Information Assurance Technology Forecast 2008

STEVEN M.
BELLOVIN
*Columbia
University*

TERRY V.
BENZEL
*University
of Southern
California*

BOB BLAKELY
Burton Group

DOROTHY E.
DENNING
*Naval
Postgraduate
School*

WHITFIELD
DIFFIE
*Sun
Microsystems*

JEREMY
EPSTEIN
Software AG

PAULO
VERISSIMO
*University of
Lisboa*

I n 2005, a forward–thinking panel dared to predict the future of information assurance to help us better plan for it. Almost three years later, we assembled another fearless group of distinguished information assurance experts to give us insight into where they see our exciting field headed. Our panel gives insights into how the threat's evolving nature, the current information technology environment, and various market forces are combining to yield new security challenges and likely new technology paths for the future. I asked the panel some of the most provocative and difficult questions I could conjure, and they met the challenge admirably.—O. Sami Saydjari, Cyber Defense Agency

What do you predict will be the most significant change in information technology in the next 15 years?

Whitfield Diffie: The rise of Web services. At present, everyone I know is vulnerable to Google: you can't avoid using it, and if it were to hand your query stream to your enemies, you'd be screwed. Over the next 15 years, we'll see the rise of many trade–secret–based companies that do everything from image rendering to statistical calculations to heat flow analysis to things I know nothing about. No large corporate program will be able to run effectively without using of these services, and to do so, it will have to tell its secrets to service providers. Webs of trust will become webs of contracts, and control over security will be out of the hands of any individual organization.

The runner up and leader down the back stretch will be the introduction of WiFi.

Dorothy Denning: I posted a list of famous predictions that never happened on the bulletin board outside my office as a reminder of how ridiculous our predictions can look down the road. Be that as it may, I'm fascinated by what's happening with Google, map–ping, and virtual environments. Perhaps we'll have a live 3D Web that effectively merges physical and cyberspace with the help of zillions of devices transmitting live video, audio, and other data into and out of the Internet. The effect will be to amplify security and privacy issues.

Terry Benzel: The next 15 years will see the commoditization of ubiquitous computing. We've been talking about this for many years, and we see evidence in leading US Department of Defense and early adapter markets. Over the next 15 years, computing will move from a stand–alone conscious activity to a fully integrated/always on aspect of daily life. Architecturally, and in terms of technology, we'll complete the move from workstations and laptops connected to networks and servers to embedded tetherless computational nodes and wide–spread sensor-based systems. As a result of Internet communication moving from point-to-point communication between hosts to multiparty information dissemination models, content–centric networking and data stores will replace fixed information repositories.

Steve Bellovin: The most obvious answer is that we'll see computers everywhere, with complex interconnections and no perimeters. This will create interesting security challenges, such as how to teach the television I bought second-hand that I'm now its authorized owner, that I'm willing to delegate certain—but not all— watching and recording rights to my family, and that the previous owner should be excluded except for the right to retrieve certain stored content associated with it.

Jeremy Epstein: As computerized devices become ubiquitous (not just as laptops, BlackBerries, MP3 players, and so on, but everywhere), it will become impossible to deal with the need to manage each device individually, as typically gets done today. While large corporations have dedicated (but overloaded) IT staff to manage complex networks, home users have no such infrastructure. As a result, manageability will surpass security as a key driver, for both home and business users.

Meet our panel

Steven M. Bellovin is a professor of computer science at Columbia University. His technical interests include network security, privacy, and the social implications of computers. Bellovin has a PhD in computer science from the University of North Carolina at Chapel Hill. Contact him at smb@cs.columbia.edu.

Terry V. Benzel is the deputy director of the Computer Networks Division of the Information Sciences Institute at the University of Southern California. Her technical interests include security experimentation and test, malware analysis, and the social aspects of information sharing. Benzel has an MS in mathematics from Boston University and an executive MBA from the University of California, Los Angeles. Contact her at tbenzel@isi.edu.

Bob Blakley is principal analyst at the Burton Group. His interests include security, risk management, privacy, and identity, all of which he feels are primarily social, legal, and economic problems rather than technical problems. Blakley has a PhD in computer and communications science from the University of Michigan. Contact him at blakley@burtongroup.com.

