Responses to Questions Posed by Ms. Melissa Hathaway During Her Presentation at the National Science Foundation on March 18, 2009

March 31, 2009

On March 18, 2009, Ms. Melissa Hathaway addressed a group at the National Science Foundation (NSF) comprised of Dr. Jeannette Wing (Assistant Director for Computer & Information Science and Engineering (CISE) at NSF), several NSF CISE Program Managers and, by teleconference, approximately 30 computer scientists who are active in computer security research and education. In that presentation, Ms. Hathaway posed to the group the following eight questions:

1. How do we optimize/derive identity management and authentication procedures while balancing requirements for privacy including anonymity on the digital infrastructure?
2. Who is responsible/accountable for the protection of rights and property given the dichotomy between the conduit and the information that flows on it fall under different regulatory/ compliance regimes.
3. How do we reach a shared vision for a future (trusted, resilient, survivable…) architecture?
4. How do you get industry buy in and connect them to the research agenda?
5. What are the quick-wins (low-hanging fruit) as we build the roadmap? (what is achievable, where are the really hard problems?)
6. How do we achieve trusted transactions in the near-term as we move toward the vision?
7. There has been a great deal of research on fail safe, self healing, software networks…what is the status and can it be leveraged now for some of the transformation initiatives (e.g., smartgrid, nextgen FAA, health IT)?
8. Can you develop an index of digital maturity vis-à-vis other nation states?

This document includes preliminary responses to seven of these questions, each of which was formulated by a small, ad hoc group of the computer scientists. No response was formulated for Question 2, as the group felt that there was insufficient time to consult with the broad range of other disciplines (law, public policy, etc.) that it felt would be necessary to produce an informed response. More specifically, these responses were assembled with promptness as the first priority, in the hopes that they will be useful to Ms. Hathaway in her 60-day review (which was at the half-way point at the time of the presentation). As such, we recognize that these responses are less well constructed than they might have been with more time, and we would welcome any inquiries from Ms. Hathaway to clarify points of potential interest; Dr. Wing at NSF has volunteered to serve as a conduit for any such inquiries.

More generally, the academic research community is eager to respond to the needs that Ms. Hathaway outlined. For example, we would welcome the opportunity to establish a modern version of a Manhattan project on authentication and identity management, or on any of the other pressing topics that she discussed. To successfully address such a large-scale challenge, we would need support, in the form of funding and endorsement, so that a large-scale collaborative effort is able to attract the best people and work effectively with industry and government organizations. We believe that such an effort, with properly managed goals and expectations, could lead to innovations that dramatically improve the trustworthiness of future information technology infrastructures. We further wish to emphasize the importance of education and the proven synergy between research and education. We believe that there is a pressing need to expand the workforce and increase research and development in this important area, and stand ready to do our part toward this end.

We are grateful for the opportunity to contribute to Ms. Hathaway's review, and hope that this will be the beginning of a fruitful discussion.

(contributors listed on the last page)

Question 1: How do we optimize/derive identity management and authentication procedures while balancing requirements for privacy including anonymity on the digital infrastructure?

The identity-management problem. Identity management involves identifying and authenticating entities such as people, hardware devices, distributed sensors, and software applications that may request access to critical information. In an increasingly networked world of interdependent systems, it is critical to develop identity management solutions for federated systems that may be beyond the control of any single organization. Identity management systems involve establishment of identities, management of credentials, oversight and accountability, scalable revocation, establishment and enforcement of relevant policies, and resolution of potential conflicts. Identity management is much broader than simply identifying known individuals by an identity card, identifying number, or stored biometric signature. Today's and tomorrow's systems require flexible, secure, and understandable identity management approaches that scale to enormous numbers of users, computer systems, hardware platforms and components, computer programs and processes, and other entities.

