1

Strategic Directions in Computer Security Research

Teresa F. Lunt
U.S. Defense Advanced Research Projects Agency
Information Technology Office

Introduction

As a society we are becoming ever more dependent on computing and communications technologies. Yet the technologies we are using are fragile. We need to develop technologies for a robust and secure information infrastructure. These technologies should be affordable, verifiable, and scalable. In addition to fundamentally new technologies, we also need to develop strategies for strengthening legacy systems that form part of the critical information infrastructure.

Directions for New Security Technology

Affordability

Affordability is essential for security solutions that must be ubiquitous and involved in routine, day-to-day communications and information processing. Affordability is achievable only if security technologies that are shown to be effective in our research laboratories transition to commercial use.

Generic Security Over the past two decades, much of the research in computer security has been sponsored by the U.S. Department of Defense. Much of this work has focused on multilevel security (MLS). Several laboratory prototypes were built to demonstrate the feasibility of high-assurance MLS operating systems and database systems. Very little of this work has transitioned to commercial products, however. Many vendors did produce MLS versions of their standard operating systems and

P. Samarati et al. (eds.), *Database Security*

database systems; however, these were generally MLS versions of their standard product, and the two (MLS and standard) product lines quickly diverged. This divergence leads users to prefer the standard versions, since most of the popular applications will not be available or may not work correctly on the lesser-known MLS versions. And this, in turn, means that those users who need MLS still do not have an affordable solution; much customization and special-purpose applications and integration code must be written.

What is desired, instead of a separate "secure" product line, is for vendors to build security into their mainstream products. This is feasible only if the vendors perceive that a large segment of their users want the security. With traditional MLS security, vendors do not see a good match with the needs of their users. We need a broad set of interests to be met in the security features of general-purpose products. Rather than build special-purpose products to a fragmented security marketplace, we should seek ways to unify these needs and satisfy them in a more general way. Policy-neutral security mechanisms are preferable to simply MLS mechanisms. Such mechanisms would be capable of enforcing any number of organization-specific policies, including MLS, but would not have any single policy "wired in." These mechanisms should allow a broad enough set of policies to be specified and enforced so as to appeal to a wide set of user communities, such as finance, health care, and commerce, as well as defense. One can envision a future in which national-security-"blessed" policies will be available from third-party vendors for use with these generic, but specializable, products.

Richer Policies Most organizations have more complicated information protection needs that simple mandatory and discretionary access control matrix-oriented policies are capable of expressing. In addition to the familiar mandatory and discretionary access control policies, we should also explore richer policies such as role-based, task-based, and workflow-based policies. This is necessary if our work is to appeal to the broadest possible constituency. To do this, however, we need to identify a desirable range or class of policies, investigate natural ways of expressing such policies, identify and develop a common set of mechanisms capable of enforcing the desired range or class of policies, develop policy "compilers" to map user-specified policies into the base mechanisms, and address the related assurance issues.

Easier Administration Along with much more expressive policies comes an increased administrative burden. So, we also need to do research aimed at increasing the ease of security administration. This will be expecially important when we consider using many of these products together, and then using many instances of these different products in a networked environment.

Modular Systems Based on current research trends in operating systems, we expect future systems to be more modular. This may also be true someday of database systems. This will give us the opportunity to make security a modular and reusable component of systems. This has the advantage that the end user need only use the security modules if they need and are willing to pay for the security. It also means that various degrees of security can be made available for use with the same products. Moreover, it may be possible for several different systems to share the same security "modules," so that a common security policy can be enforced across diverse system components. This will allow us to move away from the confusing situation we are in today, in which multiple access control mechanisms are layered on top of each other, which interactions that are not well understood, because the system uses many different products layered on top of each other. There is the additional advantage that security modules can be replaced by high-assurance national-policy-enforcing modules when the systems are used in certain defense applications.

Of course such extensible and customizable systems bring with them new security vulnerabilities. This is the topic of the next section.

Assurance

New approaches to building systems are motivated by the need to serve a diverse set of needs in a general-purpose product while at the same time providing maximum performance for any given use of the system. To meet these conflicting goals, various approaches are being explored that allow an application to specialize the operating system to its needs, remove the application/operating system boundary, co-locate tasks in a single address space, or to provide its own extensions to the operating system kernel. This new thinking diverges with the traditional security engineering practice, which dictates that strict (preferably hardware- enforced) separation must be maintained between the system or OS kernel and the applications, that security mechanisms be encapsulated in a trusted computing base or security kernel, and that strict domain isolation be maintained among processes. New research opportunities arise for addressing the vulnerabilities that come with these new approaches, and for developing new approaches to system security. In addition, it may be possible to take advantage of these new system "features" for introducing security into the system in novel ways, say, by including a variety of security "specializations" which can be used only when needed. Different applications could use different security specializations depending on the need, by choosing from among a variety of mechanisms (including no mechanism) of different strengths and costs. Another approach to security in such systems may be to transfer part of the security enforcement responsibility to the application. Other approaches may include type-safe languages and compilers. New methods for establishing the assurance of such systems might be investigated,

for example run-time checks of correct policy enforcement.

