# conference reports

This issue's reports on the Cybersecurity, Research, and Disclosure Conference, and on LISA '03.

Our thanks to the summarizers, Cedric Bennett for Cybersecurity and the following for LISA '03:

**Josh Simon**, for once again coordinating the summarizing of the Conference events
Venkata Phani Kiran Achanta
Siddharth Aggarwal
Kenytt Avery
Emma Buneci
Marko Bukovac
Carrie Gates
Robert W. Gill
der.hans
Hernan Lafitte
Jarrod Millman
Bryan Palmo
Ari Pollack
William Reading
Jason Rouse
Kevin Sullivan
Aaron Teche
Steve Wormley

## Cybersecurity, Research, and Disclosure Conference

### Stanford University Law School, Center for Internet and Society, Stanford, California
### November 22, 2003

**Summarized by Cedric Bennett**

Cedric Bennett is an independent consultant specializing in the management of information security in higher education. Most recently, he served as the director for information security services at Stanford University.

*Ced.Bennett@Stanford.edu*

The Center for Internet and Society (CIS) is a public interest technology law and policy program at Stanford Law School and a part of the Law, Science, and Technology Program. The CIS brings together scholars, academics, legislators, students, programmers, security researchers, and scientists to study the interaction of new technologies and the law and to examine how the synergy between the two can either promote or harm public goods such as free speech, privacy, public commons, diversity, and scientific inquiry. The CIS strives as well to improve both technology and law, encouraging decision-makers to design both as a means to further democratic values.

The Web site describing the purpose of this conference at *http://cyberlaw.stanford.edu/security/* as of 12/9/03 read like this:

"This conference explores the relationship between computer security, privacy, and disclosure of information about security vulnerabilities.

September 11th gave new urgency to the debate over whether information collection and dissemination is dangerous or empowering. One view is that vulnerability information should be kept secret and out of the hands of potential criminals and foreign agents. Another view is that the public needs to be informed about security weaknesses, so that people can take appropriate precautions and so that there will be a constituency to pressure for the rapid repair of vulnerabilities. Meanwhile, policy makers struggle to find a balance between promoting security research, constructive information sharing, remediation and protecting commercial interests. Industry has tried to develop 'best practices' for reporting and repairing vulnerabilities, but major disagreements – over how much information to disclose, to whom, and when – persist.

The federal government has tried to both establish standards for commercial entities to share information about vulnerabilities and to pass laws to deter the distribution of information that may enable cyberattacks. However, critics say these initiatives help only a select few, threaten proprietary information, deter legitimate security research and are overly expensive. During the course of this day-long conference, featured speakers and participants will work towards a solution for both industry and government that promotes computer security and addresses the economic, governmental, and social issues that arise under current research and reporting practices."

The format of this conference was a series of brief panel presentations, each considering a particular question related to the conference subject. Discussion followed each panel, facilitated by the session moderator. The summaries represent my own observations and notes. However, some details were gleaned from the blog maintained by the confer-

ence organizers at *http://cyberlaw.stanford.edu/blogs/* as of 11/27/03.

### WELCOME AND INTRODUCTION

Jennifer Granick, Center for Internet and Society (CIS)

In her opening remarks, Ms. Granick talked about the critical nature of computer and network systems and the subsequent importance of security. She touched on a few of the relevant and legal issues, including some of her own experiences, and set the stage for the remainder of the day's discussions:

- Security vulnerability information treated as trade secrets
- Internal emails showing security vulnerabilities in voting systems treated as DMCA copyright violation cases
- Individuals prosecuted and imprisoned for disclosing security vulnerabilities

The "full disclosure" faction argues that disclosing vulnerabilities assists everyone in patching those vulnerabilities and that it facilitates effective risk management. The other side argues that such information out in the open assists those who would do wrong things. She asserted that there was general agreement that responsible disclosure helps security more than it harms it. Of concern around that proposition, however, are the following questions: Who makes that calculation? What factors are considered? What are the long- and short-term costs? Are they worth it?

