



Dejan Milojicic
Hewlett Packard Laboratories
1501 Page Mill Rd., MS 3U-18
Palo Alto, CA 94304
dejan@hpl.hp.com

Security and privacy

Security and privacy have received a lot of attention recently. Fueled by the worldwide availability of information, restricting access to your own data becomes an issue relevant to both business and personal privacy. At the same time, the variety of attacks on and breaches of this privacy has become a real and common danger to everyone.

The impact of being able to access any information from anywhere in the world is revolutionary and creates many opportunities for mankind. Unfortunately, it has also opened up opportunities for criminal and unethical individuals and groups. It lets people from countries with unregulated cyberlaws attack businesses and intrude on other people's privacy.

Early attacks on government and bank data were rather limited geographically or through the use of controlled channels, such as leased lines. They were also less visible and, once detected, easy to counteract. Today, attacks on any prominent Internet service can easily impact your needs, making them more noticeable.

The impact on businesses can be even more dramatic. Recent denial of service attacks on Yahoo and other Internet companies cost millions of dollars. The

companies that rely on the Internet for everyday business can not afford to go offline, let alone for long periods of time due to security breaches. Even as I was writing this introduction another virus—actually a worm—appeared on the Internet: ILOVEYOU. If started, it will, among other things, overwrite some system configuration files and then send e-mails to all the e-mail aliases in your Outlook address book, spreading itself further. This kind of virus both breaches security (it can destroy your configuration) and privacy (it uses your e-mail aliases and can hence send any other information).

An even bigger problem is related to privacy. The Internet lets you easily access your own data but also lets others track your behavior, movement, and habits. The implications are numerous. Many companies realize that this is valuable information that they can trade to track what you buy and browse, from where you connect to the Internet, and so on. The world of wireless communication with all kinds of gadgets and innovative toys will open up even more opportunities for attacks.

All of this creates many challenges for security and privacy experts today. Some

are technical, such as scalability, more efficient cryptographic techniques, safe languages, and better secret key distribution. Social and legal aspects pose more challenges. People must be educated about what is socially acceptable behavior on the Internet, what is criminal, and what is legally binding. We need some strictly enforced regulations as well as some commonly accepted unwritten rules as soon as possible. Unfortunately, not all governments necessarily have the same interests as users.

I've raised only some of the challenges of the security and privacy field. These interviews bring up many more interesting examples from the top experts in the field. Dan Geer is the CTO of @stake and the president of Usenix. Li Gong is the Director of Server Products at the Consumer and Embedded division of Sun Microsystems. Clifford Neuman is a professor at the University of Southern California and chief scientist at Cyber-Safe. Marcus Ranum is the CTO of Network Flight Recorder. Mary Ellen Zurko is a security architect at Iris Associates and is currently responsible for Lotus Notes' active content security.

—Dejan Milojicic



Dan Geer

What are the crucial innovations in the history of security and privacy?

The crucial innovations are related to our ability to make boundaries where there are none. We understand this now, but we

haven't for very long. In the electronic world, there are no natural boundaries, whether we are talking about national boundaries, corporate boundaries, or even the difference between my e-mail and me and my life and my privacy. All of those things are, at least at the outset, undifferentiated, because there are almost no costs to reproducing electronic information, and there are almost no boundaries to contain it. Clearly, you are aware of when you cross the border between Germany and Poland, walk

out your front door, or retrieve a document from a filing cabinet inside a law office. Those physical boundaries are clear and we base our expectations on them—what we know about the physical world and its boundaries. There aren't such distinctions in the electronic world. In effect, recreating those boundaries is the job of electronic security. Unlike in the physical world, we add value in the electronic sphere when we make boundaries around information by prohibiting it from being reproduced.

Everything we are used to thinking about in the physical world is reversed in the electronic world. You win by creating boundaries and increasing reproduction, or at least access, costs—that is what develops a sense of security. When I log into a machine, I want to know that my information is not being inadvertently shared with others. When I send an e-mail, I want to know that it is not being stolen, copied, or intercepted during transmission. When I trade stock online, I need to know that when I say “Buy 100” that is what happens and not some other transaction.

All the innovations here are based on some kind of cryptographic-like function. Only by recreating “brick-wall secrets” are we going to be able to produce these necessary adjuncts to what otherwise is a seamless sphere of information.

What are the most important remaining problems?

The number of places in which we can collect information is increasing. This is natural because the electronic sphere touches more and more things. For example, within five years, the majority of computers will have no keyboard, because systems are going to be embedded in everything. Every appliance of any size will have some sort of information gizmo in it and by 2003, high-end cars will come with a network address.

I can go on and on about the number of data sources—it's just going to skyrocket. The result is that the necessity of security technology is going to rise, probably faster than we (security experts) can stimulate a broad understanding of the risks and trade-offs. Our major challenge will be taking what we know now and applying it to the world with an order for magnitude more data handlers connected to an order of magnitude more sources of information, most of which are automated collection.

Will there be a whole paradigm shift?

