National Cyber Defense Financial Services Workshop
Report
"Helping Form a Sound Investment Strategy to Defend against Strategic Attack on Financial Services"
October 28-29, 2009

Hosted by: BITS, FSTC, and Financial Services Roundtable (on behalf of the Research and Development Committee of the Financial Services Sector Coordinating Council)
1001 Pennsylvania Avenue NW, Suite 500 South, Washington, DC

Sponsored by: National Science Foundation and Department of Homeland Security Science and Technology[1]

Organized by: National Cyber Defense Initiative

Edited by:
O. Sami Saydjari, Cyber Defense Agency, Inc., ssaydjari@CyberDefenseAgency.com
Salvatore J. Stolfo, Columbia Univ. Dept. of Computer Science, sal@cs.columbia.edu
Dan Schutzer, FSTC, dan.schutzer@fstc.org

4 February 2010

[1] Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the sponsoring or participating government organizations."

Executive Summary

The National Cyber Defense Initiative ad hoc group organized this workshop to better understand the nature of high-impact, large-scale attacks on the banking and finance sector, approaches to addressing those classes of attacks, and ways that industry, academia, and government can work together on such approaches. The National Science Foundation and Department of Homeland Security Science and Technology co-sponsored the event in cooperation with the Department of the Treasury and experts from the Financial Services Sector Coordinating Council for Critical Infrastructure Protection and Homeland Security (FSSCC).

The group concluded that high-impact, large-scale attacks that target the entire sector are theoretically possible and underanalyzed. A continuing dialogue on defending against such attacks and how to effectively address them in cooperation with government would be productive and useful. The group also concluded that banking and finance sector problems are unique and important and require basic research in modeling and analyzing large-scale interdependent financial systems and in constructing inherently recoverable distributed computation.

Summary of Key Recommended Research Directions

Analytical Models. Model large-scale banking and finance sector systems and their transaction flows and interdependencies using a new monitoring infrastructure that measures and analyzes the properties of these systems, including instability, security, hidden dependency, and cascading effects. Use such models to understand the infrastructure at a network and transaction level and the mapping between the two to inform the dependency analysis. Conduct simulations and exercises to better understand the impact of events. Initially focus on high-impact transaction systems to simplify the problem.

Collaborative Situational Awareness/Understanding. Define information-sharing requirements and methods between organizations (government, intra-industry, academic), taking into account the need for recovery and privacy (through anonymization techniques, for example) as well as the risk-reward calculation in sharing certain types of information. Specifically, define real-time information sharing with academia to obtain additional help and with intelligence agencies to improve situation understanding so that both communities can understand enough of what is happening that they can effectively respond to high-impact threats.

Resiliency. Develop methods to create more secure and resilient transactions under the load of high-impact attacks. Create inherently resilient architectures. Collaborate with academics in select spots within the secure transaction space. Develop data-centric protection strategies such that the data's integrity and provenance are preserved despite attacks.

Authentication. Improve identification and authentication capabilities for people, devices, and digital objects in such a way to make high-impact, large-scale attacks more difficult to accomplish.

Leverage Research. Leverage existing research for possible application to banking-and-finance-sector-critical problems such as those above and those identified in the FSSCC Research Agenda. Develop models for more effective interactions between banking and finance sector experts and academics interested in these classes of problems. Leverage existing models with government as a catalyst. Establish a standing "semiformal" advisory council of academic researchers who would meet regularly with and advise and assist the FSSCC R&D Committee to set research priorities and identify projects that may be conducted in collaboration with academia and government.

Table of Contents

1. Introduction 1
   1.1. Background 1
   1.2. Purpose and Goals 2
   1.3. Participation 2
   1.4. Intended Audience 2
   1.5. Opening Problem Statement 3
2. Establishing Common Ground 3
   2.1. Understanding the Banking and Finance Sector 3
   2.2. A Research and Development Agenda Foundation 5
   2.3. Unique Characteristics of Banking and Finance Sector 7
   2.4. Importance to National Leadership 7
3. Problem Elaboration 8
   3.1. Focus 8
   3.2. Changing Threat Environment 9
   3.3. Results from a Banking and Financial Sector Exercise 12
   3.4. Hypothetical Attack Scenarios 13
4. Challenges and Approaches 14
   4.1. Prevention 15
   4.2. Detection and Response 15
   4.3. Recovery and Reconstitution 16
5. Innovative Industry-Government Partnership Models 17
   5.1. Exemplary Partnership Models 17
   5.2. Discussion 20
6. Next Steps 20

Appendix 1. National Cyber Defense Financial Services Workshop Agenda 22
Appendix 2. TRUST Program—An Example of Related Research 24
Appendix 3. Prevention Details 25
Appendix 4. Detection and Response Details 28
Appendix 5. Recovery and Reconstitution Details 32
Appendix 6. Organization and Attendees 39

1. Introduction

This report documents a one-and-a-half-day public-private financial services industry workshop on national cyber defense sponsored by the National Science Foundation (NSF) and the Department of Homeland Security Science and Technology Directorate (DHS S&T). The workshop is one in a series organized by the National Cyber Defense Initiative[2] steering committee. The workshop took place 28-29 October 2009 at BITS headquarters in Washington, DC.

The report is organized as follows. This introduction section represents materials prepared by the National Cyber Defense Initiative ad hoc steering committee and the organizers of the workshop. It is intended to set the stage for the workshop, not to represent workshop consensus. Also, some items on the original goal list were discussed more than others based on available time and attendee interests. Section 2 on establishing a common ground captures presentations and discussions to help participants understand the context of the sector and its issues. Section 3 on problem elaboration provides more detail on the nature of the threat. Section 4 summarizes the results of three working sessions on preventing, detecting and responding, and recovery and reconstituting systems in the face of a high-impact attack; details are provided in the appendices. Section 5 discusses potential partnership models between the financial industry, the government, and academia. Section 6 covers next steps and potential ways forward.

1.1. Background

Large portions of our country's economic, industrial, social, and governmental functions now depend on a cyber infrastructure assembled from readily available commercial information system components[3]. Much of this infrastructure is organized to tolerate random failures and outages, including targeted attacks against specific institutions, but could fail under concerted attack against many of our country's economic, industrial, social, and governmental functions. Leadership is needed to substantially reduce this serious vulnerability. Many efforts are currently underway to begin to address these issues at the national level. Representatives from leading banking and finance sector firms participated in this workshop, which was intended as a forum for members of the financial industry to contribute input to multiple planning and strategic efforts and to define actionable plans to take back to their organizations.

---

[2] NCDI is an ad hoc group of security professionals dedicated to helping government formulate sound research strategies for the key problems of cybersecurity. See http://ncdi.nps.edu/ for details.

[3] See, for example, "Trust in Cyberspace," The National Academies Press. 1999, p12. The President's Commission on Critical Infrastructure Protection (PCCIP) study reached a similar conclusion in 1997 in their report "Critical Foundations: Protecting America's Infrastructure," October 1997.

1.2. Purpose and Goals

The workshop's goal was to develop a shared view of an attack-resistant and attack-tolerant cyber infrastructure and the specific steps needed to reach that vision. Specifically, participants sought to

1. Understand the changing threat environment, the increasing possibility of extraordinary attacks mounted by nation-state adversaries for strategic gain, and the growing sophistication of criminal organizations for financial gain.

2. Review FSSCC R&D priorities and National Cyber Leap Year priorities. Discuss research that supports these priorities and specific game-changing technologies and processes that may apply to the banking and finance sector.

3. Discuss ways to incorporate new innovation partnership models. Create processes whereby the banking and finance sector, the research community, and the U.S. government can produce relevant solutions to current and long-term challenges facing the sector and plan for a more efficient transfer of these research products to industry.

4. Produce a public report and plan to help inform government of needed R&D resources to defend the financial services infrastructure[4].

1.3. Participation

Technical and business leaders in the banking and finance sector, government (including representatives from the White House Office of Science and Technology, DHS S&T, NSF, and the Department of the Treasury), and academic researchers participated in the workshop. Attendees had a deep understanding of the technologies and operations and of the significant failures that have happened to date. To assure focus and productivity, the meeting was limited to approximately 40 participants. The workshop was chaired by Sal Stolfo (Columbia University and representing the National Cyber Defense Initiative steering committee), Dan Schutzer (President of the Financial Service Technology Consortium and representing the FSSCC R&D Committee), and Brian Peretti (the U.S. Department of the Treasury's Financial Services Critical Infrastructure Program Manager).

1.4. Intended Audience

This report has three main audiences: (1) government R&D planners, (2) financial services industry strategic leaders, and (3) academic leaders with the capability to help the banking and finance sector. Our goal is to provide input into the government planning process and to identify key strategic problem areas where additional effort may be placed for the nation's benefit. For the financial services industry, the report is intended to give some insights and approaches for planning a defense against the large-scale attacks that are within their purview. For the

[4] Background and related reports are provided at http://ncdi.nps.edu/. The site provides a parallel report done with the information technology industry on which the banking and finance sector depends.

academic community, we expect to help them better understand the nature of the challenges facing the financial services industry and thereby help them direct their research at those challenges.

1.5. Opening Problem Statement

Financial institutions have done an effective job in managing operational risks and in responding to increasing threats from cyber crime. While losses for some types of payment channels are growing, banking and finance sector participants believe they currently are tolerable. However, banking and finance sector and government experts are increasingly concerned with cyber threats and want to better understand the interdependencies and strategies for mitigating these changing risks. There is particular concern with the extraordinary attacks that nation-state adversaries could mount for strategic gain.

2. Establishing Common Ground

The materials presented and discussed at the workshop provided common ground in understanding the nature of the operations of the banking and finance sector.

2.1. Understanding the Banking and Finance Sector

This section is intended to help readers better understand the complexity and nature of the banking and finance sector. The actual material below is extracted from a portion of a report intended to provide similar context[5].