### PANEL

#### When does disclosure best promote security and minimize exploitations, and how much information should be disclosed at a given point in time, and to whom?

Jennifer Granick, Stanford CIS, moderator; David Litchfield, NGS Software; Tiina Havana, Department of Electrical and Information Engineering, University of Oulu, Finland; Gerhard Eschelbeck, Qualys

Ms. Havana focused on what she called a checklist for designing a vulnerability disclosure policy and considered the political aspects of security and disclosure. "Can we manage to get the process such that there is no need for public regulation?" In her somewhat philosophical presentation, she spoke about the complexity of communicating security discoveries and also about the timing of an information release strategy.

Mr. Litchfield self-identified as an individual who has published proof-of-concept code eventually used in the Slammer worm (because "code is a better way to get an idea across than English"). But he also believes that we must not use the "real" code to illustrate the problem. He believes strongly in the value of responsible disclosure and promotes the guidelines for security vulnerability published by the Organization for Internet Safety. He is concerned that many researchers and vendors do not adhere to those guidelines even while claiming they are part of organizational policy.

He believes it is important to stick to those guidelines not only because they promote responsible disclosure, but also because by doing so an example is set for others. Researchers and vendors have a responsibility to stick to their words about disclosure and repair, respectively. Because of his experience with the Slammer worm, he has publicly declared that he will no longer publish proof-of-concept code, because something he wrote intended for educational use was misused for nefarious purposes – he doesn't want to be part of that. When asked if his code made it that much easier for the black-hat hackers, he said that they are smart enough to write their own code: "If anything, I saved him about 20 minutes." When asked why, if it only saved 20 minutes, he made the moral decision to stop writing such proof-of-concept code, he replied, "Because 20 minutes is still 20 minutes."

Mr. Eschelbeck, a researcher who is developing a database of vulnerabilities and exploits, reported on some of his findings:

- Although there are thousands of vulnerabilities, we only need to worry about approximately 10–15 high-profile ones that cause all the trouble, but those very prevalent vulnerabilities change over time. That is because systems are constantly being modified, installed, and reinstalled, which causes many vulnerabilities to reappear.
- Many, many vulnerabilities remain unpatched for extended periods of time.
- The half-life of a critical vulnerability is 30 days (i.e., it takes about 30 days to clear a critical vulnerability by 50%).
- Some vulnerabilities don't go away (for keeps) – they keep coming back.
- According to the data he has collected, coordinated disclosure is better than uncoordinated disclosure in fighting exploits.

In discussion there was significant agreement among the speakers that vulnerability information should only be publicly released when a patch is available. However, if the vulnerability is discovered in the wild, that is a different situation. There was also some agreement that the threat of vulnerability disclosure has some impact on patch production by vendors.

### PANEL

#### How can independent researchers be adequately compensated for the valuable service they provide to vendors and customers while encouraging responsible reporting?

Chris Sprigman, Stanford CIS Fellow, moderator; Len Sassaman, Anonymizer, Inc.; Chris Wysopal, @Stake

Mr. Wysopal contrasted two kinds of incentives, the academic model and the commercial model. In the academic model, recognition is the reward, as is the sense of contributing to the "com-

mon good." Recognition also leads to gaining a reputation as an expert, which can lead to job offers, book publications, etc. In the commercial model, on the other hand, the rewards can be jobs, time-based value for software vendors (i.e., first-to-market), direct selling of vulnerability information, "bug bounties" by vendors (most recently), and, possibly, government-sponsored research (although this may come with strings attached).

Mr. Sassaman believes that vendors are not motivated to release secure products unless it can be shown to affect their bottom line. He also believes that black-hat malware creators are doing us a favor by bringing vulnerabilities out into the open (and that this is okay because they don't target but just "blast"). Fortunately, although zero-day, targeted exploits are possible, they are not yet common. Motivations of researchers must be considered; what does a researcher gain (or lose) by adhering to vulnerability publication guidelines? We need to structure an environment in which good behavior is rewarded and bad behavior has consequences.

