment in security, thus raising the appeal of cyber-insurance.

2.5.4 Action Plan

What are the reasonable paths towards bringing about that change?

Increased availability of incident and impact data could enable the market drive this change (see Cyber NTSB and Data & Metrics for Cyber-Security Analysis ideas). Precise definitions of who is being insured (user, vendor), what classes of incidents being insured against are necessary, what would constitute an insurable event, and where insurance might actually be attractive (as opposed to self-insurance used today). Perhaps a task force or a workshop could be organized to address these issues. Ultimately, a secondary insurance market for these new instruments will be needed.

What would accelerate this change?

The ability to accurately assign clear liabilities would give incentives to this market for this kind of insurance. Better data relating to incidents and their causes would also help (see Cyber NTSB and Data & Metrics ideas). Government incentives and economic forces (once the market begins to take hold) will drive the adoption. A drive for accountability for cyber security exposures and incidents would also encourage the adoption of cyber insurance.

What are the missing technical pieces?

Insurers need a breadth of data around the things and events that they cover. The identification of the necessary data and the formal means for collecting and vetting it are key. The automatic collection and processing of the data, which will be of great volume (at least initially), will also be key.

2.6 Idea - Empowering ISPs, Registrars, and Registries

2.6.1 Description

Social welfare increases when the party who is in the best position to secure a system (for instance, because its costs of improving security would arguably be lower than the costs for other parties) is also given the responsibility of securing that system. A technological and legislative framework that empowered (but also required) ISPs, registrars, and registries to halt clearly abusive or criminal behavior may offer the tools to prevent various cyber-crimes to those in the best position to help (for instance, after McColo was disconnected, spam traffic declined dramatically – albeit only temporarily so). Consider the following scenarios: ISPs temporarily disconnecting compromised users originating spam and DoS attacks from the network, registries blocking malicious domain name registrations to disrupt botnet communications, or registrars preventing miscreants and criminals from registering domains using false credentials and fraudulently obtained credit cards or payment accounts.

2.6.2 Inertia

Why have we not done this before?

While the issue of empowering ISPs and registries has been discussed in the literature, and some experiences outside the United States have already been observed, the idea still faces numerous legal and economic challenges. Among those challenges: firms may not want this kind of

empowerment, as it may expose them to other forms of liability (for instance, the threat of an ISP becoming responsible for child pornography sent through its network, or the threat of litigation should a registrar suspend a domain name that results in loss of internet presence or commerce to a legitimate (innocent) part); in a highly competitive market, ISPs do not want to risk irking their current paying customers (or disincentivizing their potential future ones). The definition of what constitutes clearly abusive or criminal behavior may not always be so clear-cut (in absence of rigorous definitions and guidelines for the operators); any such initiative faces significant IP complications.

What would derail this change?

In addition to the reasons why the change has not been possible so far, new infrastructure may be required. In other circumstances, service delivery models and automation commonly employed by ISPs and registrars are affected. Simply put, certain checks and balances may result in a delay or temporary blocking of service to customers, altering customer experience in a seemingly adverse way. Furthermore, given the added costs and risks to ISPs (and, possibly, registries, and registrars), strong incentives would have to be provided to the operators to make this change amenable.

Furthermore, privacy considerations and the potential threat this empowerment may constitute to net neutrality could derail the initiative.

2.6.3 Progress

What technologies are emerging that makes this change look doable now?

Technologies for admission control, verification, and abuse monitoring are either available today or existing implementations could be extended to satisfy a need in a relatively short time frame at reasonable cost – in fact, the feasibility of similar programs has already been demonstrated on smaller scales. Consider the following examples: college campus monitoring for illegal distribution of copyright-protected software or music, admission controls implemented by enterprise network operators to prevent systems infected by malware from joining local or campus networks, and identity verification measures employed by financial institutions to detect and block impersonation attempts. Such implementations provide ISPs, registries, and registrars the means to better defend services and systems against common abuses and exploitation.

What environmental (business, political) changes are pointing in this direction?

We already observe a move towards ISP “pushing” security solutions to their own users – albeit this transition does not seem to be happening fast enough. Certain ISPs offer security solutions such as anti-spam, anti-virus gateways to residential broadband and dialup customers. Others offer traffic filtering, intrusion detection, and denial of service (DoS) attack abatement solutions to corporate customers. Certain registrars provide measures to protect businesses and individuals from domain registration abuse and DNS misuse. Other domain name registrars and registries have highly proactive anti-abuse programs. The success of such service offerings encourages adoption by other operators and providers and ongoing innovation.

Furthermore, decreasing traffic due to spam and other attacks (such as DoS) would benefit ISPs themselves.

2.6.4 Action Plan

What are the reasonable paths towards bringing about that change?

Define a clear legal framework for what ISPs, registrars, and registries can or can't do, including common carrier-like exemptions, good faith safe-harbors, and – possibly – mitigations of liability costs (similar to, for instance, those established in the Patriot Act) in order to make the change amenable to the stakeholders.

Positive reinforcement could also be used, instead of liability: for instance, ISPs, registrars, and registries could be monetarily incentivized to identify and halt compromised hosts, or to bundle security services for end-users into ISP, registrars, and registries subscriptions.

Internet registries could be required to verify identities of registrants.

A pilot government program could be enacted to reward ISP's who discover, repair, and clean computers infected with crimeware – in order to realign the incentives of ISP's to keep their networks from harboring crimeware; this could reduce the overall costs of bandwidth, spam and fraud. Such a pilot program would have a useful side effect of providing more information about infections. The results of the pilot could be used to evaluate whether such interventions might be helpful on a large scale.

What would accelerate this change?

Learning from the experiences with ISPs that are already engaging in similar processes.

Encouraging and leveraging existing mechanisms on reporting security problems (e.g. botnet reports).

What are the missing technical pieces?

Many technologies that could be employed by ISPs have demonstrated success in LAN but not WAN deployments. For example, certain technologies, e.g., network admission controls, may need to be modified or complemented with additional functionality to provide equivalent protective measures for users of residential broadband, dialup, and dedicated access services. Similarly, certain abuse and attack monitoring techniques may require innovation or modification to scale and perform effectively when deployed over wide area networks.

2.7 Idea - Property Rights of Personal Information

2.7.1 Description

A “property rights” approach to the protection of personal data would be established, explicitly assigning clear and enforceable rights to data subjects and data holders.

Such an approach may be beneficial, because a substantial fraction of cyber-security costs do not derive from malicious intent but simply carelessness and misunderstandings between data subjects and data holder. Furthermore, it would decrease firms’ uncertainties regarding their actual ownership of, and obligation towards, the personal information of their consumers. Given the considerable heterogeneity in the valuation of the worth of individual data, significant potential gains from trade could also be achieved.

2.7.2 Inertia

Why have we not done this before?

In a sense, implicit markets for personal information already exists – as consumers we routinely trade-off personal data for tangible and intangible bargains, often as a secondary aspect of a

different primary transaction. However, explicit property rights on personal data, albeit often discussed in the legal literature on privacy, have not yet appeared.

First, rather than a strong regulatory framework, the approach in the United States has focused on self-regulatory efforts and market-based solutions. As a consequence, under the current regulatory regime, the concept of “ownership” of personal data is not well defined – the very concept of personal information as “property” may sound novel to most people.

Second, enforcing data ownership even in the presence of legislative protection is difficult (consider the challenges associated with controlling the secondary use of data).

Third, transaction costs for contracts involving personal data are high, and individual decision making in this area is likely to be affected by cognitive and behavioral biases: consumers often lack the understanding as to the ramifications of ceding control over their personal data is, and can’t assess the long-range implications of such decisions.

Fourth, progresses in data mining have increased the economic value of personal information for data holders, trumping the economic interests of data subjects.

Fifth, there exists a legitimate doubt that a property right approach may disrupt flows of personal data that are beneficial not just to data holders, but to the data subjects themselves.

What would derail this change?

Even in the presence of a regulatory framework, concerns that contracting costs will be too high and the difficulty of enforcing the rights may derail this change. Considering these challenges, law scholars such as Pamela Samuelson have suggested alternative approaches based on models akin to “trade secrets” for personal information, rather than formal ownership of that data.

Personal information has enormous value to organizations for business purposes. This may cause organizations to resist laws changing how personal information is collected and monetized.

2.7.3 Progress

What technologies are emerging that makes this change look doable now?

Possibly, progresses in the areas of DRM and Access Control technologies may help making property rights on personal data enforceable.

What environmental (business, political) changes are pointing in this direction?

Judging from surveys, interviews, and reports, consumers’ dissatisfaction with the current status of protection of their personal data may spur support for such an initiative.

Furthermore, progresses in research on digital provenance, and lessons learnt from the management of IP rights in other areas, may be applied to this area.

2.7.4 Action Plan

What are the reasonable paths towards bringing about that change?

First, an analysis of why the market has not delivered this solution (notwithstanding several similar proposals in the past two decades), and why, instead, in the current equilibrium, it is firms that take complete ownership of consumer data.

Second, develop an understanding how existing DRM and digital goods licensing technologies may be leveraged to allow for such granting and division of rights.

Third, and more importantly: this is a change that has been often discussed, but that the marketplace alone has not delivered; a true property rights approach would not be possible without actual government regulatory intervention.

What would accelerate this change?

Leverage the body of existing work on defining IP rights for personal data.

Do an economic analysis that showed that clear definition of rights may be beneficial also to data holders (such as firms), since they would decrease their uncertainty in terms of the appropriate policies to apply to personal data.

Develop a short-term (e.g., 60-90 day) proposal of what a feasible and efficient division and assignment of rights to data subjects and holder would look like.

What are the missing technical pieces?

Among others, proof of the ability to enforce rights on the secondary use of personal information through technology is missing.

2.8 Idea – Infrastructure Diversity

Currently, most large organizations are trying to transform their IT infrastructure towards a standard set of components. The goals of standardization are to drive down the cost of managing this infrastructure, the cost to train users to use the technology, and to simplify their supply chain. Moreover, procurement managers are generally wary of purchasing diverse components, especially when the market leader is perceived as a safe investment.

However, this homogeneity is dangerous from a security perspective. It lowers the attackers' costs, increasing the probability that his attack could compromise a large number of machines: an attack on any one component, which is pervasive throughout the organization, could potentially be leveraged into a catastrophic attack against the entire infrastructure.

2.8.1 Description

What does the change look like?

If firms were incentivized to have a diversity of infrastructure components instead of a monolithic infrastructure, it would be much more difficult for any one attack to bring down the entire infrastructure. Indeed, a heterogeneous infrastructure should in principle be more resilient than a homogenous one.

2.8.2 Inertia

Why have we not done this before?

To a certain extent this idea has been done before. It is common practice for large organizations to diversify their supply chain so that if a particular supplier fails, they have alternative sources. In addition, government procurement policies already dictate that a diversity of vendors must be able to participate in government contracts.

In fact, the notion of diversity as a way of improving enterprise security has been debated for a number of years. Reduced diversity at the product level can bring not only reduced costs but also improved security resulting from ease of management and configuration. While interoperability

may simplify the task of introducing diverse components, it does not yet simplify the task of managing such components – and it seems likely that standardizing management would reduce the diversity that brings putative benefits. Diversity resulting from such techniques as address space randomization, in contrast, may improve resistance to attack without requiring error prone customization of infrastructure management.

However, other market forces push firms in the direction of a monoculture. For instance, first mover advantages and economies of scale make it difficult to create a market with a significant diversity in any one technology area. Economies of scale and network effects internal to the firm also explain why firms may resist diversifying their IT infrastructure.

Finally, while governments have the ability to mandate diversity in their own ecosystems, it may be difficult to impose such a constraint on the private sector.

What would derail this change?

The cost of implementing infrastructure diversity may outweigh the expected loss of security incidents associated with standardized enterprise architectures. Emerging technologies, such as cloud computing, have the potential to make security problems associated with standardized infrastructures less of an issue for organizations in the future.

2.8.3 Progress

What technologies are emerging that makes this change look doable now?

As systems become more interoperable, heterogeneity becomes less of an issue. Moving forward, the continued standardization of infrastructure components may make this idea much more feasible. Indeed, we already have a diversity of hardware components from multiple manufacturers that can run identical software. Perhaps this idea is just a natural evolution up the technology stack.

What environmental (business, political) changes are pointing in this direction?

If anything, as pointed out above, there is significant momentum in the opposite direction -- toward ruthless standardization in enterprise architectures.

2.8.4 Action Plan

What are the reasonable paths towards bringing about that change?

There are both policy and technical approaches to bringing about this change.

One could imagine limiting the scope of this change to government systems. A procurement policy could be instantiated that dictates that a certain percentage of components of a particular type must come from multiple vendors. For example, instead of standardizing on one type of web server, the procurement policy would dictate that a certain percentage of web servers must come from alternative sources.

But diversity can be achieved by other means than diversifying vendors. For example, we can customize individual instances of infrastructure components to eliminate certain classes of attacks. Techniques such as memory address randomization and basic block shuffling have already been employed to realize such a vision.

Finally, this approach could be done incrementally. Instead of trying to enforce heterogeneity within every organization, we could begin by having incentives to have heterogeneity between

organizations. Therefore, a catastrophic attack against a single infrastructure component would not be likely to bring down an entire industry.

What would accelerate this change?

We still do not fully understand the cost/benefit tradeoff of diversity as an approach to security. We do not have sufficient data to say whether or not the additional costs of purchasing and maintaining a diverse infrastructure is worth the marginal risk reduction one would achieve by implementing such a strategy. If we could extend the economic analysis that is emerging in this area, and the analysis showed a clear advantage to diversification, then clearly adoption of this approach would be significant.

We need to have economic mechanisms to facilitate market entry for competition.

What are the missing technical pieces?

This is largely a non-technical issue. But there are a few areas that would help. We could use better tools to manage large diverse environments. And, additional research could be performed on ideas like memory address randomization which would allow diversity on a component by component basis.

