New Security Paradigms Workshop

White Mountain Hotel and Resort, NH, USA
September 18–21, 2007

Summarized by Matt Bishop (bishop@cs.ucdavis.edu)

The 2007 New Security Paradigms Workshop (see http://www.nspw.org) began with a reception and dinner on Sept. 18 and ended at noon on September 21. The workshop was highly interactive, with participation limited to about 30 people. It encourages authors "to present ideas that might be considered risky in some other forum," and all participants were charged with providing feedback in a constructive manner. The resulting intensive brainstorming proved to be an excellent medium for furthering the development of these ideas.

■ ***Security and Usability: The Gap in Real-World Online Banking***

Mohammad Mannan (presenter) and Paul van Oorschot

This paper examined what banks expected their online customers to do, and how that matched what customers knew they had to do and whether they could do it. The notice that banks give users (typically on the bank's Web site) is small, often overlooked, and contains fine print. As a result, many users are unaware of these expectations. For example, when the researchers asked a group of computer science students, researchers, and professionals how many of the requirements they met, most did not meet them all—and the researchers thought this group would be most likely to know, and meet, those expectations.

Banks expect online customers to have firewalls and antivirus software, and to keep up to date with security patches. But many users are not aware of security problems. The banks also gave misleading information. For example, one bank instructed users to ignore a message about an SSL certificate that failed to verify for its intended purpose. Banks often contracted with third-party firms for marketing purposes, and the resulting URLs looked suspiciously like phishing URLs. Finally, the banking Web sites failed to authenticate themselves to online customers, which contributed to the problem.

The researchers concluded that expecting users to follow the "shared responsibilities" or protecting their banking information was unreasonable given the lack of clarity and the nature of those expectations.

■ ***A Privacy and Security Assurance Offer System***

Jeffrey Hunker (presenter)

Currently, when a provider fails to protect a consumer's private information given to it for a limited purpose, the consumer has to take extensive action to protect him- or herself, while the provider usually faces only the consequences of reputation loss. To better link the responsibility and accountability for security of privacy-related information, this talk suggested an alternative approach, in which the consumer can opt in to one of several privacy guarantees (contracts) for a fee. The provider would have insurance policies supporting these guarantees. If the provider violates the guarantee, the consumer would have appropriate redress (e.g., be financially compensated or receive some other form of restitution). This scheme is a risk management scheme with insurance providing much of the incentive.

Pricing insurance premiums is not an exact science. Some markets do not support pricing risk (e.g., insurance for rock concerts), but insurance companies provide insurance for them. Two approaches enable violations to be detected.

The first is to write the privacy guarantees (contracts) in such a way that violations become clear. The trial bar also has an incentive to detect these problems, because its members can sue for them.

This in many ways resembles an architecture that provides software services rather than software. Finally, if the different privacy guarantees could be structured as a lattice, much of the work done on multilevel security policies may be applicable.

■ ***Authenticated Names***

Stanley Chow (presenter), Christophe Gustav, and Dimitri Vinokurov

This paper tackles the problem of authenticating identity to prevent phishing. For example, if you get a telephone call and have caller ID, the name of the caller is displayed. How can you be sure that it is accurate?

The authors propose a scheme based on the way trademarks are handled.

The RealName scheme defines "local jurisdictions" as geographical or professional groupings of brand names. Each jurisdiction runs a RealName registry that registers brand names. Each registry has its own name space and is authoritative over that name space. When a company registers, the registry gives it a certificate.

A user who wants to verify that a site's claim to belong to a particular brand is true requests the certificate from the site. The user then validates the certificate as coming from a trusted RealName registry. The user will normally trust a small set of registries.

Suppose a user is looking for a particularly unusual item, and a search engine says it is available at the XYZ company. The user finds a Web site for the XYZ company. If the company is in a trusted registry, the user can authenticate the identity to be sure it is the right XYZ company. If not, the user must establish trust in the company and then import its certificate, or establish trust in the company's registry and import its certificate, and then proceed as before.

Various extensions to handle delegation were presented.

■ ***Security Automation Considered Harmful?***

Keith Edwards (presenter), Erika Shehan, and Jennifer Stoll

Conventional wisdom holds that users comprise the weakest link in the security chain, so the system should do as much security management as possible to eliminate this link. The authors' thesis disputes this approach, holding that inappropriate automation is a direct cause of many of the problems associated with "usable security."