Dorothy E. Denning is professor of defense analysis at the Naval Postgraduate School. Her research interests include conflict in cyberspace, information operations, and information assurance. Denning has a PhD in computer science from Purdue University. Contact her at dedennin@nps.edu.

Whitfield Diffie is chief security officer at Sun Microsystems. His technical interests include cryptography, network security, and signals intelligence. Diffie has a PhD in technical sciences (honorus causa) from the Swiss Federal Institute of Technology. Contact him at whitfield.diffie@sun.com.

Jeremy Epstein is senior director of product security at Software AG. His technical interests include security of SOA and voting systems. Epstein has an MS in computer sciences from Purdue University. Contact him at jeremy.epstein@cox.net.

Paulo Verissimo a professor in the Department of Informatics at the University of Lisboa and director of LASIGE, a research laboratory within the department (http://lasige.di.fc.ul.pt). His research interests include architecture, middleware, and protocols for distributed, pervasive, and embedded systems. Verissimo is a fellow of the IEEE. Contact him at pjv@di.fc.ul.pt.

Paulo Verissimo: First, computer systems will become component-oriented, a concept that entered the car-making industry and started a revolution. Virtualization, multicore, and multichip are key factors of success, but the final enablers will be new business models from software vendors who should recognize this trend and won't dare to charge as many licenses as modules or else risk being on the losing end to open source and free software; this is already happening in the virtualization market. Consequently, the current trend to use a single physical-logical PC for everything will reverse—for example, I'll eventually use my "office work" computer, my "Web browsing" computer, my "e-banking" computer, my "personal theatre" computer, and so forth. The fact that they may all exist in the same box is irrelevant because what matters is that they'll be self-contained and noninterfering, yet will also have different levels of security and dependability. The benefit on assurance will be overwhelming.

Whitfield Diffie: Over the next 15 years, virtual realities will move from a playland of gamers and early adopters to the preferred venue for many (maybe most) business transactions. The security of virtual reality is a barely touched subject. Online chat has destroyed the virtually impenetrable security of a 10-year-old's playground (the cost of having an adult impersonate a 10-year-old is pretty high), and virtual reality has done the same to the relative security of the boardroom, the cafe, and the strolling conversation.

Bob Blakley: The devices we use to connect to information systems will become smaller and smarter, but also more specialized and more diverse. Most of the functionality we're used to on our desktops today will move "into the cloud" as connectivity gets more ubiquitous, more robust, and cheaper. More information about everything, including us individually, will move into the cloud along with this functionality, with the result that a virtual world of information will be increasingly commonly overlaid (via our tiny, portable, always-connected access devices) on the real world we live in. This has profound implications for our view of security; availability will be both vital and assumed. Privacy will become an even thornier issue than it is today. And information integrity will be essential to our safety and quality of life.

Looking back 15 years, what were the big surprises in information technology that significantly affected the information assurance problem?

Jeremy Epstein: The big surprise wasn't that we've made minimal progress in improving the security of the software we rely on for everything from e-commerce to entertainment. Rather, the surprises were

- Consumers continued to purchase systems knowing that they're flawed or at risk. It's the "boiling frog" syndrome—if consumers and businesses knew 15 years ago that they'd have to do regular patching and still would suffer periodic bouts of vulnerabilities, they might not have adopted IT so willingly, but because it happened slowly, they've gotten used to being constantly at risk.
- We've witnessed the wholesale move to online financial management without an infrastructure allowing it to be done securely on either the client or server side.
- The dramatic increase in software size hasn't slowed, despite the information assurance problem. Fifteen years ago, we struggled with security for Windows 3.1, which was at most a few million lines of code, whereas today, businesses and most consumers use either Windows XP or Vista, which have tens of millions of lines of code.

Fifteen years ago we didn't think we could secure a million lines of code; today, we're sure we can't secure tens of millions.

