National Cyber Leap Year Summit 2009:
Exploring Paths to New Cyber Security Paradigms
Draft Report of Participants' Ideas

August 24, 2009

New Game: Crime doesn't pay

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

The following ideas were captured in unedited form at the National Cyber Leap Year Summit. The ideas are a summary of the discussion of the participants in the Cyber Economics session. They do not necessarily represent the opinions of the co-editors or the organizations they represent. The Summit is 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.

Please **provide your comments**, if any, **by September 3, 2009** for utilization by the Summit's program co-chairs at http://www.co-ment.net/text/1446. To add a comment, select the "Add" tab in the left navigation menu, select (highlight) the portion of the document you are commenting on, and provide your comment. If commenting on an entire section, you may select the section heading to anchor your comment.

If you have any further questions or comments, please visit the National Cyber Leap Year Web site at the following address: http://www.nitrd.gov/NCLYSummit.aspx, or send email to leapyear@nitrd.gov.

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.

1 Introduction

This report summarizes ideas proposed and debated by the participants in the Cyber Economics session at the 2009 Cyber Leap Year Summit Workshop. They do not necessarily represent the opinions of any individual participant or co-chair, or the organizations they represent.

2 Idea – Data & Metrics for Cybersecurity Analysis

2.1 Description

Markets do not work efficiently under incomplete information. Such is the case of the market for cybersecurity. Notwithstanding the recent progresses in the economics of cybersecurity, we still lack reliable and exhaustive data and rigorous metrics on cybersecurity incidents, attacks, and infection rates. 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?

What does the change look like?

The game change consists of incentivizing (through government subsidies, mandated best practices, or other public initiatives) information sharing among private and public sector entities, in order to create a public repository of data on incidents, attacks, and infection rates. This data would enable a variety of applications and more finely tuned policy making - such as more accurate cyber risk management or cyber insurance.

2.2 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 poorly investigated. As a community, we don’t understand what information should be collected, how that information could be used, or how to deal with uncertainty and inaccuracy associated with such information.

What would derail this change?

Firms may continue to be wary of sharing their information publicly. While data could be made anonymous to protect the confidentiality of the firm, the type and depth of data 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.

2.3 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. 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 event may not be as damaging as once believed.

2.4 Action Plan

What are the reasonable paths towards bringing about that change?

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?

We could explore data collection in a limited context – for instance, at a university -- as a test bed for the approach.

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.

What are the missing technical pieces?

We should be collecting an articulation of the right set of security metrics. 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.

2.5 Jump-Start Plan

- Coordinating with DoD and others where work has already begin, plan and announce a conference to unite the various multidisciplinary research communities with the goal of defining a path forward..

- Plan and announce a subsequent NSF program focused on the many research and development challenges posed by this topic.

3 Idea – Vendor Incentives and Accountability

3.1 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. 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.

3.2 Inertia

Why have we not done this before?

It is not clear that market forces alone 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). 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.

3.3 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.

3.4 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. One suggestion coming out 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, 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.
Other models exist, but all of them will require a strong focus on the research and development of enabling technologies for this idea.

What would accelerate this change?

Convening a multidisciplinary workshop, perhaps even an annual series, on the technologies and policies to support accountability in cybersecurity (TAPSAC) would provide jump start and an ongoing drive for this research and development process. . Secure development practices standards for hardware and software could be established and incentives put in place for their adoption.

Furthermore, 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.

What are the missing technical pieces?

See the answer to "What technologies are emerging …"

3.5 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.

4 Idea – Cyber "NTSB"

4.1 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. NTSB is also charged with issuing safety recommendations aimed at preventing future accidents. 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.

4.2 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.

4.3 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.

4.4 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.

4.5 Jump-Start Plan

5 Idea – Cyber “Interpol”

5.1 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.

5.2 Inertia

Why have we not done this before?

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.

Existing organizations may already be addressing this problem, but it was unclear to the participants in the Summit exactly which organizations currently do or do not exist for this purpose.

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.

5.3 Progress

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

None were identified.

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.

5.4 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.

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.