2.9 Idea - Multiple Networks

The success of the Internet is largely due to its openness. Anyone can get on and participate, from the individual, to large, complex organizations. However, the openness of the Internet also creates security problems. Attackers can use the Internet just as easily as anyone else. Legitimate activity is commingled with illegitimate activity.

(Note: This proposal inspired aspects of the "A new virtualisable network architecture" idea listed in the Additional Ideas section of this document.)

2.9.1 Description

The game changing idea is to enable communities of interest with dedicated, isolated and virtual networks that are secure from end to end. For example, one could imagine a network dedicated to financial transactions and another dedicated to online gaming. These networks could be implemented as secure overlay networks on top of the existing internet.

We could define policies associated with each network about the types of traffic allowed, who can participate in those networks, the level of anonymity permitted to participate, what actions are permitted, and what will be monitored and logged.

The challenge is that the end point (i.e. user machines), would have to connect to multiple of these networks to be functional. There must be strong guarantees that those endpoints do not act as a conduit to allow information to flow between these dedicated networks.

From an economic perspective, the goal is to decrease the revenues of the attacker, since the networks that are likely to be easier to access are also those less likely to carry valuable information, such as financial and personally identifiable information. Therefore, they are less valuable to criminals.

2.9.2 Inertia

Why have we not done this before?

To some extent this has been done. Today we have multiple networks: the Internet, the phone network, cellular networks, etc. Furthermore, businesses have been using VPNs to extend their corporate networks for a long time. Research efforts are underway to implement overlay networks such as those in PlanetLab and the GENI initiative. From an economic perspective, these may all be examples of the market providing a mechanism to enable such networks where they are needed.

However, some of the recent initiatives (such as GENI) have not really focused on cybersecurity. Furthermore, Summit participants debated whether we already have acceptable solutions for securing the endpoints. According to some, we currently do not have adequate commercial-grade technology to provide strong isolation between multiple compartments on a single endpoint, although much research is happening that could provide this capability in the future. For others, the problem does not rely on the technology per se – which would be available – but the fact that it has not been widely adopted. In addition, users' acceptance of delays between virtual networks remains a challenge.

Finally, as discussed above, the openness of the Internet has been one of the great success stories of the last century and perhaps the primary reason why the internet has been so successful. Trying to change this paradigm may run counter to what is fueling its success.

What would derail this change?

The market forces behind the open internet are so strong, that this approach may not be able to compete with the way the internet works today. In fact, in most people's minds, it is likely that these types of trust solutions typically would have lower priority than having more functionality and flexibility.

Each entity, whether they are an individual or a corporation, may want to have control of how it interacts with other users on the internet. For example, VPNs seem to be a fine solution to this problem for corporations today.

As discussed above, one of the primary technical challenges is creating secure endpoints. The increasing diversity of platforms in mobile access devices also makes this a moving target.

2.9.3 Progress

What technologies are emerging that makes this change look doable now?

There are several technologies that make this change look doable. In terms of securing endpoints, virtualization is becoming increasingly popular. One could imagine having dedicated virtual machines on each endpoint for each of the networks that machine participates in. This can be done in a highly trusted way with hardware authentication approaches such as those being championed by the Trusted Computing Group (one idea discussed was the use of cheap, secure devices for dedicated use, such as online banking only).

In terms of keeping the individual networks secure, we already have technologies such as VPNs and other forms of link encryption, and network isolation technologies such as VLANs. Other standard techniques such as white lists could be helpful in this context.

What environmental (business, political) changes are pointing in this direction?

As the perimeter of organizations continues to erode, these organizations need to have some mechanism to create strong virtual networks that operate over assets they do not own, so a

solution along these lines will be necessary. Identity theft and other attacks against our financial systems are raising the incentives to create highly secure separate networks over which consumer financial transactions can take place that are strongly isolated from other, potentially risky user activities.

2.9.4 Action Plan

What are the reasonable paths towards bringing about that change?

The challenge lays in how to boot strap the process. The change could start at a small scale, within a closed environment -- for example, within a government or university network -- or by defining an isolated network specifically dedicated to financial transactions.

The DOD has unparalleled experience in this area – some of their practices may be shared and adopted.

We would need to understand the taxonomy of possible networks, how to express policies for the networks and gain a better understanding of the financial incentives and disincentives to making this work.

Finally, there will be situations in which it will be necessary to move information between these networks. This raises the issue of how to enable communications between networks with different security levels.

What would accelerate this change?

If an entity were formed that would be responsible for defining, managing, and regulating these networks, this change could be accelerated. Although one could imagine how this might be feasible to do in a completely distributed way, this change raises coordination problems that the government could help address. Obtaining ISP's support (and presenting it as a Quality of Service product) might also help to accelerate this change.

What are the missing technical pieces?

We currently do not have adequate commercial grade technology to provide strong isolation between multiple compartments on a single endpoint.

2.9.5 Jump-Start Plan

Recommend:

- The next round of GENI funding focus on computer security
- Other networking programs include security components at a minimal level in terms of calls and funding
- DARPA implement non-military open calls, working with NSF to allow larger-scale research in the academic community using peer reviews. DARPA-level funding with NSF-level review and outreach can result in a set of large centers with representatives from every state for work on logically distinct networks/cascade-free authentication.

2.10 Idea - 911 Cyber

2.10.1 Description

While large organizations have the ability to report cyber security incidents, receive assistance and advice, consumers and small to medium businesses have limited ability to get access to these resources and, when possible, redress.

The idea is to have a centralized agency that could collect and respond to large scale incidents, and potentially identify problems that are distributed across a large number of stakeholders. Data collected would be anonymized such that they would contain no identifiable information. Reports would be submitted to an independent central organization that is not a vendor or service provider. The organization would have personnel and resources in place to respond in a timely fashion, including interfacing with the appropriate vendors and law enforcement authorities where appropriate.

2.10.2 Inertia

Why have we not done this before?

To some extent, individual vendors and companies often already offer assistance for their respective products and services. Furthermore, the government – through the FTC – offers a hot-line for individuals who believe have been victim of identity theft. Websites such as 911.com have a model very similar to the one discussed here. However, considering the scale and breadth of the initiative we refer to here, personnel with the appropriate level of expertise are hard to find, and the risks of cost-duplication (vis a vis similar, distributed initiatives in the private sector) are high. The actual identification of a problem as a cyber security problem is often difficult to do, and it is often obscured by other system, software, or user problems. This increases the complexity of the job. Who would benefit from this service is unclear: is it just the consumer, or the community, company, or nation?

What would derail this change?

The average level of troubleshooting expertise of typical users of this service will be low, resulting in numerous non-security related calls, which will likely overload the service providers. The service will likely be costly to provide, which raises the question of who will pay for it. Sufficient personnel with the appropriate level of experience will be hard to find for this effort.

2.10.3 Progress

What technologies are emerging that makes this change look doable now?

Trusted computing hardware modules may aid the development of this idea by automatically reporting incidents that otherwise would be overlooked by end users.

What environmental (business, political) changes are pointing in this direction?

A sharp increase in consumer security consciousness and their frustration in attempting to improve have further accented the complexity securing our systems. The growth of botnets is also driving attempts of improvements to all classes of systems (personal, academic, enterprise, and government).

2.10.4 Action Plan

What are the reasonable paths towards bringing about that change?

Leverage the existing investigative bodies, such as the NTSB, and other reporting bodies, such as the ITAC for financial services, to design the new service. ISP's will be key to the realization of this idea, so their early engagement will be vital.

What would accelerate this change?

National and State centers to support this effort. Programs at universities and community colleges focused on producing graduates with the necessary skills will be essential.

What are the missing technical pieces?

Government-provided open-source software, which is very important for privacy concerns, would aid in the reporting process and perhaps result in automating it.

2.11 Idea - Swimming with the Sharks

This was an interesting discussion that many participants felt was absolutely fundamental to the problem, but the group as a whole struggled with how to turn into a game changing idea. Nevertheless, the co-chairs wanted to include references to that discussion in this report for completeness.

2.11.1 Description

Systems are so resilient that they can tolerate security vulnerabilities and attacks without impacting system operation. In essence, we accept the fact that security vulnerabilities are inevitable and we figure out other ways to deal with the problem. This represents a paradigm shift away from traditional thinking about security mechanisms, and instead focuses on alternative approaches such as resiliency.

2.11.2 Inertia

Why have we not done this before?

Numerous projects from a variety of agencies and research institutions have been pursuing ideas like this for decades. In a sense, we have been moving in this direction – even though unwittingly. Utilization of content delivery networks, such as Akamai, has provided a way to mitigate DDOS attacks as well as deal with problems like temporary congestion episodes. While there has been progress, the complexity of multiple software and hardware components and the challenges of system usability and human behavior have, along with other challenges, made limited progress.

What would derail this change?

Nothing was identified.

2.11.3 Progress

What technologies are emerging that makes this change look doable now?

Evolving technologies, such as virtualization and cloud computing, are making it easier to implement some of these approaches.

What environmental (business, political) changes are pointing in this direction?

Previous projects from a variety of agencies and research institutions in this area.

2.11.4 Action Plan

What are the reasonable paths towards bringing about that change?

While resiliency has been extensively studied for military applications, much of that research has not made its way into the commercial sector. Further explorations on trustworthy platforms, security usability, security metrics, and testing for reliability would help, as well as techniques for shifting risk around such as cyber-insurance.

What would accelerate this change?

One of the biggest assists would come from making this policy explicit, which would make people think in different ways. For example: rather than focusing on clean modular design, students may be taught to learn from malware creators, on how to obfuscate code, and so forth.

What are the missing technical pieces?

Nothing was identified.

2.12 Idea – Minimize and Target Authentication

Organizations reuse the same authenticating/identifying information across different domains. The end result is a small set of authenticators that is available to attackers at low costs. Pieces of data which are used (and abused) as authenticators (such as SSNs), once obtained, can be reused by criminals and never be “cleaned” by victims.

2.12.1 Description

Changing authentication requirements so that organizations do not store high-value authenticating information without high levels of accountability for how that data is protected.

2.12.2 Inertia

Why have we not done this before?

While cryptography offers many powerful tools to protect authentication processes, those tools are often not properly understood or used by most decision makers. Furthermore, a first mover disadvantage exists: the company that moves first in adopting or enforcing more sophisticated means of authentication may lose consumers who chose the instant gratification of easy authentication over the security benefits of more secure tools.

What would derail this change?

One company securing its customers information still bears the costs of other companies handling information without adequate protection (cascading failures in identity verification).

Overrated claims of costs of deployment and the choice of inappropriate technologies may also derail this change.

2.12.3 Progress

What technologies are emerging that makes this change look doable now?

Decreased processing costs make small cryptographic devices affordable. Decreased communication costs make multiple rounds of communication feasible.

What environmental (business, political) changes are pointing in this direction?

Many businesses realize that current practices provide organized cyber-crime with billions of years to invest in their criminal syndicates.

2.12.4 Action Plan

Sectored standards for endpoint authentication fit well with the concept of logically separated endpoints. These are strong complements.

2.13 Other Ideas

The following ideas were also raised initially during the summit, but were not exhaustively discussed by participants. They are included here for reference and possible further development.

- **RE-EXAMINE OLD IDEAS IN LIGHT OF NEW STRUCTURES.** Significant security ideas were generated during the development of basic computing technology during the 1960s and 70s. Some of these ideas were discarded because, at the time, the available computing power, storage, and communication requirements were not available. Some great ideas may have been lost. Unlike information produced in the last 15 years, these ideas are not readily available on the Internet, making access to those ideas more difficult. The suggestion is to systematically go back and re-examine these ideas to see if they are now feasible.
- **PRE-EMPTIVE DISCLOSURES.** Before deploying security system, vendors should disclose an analysis of the costs to various stakeholders (similar to an Environmental Impact Analysis), including forecasted users' efforts.
- **RISK ADJUSTED RATE OF RETURN.** Collecting the right security data is one step in the process. But we need models that can use that data. Specifically, the research community is still trying to develop an acceptable way of calculating risk adjusted ROI – such a metrics would help research and decision making in the area of information security.
- **ATTRIBUTION.** Improve attribution, e.g., through a directory service, in order to restrain malefactors; have Internet “driving” license.
- **STANDARDS.** Focus research effort on developing standards for what needs to be monitored: what kind of information should we monitor, how do we develop the right metrics.
- **CYBER BILL OF RIGHTS.** We need a Consumers Cyber Bill of Rights - addressing ISPs, consumer software, etc. Think NRC consumer rights being enforceable by the FTC.
- **CYBER VIGILANTES.** Victims should have the legal right to aggressively repel and/or counterattack cyber attacks. This would require a legal structure that encourages self defense if attribution can be determined.
- **DISRUPT THE ATTACKERS.** Develop various offensive tactics to raise the cost to attackers, including decoys, flooding miscreant markets, revealing their methods and tactics, and so forth.

- **CAP AND TRADE.** A cap and trade system for cyber security “pollution” (poor security = pollution). Imagine pollution credits, budget risk by industry, and an “insecurity load model” to correctly capture value GDP “health” of the country.
- **REDUCE BARRIERS TO ENTRY.** Reduce barriers to entry for security solutions by speeding up certification of security products.
- **IMMUNIZATION AS IMMIGRATION.** Endpoints are critical with respect to infection or attack, as are servers. Consider network admission control on a wide scale, i.e., an immunization check: certify “immunization” before granting access similar to immunization (like border controls). Stop the bots – limit access to specific services.

3 Digital Provenance

New Game: Basing trust decisions on verified assertions

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

What is the new game?

In today's game we have to expend considerable energy to discover whether to trust digital objects for any intended purpose. We are in the situation of a shopper who walks into the meat department of his grocery store and finds a case full of wrapped but unlabeled meat. While he might be able to determine if it is safe to eat through laborious chemical and microbiological analysis, some things he will never know: is it kosher; did the animals range free; what were they fed? Fortunately, USDA regulations ensure that each consumer does not have to invest in sophisticated laboratory equipment to analyze his beef, but in the digital world, this is often the very situation he finds himself in. Today, with no guarantees as to the source and integrity of digital content we have to check everything to be sure it is not harmful; with reliable digital provenance we can concentrate our resources instead on how we wish to handle the varieties of authorized content we receive.