Socially acceptable identity management must allow for acceptable notions of privacy. For example, presenting a digital identity to one site should not allow that site to learn additional, extraneous information about the individual. A tremendous problem with the current widespread use of passwords is that many individuals use the same password at many sites. As a result, insiders or penetrators of one site (such as common social networking sites) are immediately able to impersonate users at banks and financial sites, posing threats to their accounts and assets. It should be stressed that privacy is a form of security and robustness – when presenting digital identification reveals only the information that is required for the requested transaction, everyone benefits. In particular, privacy for government agents means that critical information about them is *not* available to malicious individuals or organizations.

Successful identity management must solve three key problems:

- **Authentication.** *How does an individual or device identify itself?* Authentication can involve demonstrating identity according to some identification scheme provided by some local authority, but it also includes methods for demonstrating properties of the agent presenting a credential, regardless of whether this process results in the unique identification of an individual. For example, in many credit-card transactions, the primary goal is not to verify that the purchaser is the owner of the card, but that the purchaser is authorized to spend the amount required. An efficient technological system enables a card issuer to assure the merchant that a card is authorized for the amount in question, and the merchant knows payment is forthcoming if the card issuer authorizes the transaction, regardless of what transpires between the card issuer and the customer.
- **Authorization.** *What is an authenticated agent authorized to assert or commit?* While many enterprises hard-code authorization rules into their systems, future systems require flexible authorization policies that can be audited and maintained reliably.
- **Cross-domain interaction.** *How do interdependent systems operate together?* In federated systems, such as several branches of the federal government drawing on the same identity management system, authentication and authorization will rest on local authorities and their interdependent policies. For example, an employee of one department may have a credential (password, certificate, token, biometric signature) issued or registered with their department, and may need to use this access to some resource at another organization. To access that resource, some cross-domain credential exchange must occur.

Successful identity management solutions must resist attack and misuse. Identification and authentication are pervasively attacked by a wide range of attackers with diverse motivations, within large-scale organizations and across multiple organizations. Insider and outsider misuses are commonplace.

Near-term applications of existing research results. Many serious problems with current piecemeal approaches to identity management *can be solved by putting known research results into practice now*. The academic research community is eager to engage with entrepreneurs and government customers, through public-private partnerships and possible new collaborative frameworks, to see results of past research put to use. A process that allows the research community to do this will not only benefit user communities, but this also will allow researchers to understand more clearly where existing research falls short. This will lead to better research, better education through the successful way that universities combine teaching and research, and better solutions. We therefore strongly advocate a framework for near-term collaboration and contribution.

Examples of research concepts and methods that are promising candidates for near-term (1-3 year) application include:

- Improved web browsers that use cryptographic authentication mechanisms. Authentication protocols that can be deployed in current browsers will sharply curtail phishing. These protocols, which include password-authenticated key exchange (PAKE), use of digital certificates, passwords automatically customized per site, and methods based on hardware tokens, all have the property that data received at a phishing site is not useful when forwarded or replayed at the intended victim site.
- Cryptographic methods and potential network protocols that combine authentication with privacy protection. There are a wide range of anonymous authentication tools, such as zero-knowledge proofs, group signatures and anonymous credentials. These have the potential to fundamentally change the way credentials are asserted; for example, instead of revealing her birth date, a customer can simply prove that she is over 18 years of age.
- A wide variety of methods for specifying and enforcing authorization polices have been developed in recent research. Existing research prototypes that are available for experimental use right now could be evaluated and developed into useful products (or discarded if experimental use reveals fatal flaws).

Future research directions and their importance. Identity management is a central problem in practical computer security whose successful solution will have game-changing impact. While specific research challenges are too numerous to articulate here, it is clear that substantial progress will require investigation of usability, human factors, legal issues, societal expectations, cryptographic methods, and design and analysis of systems that issue, maintain, revoke, present, and verify digital identities. In present approaches, someone has to understand who should be authorized to take various actions and why those agents should be authorized, with the legal authority and the technical capability to enforce the authorization structure. It is a multidisciplinary, long-term research challenge to identify the identity management requirements of various organizations, develop solutions that meet these requirements in increasingly successful ways, deploy these solutions, and develop and implement lifecycle policy management tools that allow enterprises to operate successfully while protecting their information resources from internal and external threats.