The U.S. banking and finance sector deeply affects the world economy. It is complex and diverse both in the varying sizes of participants—from the largest financial institutions with assets greater than $1 trillion to the smallest community banks and credit unions with a few million in assets—and involved in a variety of functions ranging from deposit-taking to lending to brokerage to insurance. The industry is organized by a unifying mission to ensure the sector and its institutions maintain their efficiency and continuity.

An important part of the profile of the banking and finance sector is that it is primarily owned and operated by the private sector, whose institutions are extensively regulated by federal, and in many cases, state government.

However, banking and finance sector and experts are increasingly concerned with increasing cyber threats and want to better understand the interdependencies and strategies for mitigating these changing risks.

Depository institutions of all types (banks, thrifts, and credit unions) are the primary providers of wholesale and retail payment services, such as wire transfers,

[5] See http://www.dhs.gov/xlibrary/assets/Banking_SSP_5_21_07.pdf

checking accounts, and credit and debit cards. The institutions use and/or operate the payments infrastructure, which includes electronic large value transfer systems, the Automated Clearing House (ACH), and automated teller machines (ATMs). These institutions are the primary point of contact for many individual customers. In addition, they provide extensions of credit, such as mortgages and home equity loans; collateralized and uncollateralized loans; and lines of credit, including credit cards.

The depository institution system is supported by electronic payment systems that link these institutions to one another and to their customers. Examples of these systems and networks are the many regional/national ATM networks that permit consumers to access their funds from ATM sites, credit-card sponsors, and ACH operators. Businesses and consumers use ACH payment systems to make recurring payments.

This sector is overseen by both federal and state regulatory authorities as well as self-regulatory organizations. Because they're part of one of the largest regulated sectors, financial institutions answer to various watchdogs that provide oversight, guidance, and examinations. These financial regulators work together through the Financial and Banking Information Infrastructure Committee (FBIIC) to coordinate efforts with respect to critical infrastructure protection issues.

To share information, the private sector has multiple options, including the FSSCC, the Financial Services Information Sharing and Analysis Center (FS-ISAC), regional coalitions (e.g., ChicagoFIRST), and associations (e.g., American Bankers Association, BITS, SIFMA). These associations work together to share information and build relationships among financial institutions as members of the critical infrastructure.

To meet the shared vision, the FBIIC-FSSCC collaboration has three primary goals with a primary focus on security, including to

- maintain its strong position of resilience and risk management via redundant systems in the face of myriad international, unintentional, manmade, and natural threats;
- address and manage the risks posed by the sector's dependency on the communications, information technology, energy, and transportation sectors; and
- work with the law enforcement community, the private sector, and our international counterparts to increase the amount of available resources dedicated to tracking and catching criminals responsible for crimes against the sector, including cyber attacks and other electronic crimes.

An important part of the profile of the banking and finance sector is that it is primarily owned and operated by the private sector, whose institutions are extensively regulated by federal and, in many cases, state government. In addition to these public-sector entities, self-regulatory organizations (SROs) oversee many aspects of market trading, and most portions of the financial services industry have

established stringent data security requirements for the processing of payments and the protection of sensitive customer information.

Discussion

It is important to understand the consequences of any of the major financial networks[6] being disabled for even a day. This is under continuous review by the government. One possible useful research idea would be to map the vulnerabilities of the ACH network, both to improve operations and inform research challenges. Alternatively, the banking and finance sector could share how network testing is currently designed and executed within institutions and then compare this against the state of the art. To better engage the research community, academics will need to be better educated on the financial services industry.

There is a need to apply research to better understand financial system behavior and learn ways to mitigate risk under high-stress attack conditions.

Today, the banking and finance sector invests in system resiliency and planning. For example, after 9/11, a sound practices white paper was created whereby certain organizations were encouraged to locate backup sites in different geographic locations run by different people[7]. An important research question is whether this investment can be reduced with more effective and efficient approaches that could lead to better resiliency for less.

The global interconnectivity of banking and finance sector systems has reached a level of complexity and interdependency that is extraordinarily difficult to fully analyze for predicting cascaded consequences or best means of control in the event of a complex well-orchestrated attack. There is a need to apply research to better understand financial system behavior and learn ways to mitigate risk under high-stress attack conditions.

2.2. A Research and Development Agenda Foundation

This section is based on a talk given by John Carlson and was intended to educate workshop participants on prior work done to establish a research agenda for key challenges in the banking and finance sector. The material below is quoted directly from the executive summary of the "Research Agenda for the Banking and Finance Sector"[8]. There was no attempt at the workshop or in this report to merge the

---

[6] The word "network" is used in the larger sense of systems, communications, people, and processes performing a critical financial service such as transferring funds, stocks, and bonds and settling payments.

[7] See "Interagency Paper on Sound Practices to Strengthen the Resilience of the U.S. Financial System" (2003) issued by the Board of Governors of the Federal Reserve System, the Office of the Comptroller of the Currency, and the Securities and Exchange Commission. See also the FFIEC Business Continuity Planning Booklet (which has been updated several times, with the most recent version published in 2008).

[8] See https://www.fsscc.org/fsscc/reports/2008/RD_Agenda-FINAL.pdf

agendas, so there is some overlap between the material presented here and later in the report.

The following are the top priorities of the Research and Development Committee of the Financial Services Sector Coordinating Council for Critical Infrastructure Protection and Homeland Security.

Advancing the State of the Art in Designing and Testing Secure Applications. Software applications are complex and often insecure and thus introduce vulnerabilities. Historically, acquisition requirements have favored functionality over security, which has led to a state of software development that often does not emphasize security. Financial institutions have begun demanding more secure application development. Because financial institutions often cannot be sure that their applications are secure, they must develop and implement costly and inefficient compensating controls. Financial institutions need a robust, effective, affordable, and timely security testing methodology, and practice to gain the confidence required to deploy application software into sometimes-hostile environments for purposes of practical and appropriate risk management. Research is needed to develop effective procurement standards, software developer education, and testing guidelines. In addition, research is needed to develop tools for producing, measuring, and testing secure application software.

More Secure and Resilient Financial Transaction Systems. The financial services industry is dependent upon information technology infrastructure, much of which is owned and operated by third parties outside the financial services industry. This infrastructure is constantly under attack by hackers and identity thieves who seek to exploit vulnerabilities in networks, devices, and applications for financial gain. Research is needed to better understand these threats, improve the security and resiliency of the financial transaction infrastructure, to enhance the protections available to prevent the increasingly common downloads of malware by criminal elements that bypass existing defenses such as antivirus and antispyware, and to develop metrics to evaluate the resiliency of the information technology infrastructure.

Enrollment and Identity Credential Management. The financial services industry depends on the ability of financial institutions to identify, authenticate, and authorize customers before accessing information and conducting transactions through remote channels where direct human interaction is not possible. Inadequate controls can leave financial institutions and their customers vulnerable to attacks. Research is needed to study how to make the identity management process better and less susceptible to social engineering attacks.

Understanding the Human Insider Threat. Financial institutions must trust employees who have access to sensitive personal and financial information. Current strategies for identifying trustworthy candidates rely upon historical methods such as background and credit history checks as well as identity confirmation. Such methods often do not sufficiently identify insider-fraud perpetrators ahead of time and can be costly to maintain. Research is needed to develop holistic solutions to the insider-authentication problem, including the development of a data frame to predict the likelihood of insider attacks based on differing scenarios, or the development of continuous, unobtrusive monitoring to reduce the risks posed by insiders.

Data Centric Protection Strategies. To maintain trust and the integrity of data, financial institutions must protect sensitive data but also share it with third parties, such as merchants and processors. Increasingly, devices and networks are vulnerable to malicious code or data breaches. Research is needed to develop secure data file and document tagging technologies to classify information and to enforce rules on access so that sensitive information is protected as intended by its original owner, regardless of where it traverses.

Better Measures of the Value of Security Investments. Traditionally, investment decisions surrounding security implementations have followed a “return on investment” (ROI) decision-making process. The ROI model does not always fit well into the security space because it can be difficult to quantify hypothetical losses averted through increased security. The creation of cost-benefit models for security spending might be more appropriate because they would take into account intangible benefits such as increased customer confidence and decreased brand exposure. Research is needed to quantify the costs and benefits of security investments using models that are understood by financial risk managers.

Development of Practical Standards. The financial services industry relies on numerous standards and practices but has not succeeded in developing quantifiable measures for how these standards and practices reduce risk and enhance resiliency of critical infrastructures. Research is needed to measure the impact of standards and practices.

2.3. Unique Characteristics of Banking and Finance Sector

The workshop's participants discussed the unique characteristics of the banking and finance sector and how these attributes impact the formulation of strategies and the development of technologies, tools, and techniques to counter cyber attacks. These characteristics include that it's

- digital centric (the whole industry's value is in the data);
- global in scope;
- universally regulated in all countries;
- highly interconnected;
- characterized by high transaction speed and volume;
- highly concentrated in value;
- technically sophisticated in terms of systems and sector;
- a hybrid of business-to-business and consumer-to-business activity;
- clearly driven by the bottom line;
- got a pre-existing culture of separation of duties;
- heavily relied on by all other sectors for funding, credit, and liquidity; and
- highly trusted by citizens to function correctly.

2.4. Importance to National Leadership

The Honorable Aneesh Chopra, the Chief Technology Officer and Associate Director for Technology in the White House Office of Science & Technology Policy, provided comments on the second day of the workshop in which he

- emphasized the importance of the workshop's study area to President Obama and how important innovation strategy is in general,
- recommended that the group give attention to both longer- and shorter-term jump starts, and
- identified the strong organizational commitment by the Secretary of Homeland Security to get moving on improving critical areas such as improved information sharing.

