Information Assurance Technology Forecast 2005

"It is said that the present is pregnant with the future." —*Voltaire*
"The best qualification for a prophet is to have a good memory." —*Marquis of Halifax*

Virgil D. Gligor
University of Maryland

Tom Haigh
Cyber Defense Agency

Dick Kemmerer
University of California, Santa Barbara

Carl E. Landwehr
University of Maryland

Steve Lipner
Microsoft

John McLean
Naval Research Lab

Ice hockey legend Wayne Gretsky once pointed out that the key to the game isn't in being where the puck is, but where the puck is going. We assembled some of the most distinguished information assurance experts and asked them to take the risk of predicting where the puck is going in this exciting field. 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.
—*Sami Saydjari, Cyber Defense Agency*

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

Dick Kemmerer: Ubiquitous or pervasive computing. The integration of computation into the environment, rather than having computers as distinct objects, opens up a whole new class of information assurance problems. The promoters of ubiquitous computing expect that embedding computation into the environment will enable people to move around and interact with computers more naturally than they currently do. One of their goals is to enable devices to sense changes in the environment and to automatically adapt and respond based on these changes and based on the user's needs and preferences. The increased information assurance needed for these pervasive systems will be orders of magnitude more than what we need today. This is a concern.

John McLean: Dick's right about ubiquitous, pervasive computing—although I would place the emphasis on the fact that these computational devices will be constantly exchanging information with other computational devices and databases all over the world. It is only a matter of time before our digital transactions will require stronger privacy-preserving authentication, encryption, and assurance technology. Further, as each device takes on more roles, it is conceivable that all my information—both important and unimportant—will reside on a single multipurpose device that accesses a wide variety of different databases. This will increase the need for improved separation technology on the devices.

Exacerbating this issue is the fact that software development and service functions across the board are being shipped overseas. We are already faced with issues of building trusted systems from untrusted components. We may soon be faced with an environment where we must perform trusted transactions with untrusted participants and build trusted systems out of components that are not only untrusted, but are, in fact, almost guaranteed to be adversarial. As the range of systems we must trust expands, the properties that any individual system must satisfy will also expand. Some of the software we will be interacting with will control autonomous systems, which will have built-in learning algorithms and nondeterminancy—properties that do not lend themselves easily to current assurance methods. Ironically, they may also form an integral part of future computer defenses, since autonomous response systems may very well be required in future computer network defense. Without developing the technology to assure such autonomous systems, we may find ourselves in the uncomfortable position of having our network defenses depend on autonomous systems that we cannot assure.

Virgil Gligor: The prevalence of portable, wearable computing and communication devices that will be impractical to protect physically and yet will be subject to loss, theft, and manipulation by a determined adversary. Security protocols will have to account for the presence of an adversary as a fully privileged protocol participant, and information assurance will have to account for poten-

Our panelists

Virgil D. Gligor is a professor of electrical and computer engineering at the University of Maryland, College Park. He has worked in security research and education for 30 years in a broad range of areas, including access control mechanisms, penetration analysis, denial-of-service protection, cryptographic protocols, and applied cryptography. Contact him at gligor@beckmann.eng.umd.edu.

Tom Haigh is a senior scientist at the Cyber Defense Agency and a member of the technical staff at Adventium Labs. Prior to that, he was vice president for research and then chief technology officer (CTO) at the Secure Computing Corporation. His current interests are the development and application of methods for measuring the relative risk of different information assurance architectures and the application of artificial intelligence techniques to problems in information assurance. Contact him at thaigh@cyberdefenseagency.com.

Dick Kemmerer is a professor and past chair of the Department of Computer Science at the University of California, Santa Barbara. He has worked in information assurance for more than 30 years and has written numerous papers on the subjects of computer security, formal specification and verification, software testing, programming languages, and software complexity measures. He directs the Reliable Software Group at UCSB, which is addressing the need for better languages and tools for designing, building, validating, and securing software systems. Contact him at kemm@cs.ucsb.edu.