Question 3: How do we reach a shared vision for a future (trusted, resilient, survivable) architecture?

We do not believe that there can be a single ubiquitous architecture, but we do believe that there is a shared vision that computer science should develop a discipline for engineering trustworthy system/architecture design using commonly available components, where security is treated as a required system property as critical as function and cost.

Commonality in future trustworthy, resilient, and survivable information-technology architectures is certainly desirable, and will not happen by accident. We do not believe that a single universal architecture (at the node, network, application, and human levels) can be invented that meets all needs; for example, military systems are unique in requiring support for information flow control that reflects classification levels, which drives the architecture of such systems. Nevertheless, we believe that a set of architectural

building blocks can be developed that can be customized for specific critical applications (e.g., process control, financial, telecommunications, health care, and military) to meet a variety of trustworthiness needs. Invention and identification of these building blocks would facilitate the construction of domain-specific architectures in a cost-effective manner, providing the "right" amount of trust for a given use. To realize this vision, investment must be made in the development of a "new engineering" for constructing trustworthy systems, where rational choices can be made among design alternatives, based on an estimate of the trustworthiness that the choice provides, and the cost associated with the choice.

Much work is ongoing in the industrial and academic technical community that could lead to the construction of these architectural building blocks. The blocks would be made up of existing commercial technologies (e.g., public key cryptography, Kerberos, IP networks and protocols, and firewalls), new security technologies currently under development, and technologies that are yet to be developed but motivated by major advances in computing technologies (such as the shift to multicore microprocessors). For example, multicore microprocessors will permit many security functions to execute in parallel with other tasks, and have the potential to drive the performance overhead of security functions toward zero. There is a critical need to fund development of individual technologies that will form these building blocks, but there is an even greater need to fund the development of candidate architectures for specific critical application domains.

Legal, regulatory, and social policies and approaches are also critical to the development of a shared vision for a set of common architectures, and should be pursued in close collaboration with technical developments. For example, regulation could be used to associate liability with those in a position to reduce vulnerabilities. Likewise, other forms of economic incentives could be created, e.g., the ability of ISPs to monetize the removal of bots from a customer computer, which would also require technical innovation directed toward creating that incentive. Finally, government procurement can be used as a carrot. In short, development of appropriate legal, regulatory, and social policies are as important as technical innovations to the construction of architectural building blocks and architectures, and must be pursued in conjunction with technical work.

Development of actionable trust, resiliency, and survivability metrics and metric estimation methods should be pursued, and would enable a "new engineering" for building domain-specific architectures out of the toolkit of technical and nontechnical building blocks that will be developed. These metric estimation methods would be facilitated by the gathering, aggregation, and release of data on attacks (in a sanitized manner), so as to allow the market to assess the correct level of investment in trustworthiness. Emphasis should especially be placed on the development of tools for estimation of *relative* metrics to permit choice among alternative architectures and configurations. (Absolute measures of system trustworthiness, although highly desirable, remain a "grand challenge.") Finally, tools should be created for measuring the relative trustworthiness of implementations. This is an area in which some success is already apparent (e.g., tools for static code analysis and for analysis of correctness of firewall rules).

Finally, we note that success in developing a shared vision for a family of trustworthy, resilient, and survivable architectures can be achieved by recognizing that fewer parts of the present architecture than one might think actually require global consensus. For example, the current Internet's architecture actually permits significant variation in the way networks are built and operated. This implies the need for a careful distinction between what is intrinsic to an architecture, and what is merely widely deployed. It also suggests that technological innovations that reduce the need for global consensus, such as virtualized networks, should be pursued. Determination of what architectural building blocks to construct and, correspondingly, which should be used in a particular architecture must be done with careful study by the research community, making use of the developed trust metric estimation methods to quantify the benefits of a proposed approach.