3.1 Idea - Stable Network Identity

3.1.1 Description

Remove the semantic overloading of IP addresses by disambiguating network topology location function from the host identity function.

3.1.2 Inertia

- Global IP software changes
- Institutional resistance
- Complex roll out strategy
- Non-Reversibility (reverse lookups very difficult)
- No community-wide incentive

3.1.3 Progress

- Proven technology with limited deployment
- Host Identity Protocol (HIP): a multi-year working group within the Internet Engineering Task Force (IETF) which
 - Enables mobility and multi-homing
 - Supports convergence of mobile and multi-homed devices
 - Has been used to secure previously non-securable devices (machine controllers, [e.g., Supervisory Control and Data Acquisition – SCADA])
- Has an open-domain code base

3.1.4 Action Plan

- Migration of standard into communication stack products
- International Regulatory Awareness Programs
- Pick one of the jumpstart activities to advance with funding

3.1.5 Jump-Start Plan

- Create use cases of how to use HIP to secure:
 - SCADA
 - Utility grid
 - Hive and composite communications
 - Healthcare remote telemetry
 - Location-based services
 - Increasing trustworthy micro-payments
- Finish standardizing HIP within the IETF / Have the approach verified by a government national lab in this domain

3.2 Idea – Data Provenance Security

3.2.1 Description

Managing and securing data provenance (DP) information. Authorizing and controlling access of principals to DP. (Data minimization, privacy, least privilege, confidentiality, integrity, and authenticity.) This is predicated on “DP definition and management” (see below).

3.2.2 Inertia

- Scalability
- Tendency of organizations to default to high levels of information (doesn’t sufficiently manage risk)
- International laws and policies differ
- No existing technology

3.2.3 Progress

- Availability of new policy- and attribute-based cryptographic techniques
- Recognized need for DP in a variety of circumstances

3.2.4 Action Plan

- Standardize technology in standards bodies like IRTF and W3C
- Strategic use cases in areas like the intelligence and healthcare communities

3.2.5 Jump-Start Plan

Design for secure provenance of immutable objects (e.g., issued patents)

- Extend to “append only” objects (e.g., log files, audit trails)
- Create a general model of secure provenance

3.3 Idea - Data Provenance Definition and Management

3.3.1 Description

Attaching context to data to track chain of custody, transformation (modification), and provenance of messages and attachments (for software, data at rest, or packets). Establish standard labeling system for quality (analogous to food labels).

3.3.2 Inertia

- Scalability
- No standards
- Complexity of the ontological model
- Privacy concerns

3.3.3 Progress

- Industry experience designing markup languages
- Existing means of cryptographically binding data and it's provenance
- Advancements in meta-data cataloging and search capabilities
- Existence of pervasive time and location services (e.g., GPS)

3.3.4 Action Plan

- Work with browser developers to incorporate into the browser and present to users (e.g., Chrome)
- Work with OS vendors to incorporate as file system meta-data and with GUI/explorer hooks for presentation to users
- Revise/extend existing government standards and software (government meta-data working group standards) to meet DP requirements
- Build upon, coordinate, and integrate with trusted systems work (including hardware trust group from this summit)
- Develop HW acceleration for attaching DP context data at network (or lower) layers
- Develop policy/legal framework for resolving DP disagreements or conflicts

3.3.5 Jump-Start Plan

Create a standards group (e.g., Defense Research and Development Canada (DRDC) efforts)

3.4 Idea - Reputation Engine

3.4.1 Description

Credibility quantification of principals and entities (by tracking popularity, responses, scoring, and other kinds of trust data) to establish reliability. Leverages cognitive sciences (perceptions) that build in mechanisms for both crisp logic and fuzzy logic systems. Enables claims-based (name, reputation, etc.) ID.

3.4.2 Inertia

- Scalability
- No cohesion
- No standard
- Identities lack anchors and are easily manufactured
- Reputations may be spoofed and misused

3.4.3 Progress

- Technology exists today
- Acceptance by consumers and stakeholders
- Proven value to consumers and suppliers
- Value in its ability to propagate, amplify, and degrade

3.4.4 Action Plan

- Build one or two real applications (proof of concept)
- Start down RFC path (proposal, review, standards, etc.)
- Build a community
- Build common exchange/interop format (e.g. genealogy as good example of similar model and format)
- Build a reputation common data model - include entities, attributes, and relationships
- Minimize spoofability

3.4.5 Jump-Start Plan

- Pick three or four commercially used reputation engines for analysis (e.g., eBay, site advisor, credit rating services)
- Find commonality and build rules.
- Pick use cases for test/verification (e.g., phishing and anti-phishing)

3.5 Idea - Trustworthy Systems

3.5.1 Description

Expanding trustworthy systems foundation to create trustworthiness (integrity) in how software treats DP.

3.5.2 Inertia

3.5.3 Progress

- Increased need recognition
- The SCAP, FDCC, and Software Assurance efforts have attracted new adopters and inspired new areas of investigation and investment

3.5.4 Action Plan

To be determined; depends on the outcome in the short term

3.5.5 Jump-Start Plan

DoJ pilot use of Digital Evidence attestation meta-data about chain of control providence

- Repositories of reusable code - DHS S&T Open Source project as a starting point
- ESAPI (OWASP) for JAVA - libraries of hard to do right security relevant functions
- Define "food label" attributes for trustworthiness of software and hardware

3.6 Idea - Government Role

3.6.1 Description

Government to serve as authoritative certification authority of digital identity.

3.6.2 Inertia

- Potential single point of failure
- Governments are not the originators of identity in US
- Privacy and civil liberties fears

3.6.3 Progress

- Consumer receptivity due to concerns about identity theft, phishing, etc.
- Need for health care information exchange
- Shifting economies, scale, and scope of cyber-attacks

3.6.4 Action Plan

- Address liability for reliant parties
- Full range of use cases
- Policy framework (US domestic, global/international)
- Forums to address issues
- Collaboration with private sector around CIP (e.g., SCADA and industrial control systems)
- Consumer outreach
- R&D on implementation approaches

3.6.5 Jump-Start Plan

- Identify early adopter use cases in financial services, energy/industrial control systems, health care information exchanges (regional cooperatives, PHR/EHR), ICT/internet
- Plan/establish pilots

3.7 Idea - Trusted Path (TP)

3.7.1 Description

A secure interface between user and trustworthy system entities that will permit provenance of actions at any layer of the protocol hierarchy.

3.7.2 Inertia

- Expensive
- Not supported in current architecture
- User interface

3.7.3 Progress

- SAK feasible – just need device driver to do this
- May be hardware mechanisms to support now. Host ID embedded in hardware with cryptographic protection
- The great need for interoperability is a driving force for remote TP

3.7.4 Action Plan

- Expand to other domains (financial)
- Use TP to mitigate spam and phishing by tying IP disambiguation via attribution
- Create anonymous access to high integrity information via public libraries

3.7.5 Jump-Start Plan

- Small field demos to show TP
- Investigate TP in situ/on platform for multi-core processors. Core to core, KUM to core

3.8 Idea - Global Identity-Based Cryptography

3.8.1 Description

Global encryption based on identity that is robust.

3.8.2 Inertia

- No proven technology
- Reliability
- Management
- No revocation
- Not post-quantum secure
- No global system available
- Privacy issues
- No compromise recovery
- Online servers

3.8.3 Progress

Technologies now exist to express scalable symmetric key authenticated encryption systems where no single trusted third party knows the final key.

3.8.4 Action Plan

- Development teams to integrate proposals into open source applications
- Identify and bring together identity stakeholders into a conference to refine requirements

- Independent evaluation of next generation proposed technologies

3.8.5 Jump-Start Plan

- Draft a high-level requirements document
- Create use cases
- Survey candidate technologies
- Independent evaluation of next generation proposed technologies

4 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.

What is the new game?

Today, weeks and months may elapse before successful network penetrations are detected through laborious forensic analysis. Despite their potential to function with intelligence, today's typical network components have very limited understanding of what passes through them, coupled with a correspondingly short memory. In medical terms, because we are not instrumenting for early detection of pathogens and their effects, our most common diagnoses are through autopsies of enterprises which have succumbed to attack. In the new game, network and host components have heightened ability to observe and record what is happening to and around them. With this new awareness of their health and safety they enjoy a range of options: they may take preventative measures, rejecting requests which do not fit the profile of what is good, a priori, for the network; they can build immunological responses to the malicious agents which they sense in real time; they may refine the evidence they capture for the pathologist, as a diagnosis of last resort, or to support the development of new prevention methods. In the new game, we should be able to monitor and control such dynamical cyber environments.

Introduction

We propose to change the game for protecting cyber-systems by looking to nature for inspiration. Examples in nature are the immune system, beneficial parasites, and social networks such as public health networks and social insects. The immune system protects the body remarkably well from panoply of threats that are continuously evolving in a dynamic and ever-changing environment. Natural systems are far more complex than our cyber-systems but they are extremely robust, resilient, and effective. Clearly, an investigation of these natural systems, such as the immune system, can be beneficial to changing the game for cyber-security. In this working group we explored and developed the following four potential 'Game Changing' idea proposals:

- Distributed Defense
- Centers for Cyber Disease Control (CCDC) and Prevention
- Using Attack Vectors
- Missing-Self Paradigm

These four potential game-changing ideas are described below.

4.1 Idea - Distributed Defense

4.1.1 Description

- Distributed defenses based on the resilience of natural systems
 - Multi-scale (computer, local network, global)

- Agility – new sensors, responses, etc., example: If there is an attack on the network, within one minute, there should be 99% immunity to the attack
- Re-engineer functions to be robust to asset damage
- Use diversity of limit fraction of assets affected by any given attack
- Sensing
 - Memory of health state – anomaly detection
 - Memory of characteristics of past attacks
 - Community reputation and trust measures for sensor data
- Signaling
 - Collaborative signaling at multiple levels, federations of communities
 - Communications standards
 - Collaborative/federated communication
- Response
 - Automatic response – appropriate to false positive rates
 - Human-in-the-loop response for high-consequence or early deployment
 - Symbiotic relationships – responses that influence adversary or cause-desired side-effects
 - Responses that anticipate and mitigate likely next steps

4.1.2 Inertia

- Data rates are high
- Usually driven by knee-jerk reactions instead of designing a systemic defense
- Low willingness to share raw data
- Specific targeting by the adversary can remove the benefit of communication
- Shared data may not represent invariants of attack
- Challenge to share more quickly than the adversary moves
- Sharing exposes what we know to the adversary
- Response systems can be gamed to deny service
- Reliance on network availability
- Rewriting applications or application protocols

4.1.3 Progress

- Critical systems more distributed now – drive distributed sensing
- Attackers more distributed now
- Sufficient additional CPUs required to do distributed processing
- Realization that peers have important real-time threat data to share
- Leverage new cloud computing architectures

4.1.4 Action Plan

- Develop CONOPS and requirements
- Gap analysis

- Use existing sensors opportunistically
- Identify new responses and sensors required to trigger them
- Develop new sensors and responses
- Verification and Validation (V&V)
 - Test convergence (control theory)
 - Quantify performance
 - Measure performance under specific realistic attacks
- Distributed Robustness
 - Re-engineered functions must be robust to assess damage
 - Use diversity to limit assets affected by any given attack
 - Bound outages and minimize impact on functions

4.1.5 Jump-Start Plan

- Pick high complexity, high savvy sites (e.g. research labs) to develop and deploy operationally
- Each of: Defense, Electrical Supervisory Control And Data Acquisition (SCADA), Health
- Fund multiple-threads of development and implementation simultaneously (~\$50M/yr)
- Build a self-sustaining community, similar to Internet Engineering Task Force (IETF) to be stewards for standards communication mechanisms, and formats, etc.
- Use management mechanisms to drive adoption
- Encourage vendors to add support to COTS
- Get industry to provide private clouds

4.2 Idea - Centers for Cyber Disease Control (CCDC) and Prevention

Provide similar public health system services for our national computer infrastructure. While establishing the center, rules and regulations need to be formulated to define the jurisdiction so that it does not to violate any constitutional rights.

4.2.1 Description

- Public health infrastructure - cyber equivalent to CCDC
- Indication of “I’m Sick”
- Overcome barriers to sharing data
- High fidelity data required to gain full understanding of illness
- Conduct data collection similar to World Health Organization and public health departments
- Collect and distribute health information to support active response
- Provide cost/benefit of interventions
- Models should comprehend key factors

- Cyber geographical statistics concerning topology, applications communities, shared software, end users analogous to doctor office, city, state, or a CDC
- Profit motive that leverages commercial opportunities and business case, e.g., service providers, etc.
- Global scale reports that provide the state of the Internet assessments, e.g., e-crime, fraud, data breaches, Comprehensive National Cybersecurity Initiative (CNCI), 60-day report, and threat intelligence report, etc.

4.2.1.1 *What is the Role of a Public Health System (PHS)?*

- Assessment of a community's problems, needs and resources
- Health needs assessment
- Data and surveillance
- Leadership in organizing effective public and private sector strategies to address community health problems
- Assurance that direct services necessary for meeting local health goals are available to all community residents such as screening, education, prevention, outreach

4.2.1.2 *What does a PHS do?*

- Monitor health status to identify community health problems
- Diagnose and investigate health problems and health hazards in the community
- Inform, educate and empower people about health issues
- Mobilize community partnerships to identify and solve health problems
- Develop policies and plans that support individual and community health efforts
- Enforce laws and regulations that protect health and ensure safety
- Link people to needed personal health services and assure the provision of health care when otherwise unavailable
- Assure a competent public health and personal health care workforce
- Evaluate effectiveness, accessibility and quality of personal and population-based health services
- Research for new insights and innovative solutions to health problems

4.2.1.3 *The Core Claim*

- "Surveillance" – The gathering and analysis of data on a national scale is a key enabler to providing public health services
- These functions can, and should be, automated for Cyberspace
- Multi-scale collection and reduction of cyber health data
- Represent the "ground truth" about cyber operations on the scale of the national infrastructure

4.2.2 *Inertia*

- Anti-virus companies are similar to drug companies
- Reactive "knee jerk" nature of business
- Absence of central driving force

- Data ownership and intellectual property issues
- Lack of Federal Government buy-in
- Fed has not created incentive system
- Data is disaggregated
- Lack of liability model
- Actuarial data needed for insurance
- Automated cyber-attack stress testing
- Public education

4.2.3 Progress

Catalysts for sharing data

- Incentive and legal precedence for sharing data (potentially intellectual property)
- Consortiums that encourage sharing, e.g., best practices, threats and attack signatures
- Data characteristic specifications necessary to jump start
- Utilizing the power of human intelligence by increasing public awareness, e.g., epidemic warnings and best practices

Why is this the right time?