Carl E. Landwehr is a senior research scientist at the University of Maryland's Institute for Systems Research. He recently completed a term of service at the US National Science Foundation as founding program director for the cybertrust emphasis area. For more than 20 years, he led research projects in computer security and information assurance at the Naval Research Laboratory; he is interested in all aspects of information assurance. Contact him at landwehr@isr.umd.edu.

Steve Lipner is senior director of security engineering strategy in Microsoft's Security Technology Unit. He has more than 35 years' experience in computer and network security in his roles as researcher, consultant, development manager, and business unit manager. He was one of the developers of the security development life cycle (SDL), which Microsoft applies to improve the security of its software, and he currently leads the team that manages the evolution and application of the SDL. Contact him at slipner@microsoft.com.

John McLean is superintendent of the Information Technology Division of the Naval Research Laboratory. He has been in the information assurance field for more than 20 years, publishing papers, establishing and heading the Formal Methods Section of NRL's Center for High Assurance Computer Systems, and eventually serving as director of the center while it developed several key technologies and devices. His research interests include formal methods and models for computer security. Contact him at John.McLean@nrl.navy.mil.

tial loss and corruption of information in a probabilistic manner—that is, we will have to accept less-than-perfect assurance, but good enough for a particular application.

Carl E. Landwehr: The advent of quantum computing, if it happens, could have a profound effect on the IA [information assurance] landscape. If it doesn't, I would agree with Dick and Virgil that the advent of pervasive computing will make the biggest change in the nature of the IA problem. I expect to see hardware-based controls play an increasing role in security policy enforcement. The automotive industry will be a particularly interesting place to watch. Onboard computing for engine control, braking, and vehicle stabilization is already in place, and Internet access and automated intervehicle communications for safety alerts and other purposes are coming soon. This industry's apparent preference for accepting liability in place of regulation and its well-known focus on minimizing cost, combined with consumer pressure for innovative, computing-based product features, could lead to innovative IA solutions.

Tom Haigh: The widespread use of computers to monitor and control safety-critical processes. Some of these processes are critical to the continuous and safe operation of society's critical infrastructure; others are critical to the continuous and safe operation of our personal infrastructure, such as our homes and cars and our healthcare systems. The opportunities for improved health, safety, and quality of life are phenomenal—when everything behaves the way it is supposed to. On the other hand, the potential for catastrophic failures will also increase dramatically. Besides the safety issues, there are the privacy issues, and these will often be at cross-purposes with each other. As John observed, the information assurance problem will have to address the need for personal privacy within this broader context of nearly continuously connected individuals who rely on these computing infrastructures for both safety and convenience.

Steve Lipner: My record as a prognosticator is not great, but I predict two sets of changes—one good and one bad. The good is that continuous improvements in best practices for building more secure software and market demand for their application will result in significant improvements in the security of commercial products. The bad is that targeted malware and sophisticated attacks will make the jobs of detection and response more difficult—and more important.

Looking back 15 years, what were the big surprises in

information technology that significantly affected the information assurance problem?

Lipner: Clearly, the explosive growth of the Internet. Fifteen years ago, we were focusing on the security of systems that might be connected to a private network; today, everything is connected to everything else.

Dick Kemmerer

The increased information assurance needed for these pervasive systems will be orders of magnitude more than what we need today.

Kemmerer: Most everyone having access to a computer connected to the Internet and a computer in most every home. This coupled with a lack of general knowledge of the dangers one can get into, or cause, is the source of most of our DDoS [distributed denial-of-service], identity theft, and other problems. We would never think of letting someone buy a car and drive on the interstate highways without first getting some training, which includes information about the possible dangers of driving. Yet we let anyone who can afford a laptop or workstation and the monthly Internet provider fee travel the Internet highway without any training. All Internet users need to understand that leaving their systems unprotected is like leaving their car unlocked with the keys in it. It is no surprise that the early computer attacks by kids were often viewed as harmless pranks, similar to a teenager joy riding in the family car.