Question 4: How do you get industry buy in and connect them to the research agenda?

For a cybersecurity research agenda to have impact beyond the funding agencies, government, and other researchers, the industry stakeholders must be involved from the start. Industry's participation helps to ground the research and to provide a means of moving the research from the lab to the field. Research agendas produced without industry's cooperation risk solving the wrong problems. Industry will support one R&D agenda over another when they believe their interests are most aligned with the research agenda. Some factors that represent that alignment are:

- Differentiators: If the research agenda gives them a leg up, they will naturally want to participate.
- Fear: If not participating in the research agenda is thought to likely result in falling behind a competitor, they are likely to choose to participate.
- Regulation: Organizations follow laws and regulations. To remain within the law, most organizations take measures to ensure their compliance.
- Reputation: Nobody wants that CNN moment when they are held responsible for a large-scale loss of personally identifiable information (PII), as an example.

Altruistic industrial support of a research agenda that fails to include any of these factors has steadily dropped to the point of nonexistence. Given differentiation, fear, regulation, and reputation as the raw motivators of industrial support for a research agenda, we can arrive at some schemes that leverage them.

Government-supported joint academic and industrial research. Often the timeframes required for research are not in line with industrial budgets. Government funding for research can bridge this gap. As an example, NIST's Advanced Technology Program (ATP) supported this kind of research collaboration in the past, but it was replaced in 2007 by the America Competes and TIP programs, which limited participation to small-to-medium companies. However, the long line of successes from the ATP program demonstrates the real benefit of the model it used for encouraging collaboration between academic and industrial labs.

Competitions and prizes can be strong motivators. Prizes like NIST's Malcolm Baldrige National Quality Award and DARPA's grand challenges generate prestige and differentiation. Similarly, competitions such as the one held for the Advanced Encryption Standard or those offered by the X Prize Foundation have proven that companies around the world will willingly expend large amounts of resources to win a competition that will not likely result in a direct payoff to balance their investment.

Reduction of liability for companies at the leading edge. Companies that meet a threshold of security capability or competence could be afforded liability protection for breaches in their products or services. An ingredient for a company to qualify for this benefit might be its record of diligence in tracking the latest research and incorporating relevant advances into its offerings.

Question 5: What are the quick-wins (low-hanging fruit) as we build the roadmap? (what is achievable, where are the really hard problems?)

Much of what can be fruitfully done as "quick-wins" in the near-term concerns policy measures that relate to well-established and/or widely used technology, as follows.

Domain registration hygiene. Internet miscreants take great advantage of the ease with which they can register new domains as a means to evade takedown efforts and fool users. These domains often anchor a miscreant's current activity by providing the mapping from the handle extended to the user, or employed by malicious code that has infected and end system, to a locus of control. The evidence is strong that while many registrars resist providing service to such miscreants, a few unscrupulous or subverted ones readily do so. A policy-driven approach to ensure that such registrars are held accountable or lose their

ability to register names on behalf of miscreants would have a measurable positive impact on the fight to secure the Internet.

Break the "bullet-proof hosting safe haven". Another element of infrastructure that greatly abets Internet attacks is the availability of staging sites that (for a fee) greatly resist pressures to remove malicious content or capabilities. The impact of the recent Intercage and McColo takedowns starkly illustrated the degree to which large-scale Internet malice funnels through these locations. What is required is a means by which ISPs become responsible (to a certain degree) for the content/activity they host, and/or that other actors (such as ISPs with which they peer) have clear standing to sever connectivity on the basis of complaints that meet a given set of standards. These mechanisms must be effective for isolating hosting sites outside of the US. Clearly, these mechanisms must be carefully designed to protect privacy and other individual rights, and require international oversight to prevent abuse.

