House Committee on Financial Services
Subcommittee on Financial Institutions and Consumer Credit
September 14, 2011
Testimony of Dr. Gregory E. Shannon
Software Engineering Institute
Chief Scientist for the CERT Program

Chairwoman Capito, Ranking Member Maloney, and other distinguished Members of the Subcommittee on Financial Institutions and Consumer Credit, thank you for the opportunity to come before the Subcommittee to testify, it is my pleasure to be here this morning to discuss the cyber threat to financial institutions.

About CERT®

The CERT Program is part of Carnegie Mellon University's Software Engineering Institute (SEI), a Department of Defense Federally Funded Research and Development Center (FFRDC), and is located on the Carnegie Mellon campus in Pittsburgh, Pennsylvania.

The CERT program (http://www.cert.org/) has evolved from the first computer emergency response team created by the SEI, at the request of Department of Defense Advanced Research Program Activity (DARPA), in 1988 as a direct response to the Morris worm incident. The CERT program continues to research, develop, and promote the use of appropriate technology and systems management practices to resist attacks on networked systems, limit damage, restore continuity of critical systems services, and investigate methods and root causes. CERT works to mitigate both cyber risks and facilitate local, national and international cyber incident responses. Over the past 23 years CERT has led efforts to establish over 200 CERT computer security incident response teams (CSIRTs) around the world – including the Department of Homeland Security (DHS) US-CERT. We have proven track record of success in transitioning research and technology to those who can implement it on a national scale.

Dr. Greg Shannon is the Chief Scientist for the CERT Program, where he leads technology innovation efforts to establish and enhance the CERT program's research, development and strategic policy initiatives.

Understanding today's cyber threats is more than just war stories, anecdotes and scare tactics. The threat is real, it is now, and it is evolving; CERT catalogues ~250,000 instances of candidate malware artifacts each month. At this volume is difficult to determine in real time which are malicious, let alone what their intent. Unsurprisingly, the limit in our technical abilities to provide security has brought about the steady corporatization of the cyber security threat. Cyber attacks have become big business, with extraordinary returns on investment; for example, botnets are both economical and flexible for creating cyber effects at scale, including attacks on financial institutions and their customers, which yield high payouts.

To efficiently fight the cyber threat we need realistic outcome based solutions, enabled by a data driven approach to research, development, policies and regulations. There is an emerging science of cyber security, and I encourage the Subcommittee to support practices that are both scientifically and operationally validated as part of a continuing dialogue on important policy discussions. CERT is working with both U.S. Government Agencies and the Financial Community to help limit the Nation's exposure to cyber attacks, but more can be done. The U.S. Government and the globalized Financial Sector need computing infrastructure that is more secure *and* more resilient in order to mitigate the escalating threats. Filling the technological gaps in current forensic abilities as well as augmented capabilities to locate the source of the attack and limit the damage are much needed. Finally, leadership and support from the government in policy discussions that bring focus and funding for more robust research in key areas of cyber security, is essential.

Examining the Threat – from the Inside

While there are many methods to launch a cyber attack on the financial sector, I would like to highlight the risks of insider threats. CERT is using incident case data in combination with hands-on collaboration with practitioners in industry and government to understand how best to mitigate insider threats.

The continued stress of the current economy on the workplace is impacting and exacerbating the potential for insider threat. Organizations are working hard to build walls around their network infrastructure to keep people out but are having a difficult time defending against potential menaces that are already on the inside of the fence.

I am going to highlight some areas of concern that we have been focused on at the SEI. These points present a general picture of the problem:

- According to our data, over the past 7 years malicious insider attacks have affected approximately half of all organizations.
- CERT's data also reveals that the insider threat is not limited to rogue individuals acting alone, but rather, almost half of all malicious insiders in the Financial Services industry colluded with outside conspirators, while a third recruited other insiders to carry out their crimes.
- In CERT's research, we discovered that around 10% of fraud cases involved organized crime; average losses in those cases were $4 million.
- Other complexities include trusted business partners (e.g. outsourcing), mergers and acquisitions, and branches located outside of the US. This exacerbates the problem due to cultural issues, national loyalties, and legal issues such as more stringent employee privacy laws.