Landwehr: The World Wide Web—15 years ago, it didn't exist. The advent of the Web brought with it the commercialization of the Internet on a vast scale, and, most recently, organized criminal activity of many kinds that uses the Internet as a vehicle, as well as attacks that exploit weaknesses in the Internet infrastructure for criminal purposes. These developments have brought the information assurance problem to the public's notice in a big way. I'm less enthusiastic than Dick about the desirability of Internet "driver training." Today's users need security training, but in the long term, this approach smacks of blaming the victim. The Internet needs to be made safe for ordinary users, not the reverse.

Gligor: I agree—the World Wide Web and the rapid penetration of Internet technology into all aspects of life. The Internet helped change the perception that security problems such as denial of service, worms, and viruses are merely a local-area nuisance and not large-scale, global threats. Assurance is no longer perceived as an end-to-end provided feature, but an all-pervasive quality of information.

Haigh: As Carl and Virgil say, it's the Web. As PCs have become commodities, and the Web has become pervasive in our lives, the old notions of IA have had to fade away. We now have an incredibly complex, federated information system of users who do not know how to defend themselves, vendors who are forced to market before they can build strong information assurance into their product, and service providers who think information assurance should be a premium feature. No one is taking responsibility for IA at a system level, and this means that everyone is left to do what she thinks is best for her or what she thinks she is responsible for.

McLean: I agree with my colleagues about the importance of the Web, but I'd like to emphasize the cultural changes it has brought about. We've come to rely on the Internet, only to realize how fragile it really is. Industry no longer views Orange Book-like certification as being affordable. We've also seen an explosion in wireless and had experience with infrastructure collapses—for example, the loss of wireless connectivity after 9/11. I think the most important change, however, is the fact that the information assurance community has expanded beyond the DoD [US Department of Defense] and government to include industry as a serious player. This reflects fundamental changes on several levels. Hacking scripts have made it possible for even low-skilled agents to initiate relatively sophisticated attacks, which has increased the attack rate dramatically. Information sharing among hackers has also decreased the time between flaw discovery and flaw exploitation from weeks to days or even hours. Identity theft, viruses, and the need to constantly download patches have helped bring security problems to customer awareness, resulting in an increased demand for better security technology from products. As a result, industry is supplementing its traditional security arsenal—such as increased pricing and insurance to cover the cost of theft and litigation as a deterrent—and is seeking improvements in security technology. This has led to an increase in the number of security researchers and a change in the security research problem set from DoD concerns to industry concerns.

There has also been a change in how both industry and the DoD have had to deal with the problem. The aforementioned ubiquity of computing affects both the DoD and consumers. Network-centric warfare, ad hoc networks—including sensor networks—and ad hoc coalitions have demanded changes in how the DoD has addressed security issues. Further, the brittleness of our systems to attack has demanded that we retrench from a secure fortress paradigm to one where we are constantly surveying our borders for intrusions and determining appropriate responses when intrusions are detected. Security has moved from being a static, built-in technology to a dynamic warfare area requiring its own tools for sensing, command and control, decision support, response generation, and so on.

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

Haigh: Two-factor authentication! So many problems would go away if it were not so easy for adversaries to impersonate legitimate corporate users or customers. A hacker steals the password of a software vendor and uses it to steal product code. Phishers fool Dilbert's pointy-headed boss into providing the password he uses for e-banking, and we all chuckle, but it won't happen to anyone we know, right? If we only used our passwords one time, these attacks would fail!

Landwehr: I suppose one could point to the invention of firewalls. I agree with Tom about the deployment of two-factor authentication, particularly in the form of token-based one-time password schemes. We are finally seeing banks deploy them to help secure online access. As Bruce Schneier points out, man-in-the-middle attacks can still subvert these schemes, but they definitely raise the bar for attackers. Perhaps the real point here is the slow advance of technology in this field. We seem now to be in the midst of a great rediscovery of the benefits of virtual machines as an isolation mechanism. This will no doubt be followed by a rediscovery of the need to share information among virtual machines and the difficulty of safely doing so. Something truly new in the past 15 years that may yet see widespread use is the concept of proof-carrying code.