Paulo Verissimo: That the "popular" computing practice contaminated the "serious" computing one, and not vice versa. Rebooting and weekly patching became the workhorse of problem mending and the symbol of reliable computing, not only in personal citizen computers but also in what used to be high-end systems and infrastructures working to tight specifications, such as data centers, telecom, banking, and even cars and planes. If the situation is unpleasant and inconvenient for personal users, it's dramatic for serious computing applications in operator, provider, and critical information infrastructures. Professionals bitterly discover, after a couple of decades, that they're tied to hopelessly flawed technologies and are equally hopeless, in the short term, to make them better and more trustworthy. There's not much of an assurance case to be made about information residing in non-trustworthy systems.

Steve Bellovin: The biggest surprise was the change in the threat model. We always knew there could be profit-based attacks, but no one anticipated the current alliance between hackers and other sorts of criminals, such as spammers, credit-card thieves, stock fraud artists, and so on. This—combined with the woeful security level of most desktop machines—has created a vast profit motive for computer crime.

Terry Benzel: The biggest surprises were less technology based and more sociocultural. The impact of spam, phishing, and the cybercrime black market caught technologists by surprise but shouldn't have. The surprise seems to come from the lack of connection between sociocultural communities and technologists. The rise of botnets is an excellent example of the bridge between sociocultural forces and technology. Myriad surprises emerged with the birth of the Web; the Internet's original designers didn't envision its current pervasive uses.

Dorothy Denning: One surprise was botnets. Fifteen years ago, people programmed distributed computations across networked computers—for example, to factor RSA private keys—but the owners of the computers volunteered the use of their machines. I don't recall any discussion in the early '90s about how an adversary could compromise and take over a massive number of machines, organize them into a network with a command and control infrastructure, and then deploy the botnet for spam, fraud, distributed denial-of-service (DDoS) attacks, and even money laundering. One of the first DDoS attacks of this nature was the February 2000 attack on Yahoo!, CNN, Amazon, and other commercial sites. Botnets are one of the most serious problems on the Internet today. The Estonian cyberassault wouldn't have amounted to much without them.

What do you think was the most significant information assurance advance over the past 15 years?

Steve Bellovin: The biggest advance was the development and deployment of SSL to protect Web transactions. It certainly has its flaws and limitations, but it completely blocked passive eavesdropping on credit-card numbers.

Paulo Verissimo: I think the most significant information assurance advances have been due to policy and practice factors. Important security technologies have emerged over the past decades, but many of them weren't or still aren't used. Banks know that if they all used certificates and mutual authentication, phishing would practically disappear, so why won't they?

I can single out a few things that have led to significant information assurance advances because they changed people's minds and practice: liberalization of cryptog-

raphy in most developed countries, the push for strong authentication, and trustworthy computing movements. We still need a lot of laws and regulations to put liability where it belongs.

Whitfield Diffie: This is hard to answer because there's a trade-off between solidity and relevance. AES, for example, is a very solid advance, but TCG technology, for all its political and technical problems, is surely the right direction for critical information infrastructure, whatever the potential for abuse in consumer electronics. In many ways, SSL is the most remarkable development: it's the most widely deployed crypto-security mechanism in the world, eclipsing all military systems and passing the previously most common (albeit low-grade) crypto device, the Zenith TV scrambler.

Naturally, I also have high regard for the security characteristics of Open Solaris. Computing on a network (even if the network is inside a box) because it lets you apply network security tools such as firewalls and is a major improvement over the von Neumann machine view. Socially or politically, the NSA's change of attitude removes a major obstacle to widespread secure systems.

Jeremy Epstein: I disagree with Whitfield—although AES is great science, crypto is no longer, with rare exceptions, the weak link in the security chain.

Terry Benzel: On the one hand, it's easy to say that we haven't seen advances in information assurance, but that isn't quite true. The reality is that there has been significant investment in research and development, and there's widespread use of technologies that weren't previously used. We've seen an incremental improvement in information assurance in the areas of boundary protection, antivirus

protection, and the use of cryptography in e-commerce applications. The issue at hand is that these advances are nowhere equal to the exponential growth we've witnessed in the threat environment.

Dorothy Denning: I think we reached a tipping point in terms of vendor interest in developing more robust products, perhaps motivated by the realization that it's better and cheaper to get rid of security holes during product development than to deal with the public relations mess and patches after the fact. Nowhere was this more evident than in Bill Gates's memo in January 2002, announcing Microsoft's Trustworthy Computing initiative. He declared "Trustworthy Computing is the highest priority for all the work we are doing," adding "so now, when we face a choice between adding features and resolving security issues, we need to choose security." Microsoft put its money where its mouth was, educating its developers and changing its software development practices.