Insider crimes in the financial services sector are not limited to fraud, but also include theft of intellectual property and insider IT sabotage. One former system administrator wiped out billions of files on a financial institution's servers all over the world at 9AM one morning; and recently an individual copied source code containing proprietary trading algorithms to servers outside the U.S. after submitting his letter of resignation.

Although there are a vast number of cybersecurity tools available, most are used on protection from breaches from outside. The difficulty in creating effective automated tools for detection from insider threats is that malicious insiders typically commit their illicit activity by performing the same types of online actions they do every day, only with malicious intent. In fact, insider fraud is often detected either through traditional auditing methods or discovery by external parties, such as customers.

CERT's work in this area began in 2001, through the support of the United States Secret Service. Over the past 10 years CERT has collected and rigorously analyzed over 700 actual cases of insider crimes spanning all critical infrastructure sectors. Within that sample, there were 147 cases involving the financial services sector. We use these cases for research and analysis, and produced models, best practices, and training for government and industry on prevention and detection of insider threat. A crucial part of CERT's research mission is to make the most of the data we have, to use insightful statistical analysis to understand the breadth of the problem and solutions, and to seek out additional highly-informative sources of data to improve our research, results and impact.

Within the past 3 years two different divisions of DHS sponsored the SEI to enhance our understanding of the insider threat problem and to make recommendations aimed at building new solutions that would be readily available to the Community. CERT recommended using existing technology that most organizations already have in place as the platform for our solutions; the first set of controls will be published and made available to the public this month. We have also developed an insider threat assessment framework, which CERT is piloting with federal agencies, to identify vulnerabilities to insider threat. This will enable CERT to provide recommendations for both the government and industry to use as benchmarks for insider threat defenses. In addition, we are building a training and certification program so that organizations can assess their own vulnerabilities to insider threats. This work is being sponsored by DHS National Cybersecurity Division Federal Network Security Branch.

CERT has also been working with the Secret Service, the Financial Services sector, and the Department of Treasury, sponsored by DHS Science and Technology, to build a model of insider threat specifically for fraud within the financial sector. As an FFRDC, CERT is able to work across the government and the private sector to ensure our research successfully transitions into operationally viable solutions for the nation, and then offer those best of breed solutions to the community at large. CERT has built a relationship throughout the financial community, including the BITS Fraud Steering Committee and the FS-ISAC to create a direct feedback loop of information and data sharing so that we can validate our models and then provide both industry and government with the methods and tools they need to combat the mounting problem.

Building More Secure Systems - Secure Coding

Software vulnerability reports continue to grow at an alarming rate, and a significant number of these reports produce technical security alerts. To address this growing threat to governments, corporations, educational institutions, and individuals, systems must be developed that are free of software vulnerabilities.

CERT takes a comprehensive approach to eliminating vulnerabilities and other software defects, starting with a detailed analysis of vulnerability reports originating from the U.S. Department of

Defense (DoD) and other sources. By analyzing thousands of vulnerability reports, CERT has observed that most vulnerabilities stem from a relatively small number of common programming errors. Software developers can take practical steps to eliminate known code-related vulnerabilities by identifying insecure coding practices and developing secure alternatives.

These new coding standards, developed in coordination with security researchers, language experts, and software developers using a wiki-based community process (More than 500 contributors and reviewers participated in the development of secure coding standards on the CERT® Secure Coding Standards wiki.[1]) encourage programmers to follow a uniform set of rules and guidelines determined by the requirements of the project and organization, rather than by the programmer's familiarity or preference. Moreover, they provide a metric for evaluating and contrasting software security, safety, reliability, and related properties; and when applied during software development these coding standards can create more secure systems.

Resilience

Operational Resilience is a new discipline that blends computer security with business risk management; and is the ability of a network computing system to provide essential services in the presence of attacks and failures, and recover full services in a timely manner. In 2004, CERT began a partnership with the Financial Services Technology Consortium (www.fstc.org) to examine the application of survivability concepts to the complex problem of managing operational resilience in the U. S. financial sector. Due to CERT's trusted relationship we are given unparalleled access to some of the best practitioners in the security and business continuity space.

The focus is not only to thwart computer intruders, but also to ensure that business goals are met and critical business functions are sustained despite the presence of cyber attacks. Improving survivability in the presence of cyber attacks also improves the ability of business to survive accidents and system failures that are not malicious. Resilience depends on three key capabilities: resistance, recognition, and recovery. Resistance is the capability of a system to repel attacks. Recognition is the capability to detect attacks as they occur and to evaluate the extent of damage and compromise. Recovery, a hallmark of survivability, is the capability to maintain essential services and assets during attack, limit the extent of damage, and restore full services following attack.