As a good engineer, your first responsibility is to correct problems. Your next responsibility is to provide the minimal solution, to not foliate your solution with excess baggage. We just don't understand the security problem well enough, yet the necessity for early solutions is strong. I want to use the word scalability, but I think our definition for that word will be under some stress as we consider what scale means. In particular, what are the default behaviors? Getting the minimal solution is typically a hard thing to figure out. Although Metcalf's famous law—the idea that a network has value in proportion to the square of the number of participants—is doubtlessly correct, it is correct only for the magnitude. The sign bit is whether the communications are or are not of value, whether they were intentional or not, whether you wanted that information shared or not. That's the sign bit you multiply by the magnitude pro-

Questions

1. What are the crucial innovations in the in the history of security and privacy?
2. What are the most important remaining problems?
3. What are the other fields that security and privacy have most impact on? What are the fields that impact security and privacy the most?
4. What are the controversial aspects of security and privacy?
5. What are the current trends in this field?
6. Who should govern security standards?
7. Who are the driving factors of security and privacy today (for example, startups, open-source companies, government, big companies, standard bodies, and R&D)?
8. Should or could security and privacy standards and practices be common worldwide, or should they follow state boundaries?

portional to the square of the number of participants. With the Internet node count doubling every six months, the magnitude rises that order of magnitude every 20. If you are at all worried about the security and privacy aspects of data fusion, then you'll have to factor in Moore's Law—the cost of computing drops by half every 18 months—as well.

What are the other fields that security and privacy have the most impact on?

If you were to walk up to most people and say, “I am sure you have some information that you would rather I didn't have access to. What is it?” they would answer some combination of medical or financial records. It's natural that the medical and finance fields are places where security and privacy concerns are important. I submit that active record-keeping creates risk. For example, some people put a transponder in their cars for toll roads, so they can whiz by the toll booths and pay at the end of the month; those are obviously interesting records that can be subpoenaed. The records are accumulated accidentally, as the side effect of another solution. I submit that we have already determined the market price for privacy in the grocery store and that it is precisely equal to the discount that an affinity card pays you. The time you save at the gas pump when you link your credit card to your key ring is the market price for privacy at the gas pump. Click-stream data has a hundred variations on this theme. Medical surveillance technologies will move overnight from novelty to necessity and each will create strong security issues and mountains of incidentally collected data. I tell people that if data is collected, you should assume it is forever. Ask anyone who has extensively posted on Usenet how he or she felt after discovering that all posts for all time have been archived in searchable form?

What do you think about operating systems? Do you think there is any chance there of advancing security?

We can do a lot to avoid the misappropriation of legitimate authority at the operating system level. We can also do a lot with operating systems in the sense of acquiring illegitimate authority. However, we can do much less about the misuse of data after it is outside the hands of the entity who collected it.

On the Internet, every sociopath is your next-door neighbor, and operating systems are our first defense, as a low wall against amateurs and sociopaths.

Operating systems today don't do enough, but I think the insurance industry will help fix that. It will insist on the same standards of a quality it uses for other industries. It will expect the same kind of cut-and-dry rulemaking to apply in the electronics sphere, or it will not underwrite the risks. Even though the growth rate of e-commerce is spectacular, it is not a substantial part of the world economy yet. When it is, people with billions of dollars on the line are going to want the kind of loss protection that the insurance industry can provide. I don't think that industry is going to let its underwriting standards collapse just because it seems hard or inconvenient to set up the proper security in the e-world.

What are the driving factors of security and privacy today? Start-ups, government, big companies?

The driving force is in new technology. The adoption curve is based on large telecommunication providers. What do we do for security of the whole permanent Internet connection? What do the DSL or cable modem guys do? What are sold as home devices? That is where the adoption side of it is. The production side is in the start-ups, the ones who are not unduly fettered by backward compatibility and an installed base, putting aside here the issue that IBM is still getting 10% of all patents granted.

Should or could standards and common practices of security and privacy be common worldwide or should they fall into the same boundaries?

I want the interoperability standards to be common worldwide, but I want them to be in the form of alternatives that we can adopt but that don't preclude us talking to other people. The key word is interoperability—function and not policy is what belongs in standards. I really do not want a worldwide standard for anything else. Anybody who values diversity of any sort whatsoever must recognize the value in having differences that do not preclude communication. I am not a fan of a worldwide standard in policy matters or kowtowing any monopoly. Whether it is a company or a government, having a monopolistic viewpoint about what does or does not constitute proper use guarantees trade-offs that eventually become dissatisfying. As a wise professor once told me, not every problem has a good solution but every solution has side effects. If you take home just the one solution, those side effects will eventually dominate people's thinking, such that you will end up treating the side effects rather than rethinking the original problem.

Daniel E. Geer, Jr., is the chief technologist officer at @stake, a privately-held confidential security consulting firm. He has previously worked as the vice president and senior strategist for CertCo, a director of engineering at Open Market, as chief scientist, VP of Technology, managing director for OpenVision Technologies, technical director within Digital Equipment Corporation's research division, and the manager of systems development for MIT's Project Athena. He has a BS in electrical engineering and computer science from MIT and a PhD in biostatistics from

Harvard University. He was involved in medical computing for 15 years before moving into distributed systems with the Athena Project. He has testified before Congress by invitation on several occasions and currently serves as a member of the Federal Trade Commission's Advisory Committee on Access and Security. Contact him at @Stake, One Kendall Square, Bldg. 200, 2nd Floor, Cambridge, MA 02139; geer@world.std.com.



Li Gong

What are the crucial innovations in the history of security and privacy?