Misunderstood social and environmental contexts, mismatched values, and missteps in user experience all limit the effectiveness of automation. Automation implies "one size fits all," but differences in contexts mean different security needs. Because the end user usually uses a preconfigured policy, the user's need for anonymity, for example, can conflict with the preconfigured settings ensuring accountability. Finally, if automated security mechanisms are only "mostly right," the mechanism may call upon the user to disambiguate exceptions (which most home users are not knowledgeable enough to do) or may ignore errors.

The authors recommended exposing the security infrastructure, rather than hiding it for all but exceptional cases, tying security decisions to user actions, and using approaches drawn from social networking. They agreed with a questioner that making the workflow model of the system match the user's mental model of the system would improve the ability to automate appropriately.

They concluded that automation has inherent limitations even if the technology behind it is faultless.

■ ***Self-Healing: Science, Engineering, and Fiction***

Michael Locasto (presenter)

This position paper argued that a self-healing system is a pipe dream, because computers cannot anticipate failure conditions that the programmer does not know about. The standard code for binary search demonstrates this; despite having been proved correct, it had an integer overflow flaw overlooked for 30 years! Further, systems are products of inherently flawed human processes, with constantly shifting demands, and can be physically unreliable.

This thesis distinguishes between restorative healing, which responds to symptoms rather than causes (e.g., the skin healing in response to a cut), and improvement, which repairs the underlying cause of the problem. This is essentially the difference between detecting new instances of known classes of failures (responding to symptoms) and detecting new, previously unknown classes of failures (responding to underlying causes). Which do we expect from self-healing systems?

The discussion identified some limits to self-healing. First, how does a system with inherently incorrect code "self-heal"? On a deeper level, how does the system establish that there is a problem that needs to be healed? This is the same problem as anomaly-based intrusion detection faces: establishing what "normal" means.

Another important question is whether developing self-healing systems is appropriate, given the cost of development and the impact of self-healing mechanisms on efficiency.

PANEL ON THE FUTURE OF BIOLOGICALLY INSPIRED SECURITY: IS THERE ANYTHING LEFT TO LEARN?

Chair: Paul van Oorschot

Panelists: Michael Locasto, Jan Feyereisl, and Anil Somayaji

To foster debate, initially the three panelists took simplified positions. Michael started by saying that we learned much from biological systems; for example, strategies for anomaly-based intrusion detection and response have been

learned from the immune system, and artificial diversity has been inspired by nature. But creating workable computer systems that really are analogous to biological systems has been pretty unsuccessful, because of biology's complexity. We've already learned the big concepts from biology, and so it is time to move on from biologically inspired security.

Jan pointed out that we understand relatively little of how biology works, particularly as security researchers. Medical doctors spend many years studying biology and they still don't understand very many things. There are many biological systems, such as those involved in reproduction, that have not been adequately studied for their security properties. Thus, we have just scratched the surface of the possibilities for biologically inspired security.

Anil then argued that applying biological metaphors to computer security was a mistake for three reasons: It led to poor research because people familiar with biology but not computer security tend to produce poor-quality security research; biological systems address the wrong problems from a security standpoint because they focus on availability for survival rather than integrity or confidentiality; and biological systems are too complex—if we imitate them, we'll produce computer systems that are too hard to understand.

The subsequent discussion was lively. One attendee noted that rejecting analogies to biological systems because they don't solve the security problem ignores the fact that nothing in security works. A reply pointed out that maybe biology got it right to focus on availability over confidentiality or integrity, as availability is what is most important in practice. Others from the audience argued that we do not yet have the right model for translating biology, because we need to take the "ecological context" of living systems into account (the security equivalent being a threat model). Some attendees thought, though, there was some promise in designing fear into computer systems.

A key point raised was the difference between the evolved systems of nature and the designed systems that we have in computer science. Designed systems are always different from evolved systems because they are created by different processes (purposeful design vs. random search). One attendee argued that it is important to pay attention to epistemology here, as there is a big difference between copying a system and being inspired by one. Further, biological systems and computer systems are fundamentally different, and the key question is whether looking into the commonalities is apt to be more fruitful than doing other things. Anil disagreed with this statement, arguing that there is no fundamental difference between designed and evolved systems.

The panel concluded with Michael saying that given that we will soon be engineering biology as we now engineer computers, bioengineering raises critical security issues itself!