Through our collaboration with the FSTC (FSTC has since been incorporated into BITS as part of the Financial Services Roundtable), as well as from extensive review of existing codes of practice in the areas of security, business continuity, and IT operations management, CERT codified a draft process definition for operational resilience management processes called the CERT Resilience Management Model (RMM). Recently, CERT was asked to participate in the Payment Risk Committee's Executive Seminar on Business Continuity and Resilience Planning: A Decade After 9/11 at the Federal Reserve Bank of New York, to speak about the RMM, in our speech on "Resilience Management: Ensuring Sufficient Forethought." Building on the success of previous work in the financial sector, CERT is currently conducting a year-long resilience improvement workshop series with members of the defense industrial base, financial services,

[1] https://www.securecoding.cert.org

government, and academic communities. As a result of the success of this workshop series, CERT is currently in the planning stages with major financial institutions to begin a series of workshops focused on the Financial Services Industry using the CERT Resilience Management Model and the Insider Threat Assessment Methods to craft an improvement path for managing operational risk from incidents, insiders, and external events. In addition, through these workshops CERT is making major inroads in developing and piloting measures that can be used to verify that organizations have indeed become more operationally resilient as a result of the controls, practices, and processes they are implementing and sustaining. Participants in current CERT-RMM workshops are piloting and testing these measures as a way to evolve their security and continuity programs away from mere compliance activities toward verification and validation activities.

Forensics

Computers are no longer just the targets of crime; our adversaries now use them to facilitate every aspect of their illicit activities and achieve effects at scale. Once an incident occurs the federal government and the financial community face several hurdles to recover the needed data in order to locate the source of the incidents and contain the problem.

First, computer forensic labs are constrained by a lack of resources, creating an enormous backlog rendering them unable to handle the mega-fold increases in the volumes of data that need to be examined for evidence. When an intrusion into a financial institutions network occurs, time is crucial, immediate decisions are required to deal with instant consequences as well as longer-term consequences (like prosecution of a case). While some entities may have the qualified examiners, and many do not, they lack the funds to properly equip them for the mission. For example, current examination methods rely heavily on processor power, but due to dramatically increased computer memory, examination stations often cannot keep up. Finally, the current state of the practice does not allow examiners to easily access varied levels of expertise in a timely or cost-effective way, resulting in time delays and increased costs.

To successfully respond to cyber incidents these obstacles must be overcome in a way that allows for high-quality, expedited collaborative examinations. For instance, what would happen if an adversary perpetrated an actual, severe cyber event on the financial sector with national consequences? Currently, there is no one facility or lab that could support the volume of data these kinds of events would generate. Under current conditions, data would have to be distributed, adding to the time and complexity of conducting examinations. Analysts and investigators will need flexible, secure access to high-performance systems, to increase productivity and facilitate effective distributed collaboration in a scalable and cost- effective way. Additionally, to support better response, security, privacy and resilience organizations must design and architect their processes and systems in line with scalable assurance principles developed at the SEI and in the broader software and systems engineering communities.

To enable organizations to accelerate the tempo of their investigations CERT is working on a new incident analysis framework which speeds up the velocity of investigation and allows for a faster and adaptive defense and mitigation opportunities otherwise not available in near real-time. To help augment the cyber forensic capabilities of law enforcement the CERT program created the Clustered-Computing Analysis Platform (C-CAP). C-CAP is designed to support 200 concurrent computer examinations looking at 200 terabytes of data, allowing for a massive, coordinated effort. Absent catastrophic events, the C-CAP environment can offer underequipped or overwhelmed agencies real time additional resources. C-CAP is a state-of-the-art forensics analysis environment that provides a complete suite of tools for host-based and network investigations. C-CAP augments scarce resources by allowing multiple users to view the same data, either remotely or locally; while maximizing the application of specialized computing resources to the forensic and incident response missions. Analysts and investigators enjoy flexible, secure access to high-performance systems, increasing productivity and facilitating distributed collaboration. Designed specifically for forensics and incident response analysis, this unique integration and packaging of tools, accelerates the analysis processes, maximizes performance and reduces costs. C-CAP is a flexible solution, allowing agencies to add or remove components that are relevant to their particular needs. Its unique centralized management interface allows organizations to rapidly allocate platform resources to tasks or analysts. Scalable and cost-effective, C-CAP can be customized to suit any organization, regardless of size and mission.