Jeremy Epstein: Steve and Whitfield are right about SSL being a huge advance. I'd further Dorothy's comments on vendor interest by noting that there's been a lot more energy put into (relatively) reliable automated patching systems, which allow us to survive despite the poor assurance of commercial software.

Another major advance has been automated techniques to catch and stop many of the most common security problems (such as buffer overflows). Although this

is simply a band-aid for a problem we've known about for 30 years, the widespread use of such tools and techniques (both dynamic and static protection) have had more practical impact on assurance than any theoretical advances.

The increasing use of Java and other type-safe programming languages, which have largely supplanted C and C++, have greatly reduced the types of security flaws we saw through the '90s. Unfortunately, new types of flaws have taken their places.

Bob Blakley: There were none. In fact, you couldn't have picked a more perfect date for "the end of history" in security than 15 years ago, 1992. My timeline of interesting information assurance events goes something like this:

- 1965: Gordon Moore proposes Moore's law, which tells us by inference that the security problem will get twice as hard every year.
- 1974: The TCP protocol is invented, ensuring that universal connectivity will one day be a reality.
- 1976: Whitfield Diffie publishes the first unclassified description of public-key cryptography.
- 1979: The Anderson Report codifies the standard model of information security.
- 1981: The IBM PC is released, and it becomes clear that the majority of computers will be administered by incompetents.
- 1983: The Orange Book tells us how to secure expensive, disconnected, competently administered computers.
- 1984: Fred Cohen describes the

Umpteen wart hogs grew up, but two mats tickled Paul. One wart hog grew up, however five dwarves auctioned tickets to see alcoholic celebutantes excuse their 4-line pull quote

computer virus (which had been discussed privately among security researchers for more than decade) for the first time in the public literature.

ware compliance against CERT's C and C++ Secure Coding Standard; software developers will soon be able to run their code against Fortify's Source Code Analysis tool. Also

Umpteen wart hogs grew up, but two mats tickled Paul. One wart hog grew up, however five dwarves auctioned tickets to see alcoholic celebutantes excuse their 4-line pull quote

- 1991: The World Wide Web goes live, and the network effect becomes a reality.
- 1992: Nothing interesting in security starts to happen, initiating a trend that continues until the present day.

Despite the fact that both attacks and losses have approximately doubled every year since 1992, we continue to rely on old models that are demonstrably ill-suited to the current reality and don't inhibit the ongoing march of failure.

What breakthroughs do you see as likely in information assurance technology over the next 15 years?

Steve Bellovin: I hope we'll see progress in two areas. First, we need to work on the human interface. Today, people don't understand the consequences of various security-sensitive actions or simply don't know how to do them. Second, I hope we can design breach-containment architectures in which the inevitable failures won't lead to wider system penetration.

Dorothy Denning: Given the growing interest in trustworthy software and recent efforts to develop secure coding standards and practices, it's possible we'll see a breakthrough in secure coding. CERT announced in early October 2007 that it was teaming with Fortify Software to automate the process of testing soft-

in 2007 and in collaboration with CERT, SANS began issuing GIAC Secure Software Programmer's certifications to developers who pass their C or Java exam. In addition, SANS offers training in secure coding. However, we still need colleges and universities to emphasize secure coding throughout their programming and software engineering curricula (not just in security classes), vendors to require their developers to follow secure coding standards and their suppliers to provide code that complies, and customers to demand software products that meet secure coding standards. None of these efforts will eliminate software vulnerabilities, but they're likely to make the problem more manageable and allow us to cope with the increasingly complex software environment.

Jeremy Epstein: Availability of low-cost hardware will make feasible information assurance technologies that we've known about but couldn't practically use. For example, the use of virtual machines is moderately feasible today; in the next 15 years, they'll become one of the key ways we get information assurance. Similarly, as it becomes feasible to have networks of dedicated processors in a single computer or device, we'll see partitioning of applications with well-defined boundaries, which will help us gain assurance.