Enable disruption of botnets. Currently, both network operators and researchers are on uncertain legal ground concerning steps they might take to undermine botnets. What is needed is a policy and legal framework that enables these parties to readily work with international law enforcement (or other governmental representatives) to disrupt the coordination mechanisms upon which botnets rely. One can envision this system operating, for example, via a "botnet task force" whose goal is to clearly specify which activity characterizes a botnet, and support, on a case-by-case basis, the disruption of specific botnets. Such activities might include coordination with registrars to modify DNS records, or publication of IP address / URL blacklists.

Clarify research best practices. Currently, researchers face a confusing set of legal constraints and requirements when investigating botnets, malware, and cybercrime. The legal framework governing wiretap, computer fraud & abuse, the DMCA, and human-subjects research are ambiguous on these issues. Providing clear guidelines for best practices in such undertakings can help remove barriers to research-driven investigations. One exemplar in this fashion comes from the UK regarding the handling of child pornography, per http://www.cps.gov.uk/Publications/docs/mousexoffences.pdf .

Liberate data. Progress on acquiring insight into Internet attacks and then developing defenses in response is to a degree stymied by the lack of access to realistic data. There is major benefit in having, where legally possible, those agencies working in the non-classified part of the cybercrime space (e.g., FBI, Secret Service, FTC) to provide data about their cases. Such data can still have significant benefit if aggregated or anonymized, and today occasionally appears as such in indictments that become public, though this latter form can be difficult for researchers to understand in depth. A broader and more powerful form of such data release would mandate reporting on cybercrimes. For example, from a research perspective today there is **no** comprehensive data on the true cost of identity theft or stolen bank accounts, let alone a breakdown on the different elements that contribute to those costs.

Collate data. Today, information about malware, as well as specimens themselves, is disseminated via an informal "old boys" network. Such information could be more effectively managed via a high-quality, shared archive, perhaps patterned after NIST's National Vulnerability Database. This would operate in a quasi-public fashion, like the Center for Disease Control manages specimens, and with a means to also facilitate dialog between researchers regarding analysis.

Encourage ISPs to isolate infected customers. Today's ISPs vary widely in their response to notifications that a customer of theirs has exhibited behavior indicating external subversion, in part because isolating such customers incurs costs with no direct benefit to the ISP. A policy framework that provides incentives to ISPs to impose such containment could increase responsiveness and decrease the resources available to botmasters; however, to be broadly effective such an approach must be coupled with some sort of mechanism for isolating infected hosts that reside outside of US jurisdiction. One

might consider some sort of “infection notification” requirement similar to today's data breach requirements as a means to provide economic incentives in this regard.

Develop elements of the namespace with strong basic confidence properties. As an example, consider introducing a domain ".bank.us", for which registrants must qualify as domestic banks under US legal rules. Once established, such a domain would provide an imprimatur of vetting that could (1) enable better security policy decisions based on the presence/absence of names rooted in the domain, and (2) incentive for more of those who qualify to use the namespace to reap the benefits of (1).

Spur public/private partnerships. Organizations like the National Cyber-Forensics Training Alliance can create value partnerships between law enforcement and researchers who seek a deeper understanding of how cybercrime plays out in practice.

Disseminate knowledge about how to build security into the software development process. Companies such as Microsoft, Google, Wells Fargo, Cigital, and others have developed considerable expertise in the area of developing software with strong security properties. What is now needed is a means by which to more broadly build on this knowledge base and disseminate information about software design practices.

Ensure widespread deployment of security updates. Some forms of software security updates are only provided to legitimate (licensed) copies of the software. While understandable from an anti-piracy perspective, security updates in particular provide broader public benefit if made readily available to all copies of vulnerable software, including unauthorized ones. One can frame this as an analogy with stolen cars — despite the theft, it is still in the public interest that they be safe to drive.

Increase incentives to deploy DNSSEC. This technology secures one facet of Internet infrastructure that attackers leverage in order to undermine users. One can pursue this via enforced deployment in .gov and .mil domains, at the level of state governments, and by applying pressure to businesses that work with the government to deploy it.