More Robust Research Agenda

Research is only as good as the data it is created from, and currently researchers have limited access to data, resulting in sub-par solutions and stifling innovation. To truly begin to combat the cyber threat we must gain better situational awareness, and it is the federal government’s role to generate situational awareness beyond what any private entity has the incentive to produce. However, achieving this enhanced situational awareness will require continued research on network traffic and data and the cooperation of the financial community.

Richer data needs to be shared with the research community, not only incident data itself, but also data-sets that will enable an understanding of what “normal” resembles, enabling the detection of malicious markers that are invariant, such as behavioral based indicators (e.g. insider threats). Currently, there is not a clear understanding of what this data set would look like; but if situational awareness is to develop beyond simple indicators, the financial sector must allow access to everyday data, so that researchers can begin to recognize what data sets are important.

This data sharing should start with limited access to high-fidelity datasets for researchers so that data with scientifically proven value is considered for sharing operationally. Otherwise, policymakers and experts are left to speculate what is the right data to share. Furthermore, if the research community was able to successfully determine which data sets where imperative to combating the cyber threat, then in effect less data would need to be shared to productively handle cyber dangers.

However, I realize information sharing on this scale tends to exacerbate an already contentious relationship between security and privacy. Security and privacy advocates often are at odds with one other in discussions of how security degrades privacy or privacy degrades security. This is an unhealthy condition, and our adversaries are exploiting it and degrading cyber space for us all. Privacy advocates contend without privacy there is no security. But given our ever more interconnected world the loss of anonymity is unavoidable, and I believe that without security there is no privacy.

Lastly, the government in collaboration with the financial sector needs to encourage and enable more research to be done in the areas of identity management and authentication. As banking becomes more and more mobile, at the demand of the consumer, we need better ways to secure our data while preserving a user-friendly platform.

Conclusion

I cannot end my testimony without saying something about the need for a robust cyber workforce. An educated and equipped workforce is essential to handling the cyber threat to financial institutions. However, the rapid changes and dynamic nature of cybersecurity make keeping the workforce up to date a very challenging problem. The most common workforce development training solution is the traditional classroom training model. While this training model is easy to implement and is widely used, there are a number of reasons it is not adequate for providing effective, large-scale training to a technical workforce including time, cost and scalability. CERT uses innovative platforms for the federal workforce and these models, such as CERT's Virtual Training Environment (VTE) or Exercise Network (XNET), could be emulated to successfully train cyber professionals within the financial community

In conclusion, good cyber security must be built on scientifically sound research and operationally valid data. I have shown where CERT successfully applies these approaches. I believe that such well-founded security enhances privacy. I encourage the Subcommittee to consider policies that promote other such "dual" impact cyber security R&D. If the subcommittee can foster access to data for scientifically valid research, I believe that not only will research be better, but policy makers will have better information for making regulatory decisions. I hope the Subcommittee will encourage policy and research discussions between the security and privacy communities; we all would benefit with improved cyber security and privacy.

Appendix A

Insider Threat Blog
http://www.cert.org/blogs/insider_threat/2010/12/case_trends_for_type_and_status_of_insiders.html

Insider Threat Case Trends for Employee Type and Employment Status

By Insider Threat Team on December 21, 2010 10:45 AM |

We recently met with leaders from the U.S. financial services sector, and they asked a number of questions about recent trends in insider threat activities. We are often asked these types of questions, and we can answer many of them right away. Others require more extensive data mining in our case database. In this entry, we address the following question:

Between current employees, former employees, and contractors,
is one group most likely to commit these crimes?

The answer to this question has some important implications, and not just for these particular meeting attendees. If, across all types of incidents and all sectors, the vast majority of incidents are caused by current, full-time employees, organizations may focus on that group to address the vulnerability. If, on the other hand, there are a large number of part-time contractors or former employees, there may be different controls that an organization should consider using.