If I put a filter on what is actually used on a large scale, I can think of three different things that have really improved security for commercial products. The first is public key systems, which form the basis of the SSL (Secure Sockets Layer) that is used in all major Web browsers and is the secure communication mechanism that underpins Web-based commerce. The second is one-time passwords, used often in the form of secure tokens for remote login. This is a significant advance compared to transmitting plain passwords over the Net. The third is the concept of the reference monitor, which was designed in the 1970s for securing standalone systems but is such a fundamental concept that it is still the way people think about how to build security today.

What are the most important remaining problems?

There is a technical one and a social one. The technical problem is how to handle denial of service. The recent attacks on Yahoo and other Internet sites show that there is no good solution currently available that can prevent people from attacking your system and disabling its service. The social problem is the huge gap between what is possible in computer security theory and what is realized in practice. A major problem is the user interface; security cannot be fully automated, but it is very difficult to present users with choices that are meaningful to them.

What about legacy applications and those that were made without any notion of security? Do you think it's easy for security experts to subsequently make them security aware?

No, it's almost impossible. I have never seen a case where you can retrofit security nicely and easily. If you don't have security, you start building so many dependencies into the system that the moment you start to put security in, you'll break the dependencies and thus the product's backward compatibility. In the commercial world, backward compatibility is often a critical need—people don't want to release a new version of the system that breaks existing applications, so the end result is that security stays off the product feature list.

What other fields have security and privacy impacted?

Security has a big impact on the Internet because business models (in addition to e-mail and research) are driving what we do on the Web, especially e-commerce. Without commerce, the

Web would not be what it is today. Ebay, Yahoo, E-Trade, and other such businesses make e-commerce possible: without a minimum level of security no one will use the Web to buy things. That minimum level is pretty low, but the impact is huge.

What are the controversial aspects of security and privacy?

One main problem for me is who pays for security. That point is controversial for many people. Most won't pay that much (or anything) for it. Even people in commercial companies say it is important, but in reality nobody puts enough money behind it. This translates into underfunding for technology and product development. Security features in any product tend to be the last to be put in and the first to get cut under time or budget pressure. It is a loop that security professionals have to break.

What are the current trends in this field?

One trend I see is the emphasis on the whole system's security. Secrecy, integrity, and availability are the three elements of a secure system. Traditionally, people study these areas separately and develop solutions that are not applicable to all areas. I see people starting to pay attention to overall system security, but a lot remains to be done.

Who should govern security standards?

Industry seems to be a good candidate for setting technical standards for security. Governments have a role in driving the need for security standards, but they should not be prescribing the technical specifications. Military computer security in the 1970s and 1980s was a disaster in that most of the work that came out of that area is unusable today.

What are the driving factors of security and privacy today?

I am biased toward deployment. In that regard, start-ups seem to be a major force. They decide what is needed and just do it—they have the money and motivation, and they go for it. The downside to this is that not enough security experts or security professionals are around, so many start-ups tend to get a part-time consultant or don't have a consultant at all. They tend to hack up something that is not really secure. But at least they're doing something. Government has been active in some areas, but I don't see standards bodies or R&D labs being real drivers with significant impacts.

Should standards and common practices of security and privacy be unique worldwide or should they follow state boundaries?

I tend to think they should be worldwide—state boundaries seem to be accidental. There are some natural qualities attached to boundaries, but I don't think there should be any politics behind it.

Li Gong is a distinguished engineer and the director of Server Products at the Consumer and Embedded Division of Sun Microsystems. He chairs the Java Expert Group at Open Service Gateway Initiative, and he's an associate editor of *ACM Transactions on Information and System Security*. He serves on the editorial board of *IEEE Internet Computing* and the *Journal of*

Computer Security. He was awarded the Leonard G. Abraham Prize by the IEEE Communications Society in 1994, and he published a book last year entitled *Inside Java 2 Platform Security* (Addison-Wesley, 1999). He received his BS and MS from Tsinghua University, Beijing, and his PhD from the University of Cambridge, England. Contact him at li.gong@sun.com.



Marcus Ranum

What are the crucial innovations in the history of security and privacy?

Brilliant people have produced some terrific technical innovations, such as public key encryption and opening up the art of privacy to the masses. But the main issues in security privacy, in my experience, are social. I tried to promulgate a law once called Ranum's Law: "You can't solve social problems with software." Most of what we're trying to do with computer security and privacy is handle social maladaptions using software techniques. That's difficult.

What do you expect the future to bring?

The real question is, can we come up with a technological innovation for security that is so broadly installed and so broadly installable that it begins to have some leverage? It's sobering to consider that 95% or 98% of the world's computers run Windows, and most of those don't even have a security model. I started my career in security building firewalls in the late 1980s; today, most Internet-connected sites and systems still aren't using them. Any security expert will tell you that you're an idiot to connect to the Internet without one. There's a long way to go.

We need to dramatically change people's attitude toward cyber crime. Like in the Wild West days of the US, kids want to be gunfighters instead of the town sheriff, because the media has portrayed cyber crime as cool. In today's job market, unfortunately, if you have a criminal record as a cyber criminal or hacker, you can get a top job as a consultant based solely on the fact that you were incompetent enough to get caught, tried, and convicted. This is a slap in the face to network security professionals who stay on the good side of the dividing line. Our social problems and our attitudes toward those problems will delay the uptake of technology.

So we also need to educate the whole community?

Yes. Normally, hackers treat the media as allies, but in the recent denial-of-service attacks, the bad guys interfered with CNN.com. Suddenly, the media is giving them a less favorable message. They're no longer brilliant whiz kids looking for something to do; they're cyber criminals.

