

As Simple As Possible --- But Not More So

Steven M. Bellovin, Scott O. Bradner, Whitfield Diffie, Susan Landau, Jennifer Rexford

The problem: secure the cybernetwork of a large enterprise supplying a vast array of services including disease tracking, astrophysics research, weather predictions, and veterans' health care services to a population of three hundred million. Requirements include doing so cheaply and efficiently.

The enterprise: the U.S. government.

The solution? Because providing services to the public is a fundamental role for U.S. federal civilian agencies, many agencies turned to the Internet to do so. While confidentiality, integrity, and authentication dominated early federal thinking about Internet security, agencies faced phishing, IP spoofing, botnets, denials of service (DoS), and man-in-the-middle attacks. By the early 2000s, the growing number of attacks on U.S. civilian agency systems could not be ignored. The U.S. government's solution has been to build intrusion detection systems (IDS) and intrusion prevention systems (IPS) at large scale. The project, called EINSTEIN, works at an agency-wide, and in some cases, multi-agency-wide level. Federal civilian systems have two million direct users and serve many more. While few doubt the value of IDS and IPS as part of a cybersecurity solution, can EINSTEIN really work? What attacks does EINSTEIN prevent? What does it miss? What are the privacy implications of using the interception program? We sought answers, answers that have become more important in light of talk of extending EINSTEIN to critical infrastructure [2].

The purpose of the 2004 EINSTEIN was to do real-time, or near real-time automatic collection, correlation, and analysis of computer intrusion information. IDSs were to be located at federal agency access points to the Internet. If incoming traffic appeared "anomalous," session information would go to US-CERT, the US Computer Emergency Readiness Team, a federal clearing house for cyber intrusion information¹. But information sharing did not happen in real time and EINSTEIN's voluntary nature meant that many agencies did not participate.

Part of the difficulty was Internet connections. Every small, medium, and large federal agency was connected to the network, sometimes in multiple ways, making control of incoming data and real-time information sharing extremely difficult. The government went about reducing the number of federal connections to the public Internet from a few thousand to several hundred.

¹ US-CERT collects information from federal agencies, industry, the research community, state and local governments, and sends out alerts about known malware; see <http://www.us-cert.gov/aboutus.html>.

The next program, EINSTEIN 2, uses devices located at the Internet access points to monitor traffic coming into or exiting from government networks and to alert US-CERT whenever traffic matching *signatures*, patterns of known malware (e.g., the IP address of a server known to be hosting malware or an attachment known to include a virus), were observed in incoming packets [7, p. 3.] Participation lagged, but EINSTEIN 2 is now mandatory for federal agencies.

The third effort, EINSTEIN 3, will really up the ante by using intrusion prevention systems to stop malware from reaching government sites. EINSTEIN 3 devices will be performing deep packet inspection of content, discarding suspect traffic before it reaches federal systems. (The architecture is such that only communications destined for the federal government are so inspected.) As of this writing, EINSTEIN 3 has been tested only at a single medium-sized federal agency.

Initial concerns about the EINSTEIN effort focused on privacy threats raised by the project. Because EINSTEIN IDSs and IPSs would operate on all traffic destined for federal networks, the system would intercept private communications of federal employees (e.g., if a federal employee used an agency computer to check a private email account during lunch). In this, a federal employee is not different from employees at regulated industries using company-supplied equipment for personal communications; they, and the people with whom they communicate, are also subject to company monitoring. One problem is that the monitoring techniques to be performed by EINSTEIN 3 technology are not public.

Can EINSTEIN work? That depends on what “work” means. We have the following concerns:

- **Scale: Denial-of-Service (DoS) attacks can be daunting; they have been measured at 100 Gb/s. It is unlikely that the current generation of any network device would be able to resist the DoS attacks at this rate --- let alone new attack rates likely in the near future. Indeed, it is likely that new DoS attacks will be developed using EINSTEIN monitoring functionality for the attack triggering.**
- **Ability to do Correlation: Correlation is about discovering previously unknown threats in real time as they appear. But this is impossible to do in all but very small networks. No one knows how to use a percentage of the traffic---whether compressed, diarized², or sampled---to characterize arbitrary new threats. If one is hoping to deter all threats (and not just previously known ones), *all* incoming data must be correlated and analyzed.**

² “Diarize” is used by the trade to mean making a diary of the data; in the case of a telephone call, this might be the to/from, time, and length of the call, while for IP communications, this would be the metadata of source and destination IP addresses, TCP source and destination ports, and perhaps length of packet.

One way to think about potential correlation solutions is that architectures can range from highly “centralized” to fully “decentralized” while sensors can be “smart” or “dumb,” that is, having the ability to do lots of computation locally, or not.