Question 6: How do we achieve trusted transactions in the near-term as we move toward the vision?

We interpret “transactions” to encompass a broad range of online activities that involve interaction between multiple parties, such as making a purchase on a web site or simply sending an email. We consider a transaction to be “trusted” if (i) each party can ensure that it is interacting with someone legitimate, rather than an imposter, and (ii) the transaction achieves both expected outcomes (e.g., the cash transfer happens) and no unexpected ones (e.g., the client’s credit-card information is not distributed to adversaries).

Trusted transactions can be facilitated in the near-term by both technical and nontechnical measures. Technical measures include those that enable authentication of the parties to help achieve property (i), and those that harden computers and protect data so as to help ensure property (ii). There are several technologies of each type that can be applied in the near-term.

Technologies of the first type are grounded in more pervasive use of cryptographic authentication, and in particular the introduction of public-key infrastructure to enable more effective authentication of transaction participants. A particular example of such a technology is DNSSEC, extensions to the Domain Name System (DNS) that provide authenticated name resolution via public-key infrastructure; for example, DNSSEC allows a user to trust that she is interacting with an IRS computer when connecting to a computer in the “irs.gov” domain. DNSSEC is being deployed inside the US government at the initiative of the Department of Homeland Security; ideally, this deployment would be more

pervasive. Deploying stronger authentication technologies on the client side (such as client-side certificates in SSL, or techniques to transparently customize a user's password for each site) may also help, though significant challenges remain.

Technologies of the second type include methods to build more secure software, and to affirm that the intended software is being executed. Dramatically improved commercial-grade programming languages and software analysis tools have become available in recent years; more ubiquitous adoption of these could pay large dividends. Technologies of the second type also include methods for detecting intrusions, e.g., online accounts or computers that have been "hacked", by noticing deviations from their normal behavior. Numerous examples of such technologies are ripe for deployment.

The nontechnical measures that can facilitate trusted transactions in the near-term are policy changes that better allocate risk to those in a position to improve transactions' trustworthiness, e.g., via the deployment of some of the technologies we described above. For example, because the risk of stolen credit cards is primarily borne by the credit card company, these companies have deployed methods to detect and terminate stolen cards. By contrast, individuals bear the burden of identity theft; credit rating agencies and companies collecting sensitive personal information generally have little to lose if that information is misused for identity theft, and consequently have no incentive to adopt adequate measures to protect that information. Better allocating risk will provide an appropriate incentive to deploy technical countermeasures where they are cost-effective. It may also increase the chance that technical measures are appropriately matched to the specific needs in each application area and industry.

We caution that the near-term technologies we recommend above are helpful but not sufficient to achieve the goal. In particular, we wish to call attention to two areas of near-to-medium term work that may be promising.

First, a problem with current technology for authenticating "secure" web servers is that it focuses on a notion of "identity" that often does not coincide with the user's notion of "legitimate." For example, consider an e-commerce site X that outsources its sales processing to site Y, or a university X that hires 3rd-party firm Y to collect college recommendation letters. Current technology allows the client to authenticate the identity "Y", but not X's delegation to Y. The research and development community has, however, provided sufficient building blocks that we are confident that solutions can be deployed in the near-to-medium term.

Second, "trusted computing" technology being spearheaded by the Trusted Computing Group consortium and becoming ubiquitous in commercial computing platforms has the potential to add significant assurance to the requirement that parties to a transaction behave in a trustworthy fashion. Sufficient questions remain regarding the needed software and cryptographic infrastructures to keep us from endorsing this as a near-term solution. However, in the medium term, this may help considerably.

Question 7: There has been a great deal of research on fail safe, self healing, software networks … what is the status and can it be leveraged now for some of the transformation initiatives (e.g., smartgrid, nextgen FAA, health IT)?