Before we discuss the data and its implications, there are some caveats. Our sample of incidents only involves individuals who were caught and prosecuted for their crimes. Also, we currently only have data about incidents that were reported to law enforcement, so these were examples that reached a certain threshold of damages and satisfactory evidence to furnish in a court of law. Finally, it is not entirely accurate to infer from our sample that the results and figures apply to all sectors and all organizations. We are providing these statistics as "food for thought" and to add to the discussion about an important threat that most organizations face.

To develop the answer to the question above, we used 401 cases of all crime types (i.e., IT sabotage, fraud, and theft of intellectual property) spanning all critical infrastructure sectors. Within that sample, there were 85 cases of all types of crime involving only the financial services sector.

The figure below shows the number of cases per year by employee type, constrained to either employee (current or former) or contractor. The graph on the left shows all cases, while the one on the right shows only financial services sector cases.

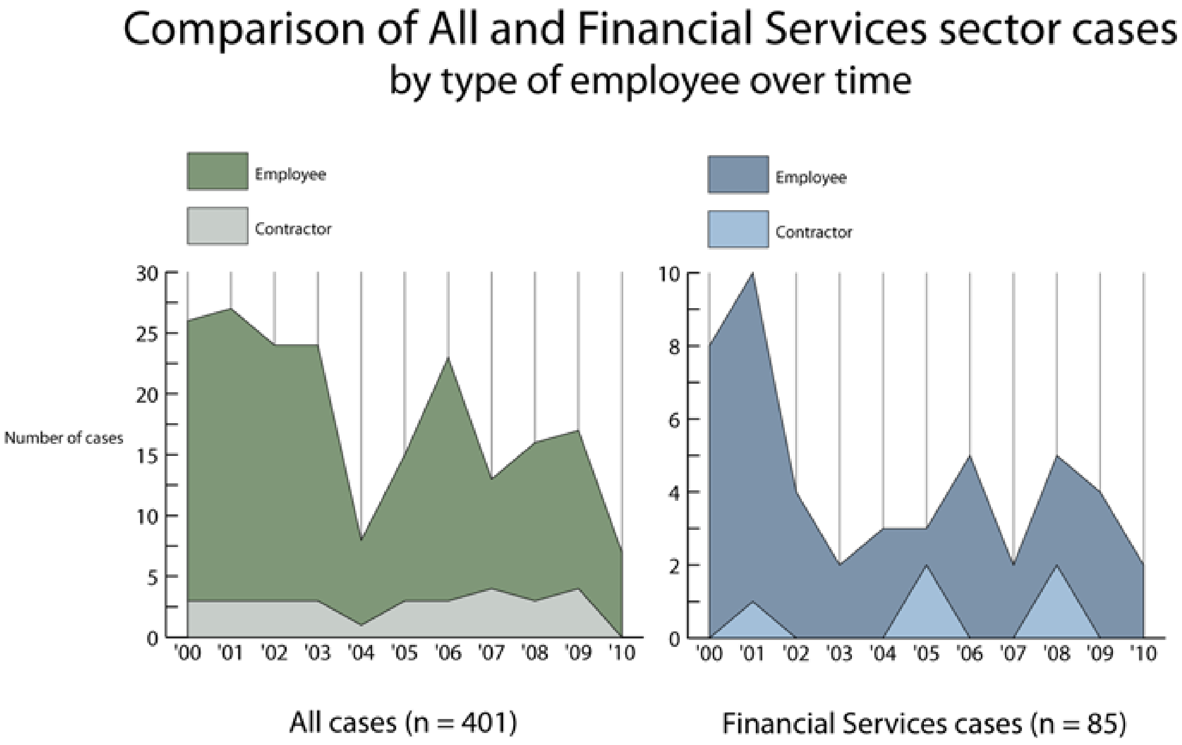

In the ten years shown in the graphs, the percentage of incidents involving a contractor hovers around 15%. Whether the number of total incidents for a particular year is higher or lower, the percentages stay roughly the same. What is most interesting about these two graphs is that this ratio has stayed the same over the course of ten years of a fairly tumultuous economic environment. This result may indicate that it isn't likely for contractor crimes to raise or lower significantly. But with almost 1 in 7 of our insider threat crimes being committed by contractors, are organizations adequately considering the risk posed by this group?

The second figure below shows the percentage of cases perpetrated by current and former employees in all cases and in only the financial services sector. The chart also shows the ratio of employment type (full-time, part-time, or contractor) depicted within each bar. In some cases, we were not certain whether the incident was committed by a current or former employee, so we indicated those incidents as unknown.