In discussion, the question of ideology as a motivator was raised. Some thought this was best, since it is unstoppable. There was also a question about a rumor that spammers are paying hackers for exploits (and some confirmation that it was true). There was a reiteration of the notion that responsible reporting gets us the greatest good with the least risk.

### PANEL

#### Does the commercialization of security information promote security, or should reporting be an academic or governmental function?

Chris Sprigman, Stanford CIS Fellow, moderator; Shawn Hernan, CERT; Simple Nomad, NMRC; Sunil James, iDEFENSE US

Mr. Sprigman started this panel with the question, "Does commercialization provide motivation sufficient to facilitate discovery, disclosure and security, or does it have the opposite effect?

Mr. Hernan believes in capitalism and free speech. But he believes that societal safety must be considered as well. He used examples of sharing vulnerability information with regard to critical infrastructures such as hospitals, power, and so on. There is lots of evidence that information has value. There is money to be made in commercialization. He objects to a model of restricted information.

He also commented that commercialization of security/vulnerability can dramatically complicate the process of identifying and fixing security problems. He believes there is a certain amount of hubris in the computer/network security community with regard to disclosure. He stated that the CERT mailing list is ~150,000 people; Bugtraq has ~50,000 subscribers;, but an episode of *Friends* draws ~30,000,000 viewers.

Mr. James believes that commercialization is good. A company such as iDEFENSE can provide incentives (payments) to researchers and other contributors and add its own value to the information. He raised the question of trust of vulnerability information which is voluntarily provided.

Mr. Nomad indicated that he was coming from an entirely different perspective and that he had a speech to make that might be considered, but was not intended to be, a rant. He has serious concerns about the role of government in such research. Today, as was mentioned earlier in this conference, the DMCA is being used to prevent disclosure of security vulnerabilities. The USA Patriot Act makes port-scanning into a country's Internet space illegal (possibly an act of war). This is creating an environment that stifles legitimate research.

On the other hand, he is aware of spammers who are paying large sums for exploit code. He has met people in Seattle who make a very comfortable living (six figures) writing spammer exploit code (but won't work for Microsoft "because they are evil").

As a result of this repressive legislation and commercialization, computer and network security is suffering. He prefers the "academic" model of research and believes information should be free. No single reporting model will fit everyone; such a model won't work.

In discussion, the question was asked when there are any situations in which it makes sense to notify the vendor and no one else. Mr. James indicated that they take that into account in their disclosure model. They notify the vendor first and then they notify their customers. He acknowledged that it is a difficult balance to maintain. This raised the question of customer certification: how do they know that their customers are not terrorists or mobsters?

Mr. Nomad asked the rhetorical question, "What if a law was passed that said that every time you discovered a vulnerability, you had to write a worm for it?" Such a law would be highly motivational in getting people to fix their systems.

### PANEL

#### What practices or policies facilitate communication between vendors and researchers? What should the researcher do? What should the vendor do? Should practices differ for small vendors, ISPs or Web site owners?

David Dill, Stanford University, moderator; Steve Lipner, Microsoft; Matt Blaze, AT&T

Mr. Blaze expressed the concern that the premise of the discussion regarding vulnerability disclosure and software patching is like the study of medicine being about the efficient disposal of corpses. His assertion is that we need to write more secure software. But, just as we don't completely understand physics, biology, or chemistry, we don't yet

understand enough about computer code.

He sees this as a research and engineering problem. Those disciplines follow certain rules (scientific methodology) , to wit: You don't just trust me because of my reputation; current work builds upon the work of others, and everything is published (except that you must argue convincingly that your publication contributes significantly to the body of knowledge). He quoted Alfred Hobbs, a figure from the 1850s whose research into and ability to pick various strong-lock mechanisms created a lot of controversy at the time, and suggested that we might learn some lessons from it if we consider it to be about Internet security rather than physical lock mechanisms.