During questions and answers, Mr. Chopra made several observations:

- The banking and finance sector is uniquely important because every sector relies on it for trusted transactions. Other sectors are both suppliers and customers of its capabilities, creating a close dependency. People have grown to deeply trust the banking and financial sector for essential services such as protecting their

identity. Lastly, the sector has the organizational capacity to address difficult problems in that it has solved hard problems in the past.

- Market incentives need to be changed to encourage adoption of new technology.
- The government needs to invest in game-changing technology to change markets.
- The government needs to help create an insurance market to encourage the adoption of advanced resilient technology.
- Internationally, cybersecurity is a top priority issue on the minds of other government leaders, so the U.S. is actively engaged in discussions with other countries on key topics in cybersecurity.

3. Problem Elaboration

This section is intended to elaborate on the nature of the problem to help identify research challenges and possible approaches.

3.1. Focus

The focus of the meeting was on high-impact, large-scale attacks, not the ordinary problems that afflict institutions on a daily basis. "High impact" means that the consequences are strategic in nature, and "large-scale" means that the high impact is judged as it affects the entire sector, not just a single institution. This focus made the meeting unique and represents a critical gap area that the sector must address. We must continue to work hard to stay on this focus.

The focus of the meeting was on high-impact, large-scale attacks, not the ordinary problems that afflict institutions on a daily basis. This focus made the meeting unique and represents a critical gap area that the sector must address.

Strategic concerns among participants included building fundamentally more secure and resilient business applications on more survivable security architectures with substantially higher software assurance and better-controlled provenance. For high-impact attacks that might succeed, there was concern about the need for dramatic improvements in fast detection of a range of sophisticated attacks including insiders, better ways of assessing cascaded impacts on the sector, and more effective responses to minimize damages. If high-impact attacks cause serious damage, there was concern for improvements in disaster recovery to maintain core sector functioning while quickly bringing full functioning back in priority order. Foundationally, participants saw a need for significant improvements in the ways of measuring systems at all levels (for example, network, operating systems, and applications) to better analyze status and more effectively manage the system under stress. At a policy level, there was interest in how national policy plays into sector cybersecurity and longer-term strategies to secure e-commerce.

3.2. Changing Threat Environment

The following material is based on a workshop talk given by Jane Carlin. The intention was to help participants understand the nature of the threat and how it is escalating.

There are an increasing number of profound cyber threats across the banking and finance sector, but data about cyber attacks on the industry is scarce, partially because there is no standard method of tracking these attacks. The FSSCC has separated cyber threats into five areas: application security, identity theft, mobile devices, supply chain, and undersea cables. Table 1 provides examples of potential attacks in each of the areas.

Table 1 Cyber Threat Examples

Application Security	Identity Theft	Mobile Devices	Supply Chain	Undersea Cables
• Unauthorized access to sensitive applications • Toxic combinations, e.g., execution and approver • System shutdowns • Software standards	• Customers' personal or private information used for financial gain • Unauthorized access to critical infrastructure • Credentials vs. access	• Paralyzing communication during a cyber attack • Access vulnerabilities exposed to smart devices • Data leakages • Internet/telecom congestion	• Embedded backdoors • Reliance on third-party vendors across borders • Quality controls at interaction points • Concentration	• Component failure • Landing station concentration risk • Physical attack resiliency (poor design)

Communication within the sector and between other sectors is crucial. Figure 1 shows response gaps and demonstrates that sector-wide coordination is important to contain infrastructure threats. The response gaps appear most prevalent in the transition/elevation from firm- or institution-specific threats to sector- or country-wide attacks. Furthermore, identification of patterns in numerous low-impact events is important for early detection and prevention, reducing the risk of escalating severity and pervasiveness.

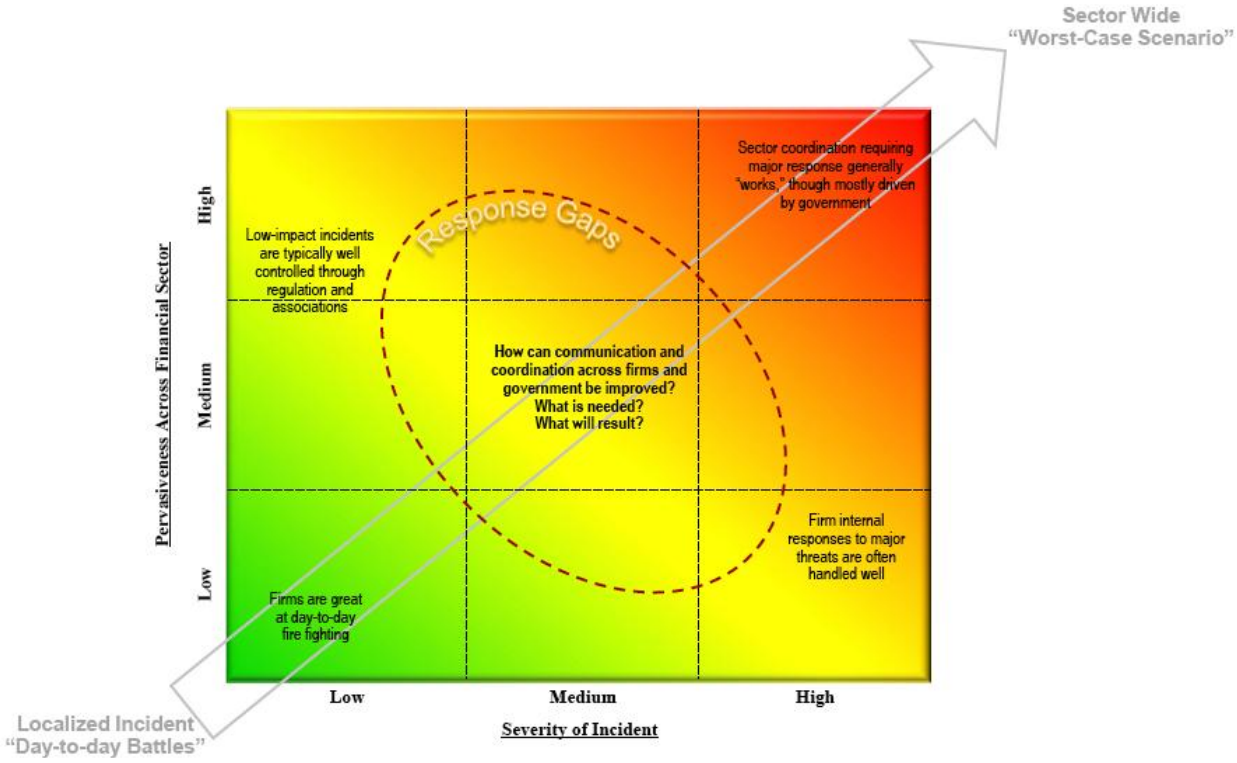

Two key items must be determined: the differences between an attack and a persistent attack and between malicious and non-malicious attacks. The Pakistan YouTube misdirection and the problem with Sweden dropping off the Internet[9] were both caused by accidental misconfigurations. Neither event was an attack, yet the impact was high. One approach is to focus on the impact (e.g., the paralyzing effect) rather than whether it was accidental or intentional.

There are two classes of the problem—internal vs. external to the system (such as undersea cables)—and approaches to each are likely to differ because of the differences in the level of control one has over the systems involved. For example, the Internet is an example of an external system with no single entity charged with ensuring its continued operation under attack because there is no central control (a point not well understood by the general public and some leaders).

Discussion

The workshop talk quickly became highly interactive with the participants contributing additional material and viewpoints. The following material reflects the conversation that took place.

One participant observed that contingency operation modes are useful but not sustainable. During Y2K preparations, the issue of backing up information was

[9] See, for example, http://www.v3.co.uk/v3/news/2251165/dns-gaffe-cripples-swedish

addressed, so that if necessary, businesses could revert to a lower-tech version for high-value transactions (and temporarily ignore lower-value transactions). There are no survivable procedures or resources to continue this on an everyday basis.

Each firm has a unique "self-help" model during times of crisis—individual guardianship. Communication among firms is limited as firms wait for the dust to settle.

As the exchanges had yet to open preceding the September 11, 2001 terrorist attacks against New York City, the impact on the markets was not as great as it might have been otherwise. . The Securities Industry and Financial Markets Association (SIFMA; at the time, the Securities Industry Association [SIA]) even ran tests to make sure that no one was excluded from market activity. Government and financial services leaders gathered people together for collaborative problem solving. The key was communication among agencies and organizations as they worked to address challenges facing the financial sector. Public confidence in the integrity and trustworthiness of the financial markets was maintained. If confidence in the system eroded, the impact of these attacks could have resulted in additional harm to both the sector and the economy as a whole.

Attacks that impact liquidity could also result in a profound impact on the general economy.

The banking and finance sector could be impacted in lots of ways as a side effect because it depends heavily on many sectors such as telecommunications, information technology and network service providers.

Attacks that impact market liquidity[10] could also result in a profound impact on the general economy. It's helpful to think about the types of attacks and how they affect important capabilities such as liquidity—for example, it might help to identify the services and infrastructures on which financial institutions depend. It would also be useful to leverage current, ongoing threat analyses to understand how they intersect with the financial services industry. Analysis of the Societe Generale attack, for instance, shows that it was actually an attack on the separation of duties and other security mechanisms, including shared passwords, poor recovery, poor administrative controls, etc.[11].

The banking and finance sector has invested a lot of energy in developing best practices and working with other sectors to address software security concerns, as evidenced by the papers that BITS, SIFMA, and FSSCC have published during the past 12 years. Still, there is a need to design more resilient and secure business

[10] Liquidity refers to the ability to sell an asset without significant loss of value. Market closures or degradations can impact liquidity because it can reduce the potential number of buyers and/or sellers for an asset.

