Making Information Safe

Our dependence on technology has made the Computer Security Technology Center's developments—electronic counterparts to guards, guns, and gates—crucial for protecting our nation's information assets.

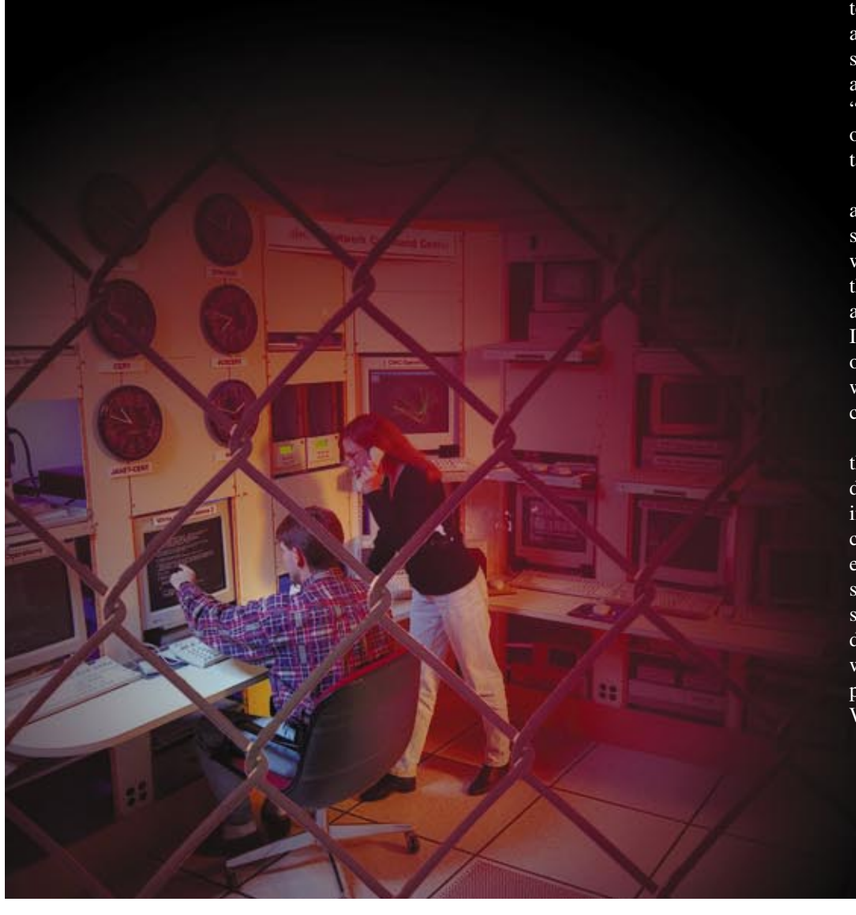

VERY late one night in November of 1988, a warning appeared over the Internet: a virus was running loose in cyberspace. As it turned out, the warning was apropos but incorrect—it wasn't a virus but something worse. A computer virus needs the help of a user to activate and spread it; whatever was attacking systems on the Internet was seemingly able to search for and infect any location without assistance. It "wormed" its way through networks, overloading machines with invisible tasks and preventing their effective use.

As word spread, system administrators frantically shut off their systems from the Internet, hoping they weren't too late in defending themselves. They rested easier only after the worm was removed from the Internet. The worm's perpetrator was one Robert Morris, a graduate student, who eventually was convicted of computer fraud and abuse.

The Morris Worm will go down in the annals of Internet history as an early demonstration of how vulnerable and interdependent network-based systems can be. Even though it specifically exploited the weaknesses of a particular subset of UNIX systems, all Internet systems suffered days of service disruption and weeks of uncertainty while costly cleanup activities took place. The likelihood of more Morris Worm–like attacks led the Department of Energy to take two important steps to safeguard information on its computer systems: it created an incident response team to contain computer intrusions and prevent their recurrence, and it increased sponsorship of projects that advance the cause of computer security.

24-Hour-a-Day Security

As a direct result of the Morris Worm attack, DOE in 1989 formed the Computer Incident Advisory Capability (CIAC), an organization based at Lawrence Livermore that provides on-call incident response and transmits security incident information throughout DOE sites. Today, it is the oldest response team in existence funded by a federal civilian agency and is a recognized institution both nationally and internationally.