"Self healing" generally has two meanings: (1) tolerating a fault or compromise so as to provide continued service; and (2) actually fixing the fault. There has been a significant amount of progress on (1) and much less on (2). Work on self-healing systems falls into several natural categories, listed below roughly in decreasing order of maturity. The more mature technologies below could certainly play a role in the mentioned transformative initiatives, although some are controversial and which components are appropriate for any given setting would require study conducted in conjunction with domain experts.

Intrusion detection: The ability to notice a problem or unusual behavior when it occurs is frequently important to self-healing systems. The field of intrusion detection is relevant here. There are two common approaches, known as signature detection, which detects known bad behaviors, and anomaly

detection, which detects deviation from known good behaviors. Intrusion detection has been an active research area for the past 15 years: there are specialized conferences on the subject, and several commercial products exist for detecting host-based, application-level, and network attacks.

Rate limiters. Rate limiters, also known as throttling, slow down a computation or communication in response to compromise or suspected attack, for example, in response to worms that generate high-volume network traffic. Rate-limiting methods have been demonstrated at many levels, ranging from system-level processes to new network connections to inter-domain routing (BGP). The technique is deployed in HP's ProCurve network Immunity Manager. A related technique is dynamic quarantining in which misbehaving hosts or regions of a network are isolated from the rest of the network to prevent further spread of attacks. There are some commercial applications, including one by Cisco, although many details of their operation are proprietary.

Automated diversity. Homogeneity across computing platforms leads to vulnerability to widespread attack — once a method is devised for compromising the security of one computer, all computers with the same configuration become similarly vulnerable. The potential danger grows with the population of interconnected and homogeneous computers. Automated diversity methods are intended to counter this vulnerability by deliberately introducing variations among computers that are likely to disrupt replicated attacks. Diversity does not exactly fit the above definition of "self healing", because there is no overt detection/repair control loop; however, most people think of it as a self-healing technique.

Automated diversity methods have been an active research area, they have been demonstrated at many levels (including instruction sets, address space assignments, DLLs, TCP time out parameters, and XML namespace prefixes), and the methods are reasonably mature. They have been deployed in several widely used operating systems, including Vista, Red Hat Linux, and one of the MacOS releases. There is some theoretical work formalizing the classes of problems that diversity techniques can address and what they might miss, and there is research analyzing the costs and benefits of adopting diversity techniques.

Malware writers have also adopted automated diversity techniques to avoid detection.

Fault-tolerance approaches. A *Byzantine fault-tolerant* service is one that employs replication to survive malicious compromises; i.e., it continues to provide correct service even if some servers comprising the service are under the control of a hostile attacker. Dramatic improvements in the performance of Byzantine fault-tolerant service implementations have been achieved in the last decade, and it is now known how to implement some types of such services that perform nearly the same as less resilient service architectures in the common case. To our knowledge, these techniques have been commercialized only in embedded systems, though we believe they are ripe for more widespread use.

Saving the state of a computation, or *checkpointing*, at regular intervals allows the computation to be restarted from a safe state, in the event of an error or security violation. Many versions of this idea have been studied in the context of self-healing systems. In systems that can detect a failure (corruption, integrity violation, etc.) there has been significant amount of work using a combination of self-healing techniques (checkpoint/replay and automated diversity) to achieve run-time repair. For example, one type of system rolls back on failure and then iteratively perturbs the execution state (with only safe changes) until the failure no longer occurs. There are many other examples and variations of these ideas, and they are a reasonable approach for some mission-critical applications.

Run-time repair of underlying problem. This concept is the least mature and potentially most useful component of self-healing systems. There is some work on memory-fault masking (called Failure Oblivious Computing) and on data-structure repair by learning invariants at runtime. Microsoft Research developed a system called Vigilante, which has a closed loop between detection (using broad taint analysis) and blocking (using targeted taint analysis). There is recent work on automatically repairing software bugs that cause security vulnerabilities. These projects are promising, but have not to our knowledge been deployed in commercial systems.