[11] See "Security Lessons Learned from Societe Generale" - http://www.computer.org/portal/web/csdl/doi/10.1109/MSP.2008.71

products that mitigate the effects of both insider and outsider attacks. A unique focus of the banking and finance sector is on managing risk to make business determinations. It is important to learn how to better model and manage the complex interplay of the risks involved to minimize the overall risk.

At the workshop, there was discussion of the "middle ground" of risk, which refers to the middle space between overt acts of war against our infrastructure and ordinary criminal hacking. The Department of Defense has clear responsibility for overt acts of war. Law enforcement has clear responsibility for dealing with criminal activities. Major strategic attacks[12], particularly those against multiple institutions within the sector, fall in the middle ground, so responsibility for protecting against them is not clear from a policy perspective. Similarly, there are areas of critical infrastructure cybersecurity for which there are no defined owners. One example mentioned earlier is that no specific body is actually responsible for security on the Internet. This is a much-neglected area worthy of further discussion, including a prioritization based on risk probability and impact.

Because the banking and finance sector is highly regulated, we have a unique opportunity for government and industry to cooperate and improve resiliency against high-impact attacks.

Currently, the banking and finance sector is not completely aware of the full range of potentially useful academic research taking place today, indicating a need for improved communication between the sector and the academic research community. Such improved communication would give the academic community an increased understanding of the needs of the sector and the practical obstacles to adoption of their ideas, and the sector would learn about how to apply the research relevant to the challenges that it faces. As part of this dialogue, it is important to identify how the sector's problems relate to classical problems, but also the uniqueness that justifies a specific, focused research investment by the government. Because the banking and finance sector is highly regulated, we have a unique opportunity for government and industry to cooperate and improve resiliency against high-impact attacks, both in the short term (by applying existing research) and in the long term (by focusing government research investments in the key challenge areas).

3.3. Results from a Banking and Financial Sector Exercise

This section is based on a workshop talk by Mark Clancy that intended to inform participants of an example of activity by the banking and finance sector to better understand the nature of the problem.

Recently, the FSSCC and some members of the FBIIC designed a voluntary exercise for individual financial service firms, government regulators, and financial utilities and exchanges. The simulated exercise focused on a scenario that spanned five days

[12] By this, we mean attacks that cause serious damage to the U.S. economy or national security interests.

of cyber attacks on systems, with the final attack coinciding with a "triple witching hour.[13]"

The objective was to understand the direct and indirect cascaded impact to both security and business operations of an attack that affects only a portion of the sector. Malicious code in a widely used financial application within that portion of the sector was posited. The malware had the effect of causing payment mismatches within the transaction system.

The result of this exercise is to have an after-action report in late January/early February 2010 that describes in detail what the exercise examined and its findings. Learning about the degree of interconnectedness within the financial industry and the resulting propagation of damages was a key lesson. This exercise and others like it are increasingly important to the financial industry to identify sector-wide risks that require mitigation.

The overall discussion during the workshop centered on "what's next"? Attendees discussed the idea of bringing the academic/research community into future exercises. Participant anonymity could be an issue, so some sort of non-disclosure agreement will likely be needed.

3.4. Hypothetical Attack Scenarios

To better motivate the challenges and potential solutions, the workshop participants felt it necessary to enumerate some hypothetical attack scenarios to illustrate what a high-impact attack might look like. The group brainstormed a list of scenarios to better understand the spectrum of attacks:

1. a pervasive failure in a major software application that handles the flow of financial messages;
2. supply-chain corruption by a disgruntled employee or insider to enable backdoor access in a widely used operating system or major application;
3. massive infrastructure attacks to some shared services such as data network providers (specifically, a denial-of-service attack on the Internet to disrupt business-to-consumer transactions or private data networks to disrupt business-to-business operations);
4. tainting of transaction data in a subtle way so that it's not detected for an extended period of time (specifically, corrupting integrity in such a way that end-of-day settlement cannot occur);
5. exploitation of an unknown dependency to cause unexpected large-scale cascaded consequences;
6. coordinated human infiltration and sabotage across many different key financial and supporting institutions;

[13] This is the last hour of the stock market trading session on the third Friday of every March, June, September, and December when three kinds of securities expire: stock market index funds, stock market index options, and stock options.

7. physical attack on an area of concentrated risk, such as taking out a backup site farm for several institutions;
8. physical and cyber attacks on backup and recovery processes (specifically, by propagating corrupted data into recovery infrastructure so that recovery from backup is useless);
9. overwhelming multi-institution recovery sites by attacking a large number of institutions that use those sites, creating a flood of recovery requests that will not be able to be serviced because of insufficient capacity;
10. a direct attack on trust to create a run on banks, such as a media attack, done in coordination with a military campaign, to try to get people to withdraw money;
11. an attack through the IT monoculture and central system management;
12. a strategic attack on clearinghouses, large financial institutions, or on service providers;
13. pervasive attacks on the systems of financial institution partners leading to a lack of trust and widespread disconnections to minimize damage and confusion over how to regain that trust, significantly suppressing commerce;
14. trust attacks from adversarial foreign nations or organized crime (by introducing large bogus transactions, for example);
15. large-scale intellectual property loss; and
16. governments around the world overreacting to certain events and causing unnecessary impact.

One thing that all these attacks have in common is that they pose systemic risk beyond risk to any single institution.

This is just a sampling of attacks. With more time, the workshop participants would have considered the motivations behind these scenarios, attacker characteristics, and the types of systems attacked, and then group them to help drive challenges and approaches. This could be a useful action for a future workshop.

One thing that all these attacks have in common is that they pose systemic risk beyond risk to any single institution.

4. Challenges and Approaches

During the workshop, experts from academic institutions, financial institutions, and government agencies were assigned to three groups on prevention, detection and response, and recovery and reconstitution. Their tasks were to define the key research challenges and to sketch some possible approaches to the key challenges. The next three subsections reflect summaries of the outputs of these three working groups. More details are contained in appendices 3, 4, and 5, respectively.

4.1. Prevention

Prevention is the collection of technologies and processes aimed at stopping attacks from being attempted and, if attempted, from succeeding. The primary focus was on 11 areas:

- improving understanding of the nature of attacks to better understand how to design systems and processes that can prevent them;
- improving identity management systems to prevent attacks by means of impersonation and to support more fine-grained least privilege authentication and access controls;
- sharing information with a focus on early discovery of planning and reconnaissance efforts prior to attack by developing better sensors and algorithms for pre-attack detection that address needs for privacy and incentives for information sharing;
- making someone accountable to worry about attacks not directed at a single institution, processor, or utility but at attacks aimed at low probability of detection, aimed at destabilization of the system, and bringing on the onslaught of systemic damage;
- building a diversified and distributed system that can better resist attacks because it represents a moving target (this also includes the ability of systems and processes to take over for other systems and processes that fail or lose their integrity);
- applying data provenance and data-centric protection strategies to allow for continued data and transaction integrity while operating in the midst of an untrusted environment;
- creating better metrics and methodologies to understand the impact of attacks on the business and to quantify return on investment for improved designs (this includes the ability to do simulation exercises to analyze the impacts of events and identify the systemic risks and social consequences of various attack vectors);
- developing new security metaphors for improved usability that can overcome the current lack of good security behavior;
- automatically validating and verifying security system configurations and the ability to restore a system to a validated system configuration when needed;
- developing deterrence strategies and technologies that could discourage attacks; and
- formalizing relationships with academia to optimize relevant research and technology transfer of prevention technologies and approaches.

4.2. Detection and Response

Detection and response are a key part of dealing with cyber attacks. However, a key intermediate step is missing: recognition, which is an understanding that a sequence of detected events represents a serious attack needing response. Having sensors in the right places is crucial, but if the situational awareness systems (both automated and human) do not have proper recognition, those sensors are worthless. As an

example, there were adequate sensors in the Societe Generale attack[14], but the humans in the loop ignored the blaring horns.

A large focus for workshop participants was the need for greater data sharing between the financial industry and academia, with the goal of being able to gain greater situational awareness, possibly using existing Department of Defense data-sharing programs as a template. The existing data sharing through the FS-ISAC and the FSSCC has been insufficient; the industry needs to share more data and on a timelier basis so as to identify and stop attacks sooner.

Academics are largely unaware of how the financial industry functions or its requirements; a key area of discussion was how to educate the educators.

The financial industry has some unique capabilities and requirements compared to other industries. Academics are largely unaware of how the financial industry functions or its requirements; a key area of discussion at the workshop was how to educate the educators, with the side effect that today's computer science graduate students trained in the functioning of the financial industry will be tomorrow's employees.

4.3. Recovery and Reconstitution

In this area, successful attacks causing significant damage are assumed, so the question is how to recover from them.

The top-level research challenges involve discovering methods of performing distributed recoverable computations, creating analytical models for better situational understanding and control, learning what and how to share information for multi-tiered multi-institutional recovery, restoring user trust in sector systems, and developing robust recovery processes that tolerate the presence of insiders.

Approaches to these top-level challenges were many and varied, including

- Create intentional redundancy in financial-transaction-recorded-information among all involved parties to create a more robust and recoverable distributed computation[15].
- Create a series of different analytical models to examine different aspects of the problem and develop libraries of analytical capability such as a spectrum of perturbations (a series of types of attack and failures, for example) that can be applied to these models to improve understanding of their operating characteristics (cascaded damages, for example) under high-impact stresses.

[14] Cited earlier.

[15] For example, have each party in the transaction record enough information such that the entire transaction can be replayed in the future as needed; have any reports that go to third parties, such as regulators, do the same.

- Develop concentric rings of criticality of information sharing for bringing up concentric rings of recovered system capabilities.
- Develop stress and validation tests as criteria to readmit institutions back into trusted financial systems and build public confidence in the system.
- Develop a heavily audited recovery processes with built-in indications and warnings of attempts to subvert the recovery process.