Kemmerer: Public-key infrastructures [PKIs], which verify and authenticate the validity of each party involved in an Internet transaction. These systems have led to widespread electronic commerce on the Internet. Most every business has its own Web page, and most of these allow users to purchase goods from the page.

Lipner: Going beyond two-factor authentication and PKI, I'd say the pervasive application of cryptography. Many of us had been predicting growth in the importance and use of cryptography through the 1970s and '80s, but as with so much else, it took the Internet to drive the integration of cryptography as an expected component of computer systems. I believe the Internet was also the critical factor in the relaxation of export restrictions on encryption, which were very much a factor 15 years ago.

McLean: The moves from protecting paper, to protecting analog transmissions, and then to protecting bits are enormous. The last 15 years saw a move from private networks to the Internet—a move of comparable magnitude. I agree with Tom and Carl that the technology that has helped us the most in this new environment is two-factor authentication, as embodied in any number of techniques and devices. Although the technology is certainly older than 15 years, it is only in the last 15 years that it has become widespread. An honorable mention goes to intrusion detection systems. Again, the technology is older than 15 years, but it has achieved widespread deployment only in the last 15 years. The reason for including such systems is that they were important advances in our way of thinking about the security problem, a problem that we now see as one of ongoing, real-time warfare rather than one-time castle building. The reason why they are only an honorable mention is that the advance they embody is more cultural than technological.

Gligor: The recognition that information assurance is a key element of any information system is perhaps the most significant advance. From now on, what is left is innovation and hard work to improve our information assurance posture. But at least we understand what we need to think about

The recognition that information assurance is a key element of any information system is perhaps the most significant advance.

Virgil Gligor

and where to focus a substantial part of our research and development resources.

John McLean

Risk is going to increase dramatically. IT is becoming too pervasive, and we are becoming too dependent on it for it not to become a point of attack.

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

Gligor: Secure operating systems and computing platforms will undoubtedly appear and will give us an opportunity to focus on the areas of primary importance—namely, applications security and usable user interfaces for secure systems. I believe that it is only now that we are beginning to realize that the "administrator in the loop" is a major security threat that is likely to persist even after secure computing platforms find their way to the marketplace. We have not made the life of users and security administrators easy with our interface designs and have not noticed that privileged security administrators, if malicious, could inflict untold damage in a so-far-undetectable manner.

Haigh: I think the big breakthroughs will be legal and regulatory, rather than technical. We have to start seeing the IA problem as being multifaceted and for a much broader system than the computer networks we have concentrated on in the past. I like Dick's analogy with automotive safety—the system includes automobiles, drivers, and roads at a minimum. Safety depends on the interplay of all these components with each other and with the environment, and we ensure this interplay with a combination of technology, laws, and economic remedies. So I think the next IA breakthroughs will be on the broader legal and economic fronts. The merging of safety with IA should accelerate these breakthroughs. I'm not saying this will be easy: it is a much harder problem than automotive safety. I don't know how, or how well, it will be done, but it has to be done. Imagine the state we would be in if we had relied solely on a "market-driven approach" to automotive safety!

Lipner: I'm not sure that I see breakthroughs, but I do see payoffs from incremental improvements. One example: in the 1980s, we were trying to make development teams change their entire operational models and training to apply formal verification techniques that may or may not have worked. Today, we encapsulate verification technology in tools that analyze code, such as PREfix and PREfast, and tell developers to incorporate assertions at a comprehensible level that can help the tools be more effective. The results are real—maybe not as much as we were promising in the 1980s, but real and not just promised. Threat modeling is another example of a technique that has led to real-world incremental improvements and will continue to improve.