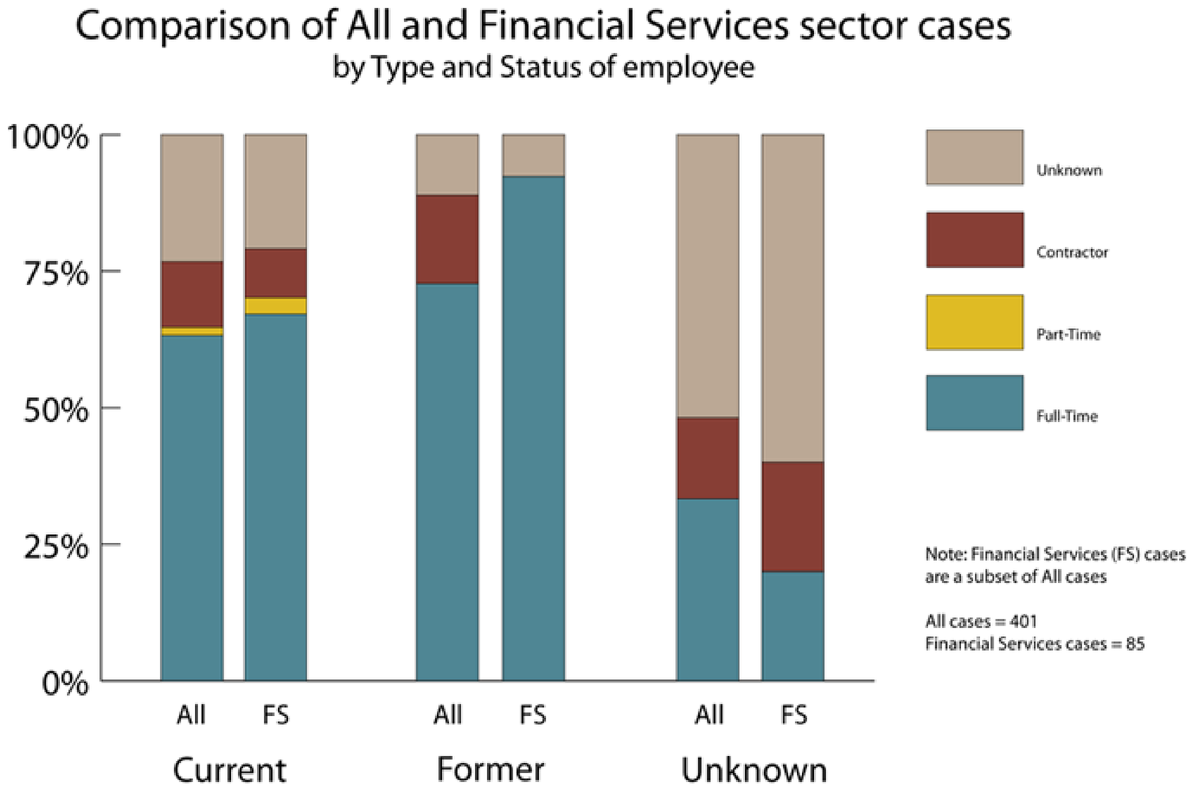

Note: Financial Services (FS) cases are a subset of All cases

All cases = 401
Financial Services cases = 85

At first glance, the financial services sector cases seem to mirror all cases. Full-time employees have the greatest percentage across all sectors for both current and former employees. Part-time employees form a small percentage of our cases across all employee status and types. The contractor results, on the other hand, reveal an interesting trend. For current employees, the percentage is about the same for financial services as all sectors. For former employees, however, 16% of all cases were contractors (indicated in burgundy in the center-left bar), and none of those were in the financial services sector.

These results may be meaningful or may be an artifact of the small number of cases (only 16) of former employees in the financial services sector. Regardless, these graphs provide some interesting data points for you to examine within the context of your own organization. Do you use the same prevention and detection controls for all employees, or are you only worried about the majority—the current, full-time employees you see on a daily basis? Use the feedback link to send us your thoughts.