If analysis is done locally at the data collection point, then the need to see all incoming data requires that *all* raw signals be sent to *all* sensors. This quickly becomes unmanageable. If there are n sensors, then each sensor must look at the data from $(n-1)$ other sensors, and there are $n(n-1)/2$ pairs of data traversing the network. This is simply unmanageable when n is at all large (EINSTEIN is designed to have between one and two hundred). And the sensors would also need protecting.

An alternative approach would be to centralize the data to perform the correlation. Because summarizing the data cannot solve the problem, all the data must travel through the system to the centralized detector. (We note that in an IP-based environment, packet summary information constitutes 20-30% of the data. Thus summarizing does not provide savings at the same scale that it would for telephone communications.) This is both enormously costly for a network of any scale, as well as unable to provide the millisecond response needed in a serious attack.

(Of course, one could try a middling solution: neither fully decentralized nor fully sharing signals. Depending on where one sets collection, the problems above will still occur.

The two alternative solutions---dumb sensors and decentralized architectures or smart sensors and centralized architectures---have the worst of both worlds: they would either miss the problems, or involve enormous investment. Neither are viable.)

In short, correlation at the scale and speed at which a system serving two million users is expected to operate is not achievable using common production technology.

- Device Management: Many EINSTEIN devices will be in non-government facilities, but will need to be remotely controlled by US-CERT. Protecting control mechanisms and pathways against intrusion, disruption, modification and monitoring will be very challenging.

- Signature Management: EINSTEIN 3 will use classified signatures developed by the government as well as unclassified signatures from commercial IDS and IPS vendors. These signatures will have to be protected from the access point operators as well as from Internet-based attackers.
- Data security: In some cases federal regulations require the use of encryption (e.g., in sharing medical records). For EINSTEIN to function as advertised, communications transiting the IDS/IPS must be decrypted. Public documents do not discuss how the system will handle encrypted traffic, and what security measures will be used to protect the data.

These complexities make it highly unlikely that EINSTEIN can achieve the job for which it is being designed.

We have concerns about cost. If we assume that the IDS/IPS function at a federal civilian agency will be similar to that in commercial network defense products built and sold by Narus, Cloudshield, a back-of-the-envelope calculation shows each router directing traffic will require 64 times as much equipment to perform EINSTEIN-type filtering [1]. This is clearly a losing battle. In addition, it means that EINSTEIN --- or at least---EINSTEIN 3 will cost roughly one billion dollars just for equipment.

EINSTEIN also raises policy concerns. Any IDS looking for long-term subtle attacks must store large amounts of traffic for non-real-time analysis. System design and configuration will determine what is stored and when. The data EINSTEIN collects will have many possible uses. History has shown that investigatory tools are often misused by those with the tools [5], [8]. There is a significant risk of mission creep for EINSTEIN, and generating detailed logs for all functions that the EINSTEIN 3 device has been configured to do is crucial. Yet current EINSTEIN 3 documentation does not describe details of the auditing system. Given the size and scope of the EINSTEIN effort, these should be public.

What EINSTEIN can accomplish is limited. EINSTEIN documentation mentions threats of phishing, IP spoofing, botnets, denials of service, distributed denials of service, man-in-the-middle attacks, or the insertion of other types of malware [6, p. 3], without noting that phishing, IP spoofing, and man-in-the-middle attacks cannot be prevented by EINSTEIN-type systems. U.S. Deputy Secretary of Defense William Lynn III has called cyberexploitation, the targeted theft of U.S. intellectual property from industry and government sites possibly “the most significant cyberthreat that the United States will face over the long term”[4, p. 3]. EINSTEIN does not protect against this except when the phishing relies on previously known malware for the attack.

While EINSTEIN is a government intrusion-detection/intrusion-prevention system designed to protect U.S. federal civilian agency systems, there is interest in extending the system to critical infrastructure, including communications and public utilities such as the energy smart grid [9]. This is contradictory---a classified U.S. federal government program for protecting widely used private-sector systems. We have great doubts about extending EINSTEIN to protect privately held critical infrastructure. The architectures and functions simply don't match.

Federal civilian systems serve two million employees, but critical-infrastructure systems in the U.S. serve over three hundred million Americans. Scale matters. Can a program that effectively protects the communications of federal agencies with a hundred thousand employees each do the same for communications giants that serve a hundred million people instead? The sheer number of communications in the commercial communications networks is dwarfed, in turn, by those of the "smart grid," the planned power network that will use digital technology to monitor and control power generation and usage.