■ ***Robustly Secure Computer Systems: A New Security Paradigm of System Discontinuity***

Jon Solworth (presenter)

The theme of this talk was that we need to stop doing what does not work. The problem is that today's systems were designed before the lack of security became a major problem. Over time, new features were introduced and others removed, without thought to the security consequences. Then came the attackers.

The speaker mentioned the usual pitfalls of nonsecure programming, emphasizing that experience shows that only a few programmers have the right mindset to write secure code. The solution is to write new operating systems and programming languages in which these pitfalls are engineered out of the system. He then described the "application trap."

The application trap is a circular trap: No one will use new systems with no applications, but neither will anyone write applications for a new system that has no users. This led to the observation that one could introduce a new operating system that is incompatible with every existing application (the "system discontinuity") but uses virtual machines to support existing application-rich systems with poor security, and new operating systems with few applications but good security. In fact, such development is underway.

Considerable discussion ensued about the nature of flaws and vulnerabilities, and whether remediating them at the operating-system level would fix them at higher levels of abstraction, e.g., in browsers. The conclusion was that a new operating system could improve things, but we need to determine how to build easier-to-use operating systems and programming languages.

■ ***Information Protection via Environmental Data Tethers***

Matt Beaumont-Gay (presenter), Kevin Eustice, and Peter Reiher

A data tether is a mechanism that makes data accessible only when in a secure environment. When a mobile computer containing data is moved out of that environment, either the data is removed or the data is encrypted and the key stored on a secure server and deleted from memory. If the data was in memory, the process must be suspended and memory encrypted, or the process must be terminated. In this way, if the mobile device were stolen, the thief could not access the data. But when the mobile device could access the secure server, the data would then become available again.

Someone made the point that even if an attacker could introduce malware onto the mobile device, the malware could not access the data, as it is either encrypted or nonexistent. There was considerable discussion about the secure environment, which was defined as one in which the secure server could be contacted. Other assumptions were

that the user is nonmalicious but not reliable, the computer is connected to a network when in the secure environment, and it was undesirable or impractical to store the data on the secure server.

The system would have to track information flow, because the policy associated with the data would determine whether the data tether needed to protect the data. Thus, the policy is associated with the data and not with its container (file, etc.). How to do this in a commodity system is one of the research challenges, as is determining the contexts of a secure environment and of the disconnected operation of the laptop. The last point caused a brief discussion of what would happen if the laptop could not reconnect to the server because the latter was unavailable. If the data on the laptop were mission-critical, this could turn the security mechanism into an effective denial of service tool.

■ ***The User Is the Enemy***

Vidyaraman Sankaranarayanan (presenter), Madhusudhanan Chandresakaran, and Shambu Upadhyaya

This position paper argued that user actions should be treated as malicious because users do not follow security best practices, through ignorance or maliciousness or because they are oriented toward immediate performance gains. It then proposed an incentive/penalty system to encourage users to abide by the security policies. Specifically, users who followed the security rules would be rewarded by (for example) allowing them to use programs such as instant messaging services and providing them with more bandwidth; those who failed to do so would be penalized by reducing their quality of service, denying use of some programs, and so forth.

The advantage to this scheme is that it directly addresses what the user does to protect, or weaken, the system. The effect of what actions the user takes is proximate and concrete. It also eliminates the problem of nagging alert boxes that users tend to close without reading or understanding.

This proposal was controversial, producing a heated discussion. Three points emerged. The first concerns the user who wants to comply but cannot; the example cited was the inability to use EndNote without triggering security alerts from Vista. The response was that a better, more stable system and software need to be designed; failure to do so would make the user feel that security is a good idea, but "not in my backyard." The second point was that the language of calling the user an "enemy," "ignorant," and so forth was probably counterproductive, driving users and security personnel farther apart; this led to the opinion that the paper really argued that user actions, and not users, are the enemy. The third point repeated an objection to the data tethers: that the consequences of penalizing the user could cause a catastrophic failure if the user were denied necessary resources because of the penalties. This led to the question of when the computer knows best and agreement that this approach assumes no false positives in detecting the user violating the security policy.

■ ***Computing Under Occupation***

Klaus Kursawe (presenter) and Stefan Katzenbeisser

The computer security battle is going badly, said the authors, and providing large-scale protection against platform compromise is becoming less and less plausible. Using a service-oriented business model, an advanced infrastructure, and high-quality attack programs, and recruiting highly skilled personnel, organized crime is outpacing the defenders, who fight compatibility issues, lack of user awareness, and slow adoption of security mechanisms that are generally not effective enough.