5. Innovative Industry-Government Partnership Models

This section documents innovative ways that industry and government can work together to address challenges and implement strategies. The primary content comes from a talk given by Doug Maughan from DHS S&T. A discussion section follows.

5.1. Exemplary Partnership Models

This section is based on a workshop presentation that was meant to help participants understand the range of possible public-private partnership models that may be useful in moving forward on working the banking and financial sector challenges.

DHS S&T funds many different types of partnership models with industry and academia, where the government acts as facilitator but lets industry and/or academia take the lead. Example models follow:

- S^2ERC – Security and Software Engineering Research Center
- I3P – Institute for Information Infrastructure Protection
- SIF – System Integrator Forum
- ITSEF – IT Security Entrepreneur Forum
- LOGIIC – Linking Oil and Gas Industry to Improve Cybersecurity
- DECIDE – Distributed Exercises
- PPISC-ES – Payment Processing Information Sharing Council – Enhance Security Working Group
- TCIP – Trustworthy Cyber Infrastructure for Power

S^2ERC is an academic consortium at Ball State University. The model here is for both industry and government to provide funding for common foundational problems.

The **I3P** academic consortium is a government-funded collection of universities working on high-priority research projects.

SIF and ITSEF were created to help match entrepreneurs and small businesses with necessary technologies and some of the key problems in industry needing solutions. They also provide a source of relevant problems for new research initiatives that industry can propose to the government.

LOGIIC was an important success story, an extensive interaction between the government and the oil and gas industry that may be highly relevant for the banking and finance sector. The project began when ChevronTexaco approached DHS S&T in March 2004 about possible opportunities to secure oil and gas cyber infrastructure. Ensuing discussions determined that this should be done sector-wide, with a workshop convened in July 2004 in Washington, DC. The outcome was to determine if government and industry could work together to (a) establish a SCADA testbed and (b) determine a working model for future research and development activities. Industry partners agreed on a technical project focus in April 2005, and the project officially started in July 2005. The group invited technology providers to show their capabilities between August and September 2005, from which industry selected the winning candidates. The project was presented to the oil and gas industry in September 2006; a LOGIIC DVD describes the program in more detail and highlights its success.

LOGIIC is a model for government-industry technology integration and demonstration efforts to address critical R&D needs. Industry contributes requirements and operational expertise, project management, and product vendor channels, and DHS S&T contributes a national security perspective on threats, access to long-term security research, independent researchers with technical expertise, and testing facilities. Chevron and DHS each invested $750,000, which was paid for by participants and includes vendor products and individuals' time.

The goal was simply to reduce the vulnerabilities of oil and gas process control environments by correlating and analyzing abnormal events to identify and prevent cybersecurity threats. The approach was to identify new types of security sensors for process control networks, adapt a best-of-breed correlation engine to this environment, integrate and demonstrate it in a testbed, and transfer the technology to industry.

Today, LOGIIC is sustained by the International Society of Automation (ISA) Automation Federation, with DHS S&T maintaining a supporting role through a Cooperative Research and Development Agreement with the federation.

Companies involved in LOGIIC and similar activities view it as pre-competitive R&D for the benefit of the entire sector.

LOGIIC was an important success story, an extensive interaction between the government and oil and gas industry that may be highly relevant for the banking and finance sector.

Similar to its work with LOGIIC, DHS S&T is now working with several commercial payment processing institutions to examine how to take the value out of data for payment processing systems.

DECIDE is building a process to do exercises at the institution, market, and sector levels, with the ability to protect information as necessary. The idea is to enable enterprise decision-makers to think through responses to operational disruptions of market-based transactions across networks: sector(s), market(s), and institution(s). The aim is to provide a dedicated exercise capability for several critical infrastructures in the U.S. Enterprises will be able to initiate their own large-scale exercises, define their own scenarios, protect their proprietary data, and learn vital lessons to enhance business continuity, all from their desktops

DECIDE is building a process to do exercises at the institution, market, and sector levels, with the ability to protect information as necessary.

The concept has been reviewed by and developed with input from experts at various banking and finance sector institutions. The FSSCC R&D Committee has organized a user group of subject-matter experts (SMART team) paid by their respective financial institutions to support the project over the next three years.

TCIP is a research center focused on security and resiliency of the power grid. It was originally funded by the NSF at $1.5 million per year for five years with additional support from the Departments of Energy and Homeland Security. More recently, the Departments of Energy and Homeland Security Science and Technology (S&T) have funded TCIP for an additional five years at $18 million. The center involves 20 senior investigators from the University of Illinois at Urbana-Champaign, Washington State University, Cornell University, and Dartmouth University, and includes an active industry advisory board of 35 owners, operators, and vendors. TCIP creates a sustained interaction that helps academics understand the industry's key problems and focus related research on the power grid. This seems similar to what has been called for by this workshop's participants.

TCIP creates a sustained interaction that helps academics understand the industry's key problems and focus related research on the power grid.

Which of these partnership models is best for industry group depends on its purpose:

- What does the group want to accomplish? Does the group want more formally organized information sharing, with a designated owner? Is this a known technology exploration and evaluation, or does it require new R&D?
- What's the government's role here, depending on the answer to the first question?
- What "formal agreements" does the group have, and does it believe it needs others?
- Are all the "stakeholders" present? If not, do they need to be? If they aren't, can the group still succeed?

- Does the group plan to put money on the table to accomplish what it wants? Will everyone contribute equally?
- Has the group considered other important issues, such as antitrust laws, liability, intellectual property, etc.?

5.2. Discussion

Since both the DHS S&T and NSF are interested in this area, is a co-investment possible such that the NSF invests in the longer-term aspects of a project and the DHS S&T focuses on the nearer-term technology application? There are rules about procurements that make coordination a challenge, but it is possible.

The sector needs both longer-term research and shorter-term applied research with deadlines and direct goals. The conversation is not yet mature enough to identify specific projects, and this report only captures the content of a single meeting during which no specific projects were selected. The group needs to identify areas for collaboration, which may best be done through a smaller focused committee, perhaps looking through the report to pick out a handful of projects. The FSSCC can then reach out to its member institutions and assess cost based on project participation.

What can be done to keep momentum going and institutionalize the process? The FSSCC R&D Committee would be a good vehicle to meet on a regular basis for follow through. The BITS and FSTC organizations can also contribute where appropriate.

In preparation for a deeper collaboration between the banking and finance sector and academia, it may be helpful to develop a boot-camp-type training session (week-long) to give information to researchers on how the sector really works.

6. Next Steps

The group concluded that high-impact, large-scale attacks that target the entire sector are theoretically possible and underanalyzed. A continuing dialogue on defending against such attacks and how to effectively address them in cooperation with government would be productive and useful. The group also concluded that banking and finance sector problems are unique and important and require basic research in modeling and analyzing large-scale interdependent financial systems and in constructing inherently recoverable distributed computations.

Some useful next steps include a very near-term follow-up meeting to continue the conversation and create the boot camp for academics and appropriate government representatives. Also, the academic community should consider forming a standing interest group that parallels the FSSCC. Until then, the FSSCC and academics can work informally in consultation with the NSF and DHS S&T to identify the relevant people to participate in follow-up discussions and the boot camp. As part of the exploration, it would be useful to find a way to survey current/previous research

that is directly relevant to the key challenges identified here, with a focus on systemic challenges to the sector. The group will work hard to sustain the workshop's valuable industry-academia-government partnerships by leveraging all mechanisms that all three groups make available to address these important national challenges.

Appendix 1. National Cyber Defense Financial Services Workshop Agenda

"Helping Government Form a Sound Investment Strategy to Defend Against Strategic Attack on Financial Services"

Wednesday, October 28

8:00–8:30	Breakfast	
8:30–8:45	Welcome /opening remarks	Brian Peretti (Treasury), Dan Schutzer (FSTC) and Sal Stolfo (Columbia University)
8:45–9:00	Workshop context	Douglas Maughan (DHS) and Lenore Zuck (NSF)
9:00–9:45	Changing threat environment and needs of the financial services industry; scenarios of likely outcomes if FI infrastructure is disabled, timing of downtime, and requirements for reconstitution; learn what we need to know	Dan Schutzer introduce discussion facilitated by Jane Carlin (Morgan Stanley and chair of the FSSCC Cybersecurity Committee)
9:45–9:55	Brainstorming research challenges for nation-state and organized crime threat to FI infrastructure; prevention strategies, reconstitution strategies Breakout into three parallel working groups	Dan Schutzer (FSTC) and Sami Saydjari (CyberDefense)
9:55–11:00	Working group 1: detection and response - how we would know that a strategic banking and finance sector attack was unfolding and perhaps what top-level actions might mitigate damages - what kind of attacks we need to worry about (service denial, attack on integrity of services or financial health)	Chair: Dan Schutzer Scribe: Jeri Hessman
9:55–11:00	Working group 2: prevention - how we increase the adversary's work factor to make such attacks much harder	Chair: Craig Froelich Scribe: Jeremy Epstein
9:55–11:00	Working group 3: recover and reconstitution - how we maintain the largest possible core of a system and recover the rest as quickly as possible	Chair: Sami Saydjari Scribe: Jenny McNeil
11:00–11:15	Break	
11:15–12:00	Interim report out of working groups to plenary session	Working group chairs
12:00–12:30	Discussion of R&D and game-changing technologies that support the needs of the banking and finance sector; review relevant recommendations from Cyber Leap Year	John Mitchell (Stanford University)
12:30–1:15	Working lunch	Dan Schutzer to introduce Mark Clancy to discuss recent cyber exercise for financial services
1:15–1:25	Continuation of discussion of R&D and game-changing technologies that support the needs of the banking and finance sector; charge to working groups	Sami Saydjari (Cyber Defense Agency)
1:25–2:30	Breakout session 2: Working groups 1, 2, and 3 continue to meet to finish and prioritize their lists of needed technologies	
2:30–2:45	Break	
2:45–4:15	Breakout Session 3: Continue discussing game-changing technologies; rank key ideas, develop strategies to address	
4:15–4:45	Reports by three working groups on top ideas	Working group chairs to workshop
4:45–5:30	Recap of the day's discussion	Led by Sal Stolfo and Dan Schutzer
6:00	Dinner	

Thursday, October 29

8:00–8:15	Breakfast	
8:15–8:45	White House's view of defending the FI	Aneesh Chopra, White House CTO
8:45–9:00	Break	
9:00-11:15	Discussion of the innovation partnership model and efficient technology transfer mechanisms	John Carlson, Doug Maughan
11:15–12:30	Review outline of report, assign authors to draft portions, and establish target completion date	Sami Saydjari, Jeremy Epstein (SRI)
12:30	Box lunches and adjourn	

Appendix 2. TRUST Program—An Example of Related Research

There are many government-funded research activities whose results could be of value to the banking and finance sector's problems as academics attempt to better understand the sector's unique challenges. One example is TRUST. The NSF funds the TRUST Science and Technology Center, which involves multiple academic institutions (Berkeley, Carnegie Mellon, Cornell, Stanford, and Vanderbilt) with a wide range of research lines.