McLean: I don't see any nascent technology that will cause a breakthrough in the next 15 years. Such breakthroughs tend to come from some unsuspected corner and are impossible to predict. Further, our failures in IA tend not to be technological, but rather human failures: failures to use technology or, at least, failures to use it correctly. I think that based on the field's progress over the last 15 years, it is more likely that instead of breakthroughs, we will see incremental improvements, realignments, and changes in direction. Given the change of landscape discussed earlier, we will certainly see incremental improvements in such things as single sign-on to help us deal with all the trust relations that will be forced upon us. On a bigger scale, we will see one of two things happening: either, as Tom suggests, we will start seeing mandated standards for IT products, or we will see a shift away from expensive, all-encompassing solutions to cheaper, piecemeal component solutions that can hit the marketplace quickly in an attempt to set de facto standards. In the latter case, market forces will create an environment where inexpensive 90 percent solutions can thrive. In the former, we will see a dramatic increase in jobs for logicians.

Kemmerer: I don't see any either. I expect that we will improve on the things that we are doing today. However, as discussed in the earlier questions, the playing field is going to be much more complex and the attackers will likely be more sophisticated. I would love to be proved wrong.

Landwehr: I'm hopeful that we will figure out how to specify and support security policies that users can easily understand and manipulate. It's true that the history of remote controls—such as for television sets and video recorders—limits one's optimism on this point. I think there is a chance to make significant progress in reducing software vulnerabilities. We should have better accountability mechanisms in place, and we should have a better means of assuring that published software is free of large classes of vulnerabilities.

What is 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.

McLean: Risk is going to increase dramatically. IT is becoming too pervasive, and we are becoming too dependent on it for it not to become a point of attack. There will be failures as well as intentional breaks caused by pranks, criminal behavior, and malicious attacks from nation states. Until now, we have operated on the assumption that an IT attack would probably be something that would be used to supplement a more traditional attack. The reasoning is based on the assumption that physical destruction provides a greater payoff for an adversary than cyberdestruction, but that cyberdestruction could seriously limit our response to a physical attack. As our dependency on IT in our critical infrastructure grows, we will have to start taking more seriously the possibility of an IT attack being an end in itself.

Further, if we continue to push for cost savings by insisting on offshore developed software and consolidating infrastructures, we will expose ourselves to even greater risk. Untrusted software is only cheaper until the point where it fails or introduces a system compromise. The current trend of consolidating infrastructures or, at the very least, using identical components across infrastructures in the name of simplification and cost savings makes it easier to bring down several infrastructures at once. The systems that will possibly be most vulnerable here will be supervisory control and data acquisition [SCADA] systems as we try to integrate control over all our different critical infrastructures.

Landwehr: The CII will grow substantially as it becomes increasingly intertwined with conventional infrastructures for transporting energy and resources. In my view, the biggest risk is that the evolving CII will at the same time become increasingly vulnerable to major failures and outages either from natural phenomena or from provocateurs of one sort or another. Unless appropriate incentives are put in place to influence corporate behavior, this future seems likely. One of the side effects of the construction of the Aswan Dam in Egypt, it seems to me, is that it created the possibility that the whole country could be flooded if that piece of critical infrastructure were destroyed by an attack. We need to avoid putting ourselves in that position with respect to our computer-controlled infrastructures.

Haigh: Given the trend toward more and more critical information infrastructure, I expect the magnitude of the risk to increase dramatically. In particular, there will be increased risk to our personal and societal safety as a result of the application of computing technologies to sensing and control for our critical infrastructures. This will make cyberattacks, or blended cyber- and physical attacks, increasingly attractive to criminals, terrorist of all sorts, and hostile nation states. Even the convincing threat of such an attack would be a powerful form of blackmail. How can one provide assurance that such attacks cannot occur?

Kemmerer: The risk is going to increase dramatically. Also, if we have systems that react to their environment automatically, then information assurance is going to need to concentrate more on the design and development end of the software life cycle. We can't afford to release systems that do not perform correctly, whatever that may mean.