When CIAC receives notice of an incident, it assesses its extent, and determines if catching the intruder is possible. If the site where the incident occurred chooses to try to capture the intruder, CIAC monitors the break-in and coordinates with other sites and law enforcement to trace the intrusion back to its origin. After the intruder is caught or if the investigation determines that the intrusion cannot be traced, CIAC provides appropriate technical resources to contain the incident and fix the system's vulnerabilities. It collects and verifies information related to the incident and disseminates information about new vulnerabilities and patches (fixes for vulnerabilities) to the DOE community. CIAC's services are funded by the DOE and are available 24 hours a day, 7 days a week.

CIAC's incident handling capability is the central, reactive component of a larger security service that also provides awareness training and education. It does so through comprehensive, customized workshops tailored to a user group's specific information-security needs. Workshop subjects include threats and countermeasures, firewalls, connecting to the Internet securely, legal issues, and even briefs on how to use CIAC effectively.

As part of its work, CIAC keeps close ties with other response teams, commercial vendors, law enforcement agencies, and other government agencies to track the latest technology trends and the latest known information about network security threats and vulnerabilities. It publishes a well-recognized security Web site on the Internet (http://ciac.llnl.gov/).

To extend CIAC services to all other federal civilian agencies, the U.S government funded a new joint effort with a sister team called the Computer Emergency Response Team Coordination Center (CERT/CC) at Carnegie Mellon University in Pennsylvania. This new virtual team is called FedCIRC (Federal Computer Incident Response Capabilities).

Integrated Protection

The Morris Worm incident occurred at a time when awareness of computer security issues was beginning to grow. In 1986, Congress enacted a Computer Fraud and Abuse Law, following it in 1987 with the Computer Security Act that established a national framework for addressing computer security issues and required federal agencies to plan and train for security incidents. Since then, awareness of computer security has increased because worldwide connectivity is increasing at exponential rates (Figure 1), and computer security compromises are increasing in parallel with it. In 1995, for instance, an estimated quarter of a million computer intrusions occurred on Department of Defense computers alone. Trends indicate that the number of intrusions doubles each year, so that by the end of 1997, it is estimated that DoD computers were attacked one million times.

Computer intrusions into DOE and other computers can range from annoyances such as chain letters (make your lucky day luckier by sending this message to a dozen friends) and hoaxes (don't open this file or read this e-mail message because it will destroy your system) to malicious attacks that deprive computer users of service, destroy files

(a) International Connectivity, June 15, 1995

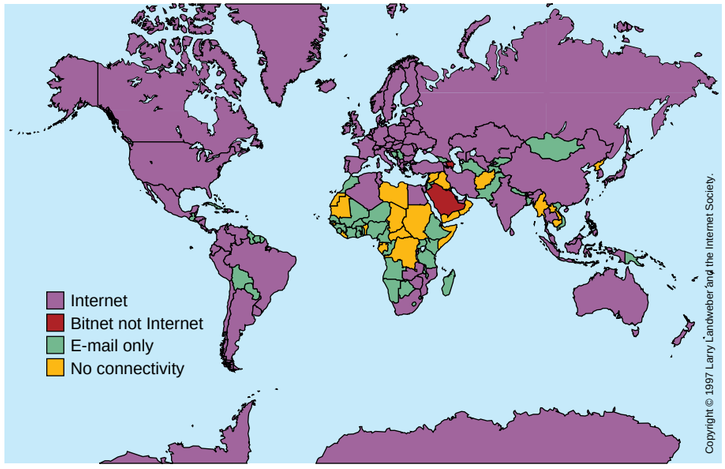

(b) International Connectivity, June 15, 1997

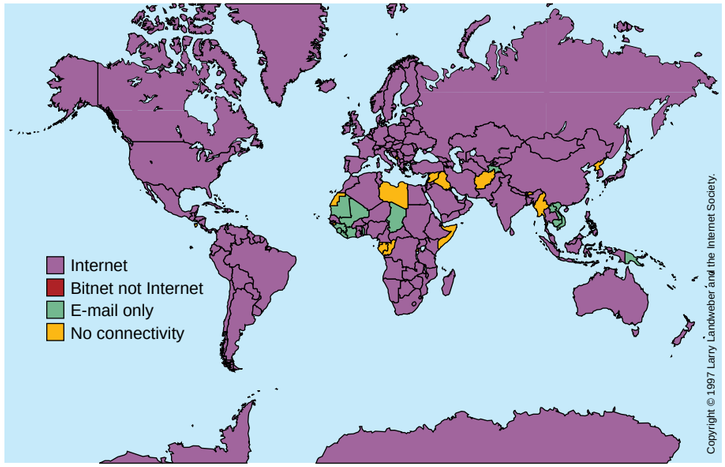

Figure 1. Comparison of (a) and (b) reveals an exponential increase in worldwide computer connectivity in just two years. Computer security issues have increased at a similar rate.

and hard drives, or steal sensitive or proprietary information.