Researchers are also trying to improve understanding of network security by learning the details about how protocols interact and by looking at a wide range of possible attacks against a system.

Some concepts being investigated include the following:

- use of mobile phones as a point of authentication;
- authenticating to a website in such a way as to provide enough information to verify identity but not so much that an opponent intercepting it could use it to masquerade as that particular user in future interactions;
- having a single site (perhaps sponsored by the banking and finance sector) act as a certified source of identity, from which all other logins are done;
- ongoing studies in risk management, particularly for systems built without adequate consideration of failure modes and intentional attack modes (at the workshop, a member of the banking and finance sector observed that in looking at the efficacy of system solutions, one must consider differences in required customer behavior change—customers do not change easily or tolerate additional burdens placed on them); and
- tracking data provenance in terms of its origins and processing history (the research community is investigating a technology called tagging that has previously been used by the Department of Defense for high-assurance systems).

Several workshop participants suggested that the FSSCC R&D Committee might benefit by establishing an ongoing relationship to have the financial industry briefed by academics on potentially relevant studies like these on an ongoing basis. Periodically, researchers have come to talk to the committee about projects going forward, so there is reasonable precedence for such a sustained interaction.

Appendix 3. Prevention Details

This appendix is an elaboration of the details of the workshop working sessions summarized in the main body of the report in section 4.2.

Research Challenges

The responses to what sort of solutions/fixes are required to meet the banking and finance sector's unique challenges in the event of a high-impact, large-scale attack were all based on the common agreement that there are concentration points within the financial industry where attacks could be catastrophic.

Everyone acknowledged that the intersections of these systems is the most serious possible attack point, but also the point at which the problems can be best addressed; despite different applications, there is system commonality. Because there are a finite number of critical choke points crucial to the entire industry, research is needed to understand how to build resiliency across systems—that is, to design systems and processes so that they can take over the function of other systems and processes when they fail or lose integrity.

The intersection of these systems is the most serious possible attack point, but also the point at which the problems can be best addressed; despite different applications, there is system commonality.

The suggestion that there be an additional focus on damage avoidance and deterrence as well as prevention received solid agreement. There was also discussion about the need to study the impact of remedial solutions on system functionality and how these solutions might create vulnerabilities.

It was also widely agreed that given the industry's competitiveness, there has to be some incentive to aggregate data to meet common security goals.

Proposed challenges included the need for

- more models on redundancy and diversity,
- research on load sharing,
- better understanding of data provenance,
- better understanding of the consequence of incentives and regulations,
- increased intelligence gathering function,
- more flexibility in complex systems,
- metrics to understand the impact of problems, as well as attacks and fixes, and
- more standards on how to respond better in crisis.

To properly protect the banking and finance sector from high-impact attacks, one must understand the big systems that comprise it and how they work together. One

must also understand the connections, vulnerabilities, resiliencies, and possible attacks, and how to design our systems to better prevent and counter those attacks. This suggests the following courses of action:

1. Design secure and resilient business practices; the sector must organize to reduce systemic risk.
2. Focus on validating the information embedded in the supply chain; it needs more internal controls, but a determination must be made as to where in the chain the protection is most needed. There must be robust actions in enforcing credentials to enter the system at that point.
3. Work on modeling the system so that the banking and finance sector can identify the nature of the problems. Use the DHS S&T's DECIDE[16] as an example.
4. Research how to build resiliency across systems. A finite number of critical choke points are crucial to the entire industry, and there are exclusive flows of data but no cross-flow or contingency. The FFIEC has put out guidance in industry resilience[17].

Consideration must also be given to the fact that the financial industry is complex, thus the implications of even a simple attack are difficult to understand, thus the impact of remedial solutions are also hard to understand or evaluate.

Strategies for Addressing Challenges

The workshop attendees offered some solutions as to how to counter the common denominator needs:

1. Create a methodology for **relating computer systems to the operational processes** that support the business mission. This, in turn, enables assessing the business impact of an attack on the information infrastructure.
2. Create a **risk assessment process** that enables the banking and finance sector to prioritize risks to single institutions and to the sector as a whole, including a better analytical techniques for economic and social consequences of high-impact attacks, their likelihood, and complexities such as delayed impact.
3. Develop more **secure and inherently resilient financial transaction systems** and architectures, moving away from single points of failure such as transaction choke points and dependency on a single monoculture system.
4. Invent better **protection strategies for data-centric systems** including the notion of protected data provenance (tracing where data originated, what systems processed it, what was done to it, and where it went from there) and

[16] See the section on Innovative Industry-Government Partnership Models earlier in this paper for a description of DECIDE.

[17] http://www.ffiec.gov/ffiecinfobase/resources/bus_continuity07/EX_NSTAC_Fin_Ser_Task_For.pdf

systems provenance (controls on software lifecycle to reduce the risk of technology supply-chain corruption).

5. Design systems with **runtime verification test and validation methods** and process control mechanisms to discover vulnerabilities (such as malware detection) and immunize against them, with the added variable that some of the systems testing or being tested may have been co-opted.
6. Improve **complex system modeling capabilities** and large-scale systemic measurement to inform and validate the models so that they can be used to better under existing sector-wide systems, new designs under consideration, and ways of designing systems with better resilience.
7. Develop technology (such as anonymization), policy, and incentives for **better information sharing** between institutions and with the government.
8. Create new security metaphors for **improved usability** so that users, administrators, and risk managers better understand what to do under specific circumstances and with what risks.

Most participants agreed that more cooperation is needed among financial institutions for a wide number of reasons.

Most participants agreed that more cooperation is needed among financial institutions for a wide number of reasons: increased levels of intelligence gathering and analysis to better design preventative measures; increased system diversity and redundancy; and more incentives for and regulation of cooperation.

Appendix 4. Detection and Response Details

This appendix is an elaboration of the details of the workshop working sessions summarized in the main body of the report in section 4.2.

Research Challenges

The following research challenges were generated through a round-robin brainstorming process in which idea breadth took priority over depth. Some combining of similar ideas was done during the editing process for coherency and flow. Prioritization was achieved by loosely grouping ideas and having panel members vote for their top three choices.

Data sharing was this workshop group's top issue. Financial institutions have vast quantities of data that might be useful for academics, but they are highly concerned about data scrubbing, both to protect their proprietary interests from competitors and because of customer privacy concerns. A research challenge is to come up with better methods for data scrubbing that will provide useful data for researchers while still protecting the banks' interests.

A research challenge is to come up with better methods for data scrubbing that will provide useful data for researchers while still protecting the banks' interests.

Once data-sharing concerns are addressed, the next step is **better analysis and situational awareness**. One research area is how to use the data to build better detectors, how to present findings, how to learn what types of additional sensors would be useful, and how to deploy them in real time.

Related to the issue of data sharing is the need for **centralized "war rooms"** in times of crisis where participants from major organizations can share data in real time, just as intelligence community command centers allow sharing by having assignees from each organization physically in one place with links back to their home organizations. Doing this effectively will require government participation because real-time intelligence relevant to the financial industry may be classified.

Effective situational awareness includes **understanding the security stance of critical suppliers**. Any war-room solution needs to include those organizations in addition to the financial institutions themselves.

Partitioning data into sector-specific attacks and infrastructure attacks is problematic. For example, prior to an attack, there may be anticipatory activity on the telecommunications system (just as there was anticipatory "chatter" before 9/11), but little activity precursors on the financial systems themselves: attackers will presumably have tested their attacks in tightly controlled environments. Thus,

any detection system must integrate data at all levels, so that it can recognize attacks even if they do not show up directly at the applications level.

The goal of data sharing is to allow organizations to find "abnormal" activity. However, to find abnormal results, systems must know what "normal" is. Some activity that might appear abnormal (e.g., rapid transactions) may actually indicate a previously unseen software package, rather than an attack. One research area is to try to **characterize normal customer behavior**, so that abnormal activity can be distinguished. Banks have significant experience in finding fraud, which perhaps could be leveraged—interesting problems show up as financial anomalies, not technical problems.

There are dependencies between the financial industry and other industries (e.g., telecommunications, information technology, power, transportation); banks need to **understand this dependency and its relationship to their institutions**. They also need to be able to subdivide the dependency, so rather than simply being dependent on "the Internet," they're dependent on various vendors for authentication, network connectivity, etc.. One research topic to consider is whether prioritization of financial traffic is a good solution (but this may be contradictory to net neutrality).