5.5 Jump-Start Plan

6 Idea – Cyber Insurance

6.1 Description

What does the change look like?

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.

6.2 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.

6.3 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.

6.4 Action Plan

What are the reasonable paths towards bringing about that change?

Increased availability of incident and impact data will 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.

6.5 Jump-Start Plan

7 Empowering ISPs and Registries

7.1 Description

What does the change look like?

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 and Registries to halt clearly abusive or criminal behavior (such as ISPs temporarily disconnecting compromised users originating DoS attacks from the network, or Registries stopping consumers from registering obvious phishing domains) may offer the tools to prevent various cyber-crimes to those in the best position to help.

7.2 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); in a highly competitive

market, ISPs do not want to risk irking their current paying customers (or disincentivizing their potential future ones); 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. Furthermore, given the added costs and risks to ISPs (and, possibly, Registries), 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.

7.3 Progress

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

The technology to implement this kind of monitoring is now cheaper and more available – in fact, the feasibility of similar programs has already been demonstrated on smaller scales (consider the example of college campuses monitoring traffic involving MP3s).

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

We already observe a move towards IPS "pushing" security solutions to their own users – albeit this transition does not seem to be happening fast enough.

7.4 Action Plan

What are the reasonable paths towards bringing about that change?

Define a clear legal framework for what ISPs and Registries can or can't do, including common carrier-like exemptions, good faith safe-harbors, and -- possibly -- mitigations of liability costs (similarly 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 could be monetarily incentivized to identity and halt compromised hosts, or to bundle security services for end-users into ISP subscriptions.

Better authentication of users of registries would help.

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?

None was identified.

7.5 Jump-Start Plan

8 Property Rights of Personal Information

8.1 Description

What does the change look like?

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.

8.2 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 asses 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.

8.3 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.

8.4 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.

8.5 Jump-Start Plan

9 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.

9.1 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.

9.2 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.

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, may make security problems associated with standardized infrastructures less of an issue for organizations in the future.

9.3 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.

9.4 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 or 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.

9.5 Jump-Start Plan

10 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.

10.1 Description

What does the change look like?

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.

10.2 Inertia –

Why have we not done this before?

To some extent it 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, we still do not have acceptable solutions for securing the endpoints. 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.

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. It may take a long time, and be very costly to push out secure endpoints into the market.

10.3 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.

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.

One idea discussed was the use of cheap, secure devices used only for online banking. At scale, one could probably construct a simple one for $100-200 that would use the cell phone networks and connect directly to your bank. A hacker would have to spoof your device in order to fraudulently access your account.

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.

10.4 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.

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, which raises the issue of how to enable communications between networks this

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. Getting ISPs on board 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.

10.5 Jump-Start Plan

11 911 Cyber

11.1 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.

What does the change look like?

The idea was 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.

11.2 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.

11.3 Progress

What technologies are emerging that make 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).

11.4 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.

11.5 Jump-Start Plan

12 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.

12.1 Description

What does the change look like?

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.

12.2 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. 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.

12.3 Progress

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

None were identified.

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

None were identified.

12.4 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?

Nothing was identified.

What are the missing technical pieces?

Nothing was identified.

12.5 Jump-Start Plan

13 Other Ideas

The following ideas were also raised initially in the meeting, but either they were discarded, or the participants felt they were lower priority than the other ideas. We do however, include them here for reference.

- 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.

- FROM IDENTITY AUTHENTICATION TO BEHAVIOR AUTHENTICATION. We should stop trying to authenticate people and focus on authenticating behavior (stop confusing offline identity with online agency), at least where possible. Often it is not the people who are the problem, but rather their actions. And, while it is sometimes easy to obtain someone else's identity, performing an illicit act may be easier to detect.

- 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.

APPENDIX A: Acronyms

Acronym	Description
DRM	Digital Rights Management
GDP	Gross Domestic Product
ISP	Internet Service Provider
ITAC	Identity Theft Assistance Center
NRC	National Research Council
PKI	Public Key Encryption
ROI	Return on Investment
VLAN	Virtual Local Area Networks
VPN	Virtual Private Network