What has particularly worried computer security specialists is the growing number of hackers, the growing technical sophistication of their attack tools, and the leveraging of their expertise. Hackers have begun sharing automated hacking tools with each other, enabling many more hackers, including less-experienced ones, to attack computer systems with impunity, exploit arcane system flaws while fully covering their own tracks. And they can do all this without necessarily understanding how the tools work (Figure 2).

In this context, the second response DOE had to the Morris Worm attack was to sponsor the establishment of the Computer Security Technology Center, or CSTC, at Lawrence Livermore. Kernels of CSTC had existed at the Laboratory since the 1970s, when prescient computer specialists such as Chuck Cole and, later, Doug Mansur (now the program manager of CSTC) began working on computer security research and development projects. Cole, who recently retired as Deputy Associate Director of Operations in Livermore's Computation Directorate, was such a strong champion of computer security that he was as much a factor as the Morris Worm attack in convincing DOE to create a formal entity dedicated to information security. Once formed, the CSTC combined the incident response work of CIAC with two other important components: advanced security research and development projects, and outreach consulting services. This integration of capabilities has proven to be powerful, and the CSTC has become an increasingly influential focal point for

information protection throughout the federal community.

Security through Penetration

Among the consulting and professional services that CSTC staff provide is one they dub the White Hat review, a friendly attack of a client's information systems. These systems are likely to be complex, with global computing functions, telecommunications, open architectures, and diverse platforms and protocols that span geographic boundaries and time zones. Their interdependencies put all components at risk if any one fails, thereby jeopardizing the security of the total system. At the same time, system complexity exceeds the protection capabilities of most safeguard mechanisms.

As a way to actively manage the risks of complex systems and improve information protection, a client can request that a White Hat team perform system and network penetration tests and acquire a snapshot of security strengths and weaknesses. Members of the White Hat team are Top Secret–cleared, information security specialists, armed with current intruder techniques and tools, who attempt to penetrate an information network and learn the state of protection measures in the system. They really are just the other side of the coin of CIAC response personnel—generalists who use their computer skills to root out security problems.

White Hat activities generally comprise three phases and use methods previously negotiated with client management: scan and map a network to determine its topology and identify its vulnerabilities, intrude and compromise systems by exploiting the discovered weaknesses, and analyze results to recommend protection improvements. Unlike organizations whose systems suffer hostile attacks, the clients requesting a White Hat team always maintain complete and continual control of their systems and the intrusion process.

Advanced Security Tools

The specialized research and development work performed by CSTC staff has led to the development of security tools now in use in DOE and other federal environments. A number of the tools have been used to catch intruders, and one of them made national news while doing so.

Detection Sets Court Precedent

In early 1996, federal investigators charged an Argentinean student with illegally accessing U.S. military computers. The student apparently had broken into his university's Internet-linked computers to steal passwords and then used the network to penetrate computer systems at the National Aeronautics and Space Administration, the U.S. Navy, the U.S. Army, and systems in Taiwan, Mexico, the United Kingdom, and South America. He had managed to get access to a variety of sensitive government information before the U.S. Navy traced the culprit and nabbed him. To apprehend this hacker, the Navy used the Network Intrusion Detector (NID) software developed by Lawrence Livermore computer scientists and based on earlier work with the University of California at Davis.

NID is a suite of tools that detects and analyzes unauthorized computer access. Working within a network of host computers called a security domain, NID runs undetected by the intruder as it collects information packets (data packaged for transmission) and statistics across the domain. First, it uses a tool called