Computer science academics do not have enough background to understand many of the financial industry's challenges. A **week-long "boot camp"** may be an effective way to teach academics (including graduate students) how the financial system actually works (as compared to how it's *perceived* to work). This would yield better results in the long run and also help researchers understand the differences between how very large banks differ from smaller banks, which typically have much less sophisticated security capabilities.

In addition, the issues identified are not solely within the solution space that computer science academics can provide. To develop solutions, researcher from other disciplines, including the social sciences, must also be brought to the table to provide guidance on how employees and customers can be educated to take maximum advantage of the solutions developed by computer science researchers.

A nontechnical but extremely critical research topic is the role of **bank responsibilities to retail customers**. Some banks would like to require that their customers have up-to-date antimalware software before allowing them to perform online banking but are concerned with the liability issues. Similarly, if a bank detects malware on a customer's personal computer, it is hesitant to make changes, lest such a change has a negative effect (e.g., "breaking" the customer's machine). Can banks, as the most security-sensitive industry, do more to help protect their customers? Since the largest handful of banks have relationships with the vast majority of Americans, solving this problem with the largest institutions can "trickle down" to protection for nearly everyone. However, financial institutions would

prefer that these controls be implemented at the service provider level, as service providers have the technology relationship with the customer.

Most retail banking customers are unwilling to deal with inconveniences to access their online banking site. A small fraction (estimated at 10 to 15%) is willing, most likely because these people have experienced financial fraud or identity theft. A useful research topic would be **incentives to encourage customers to treat security as part of their responsibility**, rather than relying on the bank to do it for them.

On a related point, authentication is a key to improving security, but many retail customers are resistant to replacing username/password authentication with something stronger. Unfortunately, replacing existing password schemes with stronger authentication systems has side effects, including the problem that many consumers deal with five or ten financial institutions, not just one, and they access accounts from their home, work, and mobile devices. Stronger **authentication solutions must be scalable for consumers** as well as financial institutions.

As is well understood, many users choose the same password for their low-risk, high-value sites (e.g., a bank) and their high-risk, low-value sites (e.g., YouTube), and many of the low-value sites lack adequate security controls. Thus, a compromise of the low-value site can indirectly compromise the user's bank account, since the password is the same. As a research topic, participants brainstormed over whether **banks should provide free security services to non-banking websites** to reduce the risk of compromise. Or are there ways to discourage users from using the same password, either automatically or through education?

Response to detected attack was given much less attention by workshop participants. The operational model was described as "being a turtle," pulling arms and legs in until the storm passes. However, research problems in recovery include distinguishing legitimate from illegitimate data and transactions since a bank can't recover to an empty database. This is one of the disadvantages of T+1 trading windows—less time to roll back any incorrect transactions before they are completed.

The above list of challenges were further combined and prioritized into the following high-level challenges:

1. **Accurate and comprehensive data gathering and situation awareness through a real-time information sharing and fusion center.** Financial institutions need broader situational awareness (including attack data at all levels) from across the industry as well as nonindustry input (e.g., from the Department of Defense), and better methods for using that data to gain situational awareness.
2. **Consumer computer system protection policies.** Financial institutions need policy clarity as to what they can do to protect consumers from

malware and/or block access to their systems by consumers whose computers contain malware.

Strategies for Addressing Challenges

As noted above, *improved data sharing and situational awareness* was the highest priority item. Key challenges for this topic include investigating whether existing governmental data-sharing procedures can provide the needed safeguards. There has been substantial progress in developing data-sharing agreements with the Department of Defense; perhaps this can be used as a vehicle to avoid reinventing the sharing infrastructure. Additionally, the DHS National Cyber Security Division could provide some of the infrastructure to facilitate sharing, although thus far, there has been little support in this area. Critical to progress is developing generic data-sharing agreements to replace the point-to-point sharing that happens today (and is not scalable).

A key concern was the mismatch between the goals and timelines of financial industry practitioners and academic researchers. Some participants felt that the financial industry historically does not look more than a year or two out for information security, while researchers perform research that may not pay off for five years or more. As noted earlier, a "boot camp" program to help academics understand the financial industry would be useful for shortening the learning curve and thereby possibly making research results more timely and useful.

For policy clarity, the dialogue between the Department of the Treasury and the financial industry must continue to establish policies that may be used by financial institutions relative to consumer systems (e.g., for online banking, mobile banking, social networking). This is not fundamentally a research problem, but once the policy is determined, research on effective methods to protect consumer systems will be needed.

Appendix 5. Recovery and Reconstitution Details

This appendix is an elaboration of the details of the workshop working sessions summarized in the main body of the report in section 4.3.

Research Challenges

The following research challenges were generated through a round-robin brainstorming process where idea breadth took priority over depth. Some combining of similar ideas was done during the editing process for coherency and flow. Prioritization was achieved by loosely grouping categories and having panel members cast five votes each:

1. **Communicate and share information reliably** across financial firms, service providers, and with governments (law enforcement, intelligence community, cyber command, DHS) in the face of an attack. Build resilient, critical, shared, communication services.
2. Build **resilience** into the network itself with automated attack recognition and response, allowing the system to defend itself.
3. Discover how to do transaction commitment and recovery in the context of complex banking and finance sector processes. Identify and recover to a **good state** after pervasive attack, such as a data integrity attack. Given that a system has multiple dependencies, how does one know the system has recovered to a collective known good state? How does one manage output commitment of transactions? In summary, discover how to structure **distributed computation so that it is inherently recoverable.**
4. Understand system **interdependencies and consequence propagation** and then structure business process interconnections in chains of clearing houses and suppliers.
5. Build **behavior models** and monitoring capabilities to characterize reasonable behavior for the entire banking and finance sector (such as what intrusion detection systems do for enterprise behavior models but on the much larger scale of a sector-wide behavioral model). Develop behavioral models that enable detection of anomalies within financial transaction systems, including money-moving models for transfer systems.
6. Create a **common operating picture** (COP) of the banking and finance sector with multi-infrastructure views that tie in communication and intrusion detection systems to fuse analysis in real time.
7. Create methods to **regain trust in** a machine or capability once it has been lost due to a pervasive attack. At the same time, create new architectures that do not depend on trusting so many entities.
8. Develop tools and methodologies for **proving programs** (such as critical transaction processing) correct. Narrowing properties may make it feasible.
9. Define what **information to share** and the processes and mechanisms for sharing to enable government and private institutions to effectively perform their mission, including attack trace-back and attribution.

10. Make **data self-aware** and security-conscious so that data objects participate in their own defense, perhaps in cooperation with the underlying system. Develop access and functional controls that travel with the data similar in concept to data rights management but in a vendor-neutral way.
11. Develop tools and techniques to automatically and dynamically assess the most **significant risks to the banking and finance sector** mission as it relates to information technology and organization dependence. The impact of scale, which is pervasive throughout the industry, needs to be a major focus.
12. Create a technology to allow **continuous verification and validation of recovery capability** (running multiple resilient sites, for example) to function when called upon to do so, especially in the face of an attack that may attempt to disable that capability first.
13. Discover a means to **recover without a rewind button** on pending and partially completed unreconciled transactions in the middle of a trading day.
14. Investigate how the coming introduction of **pervasive mobile devices** affects approaches to these problems. This is primarily a placeholder for more specific discussion on this topic at a future workshop.
15. Learn how to perform **synchronized, coordinated, multi-layered recovery** so that basic capability can be restored quickly and increasing capability can be brought back in incremental phases.
16. Investigate how end users participate in the recovery process. Ideally, the recovery is transparent to them, but part of the recovery process may be **restoring the user's trust** in the system—no small task.
17. Develop a means to recover from high-impact attacks involving **insiders,** particularly if they may be on the inside of the recovery process as well. How can separation of duties be arranged so that the people who do the recovery are not likely to be the same people who caused the incident?

The impact of scale, which is pervasive throughout the industry, needs to be a major focus.

The above list of challenges was further combined and prioritized into the following short list of five top-level challenges.

Constructing recoverable distributed computation. The challenge here is to define distributed system states with properties that need to be true of those states for them to be considered good. The sector needs a means of identifying good states and transitioning to them quickly and without creating instabilities or inconsistencies internally or with external entities (clients) that use the system. In rewinding to a previous state, methods are needed for dealing with portions of the state that affect the state of systems external to the system being recovered (for example, if an ATM distributes cash to an account holder, it is hard to back up to a state prior to that because it is difficult to "undistribute" the cash to the person).

Constructing analytical domain models of banking and finance sector. The group observed that banking and finance sector systems are very large and complex and inadequately understood in terms of what they do and how to exercise control under duress. The sector would benefit from sophisticated models of aspects like information flow, interdependency, and cascading impacts. The sector needs to be able to perform both static and runtime analysis (both formal and informal) of these models to improve their situational understanding as well as their ability to exercise control.

Organizational information-sharing requirements for recovery. As the sector begins to better understand how to do multi-institutional multi-phased recovery, it needs to understand the information that must be shared between and among institutions to achieve sector-wide recovery. Finer-grained reports and structure may well be needed. Shared data must be useful and actionable for the intended purpose, and the risk of sharing versus benefit must be better understood. Multi-phased recovery processes must include a minimum essential base, such as robust communication used to enter a guardianship mode from which to recover the system. Other capabilities would then need to be brought back in order of criticality.

The sector needs to understand the information that must be shared between and among institutions to achieve sector-wide recovery.

Recovering human trust. The banking and finance sector runs on bits of data, but those bits only have value because of the humans involved at the end of the transaction. A strategic attack on the banking and finance sector would likely result in an erosion of human trust in the system that could lead to serious consequences in the suppression of commerce. We must therefore model human and stakeholder involvement and behavior during the recovery process and develop ways to restore human trust and confidence in the system—including banks and brokers who operate closer to end users.