Mr. Lipner described some of the problems of trying to develop and maintain secure code. He works in three time frames:

- Response – to the next bug. Follow responsible rules of telling the vendor and allow them to fix the problem before telling others.
- Release – to cut the vulnerability rate down. He is against the release of concept code, since there is good evidence that it is used by others in exploits.
- New technology – avoid reintroducing old vulnerabilities as well as new ones.

About 10% of the thousands of bug reports they receive reflect real problems, with about 1% actually exposing vulnerabilities. They must look at all of the reports.

In discussion it was noted that no vendor wants to release code with security vulnerabilities. Many more people are harmed by the release of exploit code than benefit from using it for responsible testing. On the other hand, if one person thinks of a clever idea (an exploit), the chances are good that someone else has also.

## PANEL

### How do you motivate the vendor to release more secure software without crippling innovation?

Scott Blake, BindView, moderator; Mary Ann Davidson, Oracle; Bruce Schneier, Counterpane

Mr. Schneier said that transparency is critical (where "transparency" means that we understand the vendor's process for dealing with bugs). He wondered if we need a threat of transparency to make vendors do a better job. He feels it is better to get the top software vendors to introduce effective security ideas than to get hundreds of researchers to do so. Vendors are not stupid and they are not charities, but they need incentives. He believes that society has several "knobs" it can tweak to create those incentives:

- Public exposure
- Competition
- Law (criminal, statute, tort)
- Technology/economics (cheaper to develop more secure code than to fix it later)
- Society (what's "OK")

Ms. Davidson declared that security is not antithetical to innovation (security is not the enemy of feature sets). But it must be built in from the beginning. No one really knows what the cost of a secure system is (yet). Software needs to be better even though it will never be perfect. Moreover, security is always someone else's job (although that is beginning to change). She is concerned about the "L" words (legislation and liability); Congress will do something if industry doesn't.

In motivating vendors, she sees:

- Use of security as one of the purchasing criteria (the Department of Defense does this today)
- Requiring software to have secure conditions set as a default
- Security as becoming a market discriminator
- Big cost avoidance (doing it right the first time – she admits to having trouble convincing her management of that proposition)

She also wonders about:

- A "UL" approach to software security
- A required licensing scheme for programmers (we don't let just anyone build a bridge and test it by letting people drive over it to see if it stays up)
- More education on writing secure software from higher education.
- The development of better tools to automate best practices

In discussion Mr. Schneier said that he doesn't like to see more regulation but that it nevertheless may be a part of the solution. Ms. Davidson also pointed out that the government is a very large customer of software and that regulation isn't the only way they can influence vendor behavior.

One participant suggested that computing has become too "everyman"; marketing has convinced people that they need computers, but those people do not know how to maintain them [ignoring, I thought, that most people do not know how to maintain their cars either, but we don't suggest that's a reason for them not to have them – cb].

It was suggested that developers need to internalize security as a key part of the development. When someone else asked about incentives employers can offer employees to act in that way, the response was: (positively) salary, bonus, stock option, or (negatively) job loss.

## PANEL

### What policies or practices encourage the installation of patches?

Lauren Gelman, Stanford CIS, moderator; Stephanie Fohn, Security Consultant; Vincent Weafer, Symantec

Mr. Weafer feels that patching is a big issue – it is often just a matter of pure numbers (which are large). Given a fixed set of resources, where should one be

allocating energy? More than just patching needs to be considered; setup and other security measures (e.g., firewalls) need to be used. How does the industry deal with the home users with regard to patching? This may be solved either through education or automation.

Patching is very complex, says Ms. Fohn. For example, many companies don't trust patches, others can't find all the computers that need patching, and others believe that they don't need to patch as long as they have a firewall. Until recently, the risk-adjusted cost of patching has been higher than the risk-adjusted cost of not patching. These costs (of patching) have included not only the people-time and tools they use but the risk of making things worse with patches. This comparison has started to shift toward patching because the risk-adjusted cost of not patching is going up. Vendors could helpfully work on reducing the risk of patching (as another incentive to encourage patching).