Insider Threat Case Trends of Technical and Non-Technical Employees

By Insider Threat Team on January 26, 2011 10:08 AM |
http://www.cert.org/blogs/insider_threat/2011/01/insider_threat_case_trends_of_technical_and_non-technical_employees.html

This is the second of two blog entries that explore questions we were asked during a recent meeting with leaders from the U.S. financial services sector. In this entry, we focus on what role malicious insiders typically hold in an organization: a non-technical position, a technical position, or both. "Non-technical" includes positions such as management, sales, and auditors. "Technical" includes positions such as system or database administrators, programmers, and helpdesk employees. "Both" includes overlapping jobs such as IT managers.

The statistics in this entry were generated from the cases that we have collected and observed. Your organization may see a very different breakdown of the positions held by malicious insiders, especially if you have a different allocation of technical and non-technical positions.

In our repository, we have data about the organizational position for perpetrators of 355 malicious insider incidents. Of those cases, 54% held non-technical positions, 41% held technical positions, and 5% held both. Looking specifically at the banking and finance sector, we had employment data for 73 incidents. Of those cases, 66% held non-technical positions, 33% held technical positions, and only 1% held both. It is interesting to note that we did not observe a drastic difference between the position breakdown in the banking and finance incidents and in our larger sample of cases. However, the results do seem to indicate that the majority of crimes that we have observed in banking and finance involve insiders in non-technical positions. If we examine the type of crime for all malicious insider incidents, 40% of the cases are fraud. Within only the banking and finance incidents, the percent of fraud cases increases to 70%. Our research indicates that non-technical employees perpetrate the majority of insider fraud crimes, so the difference in number of fraud cases may account for the increased percentage of non-technical positions within the banking and finance sector.

We also collect data on when the crimes occur, so we can compare technical versus non-technical crimes over the last ten years. Incidents that occurred in 2010 may still be reported, so we did not include 2010 in these graphs. The first graph includes all incidents where we knew the start date of the incident.

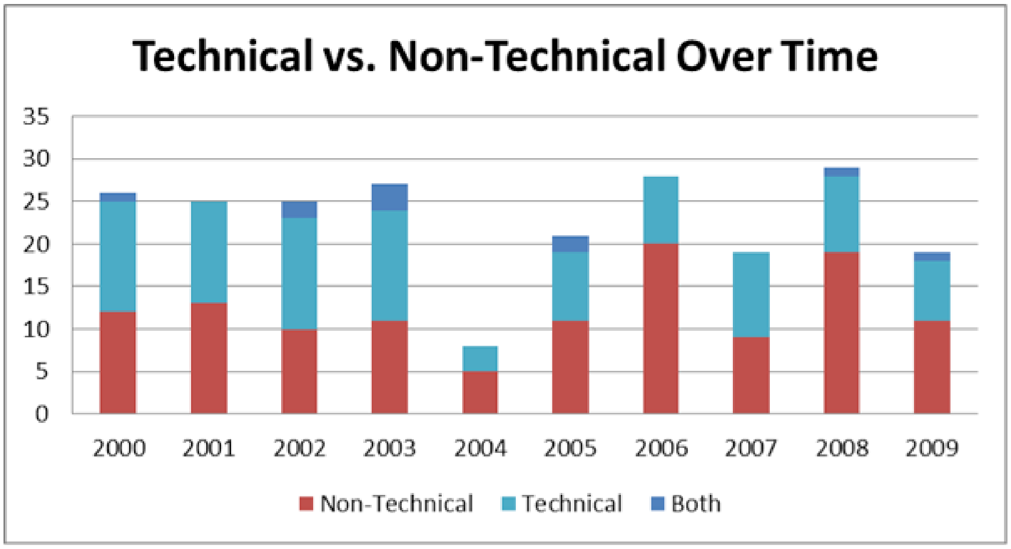

The next graph only includes financial sector incidents where we knew the start date of the incident.