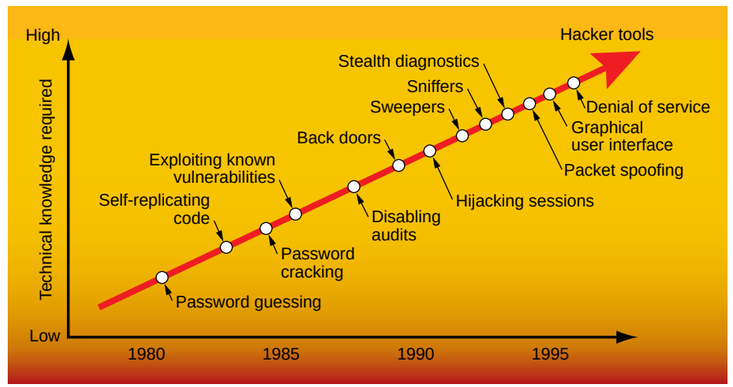

Figure 2. As technical knowledge increases, so do the number and sophistication of hackers' tools. Alarmingly, hackers without detailed technical understanding of those tools are still able to use them.

iDetect to look for evidence of an intrusion by examining information packets for intrusion signatures (that is, a string of characters known to be used for attacks). Collected evidence is presented to an authorizer that approves the transition to iWatch, NID's second evidence-gathering phase. iWatch scans a network for connections that contain the same signatures found by iDetect. If iWatch provides compelling evidence, then a third subset, iScript, is used to convert the packets of data into a transcript suitable for use in court.

Before NID software could be used against the Argentinean student, the FBI had to convince a judge that NID would not violate privacy standards such as those imposed on wiretaps. Accordingly, NID was modified to address the issue of civilan computer privacy. The modifications took into account the conflicting values of information protection versus privacy and made use of an evidence-gathering model that searches for specific patterns. If the data search detects an apparent specific pattern, permission could be obtained to continue with specific data collection.

On March 29, 1996, Attorney General Janet Reno announced on national television that an arrest warrant had been issued for the student. "We are using a traditional court order and new technology to defeat a criminal, while protecting individual rights and constitutional principles that are important to all Americans," Reno said. The case set a precedent for evidence gathering on the Internet.

Detection in Near-Real Time

Had Automated Information System (AIS) Alarms been available when the Argentinean hacker was breaking into the network, he might not have gotten as far as he did. An intrusion detection program that is in development, AIS Alarms works much like a building security system connected to a police or security station. It uses sensors distributed throughout a network to detect specific suspicious events. Sensor information is fed to a central assessment module (CAM), which is outfitted with a set of rules for interpreting the information and determining the state of system security. The assessment triggers a number of possible system responses, such as turning on more sensors to get more security data; notifying a system administrator of abnormal or improper activity on the network; or reconfiguring a firewall, router, or other network protection device to isolate particular users, addresses, or network services (Figure 3).

AIS Alarms recognizes a security incident in near-real time and with great flexibility. Its three parts—the sensors, central assessment engine, and response agents—are all planned as "plug-and-play" elements that can be configured and reconfigured easily (even "on the fly") in the computer architecture. This feature allows users to tailor the system for different networks, local policies, and threat environments. Sensors can be ramped up when a threat has been detected (the response agents can turn more sensors on) or are disabled to conserve computing resources. The rules used by the CAM can be changed to redefine what constitutes a computer attack, thus giving system administrators great leeway in specifying what should be detected and how responses should be formulated. The CAM may be made to merge information from any number of sensors, and it may be linked into hierarchical systems to protect local, regional, and national computer networks. Whatever the configuration, the AIS Alarms remains automatic: it can run unattended and will, on its own, take evasive action against attacks.

The AIS Alarms project is a collaboration of the Lawrence Livermore, Los Alamos, and Sandia laboratories. The tri-lab team has designed the software as a continually evolving system. Because there is a constant leapfrogging of security solutions and new attack methods, the team's approach has been to develop a prototype system quickly and then fine-tune it through real application and experience. The result is ever-improved security, better understanding of risks, and minimized computing resource overhead.

A Network SPI

DOE commissioned the Security Profile Inspector (SPI) analysis program specifically to counter attacks like the Morris Worm and was joined by DoD's Defense Information Systems Agency in sponsoring its development. Developed at Livermore, the program is now being used throughout DOE and DoD; the transfer of its technology to the private sector is being pursued.

SPI simultaneously assesses the security of all machines in a designated security domain. Users and system administrators can run SPI on demand or on a set schedule. Either way, they are actively defending their systems from hackers and even from insiders trying to escalate an attack to more sensitive parts of the system.

SPI has six modules that are used to collect and report system security information. They are installed on every host computer in the security domain. The modules query the status of a system's files, users, and groups; look for common security problems and known vulnerabilities (the list of which is constantly updated); uncover poorly chosen passwords; create a database snapshot of important user, group, and file information that can be used to detect unauthorized changes or additions; test the access controls; and ensure that the system contains only up-to-date, authentic software (that is, no Trojan horses) with the latest patches for detected flaws.