Recovering from insider attacks with insiders involved in recovery. Trusted insiders may well be co-opted as part of a strategic attack scenario. In such scenarios, insiders involved in instigating the attack may be part of the recovery process as well. Further, sophisticated attackers are likely to target the backup systems before attacking the primary system. Therefore, we need to design system architectures such that recovery is robust, even in the face of some insiders involved in portions of the recovery process.

Strategies for Addressing Challenges: A Scenario

In this section, we assume that a major strategic attack succeeds, such as those described earlier. For example, we might assume a scenario in which the supply chain for all financial transactions done over the last week has been corrupted—the data looks like it works, but it's not correct data. We focus possible approaches by

addressing the prioritized challenges in the previous section. Workshop participants developed the approaches using a round-robin brainstorming process.

Constructing recoverable distributed computation

1. Develop methods **of continuous resilient operations in parallel environments**—so-called "hot-hot" operations.
2. Take advantage of the **replication of transactions and reports of transactions throughout the system** on both sides of the transactions and with regulators. Develop methods to recover using these distributed redundancies from reports of transactions. Perhaps restructure the redundancies to make recovery using this method more plausible—for example, require redundant recoverable reporting at financial exchanges and clearing houses. Extend existing data standards to a fuller list of participants, including the "other side of the trade."
3. Develop a **means to identify bad data** and the data derived from it, and exclude that data from recovery processes.
4. Develop highly **robust diagnostic data** to determine what went wrong, where, when, and how. Protect this diagnostic data from attack.
5. Develop **probabilistic trust models** of which areas are least likely to have been corrupted and build trust from those.
6. Develop **reversible change control processes** that would incorporate change control libraries.
7. Develop **pre-thought-out recovery strategies** that depend on the nature of the incident, such as whether the system experienced a confidentiality, integrity, or availability attack.
8. Instead of trying to figure out what happened, recover what might have happened. Have more leeway in defining a consistent state. Use the notion of **compensating transactions** that do not necessarily reconstruct the entire transaction history, but yield an end result that is the same or close to what would have been the outcome.
9. Create **pre-planned recovery actions** and states that are standards to allow institutions to get back on the financial grid.
10. Figure out how to securely **use the "cloud" to store encrypted transactions** and create robust storage of recovery data.
11. Define **recovery processes that include corrupt data** and back that corrupt data out of transactions by removing them and all data they affected.

Constructing analytical domain models of banking and finance sector

1. Develop **network diagrams** of "who talks to whom" for various attack scenarios and identify the major transactions of highest interest.
2. Develop some kind of **model simulation with libraries of a spectrum of perturbations** representing attacks or attack components. Use this to better understand dynamic behaviors, especially the unexpected ones. Develop a functional model of the banking and finance sector that includes service providers and uses game theory to explore failures of all kinds.

3. Create the equivalent of **"radioactive" tracer data**. Inject data that is flagged in a certain way so that every component that touches it has to report that it has seen it and to where it propagated it. Some basic form of this already exists and is called "salting and seeding."
4. Create a **variety of different models** and differentiate between them at functional, transactional, and infrastructure/technology levels. This structure will help define the approaches to take. When modeling systems, for simplification, pick hubs and critical nodes instead of the entire system or network.
5. Have all data in the system **carry provenance metadata** regarding its origins and the processing it has been through, with authentication of each step.
6. **Map out logical transaction paths to physical network paths** (at least on some crucial links, like the hubs) for recovery planning at the physical level.
7. Understand technical as well as nontechnical sources of trust. Have an understanding of how trust relates to the trust placed in the data source (**data supply chain**) and trust of where the system services come from (**technology supply chain**).
8. Create ways to announce each banking and finance sector participant's **continuous trust model state** and why it holds the beliefs that it does. This will enable each party to see who trusts or distrusts that participant, why, and possible actions it can take to recover that trust.

Organizational information-sharing requirements for recovery

1. Build **concentric ring communication models** where the center is the most important, most protected, and most robust against high-impact attacks. Base the concentric circles on a natural unit of work that must be done and who must contact whom to accomplish that work. Pre-play scenarios with emergency teams representing major recovery accomplishments within these rings.
2. Strengthen financial services information sharing and analysis centers (FS-ISACs) as **standard universal anonymized repositories of data**; extend to multiple ISACs for multi-infrastructure scenarios. Work out the private- and public-liability issues associated with such data repositories.
3. **Identify minimal essential data** that needs to be shared for recovery.
4. Develop peer-to-peer secure ad hoc communication network overlay on phones and computers with radios (wireless) using applications such as Skype on top of the ad hoc network for services. This is known in some circles as "direct Wi-Fi" or more informally as "LANS across America."
5. Build critical communication channels and sharing channels and **verify their operation in peacetime** when strategic attacks are not occurring. Develop alternate channels and addressing schemes. Support social channels for informal communications (similar to those formed by the North American Network Operators' Group) as well as formal communications.

6. Set up emergency teams that represent all the major players and give them a **physical place to go** (like a command recovery center) that has limited connectivity and the ability to share on a real-time basis.
7. Develop checks and balances within the system such that the participants can **verify trust of all other members**.

Recovering human trust

1. Develop a **public relations response plan** to reassure the public and restore trust in the system. Pre-position public service announcements and target honest levels of confidence. Overselling unwarranted trustworthiness of the system could lead to a total collapse in trust later. Be accurate; once you declare that it's right, it must be right…or don't provide information until you're sure it is. Ideally, there should be a **minimal amount that the end user has to do** as part of the recovery process. Mass participation in recovery could cause confusion and further undermine trust.
2. Develop **stress test suites and validation procedures** that the government can run, with passing grades to restore trust on an institution-by-institution basis, analogous to what the government did with banks during the Great Depression to restore public faith. During the 1930s, when the government allowed an institution to reopen, it meant that the institution was okay.
3. Build in **accountability methods in advance** to establish confidence at the beginning.
4. Develop **analogues to electronic voting**. Credentials will establish transaction veracity on the consumer side.
5. Develop **anti-masquerading techniques** for all sides of transactions, including end users.
6. Plan for the **equivalent of "bank holidays**" and other time-outs to allow recovery and restablization of the system.
7. Develop **fair-exchange cryptographic protocols** to assure both sides of a transaction that they receive their expected value before approving the transaction.

Recovering from insider attacks with insiders involved in recovery

1. Develop **cryptographic protocols for Byzantine recovery—**named for the so-called Byzantine generals problem that assumes some unknown but perhaps large percentage of insiders is corrupt. You need a mechanism that will produce the desired result—in this case, recovery to a good state—despite the existence of a fraction of corrupt insiders.
2. Create **internal redundancy checks**, applying game theory and anti-cheating mechanisms.
3. Make the system as **automatic as possible** so there is little dependency on humans who can, in turn, be co-opted.
4. Invent **over-the-shoulder monitoring systems** and hold observations up to a model of what sorts of behaviors are allowable under which circumstances and in what contexts.

5. Develop **separation-of-duty** procedures of critical operations and nuclear-control-style two-man control for the most critical operations.
6. Apply **indications and warning controls** in the recovery process to flag any events that might suggest an action that is inconsistent with the recovery process.
7. Create a **means to reconstitute multiple versions** of systems with independent teams and then compare them at the system level before going operational. This is a system-level equivalent of the concept of N-version programming.
8. Create **extensive audit data during** the recovery process to identify corruption.
9. Develop **self-monitoring monitors** to prevent the monitors themselves from being corrupted during recovery.
10. Have clarity on **how many insiders** are assumed (and where they are located) in the attack and make approaches robust under those assumptions.

Appendix 6. Organization and Attendees

Co-Chairs: Sal Stolfo, Columbia; Dan Schutzer, FSTC; Brian J. Peretti, Treasury
Chartering: O. Sami Saydjari, NCDI Executive
Organizing Committee: John Mitchell, Stanford; John Carlson, BITS (and chairman of the FSSCC R&D Committee)
Sponsor Representatives: Douglas Maughan, DHS S&T; Karl Levitt, Lenore Zuck, and Sylvia Spengler, NSF
Speakers: Jane Carlin, Morgan Stanley; Mark Clancy, DTCC; John Carlson, BITS; Douglas Maughan, DHS S&T; John Mitchell, Stanford; and Aneesh Chopra, White House CTO
Participants: See Table

Financial Industry			Researchers			Government	
Dan Schutzer	FSTC		Sal Stolfo	Columbia		Brian J. Peretti	Treasury
Leigh Williams	BITS		John Mitchell	Stanford		Carl Landwehr	IARPA
John Carlson	BITS		Steven Bellovin	Columbia		Lenore Zuck	NSF
Michael McCormick	Wells Fargo		Howie Shrobe	MIT CSAIL		Sylvia Spengler	NSF
Warren Axelrod	FSTC		Ernie Drew	Norwich U.		Douglas Maughan	DHS S&T
Jane Carlin	Morgan Stanley		Farnam Jahanian	U. of Mich.		Jim Devlin	OCC
Ken Schaeffler	Comerica		Peng Liu	Penn State		Bill Newhouse	NITRD/NSA
Bruce Lane	State Farm		Sushil Jajodia	GMU		Tom Donahue	NSC
Phil Venables	Goldman Sachs		Gail-Joon Ahn	Ariz. State		Sampath Kannan	NSF
Mark Clancy	DTCC		Fred Schneider	Cornell		Kevin Walsh	GAO
Eric Guerrino	BNY Mellon		Mike Reiter	UNC		Chris Greer	OSTP
Rick Lacafta	Travelers		Mike Schroeder	Microsoft			
Craig Froelich	BOA		Robert Schmidt	Delta Risk			
			Sami Saydjari	CDA/NCDI			