What are the other fields that security and privacy have the most impact on? And vice versa?

Anything to do with e-commerce. One of the crimes of the

future will be identity theft; it's already happening. Network users feel anonymous, so they are much more comfortable misbehaving.

Could other fields such as operating systems, programming languages, or networking help you do your job better?

A lot of networking guys had great hopes for IPv6, so we'd know who we're talking to and we'd be able to keep our transactions secure. But it's not getting deployed fast enough, and even if it does, there will be too many flaws in the software running on top.

New technologies are always narrowly applied. The trick is twofold: getting it applied at all, and in a way that it functions favorably with other pieces of the Internet. It's like buying a Ferrari and then using it to tow a trailer. Our whole industry has this problem: we keep carrying our baggage from years past because it's too expensive to reengineer everything. From a security standpoint, it is essential to reengineer everything or at least make a major break with backwards compatibility.

Can regulation help in that regard?

Someday regulation will come in and begin to be effective. I've been dismayed because, for example, US organizations responsible for overseeing the buying and selling of stocks have been incredibly laggard about requiring brokerages to have a certain level of security in online trading. The question is not whether there will be more regulation; it's whether the regulation will wind up being cost-justified.

What are the controversial aspects of security and privacy?

The most controversial question is the government's role. Some sectors of the US federal government have decided that Internet security is a critical issue. Sooner or later, terrorists will use the Internet as a tool. But what about people's privacy? Will we have an environment where people can do the equivalent of electronically strip-searching you? My guess is that's going to happen.

Inevitably, in the next five to 10 years, the level of accountability on the Internet will increase by a factor of 10 or more. You've got two extremely powerful forces, commerce and government, and they're both going to take steps to increase accountability. By definition, that will erode privacy. Last but not least, there's a third force in play. Marketing has realized that the Internet is a tremendous tool for collating information about customer buying patterns. These three things taken together will be controversial, because people will give up an awful lot of their privacy when they sign up for Internet access. What's frustrating to me as a privacy advocate and security guy is that most people won't even think about it or notice it.

If I were contracted to improve Internet security, had a bud-

get, and were unconcerned about the opposing side's viewpoint, the first thing I'd do would be to set up hacker sites. I want kids who spend their time at hacker sites to wonder who they're actually swapping information with.

What are the current trends in the field?

Obviously, we're spending more and more money—Forrester Research predicts about \$2 billion worth of security software and services in 2002—but I'm not sure how much progress is being made. Insularity will increase, with each organization fending for itself. If everybody were working together, we'd make better progress, but it's politically and technologically infeasible right now.

The growth of start-ups is one thing I do see as good. The old-world security establishment moves too slowly; it could never keep up with the rate of innovation required by the Internet time phenomenon. I've seen a lot of start-ups, including mine, say, "Our customers have a problem. We have a solution we can put out and sell. We know we can do something."

Is it going to be perfect? Absolutely not. Is it going to be effective, save our customers money, and help them with their problems? Yes. Start-ups and fast-moving established players that capitalize on technology problems will dominate security in the next five to 10 years. The downside is that the solutions will be based on what the start-ups want to produce rather than on an organized plan. There won't be any kind of central architecture. Anyone who has really delved into security

will tell you that without a centralized view of how things are supposed to work, you don't get great results.

Can you move the general technology further along by solving a particular customer's problems? Will your company have additional bandwidth, maybe not now but in the near future, that you can put aside to push the envelope further in general?

My company and everyone else in this industry will do whatever they can to push the envelope as long as it's profitable. That's one of the downsides of the start-up thing: it's driven by shareholder expectations, not philanthropy. Look at the IETF (Internet Engineering Task Force), a great group of people with a terrific tradition as technologists. They're mired down in technical arguments that a start-up can cut in half in a second by saying, "Which one can we ship faster? We've got money to make."

A couple years ago, I proposed a crazy idea that I think would actually work: Take the \$600 million then being spent on firewalls and use it to rewrite all the Internet applications we use. Do a complete open-source implementation of a secure file transfer protocol and a secure stream protocol, then layer a public key infrastructure on top of that, and then a browser, file transfer, remote terminal emulation, mail transfer, and so on. You'd get a secure, portable, free infrastructure that costs a lot less than

\$600 million, but there's no financial incentive to build it. If anyone tried to pull the necessary players together to organize it, they'd immediately be derailed by commercial interests that saw their meal ticket threatened and standards bodies who saw their roles as technologists threatened.

Who should govern security standards?

I don't believe that the government is competent enough to do it, although they probably should be, or that the start-up community is organized enough. The only standards start-ups want recognized is their own, and customers are not organized enough to mandate or create standards. Standards will likely get written piecemeal, through industry groups.

Unless they're burned like CNN was, by a couple of litigations.

Exactly. We're stuck in the cycle of getting burned, having a knee-jerk reaction, and then ignoring the problem until we get burned again. In the long run, you can't build a standard that way.

Should or could standards and common practices in security and privacy be worldwide?

They must be worldwide. Even if the US government could implement great security nationally, we would still have to do business electronically with people outside the community of our standard. Suddenly, all your guarantees go out the window. If you get a connection from someone at an ISP in another country, how do you know that the ISP didn't sell that account to somebody from a third country that you're not supposed to do business with? You can't.