Lipner: We continue to connect systems to the Internet, sometimes without adequate consideration of risk. Systems can be connected to the Internet and operate safely, but if those systems aren't adequately protected, the consequences can be very significant. We don't have good ways to quantify risks, but qualitatively, this is a real concern.

Gligor: Large-scale global attacks and users' lack of awareness of the necessary precautions as well as the ever-present threat of insider attacks are likely to remain the major risks to critical information infrastructure.

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

McLean: Even if we don't see a major change in adversary capabilities, we could be in for a very rough time. The increase in the number and sophistication of scripted attacks in recent years will

Best practices for building more secure software and market demand for their application will result in significant improvements.

Steve Lipner

probably continue, putting further strain on our defensive resources. Concurrently, the time between vulnerability reports and the appearance of attacks exploiting those vulnerabilities will continue to decrease, making our current techniques for distributing patches ineffective. Finally, the increasing percentage of junk email may simply bog down our current cyberinfrastructure to the point that people cease to use it. A possible saving grace is that adversaries may start to face the same sort of information overload that we are seeing in other communities. Confidentiality may be achieved, not from better encryption, but from the fact that the chances of an adversary coming across any particular email will be extremely remote. This will, of course, lead adversaries to try and create more irresistible methods of phishing, but as users become more aggressive in their attitude toward junk mail, most of this email may be lost as well. I can easily imagine a time in the near future, where users don't accept any email from an unknown party.

It is worthwhile pointing out that most of these changes are not really new. We are simply seeing cyber versions of known methods of attack. "Cyberization" of already known attacks will probably be a greater source of new attacks than any increased capability.

Haigh: You know, it's hard for me to see how they can get much worse. Already, attackers can hit from anywhere and hide who and where they are. There are toolkits that can automate many attacks and that can be used to create attack agents that can be distributed across the Internet. I suppose a next step for both attackers and defenders would be to apply automated reasoning techniques. Attackers could use the results of preliminary scans and known vulnerabilities as inputs and have the reasoning system generate a set of attacks to achieve a particular goal. Then they could feed their preferred attacks to the toolkit and have it generate and launch the attacks. Scary.

Tom Haigh

Information assurance will have to address the need for personal privacy within this broader context of nearly continuously connected individuals.

Lipner: Today, we see the use of automated tools for finding and exploiting vulnerabilities ranging from fuzz testing to find vulnerabilities, to reverse engineering to identify the vulnerabilities fixed by a patch, to scanning tools to identify targets. In the future, improvements and extended application of these techniques and others will shorten the interval between discovery of a vulnerability and targeted hostile attack. This trend will significantly increase the pressure on security tools and security managers to do their jobs rapidly and effectively.

Kemmerer: Over the last 15 years, the number of adversaries has increased drastically. More importantly, the attacks have gotten more sophisticated. I think this trend will continue. What I am really concerned about is the lack of conscience of attackers—that is, many attacks are carried out for bragging rights only, with no consideration of who gets harmed. There are also the attacks carried out just to harm—cyberterrorist attacks.

Landwehr: I think our adversaries will be pretty much as capable as we are—and hopefully not more so! Just because our adversaries are capable of inflicting damage to our CII in various ways does not mean they will have the incentive to do so. Even if two opposing forces could damage an asset, if the asset is more useful to each of them in its whole state rather than as pile of rubble, then they likely will choose not to break it. That doesn't mean it's unbreakable, just that nobody had sufficient incentive to do so. I have a good-news prediction for Dick: cybervandalism will decline in the next 15 years. The bad news is that economically motivated attacks will increase until adequate accountability mechanisms are established.

Gligor: Every new technology introduces new vulnerabilities and potent opportunities of attack for determined adversaries. The introduction of low-cost computing devices that contain sensitive information, and yet cannot be physically protected and can be captured by an adversary, will introduce significant new avenues of attacks on information systems. Also, new attack methods will be developed based on automated tools that discover vulnerabilities of deployed large-scale systems such as the Internet. Information warfare will become a better understood concept and, hopefully, will capture a more significant part of the national debate.