Question 8: Can you develop an index of digital maturity vis-a-vis other nation states?

Yes, it is possible to develop an index of nations' digital maturity, but we cannot answer this question completely here. The extent to which countries have deployed digital technologies and leverage Internet technology in their critical infrastructures, their digital maturity, is multi-faceted, and only partially visible through public sources of information. In general, to develop an index, one has to develop the measure (what things are being measured), the methodology (how are data collected), and then muster the resources to collect data. A full index of digital maturity of nation-states, fully populated, perhaps including groups capable of independent action within nation-states (e.g., organized crime syndicates, terrorist groups), would require significant research and intelligence gathering efforts. To provide quick feedback on this question, we sought to identify existing studies that could shed some light on this question, and to suggest methods that can be used to develop such an index.

Many of the relevant authoritative reports we could find readily available provide metrics regarding the penetration and adoption of information and communications technologies. The focus is often on the "digital divide", and is thus concerned with issues of equity, availability, and cost. The most recent such report, released only weeks ago, is from the International Telecommunication Union (ITU), in its annual report on "Measuring the Information Society − The ICT Development Index" (see http://www.itu.int/newsroom/press_releases/2009/07.html). The goal of this series of reports is to "develop a single ITU index to track the digital divide and to measure countries' progress towards becoming information societies." Indeed, "One of the objectives of this publication is to ... provide policy makers with a useful tool to benchmark and assess their information society developments, as well as to monitor progress that has been made globally to close the digital divide." This document details country-by-country trends in the penetration in both fixed and mobile Internet infrastructure, the rate of adoption of Internet technologies among consumers, and the way these trends differ across nations and regions.

This and other reports we could find are, however, focused on *consumer* or aggregate access to and adoption of digital technology, mobile telecommunications, and Internet communications. A complete index for digital maturity would also require information on the adoption of computer and Internet technologies in military, civilian government and within critical infrastructures in each country. One could consider measures of the degree to which embedded digital systems are deployed (e.g., Are the control systems digital?) and the degree to which digital systems are interconnected and interdependent (e.g., Can data move from here to there? Can the service operate in stand-alone mode?). Also, a broad view of the local factors influencing the speed with which countries are gaining digital maturity could provide insights about future trends. For example, the rate of production of trained IT staff may be a strong indicator of a nation's future digital maturity. In addition, identifying other groups of interest within countries that could be monitored for digital maturity could be useful.

For a thorough answer to this question, we recommend that a team of investigators with expertise in computer science, international business, policy, complex systems, and infrastructure development, be chartered with the tasks of designing the measure (selecting and designing index components), designing a methodology to compute it, collecting existing data, creating the first international index of digital maturity, and to the extent possible, projecting trends forward.

The following persons contributed to this document:

Prof. Annie Anton, North Carolina State University
Dr. David Clark, Massachusetts Institute of Technology
Prof. Nick Feamster, Georgia Institute of Technology
Prof. Joan Feigenbaum, Yale University
Prof. Stephanie Forrest, University of New Mexico
Prof. Susan Hohenberger, Johns Hopkins University
Prof. David Kotz, Dartmouth College
Dr. Patrick Lincoln, SRI International
Prof. John Mitchell, Stanford University
Dr. Charles Palmer, Dartmouth College
Prof. Vern Paxson, University of California at Berkeley
Prof. Michael Reiter, University of North Carolina at Chapel Hill
Prof. Avi Rubin, Johns Hopkins University
Prof. William Sanders, University of Illinois at Urbana-Champaign
Prof. John Savage, Brown University
Prof. Stefan Savage, University of California at San Diego
Prof. Scott Shenker, University of California at Berkeley
Prof. Sean Smith, Dartmouth College
Prof. Salvatore Stolfo, Columbia University
Prof. Giovanni Vigna, University of California at Santa Barbara
Prof. David Wagner, University of California at Berkeley