In discussion, the question of the location of liability was raised. Mr. Weafer feels that it is more on the user than the vendor and can actually be found to be on vendors, users, and the maintenance (IT) folks as well. There was some discussion of automatic patching coming from vendors (a growing trend) and the practicality/usability of that for someone at the other end of a slow connection.

### PANEL

#### What are the practical considerations in formulating, implementing, and enforcing vulnerability disclosure policies or best practices?

Jennifer Granick, Stanford CIS, moderator; Jim Duncan, Cisco; Hal Varian, Haas School of Business, University of California, Berkeley

According to Mr. Varian, it isn't so much the technology as the practices that can be at the root of the problem. He feels that a good model is to assign the liability to the party that is most involved with the risk. (He provided an example of law regarding ATMs: in England, the liability is assigned to the customer – a bad model in his view – but in the US, the liability is assigned to the banks, which he sees as a good model.) However, he feels that strict liability is not optimal; if one party bears all of the cost, the other parties don't have reason to be careful, because they will be compensated if something goes wrong (e.g., Microsoft). A better approach is to apply a negligence rule, where the courts establish a level of due care. If the due-care standard is set well, the parties have incentive to meet that standard as a natural part of doing business.

An alternative and possibly more practical approach is to consider insurance. Insurance companies are basically selling risk management to their customers. In order to obtain the insurance, a company must conform to minimum guidelines (e.g., so many sprinklers per square foot to qualify for fire insurance). In some ways, they are imposing the due-care standards that would be set by courts and are probably better at it because they have their own financial incentives. The problem for cyber-insurance, of course, is that the actuarial databases don't yet exist, and there are no incentives yet in place for such data to be collected.

Mr. Duncan observed that vendors often deal with customers but not the actual consumers of their products. He also discussed "need to know" as an important criterion for disclosure and agreed that there is a need for transparency. We need a way to report the information safely – all the standard (security) rules apply to these transactions. Unfortunately, crypto is hard to use and most people get it wrong.

"Scoring" vulnerabilities is very subjective; everyone does it their own way. This makes measurement impossible. Timeline is another issue; nearly everyone agrees on disclosure but not on the calendar for it (even down to the particular days of the week to avoid when disclosing problems). When the issues cross vendor lines, solving them becomes even more complex. There is a lack of case law and experience, but there is more focus on these problems and we are getting better.

In discussion, there was consideration of "need to know" and of appropriate information sharing among responsible parties as a way to build the knowledge base (e.g., the 12 Federal Reserve banks sharing operating information in a non-competitive way). This was another argument in favor of transparency.

### PANEL

#### What role should legal rules play, and how can the law help or hurt security in the area of vulnerability disclosure?

Greg Schaffer, PricewaterhouseCoopers, moderator; Peter Swire, Professor of Law at Ohio State University; Stephen Wu, InfoSec Law Group

Mr. Swire presented a model for when disclosure helps security (from a book he is writing). This model explores the paradox that there are times when disclosure can be a good thing and other times when disclosure can be a bad thing ("good" and "bad" being defined as helping the defenders and helping the attackers, respectively). In illustrating this model, he contrasted physical examples and software examples. He also asserted that we might want more disclosure just because it helps our general democracy (and might help us with privacy and confidentiality).

Mr. Wu pointed out that there might be liability questions that arise from disclosing vulnerabilities and there might be liability questions that arise from not disclosing, as well (a "damned if you do and damned if you don't" kind of issue). He also raised a question about mandatory reporting requirements. He provided a quick tutorial on the sources of liability (i.e., contract, tort, and statutory

law) for the approximately 60% non-lawyers attending the conference .