To be effective, policy must be applied globally and evenhandedly. As soon as you have one weak link, the entire chain essentially becomes untrustworthy. Of course, I would not bet on achieving this.

Do you have anything to add?

I just want to wrap up our discussion about hacking. In the past, security guys had pretty good lines of communication, and information sharing was on their side; now, thousands of hacker sites trade attack tools. Wanna-be hackers who really don't know anything can take down CNN.com with tools they downloaded without even having to understand them.

I predict we'll see some knee-jerk reaction that will shut down a few of these sites and maybe teach a few guys a lesson. Distribution of hacking tools might become criminalized, and frankly, that wouldn't bother me at all. There are a lot of precedents.

We can make a huge step forward in Internet security by not allowing anyone to connect to the Internet unless they somehow sign every piece of traffic they generate—which, of course, means completely revamping the infrastructure of all our software. This is the right way to do it, but I don't think it will happen.

Marcus Ranum is CTO of Network Flight Recorder and has been specializing in Internet security since he built the first commercial firewall product in 1989. He has acted as chief architect and implementer of several other notable security systems, including the TIS firewall toolkit, TIS

Gauntlet firewall, whitehouse.gov, and the Network Flight Recorder. Ranum frequently lectures on Internet security issues and is coauthor with Avi Rubin and Dan Geer of *Web Security Sourcebook* (John Wiley & Sons, 1997). Contact him at Network Flight Recorder, 1395 Piccard Dr., Ste. 230, Rockville, MD 20850; mjr@nfr.net; www.nfr.net.



Clifford Neuman

What are the crucial innovations in the history of security and privacy?

Cryptography was perhaps the earliest crucial innovation in security and privacy.

Dating back to ancient times, cryptography provides the means to scramble information so that no one except the intended recipient can view it. This is done using a transformation that is parameterized using an encryption key.

Most recent innovations in security and privacy apply cryptography in some way. Among them is the use of cryptography for authentication, which lets you determine who originated a message and verify that the information you received is the same as what was sent.

Public-key cryptography and digital signatures extend the capabilities of cryptography further, enabling simpler management of encryption keys and providing several new protections in the security services cryptography enables.

Security and privacy also requires process and data protection as they reside on a local computer system—rather than protecting it as it is communicated over a network. The most significant innovation in protecting information within a processor was the invention of privileged mode and address space protection through virtual memory. Privileged mode lets a processor enforce privilege limitations by redirecting security critical operations to privileged code in the system kernel. Before performing the requested operation, the kernel code checks user or process privileges to determine if the operation is authorized. Virtual memory maps addresses specified by a program to physical memory locations for which the process is authorized access. Modifying the memory mapping is a privileged operation subject to security checks by the kernel. As long as the appropriate page resides in physical memory, the processor directly handles accessing memory mapped through the table as a user-mode operation.

What other fields do security and privacy most impact? What fields impact them most?

Security has pervasively affected many aspects of computer systems. It has a profound impact on e-commerce, and it has affected the data storage and communication components of computer systems.

Mathematics has had significant impact on security and privacy. Much of cryptography is based on mathematics, and mathematics and logic have affected the ability to analyze programs and protocols to prove correctness. Computer architecture has

significant impact on local protection of computer systems. Privileged mode and virtual memory were both architectural innovations. On the negative side, as computer architectures became less centralized, new vulnerabilities were created, and these vulnerabilities were subsequently addressed through cryptographic means. In the coming years, deployment and integration of smart cards and related devices (enabled by the shrinking size of processors) will improve security for decentralized systems and e-commerce.

What are the most important remaining problems?

We need to come up with a basic policy framework for security, which must make it easier to integrate security with applications. Such integration can't require significant modifications to the applications each time we come up with a new kind of policy, and it shouldn't require changes beyond what is inherently necessary to apply the policies. The application must have the ability to ask, "Can the person I am talking to do what was just requested?" and it should be able to get a yes or no answer. We should not have to modify the application to go through the authentication exchanges.

One of the downsides of the traditional approach of integrating services such as authentication at the application layer is that there are policies that don't even depend on authentication. If you implement a policy based on authentication, you are stuck with authentication based policies for your application. If you decide later to implement an e-commerce policy based on payment—not identity—then you have to modify the application to support the new policy. Whereas if the only point of integration between the application and the security mechanisms were the simple question, "Can the user I am talking to do this?" then we can add support for new kinds of policies and new security mechanisms without further changes to the application.

If you look at most of the deployed security services on the Internet today, they tend to be implemented at a single layer. There are firewalls that block access to the network from the outside. Services such as SSL, TLS, and IPSEC provide encryption for data as it goes across the network. Although this simplifies integration, it doesn't support the kinds of policies that are needed in today's systems. For these policies to effectively reflect an organization's business models, the policies must apply to individual users or groups as they access individual application-level objects.

Do you think it will be impossible to go beyond certain levels of security in e-commerce without changes to the operating system?

As operating systems get bigger, it becomes harder to provide security. The trend in e-commerce is to reduce the dependence on the operating system for security. As we move toward smaller user possessed devices, such as Smart Cards, whose security is independent of the physical hardware to which they are connected, you can reduce dependence on OS security rather than depending on the security of the entire operating system. Even this is not fool proof.

I recall someone saying it is only a matter of how much money you have in order to break into a system. Is that still true? Have we advanced in that aspect?

It depends on the system, and it depends on your assumptions.