The computers installed with these modules communicate, via secure channels, with a command host computer that aggregates, processes, and integrates all acquired information and generates reports assessing the state of the system. The command host becomes, in effect, the "system administrator" of the security domain.

A centralized system administration is crucial for safeguarding networks. Yet, when computing resources are distributed to myriad users, tasks, and workstations, this function is usually left to end users with little or no system administration experience. SPI addresses this problem by providing for uniform, expert security management across many machines from a central workstation.

Ways to Practice Deterrence

Computer security starts with a system in which a user can place complete faith: it is "clean," is properly configured, and has had all upgrades and recommended security patches installed. These are prerequisite to effective access control, account monitoring, and appropriate network services. But

(a) The AIS Alarms setup allows hacker recognition.

(b) Recognition and assessment by AIS Alarms trigger flexible responses.

(c) Notification of intrusion and response is almost instantaneous.

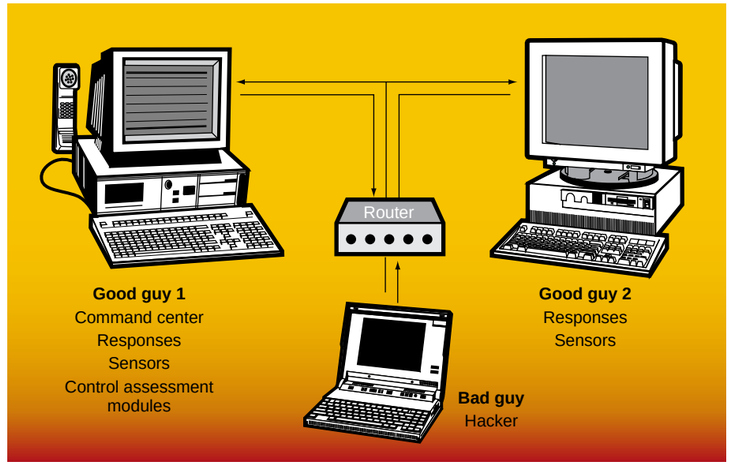

Figure 3. Currently under development by the Computer Security Technology Center, Automated Information System (AIS) Alarms allows (a and b) recognition of and (b and c) response to security incidents almost as they are happening.

keeping an electronic house in order is not as easy as it sounds, because operating system software changes constantly and system administrators must deal with many systems and heterogeneous environments. As upgrades arrive for both the core software and security patches, the versions to which they apply are difficult, if not impossible, to track. Worse, their intended applications (the system, type, version, or architecture to which the upgrade applies) are not readily apparent. Sometimes, upgrades and patches do not do what they are intended to do or arrive with incomplete or erroneous installation packages. Sometimes, they even arrive security-flawed straight out of the box. Part of the Morris Worm attack was, in fact, based on exploiting one such flaw to gain illicit access to systems.

System administrators will have some housekeeping help from

David Crawford (left), a member of the Computer Security Technology Center's incident response team, works with tools developer Stephen Wong to refine the products and services that protect Laboratory computer systems from unauthorized penetration.

Lawrence Livermore's Secure Software Distribution System (SSDS), which is currently in development. This practical, automated tool can be used to query, maintain, and upgrade the software integrity of hundreds of individual systems from a central point, through largely automated means. When completed, the SSDS tool will quickly, automatically, and regularly assess and authenticate system software, collect vendor upgrades and patches, determine the applicability of upgrades and patches to specific systems, install them and related critical system software, remove patches if for some reason a system must be restored to its previous state, detect instances of subsequent tampering, and collect sitewide software statistics or metrics.

SSDS works through two components: an SSDS server that resides in one computer and an SSDS agent in each computer being monitored. The server performs the key functions of tracking vendor upgrades and patches, converting any new ones into standard formats and storing them in a database, and comparing database information against local system files to determine what has been installed and what still needs to be installed.

Patching tools similar to SSDS attempt to keep track of the patches that they have installed by building a "patch history" file. However, because these tools do not have the capability to survey the local file system, they can be easily fooled into reporting erroneous information. In contrast, the SSDS server queries the agents about the file owner, access control list, and cryptographic "hash" and compares this information with its database to ascertain what patches are actually installed on the local file system. SSDS bypasses reliance on the local patch history file, which may be incorrect or compromised.