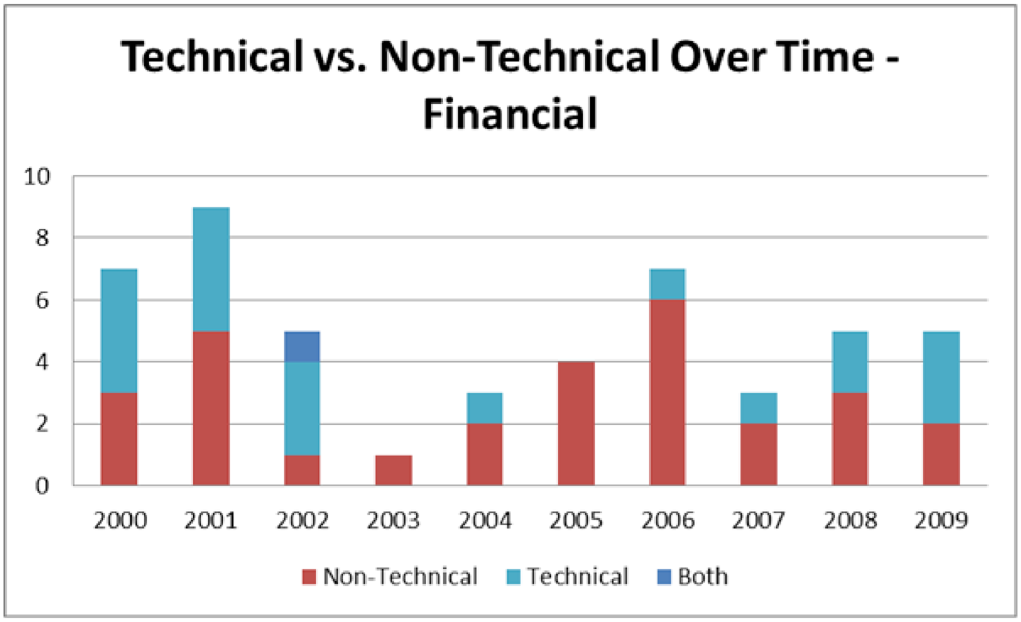

Looking at the graphs, the ebb and flow of technical versus non-technical insiders could follow U.S. economic indicators. The steady increase in non-technical crimes leading up to 2006 in both graphs may coincide with the U.S. economic downturn. Are there other possibilities? Maybe one of you in the financial service sector can compare our timeline of incidents (from this blog entry and our previous blog entry) to some meaningful measures of the U.S. economy or to other general indicators of employee well-being?

Another aspect of this issue is whether damages differ between incidents involving technical versus non-technical insiders. For example, would technical insiders have more access to IT systems and therefore be able to cause more damage? Or would non-technical insiders with much more restricted access but more knowledge of the data in the systems be able to cause more damage? Before we answer these questions, keep in mind that, for now, our case repository only includes cases that organizations report to law enforcement. Therefore, our data might exclude lower damage incidents that organizations handle internally.

In our repository, the average impact between technical and non-technical cases in the financial services sector is relatively similar. The average damages for our technical cases were more than $750,000. The average damages for our non-technical insiders were more than $800,000. (Note: The average value for non-technical incidents does exclude one outlier case of a theft that spanned several years and resulted in almost $700,000,000 worth of damages.)

How has your organization allocated resources for preventing, detecting, and responding to threats posed by technical and non-technical employees? Does your organization focus on one type of employee and not the other? Our observations indicate that there is not a substantial difference between organizational roles of malicious insiders, so organizations must consider each category of employee when implementing security controls. Insider threat could come from anyone.

As always, we welcome your feedback.

United States House of Representatives

Committee on Financial Services

"TRUTH IN TESTIMONY" DISCLOSURE FORM

Clause 2(g) of rule XI of the Rules of the House of Representatives and the Rules of the Committee on Financial Services require the disclosure of the following information. A copy of this form should be attached to your written testimony.

1. Name:	**2. Organization or organizations you are representing:**
Greg Shannon	Carnegie Mellon University's Software Engineering Institute

3. Business Address and telephone number:

[redacted]

4. Have you received any Federal grants or contracts (including any subgrants and subcontracts) since October 1, 2008 related to the subject on which you have been invited to testify?	**5. Have any of the organizations you are representing received any Federal grants or contracts (including any subgrants and subcontracts) since October 1, 2008 related to the subject on which you have been invited to testify?**
□ Yes ☑ No	☑ Yes □ No

6. If you answered .yes. to either item 4 or 5, please list the source and amount of each grant or contract, and indicate whether the recipient of such grant was you or the organization(s) you are representing. You may list additional grants or contracts on additional sheets.

The SEI is a DoD FFRDC with Contract #FA8721-05-C-0003 and an option exercise on 28 June 2010 for $583,892,230.

7. Signature:

Please attach a copy of this form to your written testimony.