- Increased public awareness
- Magnitude of problem is heightened
- Represents a business opportunity
- Technology has matured to enable collection/filtration/dissemination of information
- Government can provide stamp of good practice
- Assurance for both big and small business
- Government has expressed willingness to address cyber security issues and stimulate action

4.2.4 Action Plan

- Define taxonomy and metrics categories
- Data collection
- Current state and sensitivity analytics
- Predictive mathematical models
- Prospective studies
- Temporal data on how a “healthy network” functions
- Collect specific cohort groups of targeted populations
- Visualization of network behavior and structure
- Rapid response monitoring

Overall Recommendation Phase

- Criteria for being healthy
- Decision support, e.g., quarantine, barrier establishment, vaccination

- Synthetic cyber vaccine distribution
- Innovation center for catalyzing other health inspired innovations
- Promote continued cross-over among biological institutes and IT discipline
- Optimal sensor/actuator placement
- How much information do you need to make an optimal decision
- Control law algorithms versus machine learning on empirical data
- Given the right and/or enough data, can we machine learn the correct response

4.2.5 Jump-Start Plan

- Organize and survey
- Identify and address the required initial data
- Consider privacy issues
- Enumerate existing data sets
- Consider current taxonomies
- Detail frame and scope
- Identify other models (e.g., CDC, World Health Organization)
- Identify potential partnerships
- Identify initial stakeholders and refine data
- Establish possible business models

Establish a community of interest to further develop the concept and evolving steps to produce an RFI and establish initial pilot with seed funding.

4.3 Idea - Using Attack Vectors

We propose a set of offense techniques for cyber defense. This approach is roughly analogous to having some form of cyber pharmaceutical industry to deal with specific cyber pathogens.

Background/Motivation:

- Hordes of vulnerable computers on the internet
- Not secure because of apathy, ignorance, just don't care, etc.
- Huge problem because of botnets, etc.
- Attackers have vectors into those computers
- Same vectors used to do good, e.g., patch
- Do it without the user's consent for the greater good, e.g., Oral Polio Vaccine (OPV) - OPV transmits between individuals to provide 'passive' immunity. Passive immunity of OPV is a major reason behind the World Health Organization's choice of OPV for the world-wide Polio eradication campaign.

4.3.1 Description

Three Proposed Approaches:

- Good Worms (aka gworms)

- Piggybacking (aka ride the worm)
- Drive-By Downloads

4.3.1.1 Good Worms (aka gworms) an old idea

- Idea is to create gworms (good or benign worms) that spread love (patches, etc.)
- Been there - done that:
 - Suggested many, many times
 - Real gworms, e.g., Welchia worm (2003): detects and terminates Blaster worm, patches system and reboots

4.3.1.2 Gworm Problems

- Spreading gworms considered harmful; results in network traffic overload
- Need to move faster than a worm to catch it
- Unintended consequences from bugs
- Could harm systems that are not currently threatened and/or attacked
- Releases gworm code to the world
 - Exploit code available to blackhats
 - Transmission code available to blackhats
- Ethical and legal issues

4.3.1.3 Piggyback: Ride the Worm

- Use honey pots to catch worms
- Replace worm payload with a rider
- Rider prevents host damage
- Rider still allows network spread
- Rider goes where worm goes, possibly at the same rate the worm spreads

4.3.1.3.1 Piggyback Benefits over gworms:

- Dormant until activated, i.e. only do harm when harm is happening
- Easier to match spread rate to worm
- Rider contains no exploit or transmission code
- “More” ethical or legal than gworms
- Possibly could spread with worms even when vulnerabilities are not known a priori

4.3.1.3.2 Challenges for Piggyback

- Major technical challenges
 - Replace worm payload with rider
 - Constrain damage caused by worm
 - React fast against fast-moving worms
 - Control spread rate (if we want to)
- Legal and ethical issues need to be addressed

4.3.1.4 Drive-By Downloads

- Malicious webservers exploit client vulnerabilities to install malware

- “Good” webservers exploit same vulnerabilities to install whiteware
- Whiteware patches vulnerabilities on client, cleans off malware, etc.

4.3.1.5 The Pros and Cons of Drive-Bys

- The Pros
 - Patch vulnerabilities that can’t be fixed by gworms/piggyback
 - Address common way of spreading botnets
 - Not viral (no harmful spreading)
- The Cons
 - Penetration and auto-patching could be harmful
 - Could be useless if system is already compromised
 - Ethical and legal issues need to be addressed, but are different in subtle ways from gworms?

4.3.2 Inertia

- Why haven't we done this before?
 - Gworms previously done, but ethical and other issues remain
 - Piggyback and drive-by downloads not previously done
- What will derail this?
 - Perception, liability, legality, side-effects, lack of efficacy
 - Technical challenges, e.g., payload replacement

4.3.3 Progress

- gworms have been technically feasible in the past
- Piggyback/drive-by may have been technically feasible in the past. But now there are more technical tools available, e.g., virtual machines, more computing power
- Increased awareness of cyber-security issues may make this more palatable
- Increased problem with botnets and malware may change the cost-benefit analysis for society

4.3.4 Action Plan

- Requires research
- Technical feasibility
- Theoretical models and simulations
- Investigate non-technical aspects
- Legality, etc.

4.3.5 Jump-Start Plan

- Small workshop on using attack vectors, bringing together technologists, lawyers, government
- Early-stage research funding

GROUP UPDATE - WE WERE ABLE TO OBTAIN PRELIMINARY (AND PROMISING) SIMULATION DATA ON THE PIGGYBACK. WE HAVE FORMED A COLLABORATION BETWEEN LABS AT UCSD, LBL, AND LOS ALAMOS AND PLAN TO MEET WITHIN THE NEXT FEW WEEKS. WE ARE PLANNING TO WRITE A MANUSCRIPT ON THE PIGGYBACK APPROACH.

4.4 Idea - Missing-Self Paradigm

Background:

- Mammalian Immune System defines self in two major ways
 - Primary (Organic/Central) Self: Whatever is present at, or just before, birth, regardless of what it looks like. The only criterion is presence. This is tagged (Major Histocompatibility Complex (MHC) which is an imperfect example).
 - Secondary: What comes later is interrogated for its behavior. If it causes damage or injury or stress, the immune system is alerted to reject it. Damage is signaled to the immune system by alarm signals from the damaged cells. If it is harmless, it is not rejected. If it lasts long enough without causing harm or generating alarm signals, it becomes part of the definition of self.
- A cyber or computer system can also define self in two similar ways
 - Primary (Organic/Central) Self: Whatever is present at, or just before boot time, regardless of what it looks like. The only criterion is presence.
 - Secondary: What comes later is interrogated for its behavior or its provenance. If it comes from a trusted source, and/or if it does not cause damage, it is not rejected. If it lasts long enough without causing harm, it becomes part of the definition of self. Comprehensive sets of alarm signals in cyber systems have not yet been investigated.

4.4.1 Description

How can the Cyber system do this?

- Primary (Organic/Central) Self: whatever is present at, or just before, birth, regardless of what it looks like, is tagged. Anything that is not tagged can't run or be opened.

Examples:

- The machine generates two random numbers: one is used to tag the "self" executable entities, the other points to the "space" that the tag is inserted.
- All unlabeled executable entities that arrive later are not tagged, and cannot be "opened". This is similar to implementation of restrictive security posture, e.g. deny everything that is not explicitly permitted as in Trusted Platform Module (TPM) Management
- At shut down, all unlabelled executable entities are deleted - This is repeated at every boot up
- When a machine is cloned, all unlabelled executable entities are erased
- There is a mechanism to add to the Primary self (see behavioral self below)
- Distinguishable difference from "code signing" scheme and trusted third provided tagging

- Secondary/Behavioral Self: what comes after completion of tagging is interrogated for its behavior or its provenance. If it comes from a trusted source, and/or if it does not cause damage or generate alarm signals, it is not rejected, and if it lasts long enough without causing harm, it can become part of the definition of self
 - This is a multi-scale, collaborate behavioral pattern/model (e.g. process, host, network, user, community, enterprise, mobile device) that consists of three choices (depending on trust level or generation of alarm signals):
 - Add and tag
 - Sandbox
 - Delete
 - “Fast response” aspect and “slow response” aspect
 - Fast = Primary → self allowed to change only when you tag it
 - Slow = Behavioral → more dynamic, puts human in the loop

4.4.2 Inertia

- The amount of arbitrary code execution increased significantly, for example, malwares are getting downloaded and executed covertly
- Though there are many techniques (Vista Kernel-Module code integrity checks, Trusted Platform Module, Intel’s Trusted Execution Technology, etc.) for code and process authentication and validation, there is room for further improvement
- Different trust management systems (at process, platform and network level) are major initiators to explore tagging

4.4.3 Progress

- Potential derailers - Primary self mechanism needs change to OS, creation of sandboxes
- Technically Feasible?
 - Primary: Yes
 - Behavior: Yes, scaling is feasible, given sufficient computer power
- Environmentally Feasible?
 - Primary: Yes
 - Behavior: Yes as an overlay to existing technologies
- Mitigation of Concerns
 - Primary: none at this stage
 - Behavior: privacy concerns are mitigated because it is analysis of behavior with no knowledge of individual identity. Whatever length of time is set for a well-behaved program to be labeled as self, can be learned by the attackers and subverted.

4.4.4 Action Plan

Multi-dimensional, distributed characterization of “Primary and Secondary Self”

- Seed the research community (e.g. STTR, RFA, RFP, BAA, and SBIR) in three phases.
 - Fundamental research
 - Clinical trials (in various test environment e.g. DETER, HPC environment)
 - Deploy the system

- Standards for definition of common language & models used to signal threats & threat behaviors (e.g. threat ontology)
- Determine if self can be applied only at machine level or can it be applied to entire enterprise, cloud

4.4.5 Jump-Start Plan

- Acquire Funding (5-10-10 million dollars for three phases)
- Create a collaboration between immunologist(s) and cyber security expert(s)
- Create a group of people who care about this proposal to further it, e.g., getting help on requirements, existing capabilities and estimating dollar amounts
- Notion of tagging should be part of research

5 Hardware-Enabled Trust

New Game: Knowing when we've been had

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

What is the new game?

One of the hardest things about today's game is not being aware when we're losing. Our trusty PC has no way to notify us that it has in fact become an enemy agent or a zombie, secretly exfiltrating our financial secrets to identity thieves, or spamming our neighbors for some botmaster. Since we have no real plan for checking and restoring the integrity of our assets once we start using them, we are forced into the impossible position of having to deploy impregnable systems. In the new game we persistently monitor our assets for changes in trustworthiness by embedding tamper-resistant roots of trust in the architecture. Attacks can be stopped in their tracks if we can isolate and decontaminate their host.

Introduction

There was no attempt to provide comprehensive coverage of all the ideas in the areas of hardware-enabled trust. The list is a simplified categorization of the product of a brainstorming session. Below is a snapshot of the discussions of this group, covering most of the topics discussed during the session. Some of the ideas discussed in this report are covered in more detail in the National Cyber Leap Year Summit 2009 Co-Chairs' Report.

Seven, ten-year long-term goals were initially identified from which ideas were identified and put into seven categories. These ideas were subsequently regrouped into six and finally four ideas. The distillation process is discussed following the description of the four final game changing ideas.

The group developed action plans for the focus areas and in the process revised the focus areas to include:

- End to End Trust
- Hardware defenses for attacks
 - Hardware that does not leak
 - Hardware monitoring of normal behavior
- Resilience
- Secure Cloud Storage

A general purpose action plan strategy is:

- Institute a competition for building the best secure widget
- And a competition to break it

Common aspects of action plans:

- Develop national security standards for testing hardware
- Competition (as described above)
- Industry-academic teams are key to success

5.1 Idea - End to End (e2e) Trust

5.1.1 Description

- Need minimum (canonical) set of trust properties, protocols to exchange them, trust infrastructure to support these operations
- The canonical set may include:
 - Secure key management
 - Verifiable identity, attestable identity
 - Ability to contain or isolate
 - Research question: What are these trust properties?
 - Domain-specific abstractions are necessary for the big picture
 - Research question: Object-oriented extensible language with defined operations, supported by hardware
- Other important considerations:
 - Interoperability is key to end-to-end trust
 - Hardware based protection of audit ability
 - Heterogeneous systems, from sensors to servers, need to be able to enforce trust in a uniform fashion
- Secure the e2e trust in distributed heterogeneous environment
 - Including storage, computation and communications
 - Devices need canonical set of properties supplemented with domain specific identity
 - Privacy preserving identities
 - Identify who/what we deal with
 - Where we want to go (what devices/services/networks we want to access)
 - Set of principles describing the environment
- Observations
 - Some level of anonymity is possible in some areas but not everywhere (e.g. it is limited in cell phone networks)
 - NetBooks may be a tipping point – trust techniques may start there
 - Signatures of software components are more available than the same information on hardware components. SignaCert has over 600,000,000 signatures of software.