Providing high assurance for critical properties in operating systems goes only part of the way toward addressing the system assurance need. In critical applications, much of the software responsible for ensuring that the properties are met will be application code. We need to understand how we can achieve high assurance for critical properties in systems that are composed of components that are independently specified, developed, tested, and evaluated.

In security, an attempt was made to address this problem in the Trusted Database Interpretation of the Trusted Computing System Evaluation Criteria. There, different system architectures were considered in which a trusted operating system could provide support for application assurance. It was found that the design of the application was a crucial determiner of whether strong operating system assurances could be applied to the assurance of the overall application/OS composition. Further research is required to understand the relationship between application and operating system, and the implications for system assurance, in new operating system paridigms. Even more urgently, research is needed to understand more varied and complex system compositions, and their impact on security.

New networking approaches are being explored that will allow for customization of how network traffic is handled, in effect allowing one to "program" the network. These approaches bring with them new vulnerabilities and also new opportunities for introducing security. We must find ways to assure that the effects of can be confined according to some policy. In addition, the "code" that will form part of the communication stream must be assured to be safe. Popular new programming paradigms such as Java and software agents require similar attention to safety and assurance.

Scalability

Today's security solutions are being built for aging computing and communications technologies. Many of these solutions will not scale to the technologies of the future, and the future is just around the corner. For example, mechanisms that depend on cryptographic authentication of every packet in a data stream, that require frequent reference to distant directory servers to ascertain certificate validities, and that require lengthy appendages of signed certificates, may not be able to keep up with the speeds of new high-speed networking technologies or be at all appropriate for mutually authenticating software agents. Access control policies that were designed for a closed environment may not scale well to world-wide-web-style environments in which there are frequent interactions between unacquainted entities, or to a highly networked environment in which new alliances are quickly forged and terminated. The new phenomenon of cyberspace opens up privacy concerns that were not present

in small, closed communities where one's every computing activity was not on display to the entire world.

In many cases, the focus on security must change from the individual end system to the network. For example, in intrusion detection, we must find analysis techniques that scale to very large systems (i.e., that do not require massive amounts of data to be collected) and that can produce reasonable results with partial data (since not all portions of a network are always visible). These systems should also be refocused to monitor network activity rather than exclusively end-system activity, and they should be made to work in a variety of networking technologies. We must better understand how to instrument our systems and networks so as to give us the requisite visibility. We also need to develop both intrusion detection and system management tools that can operate across administrative domains, or that may work with a network of other autonomous detection or management systems in a cooperative or hierarchical manner. These systems should be capable of dealing with extensive heterogeneity both with respect to the systems monitored and the detection and management systems themselves.

We also need scalability in the analysis tools we use to gain assurance. Formal methods tools such as theorem provers and model checkers must be speeded up by orders of magnitude so that realistic problems, once specified, can be analyzed within minutes. In addition, their reasoning abilities must be augmented with libraries of domain-specific hardware and software design theories. This will enable these tools to have a real impact on system design.

Directions for Improving Legacy Systems

Most of the information infrastructure is going to be with us for a very long time. Telecommunications systems, electric power generation and distribution systems, financial systems, and transportation control systems will slowly evolve but will retain their legacy character through generations of technology improvements. In addition, many new critical systems, such as medical devices, defense command and control systems, and nuclear power plant control systems, are being constructed using commercial software products. We must begin work now to understand and deal with the risks of using commercial and legacy components in systems we depend on for our national well-being and personal safety.

We must develop strategies for working around the problems that are inevitably to be found in legacy and consumer-quality products. We must discover architectural "workarounds" intended to augment the strengths or compensate for the weaknesses of these components.

An approach that is being investigated is whether security can be introduced into a system by developing security "wrappers" for certain system components. With

this approach, wrappers would be used to introduce certain security functionality without altering the legacy code or the other system components that use it. The idea is to gain control over specific interfaces where a security function can be inserted. Such interfaces could be library calls, system calls, or other interfaces internal to a subsystem. For the approach to have any validity, it must be possible to ensure that all input to and output from the wrapped component can be intercepted by the wrapper; in effect, the wrapper becomes a reference monitor for the policy it enforces. This is the fundamental new assurance question for the approach.