Paulo Verissimo: I think two ma-

jor breakthroughs will be closing the coverage gap in today's systems security and achieving automatic security. Both help promote system security and information assurance to higher grounds, but they require a paradigm shift, from intrusion prevention to intrusion tolerance—after intrusion happens but before failure. Intrusion tolerance can potentially close the coverage gap—that is, the mismatch between the assumed probability of the system having a security failure and the real probability of that happening—by orders of magnitude. Furthermore, this could be achieved with mechanisms that ensure the overall system remains secure and operational, in an unattended, essentially automatic way.

Bob Blakley: Likely? Well, I'm not really a betting man. How about 'possible and beneficial'? Here are two: special-purpose security devices that are simple enough to be highly assured, and universal surveillance. Love it or hate it, the surveillance society is here. On a broader point, however, breakthroughs in information assurance technology aren't going to be as important as breakthroughs in information assurance policy. We've already seen the effects of California's SB 1386, which requires companies to notify people whose personal financial information is compromised, and we're seeing strong market effects as a result of the Payment Card Industry Data Security Standard (PCI-DSS), which requires payment processors to use security technologies that are already available if they want to continue to do business with the credit-card issuers. When our policy creates real incentives to solve security problems, the right technologies will be found and they will be used. Until our policy creates such incentives, no amount of technology is going to fix the problem.

Whitfield Diffie: Insofar as keying, infrastructure is a capital and market development problem that I expect will fix itself. The other problem with PKI has to do with the fragility of software platforms; TPM and other improvements should help reduce the plethora of compromised (or at least unaccounted for) keys that have plagued PGP and its relatives.

Terry Benzel: More incremental security technology may be marketed as breakthroughs, but, in reality, incremental changes to the landscape will only continue the arms race that we're rapidly losing. An entire shift in approach, architecture, and technologies is required. We should expect to see breakthroughs in content protection and data-centric networking. We should also look for breakthroughs in policy and practices, liability, and data/information provider accountability. But we must concede that protecting the end hosts is a lost battle.

What's the nature and magnitude of risk that critical information infrastructure (CII) faces over the next 15 years? By "critical," I mean the part whose failure would have major effects on the nation, such as economic loss or loss of life.

Whitfield Diffie: This is hard to assess. My understanding is the security in the electric grid and other such information infrastructures is a mess. I haven't done the study necessary to determine whether the mess constitutes an actual vulnerability to an attack that would damage the physical infrastructure, but my intuition is that it might.

The Protect America Act [*editor's note: see page xx for an article on the Protect America Act*] creates a vulnerability that is novel in the communication system as a whole and more serious than the analogous vulnerability in comsec monitoring. None of our "real" enemies today have the resources to get broad access to US communications. As we build in machinery for spying on our own communication system, we create the risk that the machinery might be captured by an opponent and that the cost to national security would exceed the gains.

Paulo Verissimo: Large and ever increasing. Moreover, the objective risk is amplified by the lack of perception of the risk itself existing, by citizens, policy makers, and CII manufacturers and operators. There's still a belief that the SCADA [Supervisory, Control and Data Acquisition] systems controlling these infrastructures are legacy, closed, obscure, and thus unattackable, or that it suffices to just use a firewall and an intrusion detector, but normal ICT systems protection won't be enough. To keep a long story short: Ctl-Alt-Del isn't a remedy for things that have worked continuously for more than 20 years, many security techniques hamper real-time operation, and there's still a difference between erasing a database and setting a generator on fire. This should be understood immediately or else we should get prepared for the next generation of mass hacking. Maybe all it takes for people to get serious about this is a www.scada_rootshell.com (Google the remainders of the classical www.rootshell.com to grasp the basic idea). It might be a good idea for policy makers and CII manufacturers and operators to learn the difference between crash and bang.

Bob Blakley: The biggest risk is created by technology vendors, who will continue to sell unreliable general-purpose systems for use in security- and safety-critical environments. Cascade failures will also become more common as interconnectivity increases, but targeted attacks by terrorists will be less common than economically motivated attacks.

Generally speaking, the threat will progress from "credible" to "serious" over the next 15 years, but my guess (not prediction, per Yogi Berra's wise advice that prediction is difficult, especially when it's about the future) is that it won't become critical during this period.