5.1.2 Inertia

- Infrastructure takes time to build
- A forum is not an efficient enabling mechanism
- Need legal support because of the nature of the problems - Need private sector participation
- Public does not understand the threats; outreach is necessary
- Cannot define definitive minimal set (no agreement on this)
- Application writers, system developers don't even use TPM that is widely available; will they use the new generation of trust technologies?

- Cannot test what is not specified by manufacturer, e.g., hardware backdoors/Trojans inserted by hardware designer

5.1.3 Progress

Other possibilities

- Reduce barrier to acceptance by creating (inter)national authoritative repository for whitelist component signatures (www.isitsafe.org) - software and hardware
- Create a compromised or revoked key service
- Define bottom-up abstractions – e.g., domain-specific objects with allowable operations
- Possible to define generic canonical sets of security properties, e.g., secure storage of long-term secrets (e.g., private keys), but rest is domain-specific
- Data-sheet (retrievable) of functionality provided by hardware components (Component provenance data-sheet) may be part of trust information, already exists. Data sheets need to be portable across heterogeneous test systems
- Set up environment where people will use provided hardware security features
- Need new business model to make a difference implementing suggested changes

5.1.4 Action Plan

- Determine canonical set of common health properties supplemented by domain provided information that systems should be able to request and attest
 - Supporting on demand health checks
 - Supporting dynamic measurements
- Develop a trusted unforgeable identity down to the component for devices/platforms
 - To allow for correct attribution
 - To manage connections
 - To manage and enforce trust
- Study and determine the market drivers that create demand for trust? Study why TCG concepts have yet to make traction
 - Where is the Velcro holding things back?
 - What incentives are possible to change the situation?
- Create a national trusted infrastructure test bed - Government, academia and industry participate
- Identify and develop standards for component and device identification
 - Create the DNA to describe a system top to bottom starting at the IC level
 - The information can be used for attestation (as pass/fail), but not disclosed for privacy reasons

5.1.5 Jump-Start Plan

Establish an operational pilot implementing these concepts including infrastructure to enable remote attestation (short term TCG-based; subsequently next generation of trust technologies)

- Users Group to take specifications to implementations for market segments - Verticals application domain possibilities need to be explored and facilitated Financial, Health Care, SCADA

- Sponsor a forum for verticals to collaborate and identify common infrastructure needs – Standards for inter-trustability need to be developed

5.2 Enable Hardware to Counter Attacks

- Information leakage thru side-channel, covert channel attacks
- Attacks connected with physical possession of device
- New hardware features for performance that degrade security, e.g., cache, Hyperthreading

5.2.1 Description

- Considers both hardware and software attacks
- Hardware defenses for:
 - Information leakage due to hardware -induced Side channel channels
 - Continuous measuring normal behavior
 - Robust characterization of normal behavior
 - Hardware Trojans – can develop state machine for hardware system for normal behavior
 - Hardware to protect measurements
 - Hardware to do monitoring (like IBM service processor)
 - Sanitization features for malware already present
 - What we measure
- Hardware verifies system integrity at runtime
 - Continuous biometrics, continuous monitoring
 - Dynamic measurement

5.2.2 Inertia

- Costly, chip yield is limited, severe performance degradation
- Software attacks are prevalent
- Hardware attacks had been kept as classified by governments; no good source of information is available.
- Definitions are needed:
 - Health: (need definition)
 - Integrity: hash-identity, safety (device does not blow up)
- Overhead of storage of this additional information – from monitoring and metrics-- is proportional to cache line (7-14% more storage), for those systems that already do this

5.2.3 Progress

- hardware cost decreases, while computers are being used for more important transactions, data and control of critical infrastructures
- Mobility increases the risk of hardware attacks
- Cloud computing (servers) magnifies the effect of the attacks; hardware capability would be helpful
- We look to reliability for new ideas (e.g., N-version programming)?

- Aspects of system's integrity/health that can be characterized as system evolves
- We will be able to measure integrity of high-assurance software. We will be able to determine the cause of corruption/security issues (how did software get to that corrupted state?)

5.2.4 Action Plan

Elevate the importance of security in the design of hardware performance and power features. Make security a 1st class citizen in hardware design

5.3 Sub-Idea - Enable Hardware to Counter Attacks—Hardware that does not leak, hardware defenses for information-leakage attacks

5.3.1 Description

Example: software cache-base side-channel attacks

- Memory leak problem (garbage collection)
- Software cache-based side-channel attacks

5.3.2 Inertia

- Security has not been considered important enough.
- Hard to enumerate the security properties
- Non-leaking versus alternative implementations needs to be considered
- Intellectual property issues will arise

5.3.3 Progress

- Possible if government can do something

Other discussions

- Identify a method to fingerprint hardware so it can be vetted
- Interoperability of test data

5.3.4 Action Plan

Short term

- Try hardware solution for secure and high performance cache
- Figure out metrics for side channel attacks: how do we measure the severity of side channel attacks?
- How do we quantify the risks: We need to understand how to quantify the risks associated with attacks
- Figure out metrics for evaluating the security properties of a design: What should designers be looking for? How do they evaluate design options?
- Prototype designs that have already been proposed

Long Term

- Design secure hardware subsystems that are both secure and high performance

- Establish a set of criteria that represent acceptable levels of security: The objective is to give guidelines to improve designs over time.
- Set of design principles for secure processors
- Establish a method to verify hardware integrity
- Roadmap for improving identifiable metrics

5.3.5 Jump-Start Plan

- Read proposals and prototype, and give feedback if it works
- Establish competition (open collaborative teams come up with design) design competition then break competition. Use open-cycles government has in trusted fabrication labs.

5.4 Sub-Idea - Enable Hardware to Counter Attacks—Continuous hardware monitoring of normal behavior

5.4.1 Description

- Hardware can automatically collect data and may be non-by-passable
- Hardware can protect measurement data and procedure

5.4.2 Inertia

- Serious data bandwidth is necessary to collect data
- Has been done already – in networks and in software
- Has problems in identifying legitimate behavior
 - False positives
 - False negatives – may miss problems

5.4.3 Progress

- Multicores allows parallel monitoring
- Incentivize manufacturer to join hardware fingerprinting efforts
- Run competition
- New in computing devices

5.4.4 Action Plan

- Short term
 - Identify measurement technology that can measure normal behavior of software and hardware systems
 - Measure low-entropy systems, e.g., web-server and a SCADA system - can add to SCADA system
 - See if it handles legitimate peak loads - Chron tab (irregularly scheduled jobs or activities)
 - Methods for process calibration
 - Identify what can be measured
 - Making mounds of data about program behavior available

- Long term
 - Identify what should be measured to characterize multicore behavior, user behavior
 - Methodology for applying what the parameters are for measurement
 - Hardware collection of behavior

5.5 Idea - Trustworthy Storage and Data

5.5.1 Description

- Self-protecting data is the way to proceed.
- Deployable key management solutions must be built into hardware of commodity products
- Develop Prototype Secure Storage Area Network (SAN) System
 - SAN controllers are no longer disk drives hiding behind server systems
 - Proposal is not limited to a strict definition of SAN
 - Network attached storage is also included here
 - Full fledged network nodes (using iSCSI)
 - Vulnerable to full class of network attacks
 - Need to selectively share data with multiple clients with different security properties
- User controls needed
 - Access controls need to be both mandatory and discretionary
 - Data owners need to be able to specify policies
 - Mandatory controls needed to control malware

Types of data

- User-control of data is important
- Grey data
- Self-describing data

5.5.2 Inertia

Why hasn't this been done?

- Storage designers are mostly disk designers – they view security as a problem for the server – NOT for them
- Initial version of object store didn't address limiting capability propagation
- Why bother – implement and prove object-oriented self-describing and protecting data – then let people/vendors catch up
- Unanticipated use of various technology and their confluence – perfect storm
- Cloud storage – cool – may lead to horror reality – need horror-story examples
- No incentive for industry to collect problems - Need economic motivation, but even this may not be enough
- Main motivators: FEAR and AVARICE
- Public awareness of security risks

What would derail the change?

- Controllers are not sufficiently resistant to network attacks
- SAN's security is not as good as previous disk farms, and needs built-in security
- Key-management is difficult for data at rest
- Cryptography is hard to implement correctly, especially in distributed environments
- Data leakage protection – covert channels are hard to stop

5.5.3 Progress

Technically Feasible

- Mechanisms for secure SCSI disks may also solve Side-channel attacks
- SCSI standards committee – object Store w/ capabilities
- (TPM-std) encrypted disk drives for short term, IEEE std 1667 are improvements
- Object-oriented architecture
- Context where objects can be viewed
- Crypto well developed, key-management solutions are studied; a lot of work done in this area.

Environmentally Feasible

Existing standards work ongoing in this space

- TCG has standard for disk storage
- Object Store is a capability-based standard under development by storage community
- Existing cryptographic standards for key protection

5.5.4 Action Plan

Joint Academic/Industry project to build and demonstrate a SAN controller to defend against all of these classes of attacks

Timeframe

- RFI in 60-90 days
- RFP 9 months after that
- 2-3 years contracts
- Need both academic and industrial team members
- Need to ensure competitiveness – standards need to be open – don't let one company lock itself in
- Implementations could be open or proprietary

Multiple approaches

- The RFP should permit multiple approaches and multiple contracts should be considered
- Some could be based on securing the SAN controllers
- Others could be based on encrypting the data before the SAN controllers ever see it

5.5.5 Jump-Start Plan

See 60-90 day implementation above

5.5.6 Comments

Game change by the bad guys

- Targeted spear fishing attacks to steal data
- Data Leakage Protection (DLP) products under development by various companies
- DLP products are re-discovering confinement and information flow control
 - Some developers are unaware of the extensive work in mandatory access controls dating back to the 1970s
 - Others are well aware
- As DLP products get established, commercial covert channel attacks will become common
 - Bad guys are well aware of covert channels
 - Haven't used them much YET, because other attacks were easier
 - But as DLP products become effective and widely available, the bad guys will be quite capable of using covert channels

5.6 Idea - Resilience

5.6.1 Description

Commodity hardware still executes critical services even when compromised

- Tools
 - Redundancy
 - Diversity
 - Checkpointing / roll-back
 - Reconfigurability / self-repair / evolution
- Instantiation
 - Multi-core processors

5.6.2 Inertia

- Hierarchical trust model - Full-stack attestation (TPM) – the operating system architecture has severe limitations
- To get full effect from hardware diversity, need more software diversity – extra complexity does not improve security - Diversity may add more attack vectors (want vertical rather than horizontal diversity)
- We will never get vulnerability-free software, but execution of malware is the problem

Challenges

- Industry buy-in
- Costs (area, power, complexity, design, validation, etc.)
- Lack of incentives
- Integration of different techniques

- Hard to create meaningful scoring systems for security (resilience or up-time easier) - What are the set of general properties that must be tested?
- Hard to get vendors to move up on the evaluation scale
- Done previously
- Must be unobtrusive, not hog battery, performance

5.6.3 Progress

- Processors are cheaper with multicore/manycore - Operating systems are getting better about incorporating hardware features
- We compute on many computers/devices, and attacker has to break many to get to protected data
- We can use diversity to improve security
- Open source may improve diversity techniques
- Can leverage reliability mechanisms (at some stronger level with some changes) to provide greater resilience to attacks
- NISB could apply to testing other trust properties - May need parallel board like NTSC and FAA – testing versus enforcement organization, but may be single organization
- Metrics – passes which set of sets (stars)
- Sets of tests like the EU’s randomization sets (Estream)

5.6.4 Action Plan

- Establish benchmarks / define scope
- Program to build a prototype with commodity components (1yr/\$1M) - Platform for experimentation. Develop tool kits.
- Funding for research and prototypes (5yr/\$50M) - Academic + industry teams to build a system prototype
- Establish a National Information Safety Board (NISB, a federal evaluation / test organization) Note: this entity was renamed in the Co-Chairs report.
- Security standards for federal purchases

Discussion

- Academic + industry teams are expected to integrate different tools to build a system that can meet the standards set by benchmarks for the government program
- NISB will test all commodity systems and publish the test scores. Consumers will be encouraged to buy systems with high scores just as they are encouraged to buy cars with higher crash test scores from NTSB.

5.7 History of Idea Development

As noted previously, the group first identified long-term goals and grouped them into seven categories and ultimately focused on four broad, encompassing ideas as outlined.

5.7.1 Leap-ahead, Long Term Goals – 10 year

- We will build a computer that will not execute malware

- We will be able to make a determination whether to trust a device, a network, or a software package based on dynamically acquired and exchanged standard trust information and user defined trust and security policies
- A user will be able to make an informed decision about purchasing a device or a service based, in part, on independent security scoring
- Transactions will be dynamically re-routed into an optimal trusted path, independently from their origination in terms of device, network, and application
- Distributed data objects will be able to protect themselves based on minimum security sets and user defined policies
- Security will be considered a core feature when architecting hardware
- New trust models will be introduced that are rooted in hardware instead of enforcing hierarchical interdependencies across the software stack

5.7.2 Initial Ideas

Ideas were generated and grouped into categories.