Consequently, we may have to accept the fact that most platforms are under control of some "cybermob" and learn to work under that assumption.

Thus, as defenders we need to make the attackers (ab)use of our systems resource-intensive and uneconomical, while protecting critical assets from attackers who fully control the defenders' PCs. The assumptions making this possible are that users are honest although unwilling or unable to expend resources, attacks do not target a particular individual, and the "attacker" is actually an organization with limited human resources seeking financial gain.

The discussion focused on the new paradigm of computing on systems known to have been compromised. Someone pointed out that this was to a large degree a social problem. Other suggestions revolved around mitigation techniques that would limit the gains of the attackers, but many of these also functioned as denial of service attacks against the legitimate users. The talk concluded by suggesting that the security war may be lost already, and we need to find ways to continue to use our systems by better understanding the internal structure of the attackers, our own assets, and how to use both to make the attacker's life hell.

■ ***VideoTicket: Detecting Identity Fraud Attempts via Audiovisual Certificates and Signatures***

Deholo Nali (presenter), Paul van Oorschot, and Andy Adler

This paper presented a method that helps detect identity fraud attempts by embedding audiovisual information in certificates and using audiovisual recordings in lieu of conventional user digital signatures. An av-cert is a signed audiovisual recording in which a user identifies him- or herself. An av-signature is an audiovisual recording in which the user gives consent to a particular transaction. A bank issues the user an av-cert. To purchase something, the user gives the av-cert and av-signature to the retailer, who passes both to a verifier. The verifier validates both the signature and the certificate and sends the authorization and authentication status to the retailer, who (assuming both are good) provides the services or goods to the user.

This scheme verifies identity by using biometrics and verifies the consent for the transaction. It works with both onsite and remote transactions, and it uses widely deployed tools such as Web cameras. Drawbacks include issues of privacy and questions about whether biometric identification is accurate enough to make this technique cost-effective. The automated method is inexpensive (with amortized cost of 5 cents per transaction over a three-year period), but negatives require manual intervention to determine if the negative is false, and this drives the price of a transaction up considerably (to nearly $5 per negative transaction). People pointed out thay although facial recognition mechanisms were quite accurate under laboratory conditions, when deployed in the field their accuracy was considerably more problematic.

Assuming the videoticket approach proves feasible, it shifts the risk of the transaction from the bank and retailer to the user; to compromise the transaction an attacker must coerce the user into performing the transaction. Also, given the state of the art, it is possible that an attacker could generate a human image good enough to fool current automated audio and visual biometric tools within the next five years. To address this issue, audiovisual signatures could include transaction-specific information (unpredictable by attackers). Finally, one participant pointed out that his evil twin brother Skippy might be able to impersonate him and carry out the transaction.

Statement of Ownership, Management, and Circulation, 10/1/07

Title: *;login:* Pub. No. 0008-334. Frequency: Bimonthly. Subscription price $120.
Office of publication: USENIX Association, 2560 Ninth Street, Suite 215, Berkeley, CA 94710.
Headquarters of General Business Office Of Publisher: Same. Publisher: Same.
Editor: Rik Farrow; Managing Editor: Jane-Ellen Long, located at office of publication.
Owner: USENIX Association. Mailing address: As above.
Known bondholders, mortgagees, and other security holders owning or holding 1 percent or more of total amount of bonds, mortgages, or other securities: None.
The purpose, function, and nonprofit status of this organization and the exempt status for federal income tax purposes have not changed during the preceding 12 months.

Extent and nature of circulation	*Average no. copies each issue during preceding 12 months*	*No. copies of single issue (Oct. 2007) published nearest to filing date of 10/1/07*
A. Total number of copies	7083	6825
B. Paid circulation		
Outside-county mail subscriptions	3785	3709
In-county subscriptions	0	0
Other non-USPS parcel distribution	1747	1734
Other classes	0	0
C. Total paid distribution	5532	5443
D. Free distribution by mail		
Outside-county	0	0
In-county	0	0
Other classes mailed through the USPS	57	54
E. Free distribution outside the mail	1146	678
F. Total free distribution	1203	733
G. Total distribution	6735	6175
H. Copies not distributed	348	650
I. Total	7083	6825
Percent Paid and/or Requested Circulation	83%	89%

I certify that the statements made by me above are correct and complete.
Jane-Ellen Long, Managing Editor