Size is not the only issue in transitioning EINSTEIN systems from federal civilian agencies to the private sector. While the U.S. government can mandate the types of technologies used by federal agencies, typically the types of systems used in the private sector cannot be so mandated. The biggest problem, however, in attempting to extend EINSTEIN-type technologies is the lack of applicability of the technology to privately-held critical infrastructure.

Consider commercial information and communication technologies (ICT). In the 1990s, the rate of communications transmission was sufficiently slow that the communications bits could be effectively examined and stored---at least if one did sampling. That's no longer true. Meanwhile communications technologies are in a state of constant innovation. For proper functioning, IDS and IPS should be designed to prohibit those types of communications that are not explicitly allowed. ICT use of EINSTEIN-type technologies would delay deployment of innovative communications technologies. This would have a devastating impact on U.S. innovation and competitiveness.

Or consider the power grid, which is a loosely coupled federation of many independent (sometimes competing) parties with complex trust relationships [3]. This architecture vastly complicates consolidation of the type required by EINSTEIN. Even if consolidation were possible, the need for timely delivery of real-time data and the requirement of high reliability make it undesirable to circuitously direct grid control data through a small number of consolidated access points. In the power grid, function mismatch creates another problem. IDS/IPS solutions useful for protecting U.S. federal government computer networks may not match well to the power grid. Many parties in the energy grid already have their own IDS/IPS and firewall solutions from a variety of vendors, making the EINSTEIN 3 equipment at least partially redundant. These existing IDS/IPS solutions are often integrated with other important functionality such as quality-

of-service, compression, and SCADA³ reports (which are part of Critical Infrastructure Protection requirements for the North American and Federal Energy Regulatory Commission). While these reports are generated by the same equipment that performs IDS and IPS, EINSTEIN 3 equipment cannot realistically subsume this functionality.

Putting it simply, there are deep and fundamental differences between communication networks supporting the U.S. federal government and those supporting private sector critical infrastructure. These differences create serious problems in any attempt to extend EINSTEIN-type technologies to private-sector systems controlling critical infrastructure. This is true in the United States and, depending on architecture, may be true elsewhere.

EINSTEIN sounds good in theory. In practice, even implementing EINSTEIN in the restricted environment of federal civilian agency systems is highly complex, and it is far from clear that this billion-dollar system can deliver sufficient security to be worth the cost. In the domain of privately owned critical infrastructure, the potential of EINSTEIN is much less clear. Electronic fences protecting critical infrastructure sound good, but once one examines network architecture more carefully, EINSTEIN's fit is highly questionable. In determining how to protect critical infrastructure, one should keep in mind what Einstein himself was purported to have said, "Everything should be made as simple as possible, but no simpler"---and then develop solutions accordingly.

Bibliography

[1] Bellovin, Steven M., Scott O. Bradner, Whitfield Diffie, Susan Landau, and Jennifer Rexford, "Can It Really Work? Problems with Extending EINSTEIN 3 to Critical Infrastructure."

[2] Hoover, N. "Cyber Command Director: U.S. Needs to Secure Critical Infrastructure," Retrieved from *Information Week*.
<http://www.informationweek.com/news/government/security/showArticle.jhtml?articleID=227500515>, September 23, 2010.

[3] Juniper Networks, *Smart Grid Security Solution: Comprehensive Network-Based Security for Smart Grid*, 2010.

³ SCADA (Supervisory Control And Data Acquisition) systems are used to monitor and control industrial processes.

[4] Lynn III, William, Defending a New Domain, *Foreign Affairs*, September/October 2010.

[5] United States Congress, Senate, Select Committee to Study Governmental Operations with Respect to Intelligence Activities, Final Report of the Select Committee to Study Governmental Operations with Respect to Intelligence Activities: Supplementary Detailed Staff Reports on Intelligence Activities and the Rights of Americans: Book II, Report 94-755, 1976.

[6] U.S. Department of Homeland Security, Computer Emergency Readiness Team (US-CERT). *Privacy Impact Assessment for the Initiative Three Exercise*. Washington DC., 2010.

[7] U.S. Department of Homeland Security, Chief Privacy Officer. *Privacy Impact Assessment for EINSTEIN 2*. Washington D.C., 2008.

[8] U.S. Department of Justice, Office of the Inspector General (March 2008). A Review of the FBI's Use of National Security Letters: Assessment of Corrective Actions and Examination of NSL Usage in 2006.

[9] Zetter, Kim, "Let Us Secure Your Network For You or Face the 'Wild, Wild West' Internet Alone," WIRED, <http://www.wired.com/threatlevel/2010/05/einstein-on-private-networks/>, May 27, 2010.