Steve Bellovin: I see two major risks. First, many SCADA systems are poorly protected—think of the Australian sewage spill, multiplied by a thousand. Second, I worry about attacks on the financial system. We've already seen how one rogue trader can destroy a major bank (Barings). What could a clever worm do?

Terry Benzel: Risk comes from the combination of cyber and physical attack. A concerted well orchestrated attack can disable multiple sectors of the nation's CII. We continue to be exposed to risk from incompetence as well. Too many networks, enterprises, and even home users are ill equipped and trained to manage the resources at their command.

Jeremy Epstein: The risk is a political/economic one. If the drive to low-cost production and deregulation allows connection of CII (such as electric power systems) to the Internet, then the technical problems, especially as exploited by hostile nation-states, will lead to significant economic loss. But if political/economic incentives reverse, and regulation prevents Internet connectivity (even though it increases costs), then the technical risk can be minimized.

A key question is whether attacks on the CII are subtle and slowly building (in which case, we may see the boiling frog again,

as people learn to accept periodic low-level attacks) or whether they're part of a massive attack (along the lines of 9/11 or Pearl Harbor, in which case there could be meaningful changes to how we protect the infrastructure).

Dorothy Denning: Too many variables are at play here, plus little-to-no data about any of them, to come up with any reasonably accurate estimate of the risk. The catastrophic failures people write and worry about are typically based on fictive scenarios that go beyond the cyberthreats that actually confront us. So far, our worst enemy, Al Qaeda, hasn't demonstrated a capability to conduct very sophisticated cyberattacks. But over the next 15 years, things will change.

How do you see adversary capabilities changing over the next 15 years, based on what we've seen evolve over the past 15?

Whitfield Diffie: If I'm to believe the scholars of bot nets, we might be facing an imminent crisis. Computing and communications are sinking into human culture. High tech is no longer protected by a general obscurity, but it isn't adapting fast enough to the security needs of a world in which information gets around really well. We can expect groups all over the world to explore network intelligence and information warfare. Many people worldwide are smart and well educated; as they put their minds to penetration, the playing field will change beyond recognition.

We're 19 years from the Morris worm, which—although disruptive—was a failure because of its author's lack of access to captive networks in which to debug it. In the time since, we've seen worms take over large parts of the Internet—first in hours,

then minutes. People have talked about viruses and worms since at least 1971 (when I first heard the notion), perhaps even since the 1950s, but the scale of the vulnerability took us almost completely by surprise.

In the past fifteen years, we've seen a steady attack on security—in the name of security—by people who see secure communications by our enemies as a greater hazard than secure communication by our friends as a benefit. Should cryptography and its supporting technologies become a serious problem for national intelligence, they'll once again come under attack.

Jeremy Epstein: The big change will be government organizations realizing that vendors and the public won't respond to scare tactics, especially with regard to foreign threats. The Bush administration's misrepresentation of intelligence in Iraq greatly reduced Americans' willingness to accept statements from government officials at face value, and this will carry over into the information assurance arena. As a result, we'll see more openness by government agencies in explaining specifics of international threats and the motivations and capabilities of foreign nation-states.

On the technical side, adversaries will develop more automated means of finding application vulnerabilities, which will overtake infrastructure as the primary focus of attacks.

Dorothy Denning: There has been steady progress in the development of software tools for automated attacks such as phishing and the herding, control, and application of botnets. In addition, underground markets that trade in attack tools and stolen data such as credit-card numbers are thriving. These trends will continue, mitigating some of the benefits

that would otherwise come from improved security. But even more worrisome, we may see development of and traffic in exploit tools aimed at SCADA systems and other control systems that impact physical devices, critical infrastructures, and life-critical applications. Already, we're starting to see disclosures of SCADA vulnerabilities. Fifteen years from now, we could see actual tools for attacking automobiles, robots, implants, remote surgeries, or energy systems instead of just Web sites, servers, and PCs.

Bob Blakley: Adversaries will have better education, funding, and information about the systems they attack. This isn't entirely a bad thing, however: organized crime could be more expensive but less chaotic than disorganized crime.