What question should I have asked regarding information assurance technology forecast that I did not ask, and what would your answer be?

Kemmerer: What is information assurance? *Answer:* Information assurance is the total package of assuring the confidentiality, integrity, availability, reliability, dependability, safety, and survivability of systems. That is, it applies to all aspects of safeguarding or protecting information or data, in whatever form.

Landwehr: Will IA issues be on the front pages in 15 years? *Answer*: No. Crime will still be on the front pages, and computers will be involved because they will meet the Willie Sutton criterion of being where the money is, but I think IA will be less of a story than it is now. That will be a measure of success.

Gligor: In the past, use of information assurance methods and tools for handling security threats has always trailed the introduction of new information-processing technologies. Is the existence of this gap between information assurance and new technologies an ever-present state, or is it a mere result of our less-than-stellar record in applying known assurance methods? *Answer*: Unless the basic parameters of security economics change in the future, it is likely that new vulnerabilities, invariably introduced by new technologies, will not be addressed until after visible and substantial losses are incurred. From this point of view, this gap—which was measured in years in the past—is a fundamental feature of any new technology and a consequence of rational consumer behavior. However, our ability to anticipate vulnerabilities introduced by new technologies is as good as any adversary's ability to exploit them. For this reason, I am hopeful that the parameters of security economics will change in the future, and that the proactive development and use of assurance methods and tools will improve our security posture with respect to future technologies.

McLean: Given that the threats to our systems will continue to increase, what should we be doing to ensure that our assurance methods keep up? *Answer*: Our basic assurance arsenal—security kernels, structured code, formal methods, and so on—has not changed over the last 15 years and will probably not change in the next 15, except where required by new computing technology, such as quantum computing. Even new computing technology may not bring about any significant changes in the arsenal. For example, the move from information-theoretic-based cryptography to hard-problem-based cryptography has not resulted in any new assurance methods for reasoning about the hardness of a problem.

What has changed is how we apply the tools we have now and how we supplement them. Formal methods are now applied to algorithms or protocols, rather than whole systems. Secure system design has been supplemented by firewalls, intrusion detection systems, virus detection programs, and so forth. Even if we return to system-level formal verification, we will still probably use the same tools, just in a different way.

Haigh: I like John's question. Consistent with my theme that information assurance and safety are blending together, I think the answer is that we have to expand our notion of assurance to take advantage of what the safety community has learned. An approach that intrigues me is the use of safety cases—"a documented body of evidence that provides a convincing and valid argument that a system is adequately safe for a given application in a given environment" [www.adelard.co.uk/iee_pn/index.htm]. This approach combines a number of forms of reasoning and evidence, including formal logical reasoning, statistical reasoning, and experience with best practices. The notion of building information assurance cases, similar to safety cases, is very appealing to me.

The CII will grow substantially as it becomes increasingly intertwined with conventional infrastructures for transporting energy and resources.

Carl E. Landwehr

Lipner: I don't know. Let me give you the statement of a reformed A1 systems builder: Twenty years ago, I thought we could build a secure system, verify it, deploy it, operate it according to the *Trusted Facilities Manual*, and we'd be done. Today, I think that even if we could do those things, by the time the system was completed, it would be obsolete and nobody would use it. I am also unconvinced it would be secure—to quote Earl Boebert, "security is in the weeds." I don't think we have a closed-form way to build a secure system any more than there's a closed-form way to design a crypto algorithm—software can be secure until a new attack is discovered, and then we have to change our assumptions. □

__O. Sami Saydjari__ is the CEO of the Cyber Defense Agency and chair of the Professionals for Cyber Defense. He has led information assurance research for 20 years, with key positions at the US National Security Agency, DARPA, and SRI International. His focus areas include high-assurance operating systems, network security, risk assessment, cyberstrategy and tactics, and security architecture. Contact him at saydjari@cyberdefenseagency.com.