- New trust models enabled by hardware (substituting hierarchical models with hardware-rooted models with fewer inter-dependencies)
- Resilience as a foundation for security features
- Hardware defenses for hardware attacks
- Evaluation and dynamic measurement
- End-to-end trust in a heterogeneous environment, in order to enable end-to-end communications assuring an acceptable level of security in a heterogeneous (diverse networks and devices, from sensors to servers) and distributed (e.g. cloud computing) environment
- Trustworthy storage and data rooted in hardware
- Designing crypto/randomization into core computer hardware in order to support secure execution and secure storage

5.7.2.1 Idea Development

- New trust models enabled by hardware – what does it mean? List of ideas:
- How can hardware protect trusted applications and data? - Even if the operating system is compromised
- Replace hierarchical trust models with alternative solutions
- Hardware provides essential security in systems
- Object-oriented representation of data and associated allowed operations and constraints can serve as basic architecture, permitting us to build the system bottom up
- Start from scratch with a clean slate to see what can be done - The results can potentially be retrofitted into existing systems, but often it is hard to incorporate the best ideas into existing architectures)
- Hardware decoys with low overhead can be incorporated into the standard set of security activities
- Hardware-enabled trust must be very low cost, and take into account short lifetimes of commodity hardware

- How about very thin simple hardware clients?
- Take all intelligence out of hardware to make it simpler
- Resilience as a foundation for security – what can be done? List of ideas:
- Take advantage of diversity and redundancy, e.g., for resilience and security
- Hardware built with resiliency and fast recovery mechanisms will improve overall security
- We need hardware that can run security-critical tasks even if the system is partially corrupted
- Take advantage of distributed computing platform – spread spectrum computing
- Hardware defense for hardware attacks – what can we do? List of ideas:
- Enable hardware to counter hardware attacks
- Design hardware that spots and counters malicious hardware – Must allow legitimate upgrade and replacing of hardware components
- Recognize, measure and enforce normal (not abnormal) behavior – the set of normal behavior is far smaller than the set of abnormal behavior
 - Continuous measurement, can track by user ID or other, more privacy-conscious, parameter
 - Ensure there is a way to know what is “your” hardware
- Commodity-level tamper resistance
- Evaluation – having implemented security features, we need a reliable way to evaluate them. List of ideas:
- Practical security evaluation of hardware, e.g., ability to attribute to manufacturer, ensure authenticity
- Tools that verify that the product of fabrication is correct with respect to the specifications sent to fabrication
- Evaluation process that is open and reproducible
- Both the design and each instance
- Third party and self-evaluation
- End-to-end trust in heterogeneous and distributed environments
- Need minimum set of trust properties, protocols to exchange them, trust infrastructure to support them (all networks, all devices, from sensors to servers)
 - Interoperability
 - Hardware-based protection of auditability
- Use hardware to assess identity and health of systems
 - E.g., continuous biometrics
 - Continuous measurements
 - Hardware verifies system integrity at runtime
- Secure hardware interfaces
- Operating systems leverage TPM and future trust features rooted in hardware
- Considerations raised by cross-group synthesis discussions with other groups:

- Community situation awareness and a shared ontology are necessary to convey and implement a shared version of trust
- Distributed defense based on behavior-based models will be helpful
- Transient dynamic communities of trust will emerge; need to be considered as systems are built
- How can hardware support whitelisting of software that runs on it?
- Hardware can provide acceleration of world-switching (VMs)
- Hardware-enhanced accountability is a useful notion (e.g., ensure attribution and non-repudiation)
- How can hardware help detect an insider attack? Possibilities exist.
- Can hardware provide a set of data about suspicious activity?
- Automatic reporting for national repository of suspicious activity
- Incentives for active defense (e.g., in hardware)
- A repository of patterns of communication (assisted by hardware collection)
- Multi-layered defense
- Hardware-assisted continuous ground-truth evaluation
- Different levels of service provision
- Trustworthy storage and data – what can be done? List of ideas:
- Secure cloud storage needs to be controlled by user
- Do not restrict the flow of data, restrict the interpretation of data instead
- Use hardware to provide auto-redaction (minimization, anonymity, sanitization) to handle flow of information to protect lives, privacy, etc.; ensure removal of sensitive data
- Hardware performs provenance checking, e.g., to recognize good code, even if the operating system is compromised (feedback from data provenance group)
- Need attributes defined in order to perform attribution and checking
- Protection of provenance information itself can be rooted in hardware
- A research question: How to architect a coherent secure data and storage system, e.g., Cloud, using developed ideas to achieve practical and resilient design. Combine industry and academia in a single team.
- Designing crypto/randomization into core computer hardware for secure storage and secure execution – List of Ideas:
- Can crypto improve availability? Crypto is a great tool for improving confidentiality and integrity. But are there new crypto techniques that can help improve availability, resilience?
- Is there a field of math, e.g., randomization theory, which can help improve both security and performance? Such an approach will help: thwart attacks while improving performance
- We need to bring mathematicians (crypto, randomization) and computer architects together? (process)
 - How to get effectively “no-overhead” crypto?

- How to use crypto and randomization to advantage in new environments, e.g., in processor pipelines, multicore

5.7.3 Focus Areas

A set of ideas was selected by the co-chairs for more detailed examination.

- End-to-end trust
- Use hardware to assess identity and health
- Enable hardware to counter hardware attacks
- Resilience as a foundation for security
- Trustworthy storage and data
- Crypto and randomization in processors and memory

5.7.4 Game Changing Ideas

The group developed detailed plans for the four game changing ideas discussed and in the process revised the focus areas to include:

- End to End Trust
- Hardware defenses for attacks
 - Hardware that does not leak
 - Hardware monitoring of normal behavior
- Resilience
- Secure Cloud Storage

6 Additional Ideas

The following ideas were contributed by participants at the end of the National Cyber Leap Year Summit as additional ideas for consideration and next-step activities.

6.1 Idea - Virtualisable Network Architecture

6.1.1 Description

A new, virtualisable network architecture (VNA) that rides on the current Internet that offers advanced identity management including but not limited to: authentication, non-repudiation, attribution and network introspection. Access to the VNA may be limited to hardened thin client running on a hardened hyper-visor complemented by a hardware token.

To enter an accountable virtual network domain, a multiple-attested federated id will be employed. The ID would be issued by a nation-state or other recognized entity (equivalent to and maybe leveraging passports ID's). For example this issuance of the electronic id could possibly be managed by the US Postal Service and/or US State Department in the United States.

There could exist multiple sub-domains for different sectors such as one for the medical establishment, defense industry, financial industry, e-commerce, etc. Each sub-domain could potentially have unique policies appropriate for that environment. For example a sub-domain could create a strictly accountable universe for all transactions.

This would largely eliminate Spam, Phishing, Identity Fraud/Spoofing, significantly raise the risks of hacking attacks by having authentication and attribution.

For particular applications, sub-domains could exist on a purpose built communications substrate based on a semi-regular lattice/mesh based communications infrastructure to create to increase availability, performance and security.

The new network architecture should be built using modern security and safety techniques so that it is fit for purpose in critical industrial systems, financial, medical, nuclear, mining, Government, e-commerce.

6.1.2 Inertia

Some of the techniques were not available / we didn't recognize the need for security and safety to extent needed / we didn't rely on technology at the same level we do now

6.1.3 Progress

- Significant research in the underlying enabling technologies
- Recognized need and appreciation of the need for this particularly in the defense, financial and commercial sectors, there is an acceptance if it was appropriately managed, there is a need for post quantum evolution of security systems, opportunity as e-medical is emerging
- What would mitigate our doubts?
- Transparency of system design; it is now technologically feasible

6.1.4 Action Plan

- Identify a first team of stake holders interested in participating

- Explore cross-cutting identity, policy and functionality requirements
- Develop action plan and secure funding
- Develop a prototype for a particular sub-domain such as for an emerging sector (e.g. medical establishment) or an critical sector (e.g. the energy sector)
- Who can help (in no order)
 - NITRD, DOE, USPS, US State Department, HHS, IBM, Naval Research Laboratory

6.2 Idea - Global Electronic Identity Management System

6.2.1 Description

A new robust (post quantum secure) global electronic identity management system that more accurately reflects the way human's reason about trust relationships. The proposed GEID system would implement a multiple-attested federated id that combines the best features of centrally managed certificate authorities, with the ability to have more than one entity attest to an identity. It should also be possible to electronically aggregate multiple issued id tokens to attest a single entity.

The hardware token managing an identity could be issued by a nation-state or other recognized entity. For example this issuance of the electronic ID could possibly be managed by the US Postal Service and/or US State Department in the United States.

More than one party can attest to the identity managed by that token, including Governments, large organizations or other individuals such as friends and family members. The information used to reason about an identity assertion should be managed in a distributed decentralized federated system. The system should ensure interactivity, data minimization, privacy, least privilege, confidentiality, integrity, authenticity and have the ability to be audited by all stake holders. Any enrolled user should be able to request appropriate levels of information to authenticate an identity, however each such request must be audited and in some cases require authorization by identity being queried.

The system should support "composite" identities, such as Corporations and Organizations, allowing operations to be attested to by an organization that is separate from the individuals. For example "Authorised by 3 out of 5 directors of company X". See work by NRL.

The system should be designed to protect against collusions of 'assertion' failure, and provide increased transparency into how an identity has been asserted. The system should include soft and hard reasoning ("I believe this is my child", "I have established this is my child using DNA tests").

Furthermore the system can be adapted so that when a high value transaction takes place, the identity of the actors and the transaction must be attested to by multiple entities, where the entities are held legally accountable for attesting to that identity/transaction. The accountability is limited only to matters of identity, and knowledge of the transaction, but not the transaction itself.

6.2.2 Inertia

Some of the techniques were not available / identity systems have traditionally been centrally managed.

6.2.3 Progress

- Significant research in the underlying enabling technologies,
- Recognized need and appreciation of the need for this particularly in the defense, financial and commercial sectors, due to international collaboration.
- Requirements of several different nations have been effectively captured by international implementations of first/second generation public key certificate authority architectures (See Transglobal Secure Collaboration Program) and European studies (see EU EID-STORK)

What would mitigate our doubts?

- It is now technologically feasible
- Transparency of system design
- Allow identity to audit who has access what information about them at what time and to provide varying level of access control to different organizations
- That assertion information should be distributed and decentralized, where information is selectively released by individual authorization, i.e. No single database store. Each attestation authority is responsible for managing accuracy of their data.
- Can leverage existing certificate authority efforts, and allows them to be integrated into new environment
- Must be capable of supporting different national/regional policies. Must support interoperable communications between different countries.

6.2.4 Action Plan

- Identify a first team of stake holders interested in participating
- Explore cross-cutting identity, policy and functionality requirements
- Develop action plan and secure funding
- Develop a prototype for a particular sub-domain such as for an emerging sector (e.g., medical establishment) or an critical sector (e.g., the energy sector)
- Related to other work group projects:
- Moving Target Defense: Resilient Cryptographic Systems. The current proposal outlines techniques for relying on multiple non-intersecting security domains to attest to an identity.
- Digital Provenance: Reputation Engine. The current proposal can be seen as a type of reputation engine.
- Digital Provenance: Data Provenance Security. The current proposal will share many requirements o the Data Provenance Security group.
- Digital Provenance: Data Provenance Definition and Management. A global electronic identity management system is required to support the DPD&M proposal.
- Digital Provenance: Government Role. The current proposal supports one or more Governments participating together with commercial organizations in the administration of a identities in a global system. Each Government can maintain their own identity assertions on an ID while taking advantage of assertions made by one or more over

Governments/institutions. This proposal addresses the concern of single point of assertion failure, and mitigates fears of a single ID document.

- Additional ideas: Virtualisable Network Architecture
- Additional Ideas: Global post quantum secure cryptography based on Identity. The current proposal can be hosted within the Global PQS CBI proposal.
- Who can help (in no order)
- NITRD, CyberSpace Sciences and Information Intelligence Research - ORNL - DoE, US State Department, HHS, PricewaterhouseCoopers, Synaptic Laboratories Limited, EU EID-STORK, and others to be identified

6.3 Idea - Global Post-Quantum Secure Cryptography Based on Identity

6.3.1 Description

Global cryptographic services (authenticated key exchange, digital signatures, etc) based on identity that is robust and secure against both classical and quantum computer attacks. The system exploits a federated architecture, where at least one organization from each of the federations participates in identifying users, assisting with key exchange operations and other related functions. This proposal describes an infrastructure suitable to implement the core functionality required on desktops and supporting public infrastructure.

6.3.2 Inertia

- Technologies exist, but have trust scalability limitations which prevent the creation of a global authentication/encryption network
- Voltage Security offer a commercial public key identity based encryption (IBE) product which is ideal for enterprises and small groups of enterprises. However this system has a central point of trust in the server which would prevent acceptance of single global IBE infrastructure being deployed.
- KERBEROS is an example of a symmetric federated Key Distribution Centre based technology that supports key negotiation by identity. Unfortunately there are security limitations in this context. See the paper [[Formal Analysis Of Kerberos 5](http://citeseer.ist.psu.edu/765675.html), <http://citeseer.ist.psu.edu/765675.html>].
- Current proposals are not considered to be post quantum secure
- Voltage's IBE system does not claim to be post quantum secure
- KERBEROS running as a federated system relies on known "at risk" classically secure public key algorithms to achieve scalability. Furthermore, user's access the system using passwords which may not be sufficiently secure.
- Previously no method for internationally managing name spaces in a way that protects against cyber-warfare by one large agent over another. See the problems that exist with today's public key infrastructure "[MD5 considered harmful today - Creating a rogue CA certificate](http://www.win.tue.nl/hashclash/rogue-ca/)", <http://www.win.tue.nl/hashclash/rogue-ca/>.
- The use of online servers has prevented up-take in some contexts, but is generally not a problem for Internet communications (which already relies on 24/7 online servers such as the Internet Domain Name Server infrastructure).

6.3.3 Progress

- Wireless ad-hoc mesh network architectures have advanced the study of multi-path key exchanges over distinct paths using symmetric techniques.
- Modern Smart cards can be used as trusted couriers for key material between an enrolled user and one or more online key translation centers.
- Synaptic Laboratories has introduced technologies to express scalable symmetric key authenticated encryption systems where no single trusted third party [or collusion of (n-1) out of n participating third parties] can discover the final key exchanged between two users. This addresses the core trust problem that spurred the design of public key technology (See [Quote](http://synaptic-labs.com/resources/security-bibliography/53-asymmetric-key-exchanges-classical/78-bib-celebrating-the-30th-anniversary-of-pkc.html) by Whitfield Diffie, <http://synaptic-labs.com/resources/security-bibliography/53-asymmetric-key-exchanges-classical/78-bib-celebrating-the-30th-anniversary-of-pkc.html>).
- Synaptic has proposed techniques for rapidly integrating the global authenticated encryption scheme into existing products based on SSL/TLS, SSH, IPsec, SSL VPN, and e-mail by "post-processing" the output of unmodified products. This allows all current infrastructures to use current public key standards and maintain FIPS 140-2 compliance and be incrementally upgraded to achieve post quantum security against known attacks.