In discussion, someone commented on Mr. Swire's model, pointing out that he was distinguishing between the physical world and the software world but that a distinction made between mechanism and instances would have been a better approach. Mr. Swire replied that when considering instances, one must often consider the first instance differently than others (since that will often educate the defenders and change the effect of subsequent instances). This led to a discussion of the ability of the law to operate in this complex arena (and the likelihood, or not, of lawyers staying out of the fray). There seemed to be some agreement that we will have some very confused judges, at least for a while.

#### PANEL

**Brief concluding remarks**

Jennifer Granick, Stanford CIS; Lauren Gelman, Stanford CIS; Scott Blake, BindView; Greg Schaffer, PricewaterhouseCoopers

No one today has argued against the idea that the market has failed to provide security. Instead of capitalism saving us, we are beginning to conclude that there may be a role for government, a conclusion that many of us find both interesting and disturbing.

There are some interesting (legal) questions to be answered with regard to disclosure, nondisclosure, and liability. What if one can become liable for knowing something and not disclosing it?

Security is about more than fixing "this one bug." It could be about democracy. We don't know enough about security to know that it ought to (or not) be considered differently from other scientific enterprises.

Some people think that the disconnect is about Republicans and Democrats, but it is really about the information-technology and legal communities. Both have well-developed models of their universes and like to be the masters of their respective domains. Neither likes the discomfort of not having a handle on important things that apply to their realms. There are lots of people who have not thought about these problems and won't until there is a crisis, and then the decisions are unlikely to be well-considered and thoughtful. There is a serious need for us to think about these problems in advance, as we have been doing today.

## 17th Large Installation Systems Administration Conference (LISA '03)

## San Diego, California October 26–31, 2003

### Keynote Address

**Inside eBay.com: The System Administrator's Perspective**

Paul Kilmartin, eBay, Inc.

*Summarized by Bryan Parno*

Kicking off the 17th annual LISA conference, Paul Kilmartin, eBay's director of availability and performance engineering, gave a spirited and engaging tour of the development of eBay's infrastructure, from a single PC in eBay founder Pierre Omidyar's bedroom to the current SAN-based system composed of hundreds of enterprise-level machines. Along the way, eBay's user population exploded from a few hundred in 1995 to over 85 million today.

Throughout the talk, Kilmartin stressed the incredible importance of availability. Since eBay averages $738 of gross merchandise sales every second, the prospect of any prolonged outage is costly indeed. This intense usage also makes eBay the world's 75th largest economic market, falling somewhere between Uzbekistan and the Dominican Republic. Kilmartin repeatedly emphasized how the magnitude of eBay's 85 million user-base impacts virtually every decision the company makes.

In the historical segment of his talk, Kilmartin highlighted eBay's transition from a system based on two-node Veritas clusters to a large-scale SAN. On the plus side, this cut down on the amount of idle hardware, always an important consideration for cost-conscious administrators. It also provided a greater degree of fault minimization and isolation, since the two-node clusters suffered from electrical issues during servicing. Unfortunately, shortly after the migration to the SAN, the co-location company hosting the site announced it would be going out of business. Kilmartin's team of system administrators built an entirely new SAN in three weeks and made the migration with only two hours of downtime in September of 2001. The bankruptcy of the Exodus storage facility in November of 2001 forced yet another move.

Even though the public perceives eBay as an industry leader, Kilmartin repeatedly emphasized his preference for remaining firmly in the mainstream of technology. On several occasions, he urged the audience to forge on ahead and aggressively report problems, so that after a few years of maturation, eBay could adopt the "new" technology. He offered several tips to the audience, encouraging system administrators to doubt everything, to make the system work hundreds of times before trusting it, and to challenge "best procedures" by at least asking for references. He also emphasized the importance of knowing one's role on the team, citing his initial resistance to eBay's foray into the car market (now, he says, a Corvette sells on eBay every 64 minutes). Kilmartin also stressed the need to constantly seek out a better understanding of the customer and how the customer uses the product. Commenting on hiring decisions, he reminded the audience that neither experience nor certification necessarily equates to competence. Concluding with a return to the theme of availability, Kil-