With enough resources, someone can break into any widely deployed system today, but some of the techniques used are not purely technical. Attacks sometimes involve social engineering: convincing an insider to provide a password.

If we apply a strictly limited definition of what constitutes a security breach, then we can develop secure systems. For example, given data that has been encrypted using a "one time pad" (XORing with a random-bit string), there is no way to decrypt this data without the key.

In practice, however, attackers breach such a system's security by breaking the key management system, by exploiting procedural deficiencies in generating, distributing, and storing the keys. The security of such a system depends on how well we protect and distribute the encryption key (the random-bit string) and whether the "random" bit string is really random, or only pseudo-random, possibly using a faulty random-number generator.

The public should have the right to apply any form of protection that it deems necessary to protect its data.

What are the controversial aspects of security and privacy?

Most of the controversy pertains to philosophical and political questions, namely who should have the ability or right to keep data private, and whether governments have the authority to impose techniques to keep strong cryptography out of the hands of the populace. There have been many proposals to

force users to register encryption keys with the government through "key-escrow" mechanisms.

These proposals are based on the premise that the ability to hide communications from the government can hinder the legitimate pursuit of criminals and terrorists. However, the technology is already out there, and it isn't reasonable to expect criminals and terrorists to follow such a law. As such, these constraints will only limit the legitimate protection of data by those following such laws.

I strongly object to such constraints on the use of cryptography. The public should have the right to apply any form of protection that it deems necessary to protect its data, and businesses have an obligation to apply the strongest practical measures possible to protect their customers' private data. Recall my earlier comment that cryptographic system security is usually compromised through weaknesses in key management. The proposed (and thankfully defeated, so far) key-escrow mechanisms force users to trust the government with these keys. There have been enough recent events demonstrating that the government isn't very good at preventing abuse of authority, or for that matter keeping secrets. I wouldn't want to base a system's security on the assumption that they can. What is to prevent a competitor or criminal from getting hold of your data by diverting the legal system to obtain the keys needed to decrypt the data? Such controls create significant vulnerabilities.

Who should govern security standards?

Security standards today are set in a de facto manner. Although we should try to unify many of these standards, it is usually not effective for them to be set by governments and then imposed on individuals. It should be the security community that ultimately decides what to accept and deploy. The government's role should be to create legal obligation regarding the privacy and protection of certain kinds of data—for example, confidentiality in medical records.

Should standards and common security and privacy practice be uniform worldwide, or should they follow state boundaries?

Ideally, we want to see uniform standards and common practices. On the Internet, state and national boundaries are blurry; this leads me to argue that they should fall within a common framework. This doesn't mean it's practical to expect laws to be the same in all places. I think there will be various trust mechanisms based on the concepts of assurance, trust, and endorsements. Parties will have credentials issued by licensing authorities, insurance companies, auditors, and product evaluators. Individuals will use these credentials to assess a business's security practices and verify its authority to conduct business with the customer.

What are the driving factors of security and privacy today?

E-commerce is the biggest driving factor in security and privacy today. As businesses and consumers become more dependent on the Internet for commerce, there is a much greater need for security. Privacy concerns regarding medical records have also generated security requirements that will create real improvements in the security of some systems. Unfortunately, much of the research in managing the policies associated with such requirements is still in its formative stages.

Much of the hype surrounding recent network attacks and consumer privacy has resulted in new products capitalizing on the fear factor, and not all of these products provide significant security improvements. You need to be careful when you evaluate security products to make sure that the protections provided match the policies of your organization and the environments within which the products are deployed.

Clifford Neuman is a senior research scientist at the Information Sciences Institute of the University of Southern California, a faculty member in the Computer Science Department at USC, and the chief scientist for Cyber-Safe. He received a BS from MIT and an MS and PhD from the University of Washington. He conducts research in distributed systems, computer security, and e-commerce and is the principal designer of the Kerberos authentication system, which among other uses provides user authentication for Microsoft's Windows 2000. He also designed the NetCheque and NetCash electronic payment systems and the Prospero Directory Service. He is a member of ACM, IEEE, IEEE Computer Society, Usenix, and the Internet Society. Contact him at Information Sciences Inst., Univ. of Southern California, 4676 Admiralty Way, Marina del Rey, CA 90292-6695; bcn@isi.edu; www.isi.edu/people/bcn.



Mary Ellen Zurko

What are the crucial innovations in the history of security and privacy?

There are two critical technical innovations in security. The first, public-key cryptography, set the stage for a lot of work in decentralizing cryptography and secure distributed protocols. It encouraged people to design protocols in systems that relied less on centralized trusted third parties. It distributed the trust in the security throughout the system, and made participants more responsible for their own security.

This is important for secure systems, because any centralized security creates an excellent target for attacks. Although centralization gives you one place to strengthen and concentrate your security energy, it also gives attackers a single place to concentrate on. The distributed protocols involving not only public key cryptography but also cryptographic techniques, such as perfect forward secrecy and multiple key sharing protocols, distribute the targets so attackers have to compromise multiple sites to get the information they want.

The other, more recently deployed technical security innovation that's been really important is safe languages. Safe-language research has been around for a long time, but when Java came on the scene, developers had access to a safe language that other developers were using too. Safe languages such as Java, which perform bounds checking and memory management, make it easier for developers to develop stronger systems that don't subvert the security that's designed into the application. Developers are under a lot of pressure, and safe languages help them do their job faster.