Paulo Verissimo: Let me start by qualifying the adversaries. To me, they're computer experts who live on the dark side—the Darth Vaders of the computer science and engineering profession. Since the advent of the Internet and the Web, the power of these adversaries—that is, their capacity to stage attacks on real ground—has been leveraged enormously by influencing hosts of script kids and less gifted hackers, by putting recipes on the Web, and by using computer hosts (zombie botnets) that let them deploy massive attack tactics. The DDoS set of attacks in February 2000 was just the beginning of a new era, from hacking *for* the masses to hacking *by* the masses.

Adversaries' business models could also extend to infrastructures that will become completely pervasive in a few years, such as passport and travel control systems, identity or medical card systems, and electronic voting systems. It's almost unthinkable to admit that these infrastructures could be vulnerable, but the problems already

faced around the world in these kinds of systems, from e-voting to electronic passports, make me believe some clouds are ahead.

Terry Benzel: Cybercrime is a growing segment of our economy. Well-funded, focused adversaries and nation-states will benefit from new technologies, tools, and organization that will help them stay one step ahead of the defensive security industries. Adversaries will also be able to exercise the increasing interconnectedness and interdependence in our systems. The interesting question is where cyberattack will focus as we move to the more ubiquitous computing model.

Steve Bellovin: I think they're going to get stealthier. Malware will be harder to find, harder to detect when in operation, and much harder to remove. Among other things, I think it likely that adversaries will exploit technologies like DRM and nominally locked-down appliances.

What question should I have asked regarding this information assurance technology forecast that I didn't ask, and what would your answer be?

Whitfield Diffie: "What advances in information security will not be made in the next 15 years?" The answer would be that we probably won't get much further in proving things, either that our code is correct or that crypto algorithms are secure.

You didn't mention forensics and counterforensics, a major trend in security to come within the social fabric. Military comsec failures don't allow appeal: if you didn't protect your transmissions, you won't get far complaining about being exploited. Commercial security failures are quite different, and if you collected the

right evidence, you could recover by suing.

You also haven't mentioned the laptop security problem—securing equipment versus securing events. If you encrypt data with your laptop and send it somewhere, an opponent who intercepts the message must confront the crypto. On the other hand, someone who has the laptop could well find the key lying around somewhere. Trying every consecutive 16 Bytes on a gigabyte disk is a workfactor of around 2^{26} to 2^{40}, a big improvement on 2^{128}, and being sure that the key isn't anywhere on your disk isn't easy.

Dorothy Denning: If you had $100 million to invest in security technology research, where would you put it? I'm not sure, so I really want to hear what others have to say, but I'd consider putting it in technologies for network service providers to facilitate network-wide defense—for example, to detect and shut down botnets and the Web sites used for phishing and malware distribution, as well as other pervasive threats.

Jeremy Epstein: I'd take Dorothy's $100 million and invest it in figuring out how to make systems that are more intuitive to manage securely. We have novices managing our systems, and that's not going to change. We have to make it easier for them.

Bob Blakley: Why, in the face of steadily escalating threats and losses, do we continue to cling to old security models that have demonstrably failed to solve the problem? Why, for example, do we think that intrusion detection, biometrics, secure software development, and quantum cryptography will solve problems that are manifestly not susceptible to their use? And why don't we recognize that the priority of security requirements has changed from "confidentiality,

integrity, availability" to "availability, accountability, integrity, confidentiality"? Or that the priority of mechanisms has changed from prevention to detection and recovery? Or that authorization doesn't scale, but accountability does? Or that general-purpose systems can't be secured, but special purpose systems might be?

Terry Benzel: We need to broaden the discussion to include sociocultural and economic issues. We can no longer be technologists in isolation. All too often, we suffer from technologists with a hammer in search of a nail. Taking the discussion in this direction might force us to get realistic about what we can reasonable protect and what we should concede as wild territory. We need to perform cost-benefit trades to ask and answer some hard questions.

Bob Blakley: Terry's dead right about that, and it's extremely important.

Steve Bellovin: You should have asked me how the threat model will change. I suspect we're going to see many more targeted attacks, both for profit and for national interest. If I may indulge in a geeky pun, letters of marque will be replaced by letters of mark and space. □