Integration

- This proposal can act as a platform for hosting the global electronic identity management proposal, and can support the global key exchange operations based on ID required for the Virtualisable Network Architecture.
- The Global electronic identity management proposal provides a platform for "describing and reasoning" about an identity and its trust relationships, where as this proposal supports the real-time authenticated key exchange operation between those identities.

6.3.4 Jumpstart Activities

- Identify and bring together interested stake holders
- Explore existing technologies (digital signatures, manage security functions, integrated risk management systems, current public key certificate authority requirements) and draft a high-level requirements document.
- Perform further independent evaluation of next generation proposed technologies (Independent cryptanalysis on Synaptic's proposal has already been performed by Prof. Jacques Patarin).

Further Action Plan

- Identify and bring together identity stakeholders into a conference to refine requirements
- Independent evaluation of next generation proposed technologies
- Begin development of key exchange technologies and infrastructure
- Related to other work group projects:
 - Moving Target Defense: Resilient Cryptographic Systems - Secret Key Compromise. The current proposal outlines techniques for relying on multiple non-intersecting security domains, where a cryptosystem remains secure against a collusion/compromise of (n-1) out of (n) security domains.

- Digital Provenance: Global identity-based cryptography. The current proposal outlines a more concrete proposal or achieving Global identity-based cryptography.
- Digital Provenance: Government Role. The current proposal supports one or more Governments participating together with commercial organizations in the administration of a global identity management system. This proposal addresses many the concern of single point of failures.
- Additional ideas : Virtualisable Network Architecture
- Additional Ideas : A global electronic identity management system
- Who can help (in no particular order)
 - NITRD, ORNL - DOE, US State Department, MITRE, Secure Systems - IBM, Boeing, Naval Research Laboratory, ICSA labs, PricewaterhouseCoopers, Terra Wi, Synaptic Laboratories Limited

6.4 Idea - Evaluating the Effectiveness of Data Depersonalization Techniques and It's Impact on the Community

6.4.1 Description

Establish if data depersonalization techniques used by the civilian industry are effective and assess the impacts of re-sale of depersonalized data in the community. Study the way consumers of depersonalized data use the information. If the depersonalization techniques are not adequate to protect identity (before or after sale), identify what techniques and parameters are appropriate for commercial data depersonalization. After adequate peer review, enforce these techniques and parameters as Government policies.

6.4.2 Inertia

Commercial interests for selling data / Poor community-wide awareness of the risks associated with sale of personal data collected by organizations.

6.4.3 Progress

Several papers have identified that it is possible to identify the persons present in some depersonalized data released by large organizations.

6.4.4 Action Plan

Identify the security and legal experts / acquire large representative data sets of the type of information sold / start a conference and advance it with funding.

Who can help:

NITRD, US State Department, Electronic Freedom Foundation, Jeff Jonas of IBM, weak signal analysis, other published researches in this field.

6.4.5 Jumpstart Activities

Collect a large representative sample of commercial exchanged depersonalized data (find data sold by a large online commercial store, and a mobile phone provider selling location data), bring together experts in the field to evaluate how easy it is to re-personalize the data, bring together legal team to evaluate the implications of data that is not effectively disassociated from the user. Compile any changes required to law.

6.5 Idea - Measuring the Impacts of Unauthorized Information Disclosure

6.5.1 Description

Methodologies for evaluating appropriate security controls based on the confidentiality, integrity and availability of IT systems now exist. However insufficient information exists to allow an organization to establish the value of information loss to stakeholders, including customers and clients. Without such information it is not possible to make an informed decision about the necessary level of security mechanisms required.

Large scale field studies are required to establish the value of information loss with respect to different classes of data including financial, medical, intellectual property, relationship information and geolocation of time for different groups including Enterprises, SME, and individuals. Such studies could be extended to assess the financial and emotional impact of down-time or availability of access to services.

A greater understanding of the value of information managed by others, and its management, by the stake holders can better inform organizations on how to manage their IT infrastructure and risks.

6.5.2 Inertia

Commercial interests for selling data / Commercial interests to maintain 'just-enough' security to protect against legal liability. There is little incentive for organizations to identify the true cost of security breaches against individuals.

6.5.3 Progress

Technologies exist which can be used to collect this information.

6.5.4 Action Plan

Identify interested financial, social sciences, security and legal experts. Develop action plan and secure funding. Perform studies in hospitals and other medical practices.

Who can help:

NITRD, CyberSpace Sciences and Information Intelligence Research - ORNL - DOE, RTI International, US Universities, EU Think Trust.

6.5.5 Jumpstart Activities

Identify the financial, social sciences, security and legal experts. Develop a set of questions to measure metrics on. Engage many universities and some organizations to perform surveys and collect the data.

6.6 Idea - Semiconductor Intellectual Property Protection

6.6.1 Description

Synaptic Laboratories has proposed a method of designing semiconductor devices with improved trust characteristics that protect the Intellectual Property rights and profits of the fabless semiconductor design house.

Combinatorial locks can be implemented in a hardware circuit by inserting or replacing hard-wired logic with programmable logic. The logic for the look up table is locked away in a private database such as a smart card until it is used to unlock the device. An attacker must select the correct value to unlock the programmable logic that ensures correct and reliable operation of the device. This value can be remotely programmed using symmetric cryptographic techniques. To improve the utility of combinatorial locks we propose splitting the circuit design across at least two teams (Yellow and Orange) such that each team is responsible for managing independent locks in their respective modules. The remaining unlocked source code can be exposed to all teams enabling more efficient development practices over other existing, more restrictive approaches. This process allows global placement and routing of performance sensitive code without risk of chip over manufacture due to unauthorized disclosure. Simulation of the chip design is efficiently achieved using an enhanced distributed chip simulator of two or more machines. The yellow and orange teams are responsible for ensuring their portions of locked code are simulated at full speed by machines they trust will not expose their locked logic. After a circuit is finalized traditional risk management techniques are recommended to prevent modification of the circuits before and/or during manufacture of the wafer masks, there by providing assurance against a wide range of attacks. Each team is responsible for securely loading their portion of the locked circuit behavior into each manufactured chip from a remote location or a tamper proof module.

6.6.2 Inertia

There are currently no split team development, synthesis, place-and route or simulation tools that can be used to compartmentalize portions of code.

6.6.3 Progress

New techniques to ensure verilog/VHDL software protection through to manufacture have been recently proposed.

6.6.4 Action Plan

Identify one or more semiconductor organizations. Perform an independent evaluation of the techniques. If validated, work with a company like Synplicity to modify EDA tools, and develop a complete process for working with fabrication facilities. Work with companies such as Certicom who offer chip programming facilities for supporting per-chip enabling.

Who can help:

NITRD, DOE, Intel, Certicom, Synplicity, Universities of Michigan and Rice (EPIC).

6.6.5 Jumpstart Activities

Identify a large semiconductor organization, such as Intel, that is sensitive to IP theft, and get them to perform an initial evaluation of the techniques.

6.7 Idea - Dynamic Distributed Key Infrastructures (DDKI)

Dynamic Distributed Key Infrastructures (DDKI) – a topology & Dynamic Identity Verification and Authentication (DIVA) – a process & Whitenoise – a cryptographic algorithm

6.7.1 Description

For 35-40 years we have relied on Public Key Infrastructures (PKI). They have always been vulnerable to man-in-the-middle attacks. They do not scale well. They are very expensive. It is a given that they will not be post quantum computing secure (PQCS).

DDKI provides a complete, new generation identity-based, cryptosystem that incorporates: Complete federated and distributed key and identity management configuration, for example:

Horizontal implementation example

- Complete identity can be aggregated at a central location like an non-government organization trusted third party that brings together the stakeholders from public-private partnerships i.e. government, law enforcement, industry, watch groups such as an international or national body comprised of privacy and security experts from all articulated stakeholders.
- Complete identity can be parsed and federated horizontally between different stakeholders within government to create checks and balances that reflect democratic societies. No one entity/department would have the complete identity of an individual/entity/device and act on a complete identity without transparency to other sectors of the government i.e.:
 - Department of Census: responsible for issuing identity
 - Department of Homeland Security: responsible to integrate sharing of identity with all levels of law enforcement, military, and intelligence
 - Privacy Commissioner: responsible for creating the transparency to all private stakeholders including citizens, commercial entities etc. to reflect the values inherent in democratic societies (this is the “sunlight is sanitizing” element). They would be mandated to enable the sharing of responsibility for cyber-security. They would enable and oversee effective information sharing/incidence response.
 - Department of Justice: legally (public liability rests here) responsible for “following the letter of the law” by ensuring there is no abuse or manipulation of legislation regarding identity and privacy
 - Department of Education: responsible for building the capacity for a digital nation
 - Department of Foreign Affairs: responsible to bring likeminded nations together on a host of issues
 - National Institute of Standards and Technology: responsible for enabling the building of the architect of the future. Building the architect of the future is a technological reality with the goal that the technology works securely, is accessible to any stakeholder, and that it integrates identity management. It reflects the values of democratic societies.

The architect of the future must be elastic enough that it inherently can adjust to historical context in terms of the appropriate balancing of privacy and security. For example, during times of war security may require greater latitude (by legislation) and during times of peace there are degrees of greater privacy. This is the inherent democratic challenge of balancing privacy and security in technology.

Note: for stakeholders frightened of “growing government” this structure can be condensed into

one department for efficiencies with the same kind of mandate as Department of Homeland Security whose task is to integrate all elements of law enforcement and military.

Vertical implementation example

Complete identity can be parsed, federated and distributed vertically between government/law enforcement/military and industry and citizenry. For example:

- Government is the repository for the abstract of universal identity – i.e. they issue master identity keys to authorized and trusted private commercial entities like telecommunications providers and private national security entities like the military etc.
- In public sectors, telecommunication providers can issue identity management keys to citizens and entities (devices/non human nodes) reflecting the degree of anonymity required by different activities. Note – this places a burden of responsibility upon this layer which creates incentives to act securely. For example, if they want to provide complete anonymity for their clients, then private commercial entities assume the same complete responsibility and liability as the users of their services to comply with the law. When the law is breached both the criminal and the facilitator of criminal activity assume the same (or proportional) liability. There are degrees of legislated opt out of liability paths by adjusting the degree of liability the criminal and provider have dependent on the amount of specific user information they share with law enforcement and government entities. This provides a disincentive to allow cyber crimes like hate speech, electronic fraud, etc. This provides an incentive for private commercial entities to monetize varying degrees of privacy.
- This is a flexible reality that can effectively be dialed in between stakeholders through legislation: it is not “all or nothing at all” liability. It can balance ‘the profit motive’ versus ‘the responsibility’ conundrum.
- Depending upon what the public commercial sector decides to provide, citizens and entities can each choose what level of identity they wish to utilize to use critical telecommunication infrastructures. Complete anonymity of “users” places equal liability upon the private commercial sector. Pseudo anonymity shares the responsibility between network infrastructure users and network infrastructure providers. Use of reasonable legislated Identity places the entire burden of liability upon the government. All stakeholders can ‘opt in’ or ‘opt out’ of varying levels of identity and privacy. This allows all stakeholders (government/public and corporate/citizen/private to have both public and private identities, as well as multiple kinds of Identities.

Note: at the ends of the liability/responsibility spectrum we have one of two realities:

1. The private commercial sector shares equal responsibility with the criminal private citizenry sector.
2. The government sector shares equal responsibility/liability with the private criminal sector and the private commercial sector has no responsibility/liability at all.

In between, degrees of liability/responsibility are directly proportional to the degree of anonymity that the commercial private sector can monetize.

6.7.2 Inertia

- Lack of interoperability
- The technology did not exist before. It exists and is available today.
- Competing political, philosophical and economic interests
- Complexities and costs of implementation such as scalability, access control, key manageability, reversibility (forensics), checks and balances, elasticity of systems, overall overhead and complexity of systems, and ‘privacy fears’ while remaining secure.
- Ease of use and understanding
- Lack of will power, vision, direction, incentives

6.7.3 Progress

Why is this feasible now?

DDKI and DIVA technically provide:

- Federated, distributed Identity Management
- Intrusion detection making the architecture real-time for legitimate forensic use and optimal system integrity
- Continuous Authentication providing a moving target defense
- Automatic revocation ensuring an attack can only happen once
- Repudiation/non-repudiation which is integral to ‘need to know’, ‘chain of command’, forensics, liability, and responsibility. This can be inherent within the design due to how DIVA manages authentication.
- Digital Rights Management which is integral to ‘need to know’, ‘chain of command’, forensics, liability, and responsibility. This can be accomplished by Digital Object Online Resource Sharing [DOORS].
- Authorization which is integral to ‘need to know’, ‘chain of command’, forensics, liability, and responsibility
- Complete and secure federated key and identity distribution capacity that allow systems to scale infinitely, allow ‘on the fly configuration’ to reflect changing political and social context

DDKI and DIVA and Whitenoise also:

- Exploits revolutionary identity based cryptography that embeds characteristics of a one-time pad (moving target defense)
- Exploits revolutionary identity based cryptography that is bit independent (immune to current and known cryptanalytic attacks and vulnerability) and which makes it indifferent to current technological limiters such as data/memory/key leakage which is the basis of current cryptanalytic attacks like “Side Channel” attacks in Hardware. It also makes it immune to “mathematical shortcut attacks” as well as ‘brute force’ attacks. It plugs the security hole in Hardware-enable trust. This swings the cost/benefit dynamic towards the greater interests of society by making illegal behavior prohibitively expensive and approaching technologically infeasibility. This plugs the Cyber Economic hole and ensures in the vast majority of user cases that ‘crime doesn’t pay.’