This is important, because the weak-link security theory tends to dominate the reference monitor theory of security. The reference monitor theory says you should make the security critical portion of your system small so that you can understand and develop it with high assurance. However, security is only as good as the weakest link in your system. That weakest link can be somewhere outside the reference monitor—if you're even lucky enough to have a reference monitor.

You mean security for dummies.

We're all dummies about this. You can't really make everybody a security expert in your development organization, just like you can't really make everybody a performance or database expert. You can have code reviews to find the mistakes and problems, but that takes time and resources.

What about the social and political aspects of privacy?

Various societies' expectations are widely divergent. US capitalists think about privacy this way: "If people could charge for their private information, then we'll all be better." In Europe, where they've had such incredibly bad experiences with peo-

ple abusing their private information, they expect and rely on the government to protect their privacy. They expect government protection as a right, whereas people in the US don't.

What are the most important remaining problems?

What technology can do to enhance privacy is still wide open. I see many more, larger problems in privacy than in security.

The World Wide Web Consortium's Platform for Privacy Preferences (W3C's P3P) is perhaps the first piece of technology being widely considered for privacy protection. It's about to finally make it as a W3C standard and be deployed in some products.

There's a lot of discussion about whether a technology developed by people interested in guarding privacy will actually do that or whether advertising agencies will use it as a tool to encourage you to type in private information that they can then

access and use. We need a lot more work on technical approaches to privacy before we develop any general theory about what does or doesn't work, in terms of how you treat your private information, how you share it, how you understand what other people are doing with it, and how to control and access it when it's gone somewhere else.

In security, there is a lot of research and experience from 10, 20, and 30 years ago that we aren't integrating into products. We need to find better ways to integrate existing security knowledge into deployed products, because designers and developers are on a tight schedule. That will mostly likely happen through tools. I believe on that approach, because that's how the usability community finally made progress in increasingly the overall level of usability in deployed products.

What about operating systems?

Operating systems are an incredibly frustrating area. We have a ton of experience, in both research and development, producing secure operating systems. I worked on an A1 rated operating system, a virtual machine monitor that we designed to be as highly secure as the government had criteria to rate it at the time. Multics was actually a good, secure operating system. Now, we have a lot of operating systems that integrate little of that knowledge—for example, Windows and NT.

We've had good solutions for 10 to 15 years that people just aren't using. There must be good reasons for that, but I don't know what they are other than time to market and not paying close attention.

Who do you think will enforce operating systems security in the future, if anyone?

It's got to be either market demand or government funding. It takes money to do anything, and in my experience, market demand is pretty variable.

The companies that are betting on their systems' security, such as Web-related e-commerce ventures, are able to take the risks, because somebody else will pay the price if things go wrong. For example, my credit card company will not charge me if my credit

card is misused on the Internet. If I buy a book or a CD on the Internet and my credit card number is stolen from the merchant I give it to, I bear no financial risk whatsoever (there is always the risk of the hassle I'll go through). The credit card companies are bearing the risks, because financially, it's a good trade-off for them.

What are the controversial aspects of security and privacy?

Names and other demographic information are hot topics in both security and privacy right now. In privacy technology, there are a variety of techniques to help you protect that information and anonymously send e-mail, browse the web, and use chat rooms. For instance, Zero Knowledge Systems just put out a suite of tools that lets you use five pseudonyms for that kind of Internet activity. That creates tension with people who are interested in security, because they traditionally use that information

for authentication, authorization, and auditing. We log what happens and hope that if something happens that we couldn't prevent, we can go back later, find the culprits, and attempt to get some recompense.

Recently, US President Bill Clinton said that there still isn't enough privacy protection for him to send his daughter e-mail—that would be dicey for him while he's president. On the other hand, members of the Clinton administration are against anonymity and pseudonyms,

to ensure better security against attacks such as denial of service. It's hard to figure out where to draw the line.

What are the current trends in this field?

You can get a lot of information on current trends in security by looking over the past proceedings of the New Security Paradigms Workshop (www.nspw.org). How to model trust, talk about trust, and handle trust with agents are issues that many e-commerce companies would like to get a handle on to encourage more Internet commerce.

The emphasis on survivability and information warfare indicates that we are using computers more for our infrastructure. There's a large concern about the threats and what we can and will be able to do to weather them.

As privacy gains importance, traffic analysis will too, because after you shield information about yourself, in terms of your identity and demographics, that will leave traffic analysis as the second-generation frontier.

Who should govern security standards?

The nice thing about standards is that there are so many to choose from; there are a variety of them from a variety of organizations. However, just having security standards doesn't mean that they are used in deployed systems—particularly on the Internet.

I participate in the Internet Engineering Task Force working groups. I think their emphasis on "rough consensus and running code" ensures a practical feedback cycle, but there's frustration in the IETF and all standards bodies about how long it takes to develop standards. Also, there are many important secu-

We need a lot more work on technical approaches to privacy before we develop any general theory about what does or doesn't work.

urity aspects that standards bodies don't traditionally address, such as usability. In fact, it's strictly taboo to talk about usability when you talk about protocols at the IETF. Working groups see that as a place where companies can add value, so you can't consider it in the protocol designs.

Who and what are the driving factors in security and privacy today?

The Web and e-commerce continue to be the big driving factors in terms of security deployment. The military also continues to be strong in terms of funding and driving certain forms of security research. They want to get a lot of their security in commercial off-the-shelf products, so they have a strong interest in seeing the research they fund appear in products, standards, or both.