The SSDS can be configured to support a variety of environments, whether small homogeneous networks or large heterogeneous ones. A simple configuration was described above: one server serving agents installed on all target systems. When, as at Lawrence Livermore, hundreds or thousands of systems running a variety of operating systems and architectures are in use, multiple servers will be used to collect patches and upgrades. The functions of evaluating and installing them are delegated to another subset of the system, with the number of systems performing these functions determined by the size of the security domain. The SSDS has great flexibility for supporting a variety of systems by distributing different workloads without duplicating effort.

Identifying Classified Information

Many government agencies and other organizations need to be sure that the electronic documents on their open computers are free of classified or other sensitive information. Also, since World War II, DOE, its predecessor agencies, and their contractors have generated billions of pages of classified materials. Various recent laws and court decisions now require DOE to swiftly declassify and release many of these documents. Declassification is not an easy task, because two authorized classifiers, at least one of whom must have additional training and authorization as a declassifier, must determine that a document no longer needs the protection of classification.

CSTC, through the Text Analysis Project (TAP) funded by the DOE Declassification Productivity Initiative, has been developing software tools to assist in identifying classified information for proper electronic or hard-copy storage, deletion, or declassification.

TAP works by reviewing documents against a rule set based on classification or other guidance. A TAP rule is a collection of words and phrases along with conditions based on proximity such as "within the same sentence" or "within eight words" and, in some cases, quantitative constraints on individual items such as "later than 1980" or "mass greater than 5 kilograms." Synonym lists induce multiple variants of most rules. The rule set leads to a table of rule words and to other tables specifying constraints and relating words to phrases and phrases to rules.

To process a document, TAP "reads" through it looking for rule words and tracking their locations. When TAP finds all the words for a particular rule and has determined that they meet that rule's conditions and constraints, it declares a match, or hit, assigns it a hit number, and specifies the applicable rule number and the precise location of the hit in the document being analyzed. The user can now display the document with the hits highlighted. Jumping from one hit to the next, an authorized classifier or declassifier will see additional information for each hit—the classification guide and topic on which the rule was based and the associated classification level.

TAP can batch-process large numbers of documents and provide a summary report to be used by a classifier to prioritize documents for review or by an administrator to assign documents to appropriate reviewers.

Classifiers and declassifiers are currently using TAP to support systematic reviews in which documents are separated into two categories (classified and unclassified), but no sanitization is done to turn classified into unclassified documents. Later, as DOE produces and refines rule sets targeted at various types of information, TAP may be able to support sanitization efforts and to replace one of the two reviewers required for declassification.

Solution Is a Moving Target

Tools to fend off attackers and safeguard our information have not, as we know, completely protected us from computer intrusions. They might never do so, because attack methods change and software flaws continually appear—they are moving targets. Nevertheless, the work of the Computer Security Technology Center is vital in protecting the Laboratory's and DOE's information assets; its staff will continually search for more and more advanced solutions. Doug Mansur says, "There's hope for containing these problems. For even the most perplexing security problems today, we can offer at least partial solutions."

—*Gloria Wilt*

Key Words: AIS (Automated Information System) Alarms, Computer Incident Advisory Capability (CIAC), Computer Security Technology Center (CSTC), computer intrusions, document classification, hacker, incident response, Internet, Morris Worm, Network Intrusion Detector (NID), Secure Software Distribution System (SSDS), Security Profile Inspector (SPI), software patches and upgrades, Text Analysis Project (TAP), virus, White Hat review.

For further information contact Douglass Mansur (510) 422-0896 (mansur1@llnl.gov).

About the Center

Lawrence Livermore's COMPUTER SECURITY TECHNOLOGY CENTER (CSTC) is composed of 32 computer scientists led by Douglass Mansur, center manager (pictured at left). He is assisted by Harry Bruestle, deputy center manager; Sandra Sparks, head of the incident response team; and John Rhodes and Lauri Dobbs, co-leaders of tools-development projects. CSTC got its start in 1989 with the Computer Incident Advisory Capability (CIAC), an organization begun by DOE at Livermore to identify and respond to breaches in computer security throughout the DOE complex. This 24-hour-a-day incident-response capability is made possible by a variety of new and evolving tools developed by CSTC personnel to monitor and protect computer systems and networks, to respond to and deter penetration of those research and development resources, and to identify and secure the unclassified and classified information stored in and handled by Laboratory, DOE, and civilian government computers.