- Exploits revolutionary identity based cryptography that is post quantum computing secure because the security strength of the architecture is exponential and inherently scalable ‘on the fly’ by the simple addition of subkeys to existing Identity Management and encryption (both cryptographic) keys to readily scale strength by exponential orders of magnitude.
- Exploits revolutionary identity based cryptography that will always stay ahead of the exponential computing processing threat curve because in software the speed of the cryptographic algorithm is limited by the existing computational power at any time because the speed of the cryptography is limited by the processing capacity of any hardware at any given time. This is because this cryptography is the first secure cryptographic technology that predominantly exploits the fastest available computer mathematical function, the X/Or process. This plugs the security hole inherent in current Hardware-enabled trust.
- Exploits revolutionary identity based cryptography that allows ‘virtual manufacturing and provisioning’ and lower costs by orders of magnitude, and increases accessibility (very democratic) because of the reality that software based critical infrastructure security is more secure and flexible because it is dynamic and not static. [Note: capitalistic profit motive systems have a natural tendency to drift towards a state of industry choosing the most expensive option with the least amount of service in order to solely enlarge profit margins at the expense of greater social responsibility i.e. systemic failures creeping into such systems as financial, insurance, and health care/provision]
- Exploits revolutionary identity based cryptography that allows analyzing of ‘communities of interest’, and then modeling of simulated systems utilizing key-stream as input to fractal models for evaluating health and nature inspired networks at either macro or micro levels.
- Exploits revolutionary identity based cryptography to ensure digital provenance across all technical layers of the Internet and critical communication infrastructures, enables interoperability across all platforms/operating-systems/domains, and all technological layers application-layer, network-layer, data-layer, physical-layer etc. It also enables interoperability between abstracted communities of interest: technological, social, political, philosophical etc.
- Exploits revolutionary identity based cryptography to ensure digital provenance by resolving the IP overload issue (the ‘IP Identity Problem’) caused by the semantic overloading of IP addresses containing both an IP address locator (network topology location) function from a node identity function. This enables networked entities to know the identity of its networking peers and to use that identity as a basis for authentication and authorization. This is resolved because DIVA is independent of the IP address and provides direct authentication regardless of the number of branches and modifications that are handled through the network. It is simply an end-to-end authentication system that is virtually impossible to access illegally without detection.
- Exploits revolutionary identity based cryptography to resolve the packet ordering issue. UDP headers have only routing information and no packet ordering information. TCP/IP is supposed to manage packets in their proper order. DIVA can be used as an alternative mechanism to not only authenticate but to order the incoming packets without adding bandwidth.

- Exploits revolutionary identity based cryptography to secure digital provenance of data at rest and data in the ‘cloud’
- Exploits revolutionary identity based cryptography which is the single common denominator and enabler that is required to achieve all articulated goals of the Leap Year 2009 Summit including allowing global encryption based on identity that is robust and enduring, attaching context to data, expanding trustworthy systems, facilitating unspoofable trusted paths/channels and securing data provenance on a ‘need to know’ basis.
- It is completely non-disruptive and allows seamless transition to Leap Ahead network cyber-security
- It is ready today. It addresses all the inertia problems.

Note on BOTS – As we move over to an identity based network system, BOTS will be able to be controlled and managed in a more effective way. In situations where they are not warranted they can be precluded.

6.7.4 Action Plan

What are reasonable paths to this change? What would accelerate this change?

- Commit to these initiatives with funding, education, resources (both public and private) and the full endorsement of the National Cyber Leap Year initiative.
- Strategic use cases in environments of stakeholders – intelligence/military/law enforcement, health care, financial and insurance, and utilities (SCADA – System Control and Data Acquisition) and critical infrastructures i.e. identifying and measuring the globalization and interoperability characteristics across all communities of interest and stakeholders.

6.7.5 Jumpstart Plan

Joint testing and certification

Immediately bring in technology for joint testing and certification involving the National Institute of Science and Technology (United States of America) and Communications Security Establishment (Canada) and any willing International Standards Boards and International Regulatory entities for complete transparency throughout the process.

Joint development and deployment

Engage in a joint development and deployment of DDKI, DIVA and Whitenoise into the Intelligent Grid at the British Columbia Institute of Technology and a project site in the United States of America simultaneously. [Apply scientific methodology by using a blind verification of reliability and validity of the technology and topology.]

Trial and measurement of the implementation

- Encourage trial and measurement of the implementation in a large commercial telecommunications carrier – one in the United States and one in Canada – with the simple deployment of DIVA in a secure network access protocol. This requires simply the addition of three data base fields in the login database of the carrier: a unique identifier field, a unique key structure field, and a dynamic offset field at the carrier

server. Electronically provision the endpoint with the DIVA utility (20kB – 150 kB) on any network enabled entity/endpoint/device.

- Note: this eliminates any needed integration with any firmware (all proprietary). The physical endpoint simply needs connectivity, memory/storage, and write back capacity for the dynamic, continuously-changing offset. This eliminates the possibility of impeding project progress because of lack of agreement between conflicting communities of interest or commercial private entities. Democratically, they are free to opt in or opt out without affecting the goal attainment framework.
- Note: this eliminates any risk to removal or bypassing of the protocol because there can be no network access without the continuous authentication verification. If the endpoint cannot provide the required authentication token there can be no network access.

Implement a DIVA/Whitenoise enabled FPGA

- Immediately implement a DIVA/Whitenoise enabled FPGA and test for vulnerabilities against Side Channel attacks.

6.8 Idea - Removing Barriers to Entry for Crypto Products into Federal Use

Streamline and expedite the approval process for Federal use of new security technologies.

6.8.1 Description

Many commercial security technologies are unavailable for Federal use even though they are well accepted and widely deployed in the private sector. These technologies often allow dramatic cost savings and efficiency gains over older technologies, but Federal agencies are unable to use them because the technologies have not received the necessary certifications and approvals. In some cases, the existence of rigorous, formal proofs of security should eliminate the need for the long certification and review process and allow Federal agencies to receive the same benefits that the private sector is now realizing. A decade or more is too long for Federal agencies to wait to realize the benefits of new security technologies. Let's find a way to get new technologies used more rapidly.

6.8.2 Inertia

This has not been done yet because the Federal agencies involved in approving new security technologies have relied on the "wait and see if it's secure" model so far. This approach usually determines which technologies are sound and which ones are not, but takes many years and leaves Federal agencies unable to use the innovative security technologies that are being invented today.

6.8.3 Progress

Provable security has made the "wait and see" model unnecessary in many cases. If there is a peer-reviewed formal proof of the security of a technology, that should be enough to get approval for Federal use. If the proof is correct then the technology is secure. Why wait ten years or more if that's the case?

6.8.4 Action Plan

NIST should determine a way to quickly approve provably-secure technologies for Federal use and should review existing regulations and identify ways to allow provably secure technologies

within them. This should involve, as a minimum, granting a blanket IATO to new encryption technologies with peer-reviewed proofs of security, and adding provably-secure public-key encryption technologies to the list of techniques that are allowed by FIPS 140-2. In the long run, standards and policies should be changed to allow the rapid adoption of new technologies that are provably secure.

6.8.5 Jumpstart Plan

Within 90 days, NIST should define and implement a way to approve provably secure technologies for Federal use. Within 180 days, a pilot of one of these technologies should be started at a Federal agency.

6.9 Idea – Real-Time Internet “MRI” (Orthogonal View)

Organizations such as the Cooperative Association for Internet Data Analysis (CAIDA) take great pains to measure aspects of the internet, such as internet topology, traffic flow and Autonomous System (AS) interactions. The data retrieved and analyzed by CAIDA and similar organizations are invaluable in attempting to understand the nature and complexities of the internet. However, the collection tools at our disposal are constrained by the internet itself. There is currently no “orthogonal view” of activity on the internet. Unlike tools within the medical profession where an outside observer can take an x-ray or MRI to see a global view of the situation, our view of the internet is very constrained. We are using the internet to observe itself, from an “inside the tube” view. It is as if we are attempting to map the human nervous system from the perspective of the synapse.

If a real-time orthogonal view of the internet were observable by all, then many benefits to global cyber health are enabled, in terms of diagnosis, prediction and defense.

6.9.1 Description

An orthogonal view of the internet is possible with a simple innovation. Placing information flow sensors at each AS could capture distilled information (such as number of packets per protocol sent to its neighboring AS’s). This information would be continually collected and sent outside of normal channels (perhaps via satellite communications) to a common collection point for consolidation and dissemination. A number of new possibilities are enabled:

- Real-time traffic pattern and “weather” data would be viewable by all
- Turbulence, anomalies and emerging problems could be observed and perhaps rectified
- If the collection mechanisms were real-time configurable, they could be commanded (by some national authority) to “drill-down” to provide more specific information concerning a particular attack pattern, tracking that particular threat
- An “over the horizon” threat detection could utilize this ability to see activity numerous “hops” away, before malicious activity arrived
- It would be virtually impossible for a coordinated attack to spoof information from all collection mechanisms to hide his activity. Network outages between and among AS elements would not affect the data collected and disseminated; it would be fault tolerant.

6.9.2 Inertia

This has been done before, on small scales. “Back channels” of communication are a common means of segregating communication for different purposes. Diagnostics or configuration control

messages can be segregated from normal network activity in a test/development network. However, this technique has not been attempted on anything as massive as the internet, or significant portions of the internet, because:

- No one takes ownership for the internet (or significant portions of it)
- There is an initial investment to be made that cannot be done by any single commercial or government entity.

There are a few forces that would be natural impediments to implementing the idea:

- **Funding:** There would be an up-front cost associated with building the infrastructure to collect, integrate and disseminate this data. Additional hardware resources (including perhaps satellite resources) would be needed.
- **Corporate Acceptance:** Additional cost and effort to install and maintain the collection equipment would be a deterrent, unless there was demonstrable offsetting benefit
- **Consumer Suspicion:** The idea that government may be involved with viewing internet traffic may not be accepted with enthusiasm by a suspicious public, unless done in a transparent manner

6.9.3 Progress

Technologically, this is already feasible. All needed components exist and could be aggregated for this purpose. Environmentally, the political and economic will may be at a tipping point to where bold, demonstrable action may be welcome, if that action seems to aid internet security

6.9.4 Action Plan

What are reasonable paths to this change? What would accelerate this change?

- Create a community of interest to devise specifications and implementation plan
- Specific funding requirements will arise from the implementation plan
- Enact legislation to subsidize the cost of the collection equipment, to improve chances of widespread (national) adoption
- **Momentum:** As the number of adopters grows, the benefits of the system increase non-linearly. If a small core group of adopters shows early success, the number of later adopters will accelerate.
- **Patriotism:** A campaign to contribute to the national cause to help secure the infrastructure within the US could encourage ISPs to participate. Similar campaigns could exist in other countries.

6.9.5 Jumpstart Plan

- Create a community of interest to devise specifications and implementation plan
- Announce X-Prize for best specifications and implementation plan

APPENDIX A: Acronyms

Acronym	Description
SCADA	Supervisory Control And Data Acquisition
AMD	Advanced Micro Devices
AMT	Active Management Technology
ARL	Army Research Laboratory
ARO	Army Research Office
ATD	Advanced Technology Demonstration
BAA	Broad Agency Announcements
BoD	Bandwidth on Demand
CAIDA	Cooperative Association for Internet Data Analysis
CCDC	Cyber Disease Control
CDN	Content Delivery Network
CNCI	Comprehensive National Cyber-Security Initiative
CONOPS	Concept of Operations
COOP	Continuity of Operations Plan
COTS	Commercial Off-the-Shelf
DARPA	Defense Advanced Research Projects
DDNS	Dynamic Domain Name Service
DDoS	Distributed Denial of Service
DETER	Defense Technology Experimental Research (testbed)
DHCP	Dynamic Host Configuration Protocol
DHS	Department of Homeland Security
DNS	Domain Name System
DOI	Digital Object Identifier
DoS	Denial of Service
DR	Disaster Recovery
DREN	Defense Research Engineering Network
DRM	Digital Rights Management
FIPS	Federal Information Processing Standards
FPGA	Field Programmable Gate Array

Acronym	Description
GDP	Gross Domestic Product
GENI	Global Environment for Network Innovations
gworms	Good Worms
HPC	High Performance Computing
HR	Human resources
HSRP	Hot Standby Router Protocol
IDS	Intrusion Detection System
IETF	Internet Engineering Task Force
IPS	Intrusion Prevention System
IPv6	Internet Protocol Version 6
ISP	Internet Service Provider
ITAC	Identity Theft Assistance Center
JIT	Just-in-Time
KVM	Kernel-based Virtual Machine
MANET	Mobile Adhoc Networks
MHC	Major Histocompatibility Complex
MMOG	Massive Multiplayer Online Games Massive Multiplayer Online Games
MPLS	Multi-protocol Label Switching
NAT	Network Address Translation
Nessus	A network scanner tool
NIST	National Institute of Standards and Technology
Nmap	Network Mapper
NRC	National Research Council
NSA	National Security Agency
NSF	National Science Foundation
OpEx	Operation Expenditure
OPV	Oral Polio Vaccine
OTP	One Time Password
PHS	Public Health System
PII	Personally Identifiable Information
PKI	Public Key Infrastructure

Acronym	Description
QoS	Quality of Service
RFA	Request for Application
RFP	Request for Proposals
ROI	Return on Investment
RPR	SONET Rapid Path Restoration (RPR)
S&T	Science and Technology
SAT	Boolean Satisfiability
SBIR	Small Business Innovation Research
SCADA	Supervisory Control And Data Acquisition
SDIO	Secure Digital Input/Output
SLA	Service Level Agreement
SOA	Service Oriented Architectures
SoD	Security on Demand
SONET	SONET Rapid Path Restoration (RPR)
SSL	Secure Sockets Layer
STTR	Small Business Technology Transfer Program
TC	Trusted Computing
TCP	Transmission Control Protocol
TPM	Trusted Platform Module
URN	Uniform Resource Name
V&V	Verification and Validation
vBNS	Very Highspeed Backbone Network Service
VLAN	Virtual Local Area Network
VoIP	Voice Over Internet Protocol
VPN	Virtual Private Network
VT	Virtualization Technology
WHO	World Health Organization
WIFI	Wireless Fidelity