Most published security research comes from academia. For instance, in the proceedings of the Usenix Security Symposium, a conference that emphasizes practical and deployed security over research, eight out of the 18 papers had industrial authors. Three of those were from industry labs, and two had academic coauthors. Academia is still generating the new ideas. There's a lot of government and military funding in academia, but you don't get real security just from research. You only get real security when something is deployed. The people making the products make the difference.

Who is making the most successful products today?

From a biased perspective, Lotus Notes has been shipping with a Public Key Infrastructure (PKI) and mail signing and sealing for 11 years. Integrating your security into an application gives you a big win in terms of usability and deployment but

makes it harder to get cross-application integration. Now that everyone is running on "Internet time," it's even harder to spend the time needed to produce secure products. Standards can help with that.

Should good standards and practices be common worldwide, or should they follow state boundaries?

I simply don't see how you can mandate common security practices across state and federal boundaries. Security is supposed to help keep bad things from happening. It's really hard to mandate security practices across legal boundaries. As security integration becomes more important, security standards will become more ubiquitous.

On the other hand, the privacy standards in Europe are very different from those in the US. I think that will drive either cultural or technical solutions toward homogeneity, because companies want to do business in the US, Europe, and other continents. They'll need to adopt solutions that either support the most strict privacy policies, take the risk of international litigation, or refuse to do business with people from particular areas. //

Mary Ellen Zurko is a security architect at Iris Associates, the home of Lotus Notes. She is currently responsible for active content security and authorization information. She was a member of the Jonah team that delivered the freeware reference implementation of the core Public Key Infrastructure X.509 (PKIX) standards. She holds a BS and MS in computer science from MIT. She is general chair of NSPW 2000 and will be program cochair of WWW10. She is associate editor of *Cipher*, the electronic newsletter of the IEEE Computer Society's Technical Committee on Security and Privacy. Contact her at Iris Associates, 5 Technology Park Dr., Westford, MA 01886; mzurko@iris.com.

Education

cont'd from p. 9

REFERENCES

1. L. Lewis et al., *Distance Education at Postsecondary Education Institutions: 1997-98*, Report NCEES 2000-013, Nat'l Center for Education Statistics, Washington, D.C., 1999.
2. D.A. Norman, *The Invisible Computer*, MIT Press, Cambridge, Mass., 1999.
3. M.J. Hannafin et al., "Grounded Practice and the Design of Constructivist Learning Environments," *Educational Technology Research and Development*, Vol. 45, No. 3, 1997, pp. 101-117.
4. J.R. Anderson, K. Koedinger, and R. Pelletier, "Cognitive Tutors: Lessons Learned," *J. Learning Sciences*, Vol. 4, No. 2, 1995, pp. 167-207.
5. S.R. Alpert, M.K. Singley, and P.G. Fairweather, "Deploying Intelligent Tutors on the Web: An Architecture and an Example," *Int'l J. Artificial Intelligence in Education*, Vol. 10, No. 2, 1999, pp. 183-197.
6. J. Roschelle et al., "Banking on Educational Software: A Wired Economy Unfolds," *Technos*, Vol. 6, No. 4, Winter 1998.
7. J.A. Kulik, "Meta-Analytic Studies of Findings on Computer-Based Instruction," *Technology Assessment in Education and Training*, E.L. Baker and H.F. O'Neil, eds., Lawrence Erlbaum Associates, Hillsdale, N.J., 1994, pp. 9-34.
8. A.L. Brown and A.S. Palinscar, "Guided, Cooperative Learning and Individual Knowledge Acquisition," *Knowing, Learning and Instruction: Essays in Honor of Robert Glaser*, L.B. Resnick, ed.,

Lawrence Erlbaum Associates, Hillsdale, N.J., 1989, pp. 393-451.

9. K. Snyder, *Asynchronous Learning Networks and Cognitive Apprenticeship Theory: A Model for Teaching Complex Problem-Solving Skills*, unpublished dissertation, New York Univ., New York, 2000. (Available from Univ. Microfilms Int'l, Ann Arbor, MI.)
10. C. Conati, J. Larkin, and K. VanLehn, "A Computer Framework to Support Self-Explanation," *Proc. AI-ED '97: Eighth World Conf. Artificial Intelligence in Education*, IOS Press, Amsterdam, 1997, pp. 279-286.

Peter G. Fairweather is the senior manager of the Applied Learning Sciences Department and the Accessibility Research Institute at the IBM T.J. Watson Research Center. He works on the application of technology to problems of individual, group, and organizational learning. After receiving his PhD at Northwestern University, he worked both in industry and universities developing simulation tools for training applications and researching human and machine models of text understanding. Contact him at the IBM T.J. Watson Research Center, Rte. 134, Yorktown Heights, NY 10598; peterf@watson.ibm.com.

Andrew S. Gibbons is a professor of instructional technology at Utah State University in Logan, Utah. He designs technology-based instructional systems, tools, and methods for instructional designers. He combines 18 years of industry design and consulting experience with academic interests in theories of product architectures, design process, and model-centered instruction. His work also explores high-volume production of innovative instructional forms and the application of advanced pedagogic, design, and production concepts by new designers. Contact him at the Dept. of Instructional Science, Utah State Univ., Logan, UT 84322; gibbons@cc.usu.edu.