This new approach requires new theories of secure composition of a system from components (including wrappers) and technologies for security integration. We must broaden the types of analysis that can be performed far beyond such narrow considerations as secure information flow for multilevel security. We must reason, for example, about how such diverse aspects of security as authentication, access control, and encryption contribute to overall system security when inserted into a system in various ways. In addition, our reasoning must allow for ignorance, empirical properties, or worst-case assumptions about legacy components. To support such reasoning, we must adequately specify the components; research is needed in order to understand what must be specified.

Security can be inserted in this manner to meet a variety of objectives. For example, it is easy to imagine how a wrapper could impose an access control policy on the wrapped component, or encrypt the outputs and decrypt the inputs of a components, or perform inter-component authentication, or perform message filtering. One could also design these wrappers to add security monitoring and intrusion detection capability. Ideally these wrappers should be designed so that the specific security solution is a modular part of the wrapper. This would allow the module to be replaced, for example, when it is desirable to use a stronger security solution. This should also allow multiple security modules, enforcing orthogonal policies, to be inserted in the same wrapper.

It has long been held by the security community that security must be designed into a system from its inception and cannot be added on later; we must investigate the feasibility of this new approach and discover how far and for what aspects of security it can be made practical.

How to Demonstrate Progress

Today there is no way to assess the security of a system or component of a system. In fact we do not have a set of agreed-upon definitions of security that would form the basis for such an assessment. Intuitively, we would expect that a measure of security would imply something about the ability of the system to withstand or survive attack. One notion that has been suggested in the time it would take for a system

to be compromised, given that it is under continuous attack. Such notions of system security could be augmented with more specific measures designed for particular kinds of components. For example, it should be possible to evaluate the effectiveness of an intrusion-detection or authentication component under a set of assumptions. Research is needed to develop security metrics and evaluation tools. The application of such metrics to new security solutions will allow us to demonstrate our progress towards increased security.

Surviving Attacks

In this age of global connectivity, we are exposed to a far greater potential threat than in our former closed working environments. Even with all the new security technology we will be developing, we must assume that a determined adversary can still penetrate our systems and manipulate them to various ends. Survivability against nontrivial threats is rarely considered. Generally security controls are non-redundant and are assumed not to fail, so that scenarios involving the failure of one or more safeguards are not designed for. We must develop new approaches to enable penetrated systems to continue to provide service to the most critical needs, and to prevent an intrusion from rapidly spreading its effects throughout a system. Research is needed to develop system design approaches that will assure increased continuity of service. Some current ideas have been inspired by biological and social analogies, such as redundancy, diversity, adaptibility, and self-healing. Much investigation is needed to see if such ideas will bear fruit.

Partnering with the Computer Science Community

As the world continues to develop new technologies for computing and communications, it also opens a Pandora's box of vulnerabilities associated with those new technologies. Typically these new vulnerabilities are not addressed or even recognized during development of the technology development. Security concerns are deferred until later, if ever. The security research community cannot continue to work in isolation of the greater computer science research community and to develop solutions primarily for existing technology. We must sensitize the world around us to the collective vulnerability due to our dependence on computing technologies for our continued economic well-being, national integrity, and personal safety. We must cultivate a sense of obligation to address security and survivability concerns in the new technologies that we are developing now and in the future. While new technology is still in its earliest conceptual stages there are unique opportunities to influence its development so as to minimize vulnerabilities and strengthen security. The community of researchers with special expertise in security and survivability must make connections

with the research groups investigating these new technologies. We cannot continue to perform security research for technologies after the technologies have developed.

Partnering with Industry

To assist industry in understanding the vulnerabilities we are faced with, and to encourage industry to develop the best possible security solutions that will meet our needs, it will be necessary for the government, particularly DoD, to enter into a partnership with industry. This partnership should make accessible to industry the information available within the government on vulnerabilities. It should also make expertise available to industry and the research community as those communities attempt to develop new countermeasures. Government experts could also provide informal evaluations of proposed approaches. Such exchanges could greatly leverage the expertise that exists within the government, particularly concentrated within a few hundred individuals, and that largely does not exist in industry.

Conclusion

We are in time of opportunity for computer security research. New developments such as economic commerce, combined with the awareness of new threats, such as information warfare and electronic terrorism, are converging to create an unprecedented demand for security. At the same time, many are realizing that there is a paucity of solutions. We must not repeat the mistakes of the past by focusing on the needs of only a specialized community and failing to partner with researchers and industry to address the vulnerabilities of the next generation of technologies.