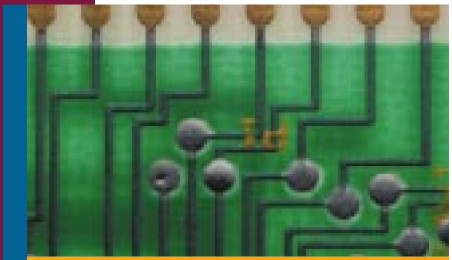

BY JOHN GRIFFITH

Protecting Corporate Data From the Hoards While Opening it to the Masses

Client/server and web access offer great opportunities to take advantage of the valuable information contained in formerly closed corporate systems. However, they also greatly increase the risks of unauthorized access. Proper planning, careful design and implementation, and active management of a security infrastructure are the only ways to minimize the risks associated with these new information-sharing paradigms.

COMPANIES that have decided to liberate the information stored in their corporate systems so that partners, customers and even competitors can access it must also consider how to protect it. Client/server and web access offer great opportunities to take advantage of the valuable information contained in formerly closed corporate systems. However, they also greatly increase the risks that unauthorized users can access or alter the information in such a way that all its value is lost.

Proper planning, careful design and implementation, and active management of a security infrastructure are the only ways to minimize the risks associated with the new information-sharing paradigms. This article will discuss the incremental risks associated with sharing data over a network, evaluate security needs through formulating application-level security policies and examine some of the means available today to reduce risk.

NEW SOURCES OF RISK — THREAT ANALYSIS

Before networks, information could be protected through relatively simple means such as locked doors, file cabinets, fences and security guards. Networking, however, opens paths to information to potentially everyone on the planet who owns a computer. Sometimes companies want this, but most of the time they don't.

In beginning a threat analysis, understand the risks involved with networked information. Placing information on the network opens it up to three primary threats:

- loss of data integrity
- destruction of data
- theft of information or money

Hackers will try to access it for fun, just to prove they can. Competitors can use it against the company. Criminals will use it to commit fraud or outright theft. Unhappy employees can misuse their access to destroy or damage the information in order to get back at the company. Mistakes can damage the information or wipe it out completely. Understanding the risks involved will help as you develop an analysis of potential weaknesses.

In general, when developing a threat or risk analysis, classify people into three groups: authorized users, users and others. Authorized users are the people you grant one or more types of access to your information. Users are people who you know have networked access to your system but not your information. Others are the people you don't know and don't want to provide any access to.

Authorized users can be the most dangerous to your information. They have the means and opportunity to do damage. They generally know the most about what information your system contains and how to obtain it. They use the application(s) that access your information on a daily basis and therefore know about any flaws that exist that might allow them access you didn't intend. Also, in many, and perhaps most, instances honest mistakes can cause the loss of your information unless you take adequate measures to protect it. Remember that people you share your information with may not be as familiar with the application used to access it as you are, so protect it accordingly.

As a group, users are the next most dangerous. Think of them as employees in other departments in the company. They have access to the network, and possibly to the system that privileged information is stored on, but the intention is that they should have either limited access or no access to that information. They often know of the information's existence, and, depending on what it is, have various degrees of motivation to access it. Losses from this class of people are generally not as innocent as those caused by authorized users, since they are only supposed to have, at best, limited access. However, because they have access to the network, they have an opportunity to eavesdrop on network traffic and steal the information.

"Others" is everyone who is not a user, authorized or not. These are people you do not want accessing your information at all. Obviously, in the case of a public Web site, there are no "others." This group of people is not supposed to have access to your network at all, but may still desire to access your information — often with criminal intent. They initially may not know anything specific about your information or applications, but can learn about them over time, given a chance. While their intentions may not be harmful initially, this can change over time, or they can share how they accessed your systems with others who have other motivations besides just breaking into your network.

In order to decide how much effort a company should take to control these groups, guidelines must be established to help decide who should be placed into which group, and also what measures may be required to help protect your information. This is a function of a security policy.

DEVELOPING A SECURITY POLICY

Developing security policies helps companies determine what information they need to protect and how they want to protect it. The eventual goal of this process is to develop application-level security policies that protect specific information repositories. The foundations of an application-level security policy may already exist in a company in the form of general security and information security policies. Companies can build on these to develop application-level policies, which in turn guide their selection of specific security mechanisms.

The first step in protecting your information is to determine what kinds of information you are sharing and the corporate policies that apply to who can use it. Larger organizations almost always have a written corporate information security policy that deals with information classifications such as these:

- **Public:** information shared with the general public, such as press releases and public access web sites
- **Proprietary:** information shared only with partners under non-disclosure agreements, but generally available to the employees or contractors, such as project information
- **Restricted:** information only available to a subset of people – such as financial data
- **Personal:** information of a personal nature, including personnel files and records
- **Confidential:** information only available on a need-to-know basis

Before networks, information could be protected through relatively simple means such as locked doors, file cabinets, fences and security guards. Networking, however, opens paths to information to potentially everyone on the planet who owns a computer.

Many such policies also classify users into basic categories:

- **The world:** everyone not in another category — the general public
- **Corporate partners:** subcontractors who are granted access to one or more internal systems for specific reasons
- **Authorized users:** the group of people who work inside corporate facilities and have access to corporate computing systems
- **Employees:** all people who are directly employed by the firm
- **Project/department members:** people who are members of a specific group who need access of a particular type as a part of their job function

A corporate security policy will generally identify classifications of data, categories of users, and the general rules and conditions for allowing access by users of each category to data in each classification. A comprehensive security policy should also identify roles and responsibilities for business managers, data owners and administrators, IS managers, and end users.

Larger organizations may also have a second level of policy beyond corporate information security — an information systems security policy. This level of policy deals with the specifics of storing and protecting information on computer systems. It may include requirements for passwords and password management, user account management, physical security, system and network level access controls, management procedures, and audit procedures.

Corporate security policies range in detail from the extremely formal, strictly enforced policies of governments down to simple policies in place at small companies with only a few computers. If an organization has a security office or a CIO, these are good places to start looking for a corporate security policy. Other sources include the finance and the personnel organizations — they have legal requirements to protect some or all of the data they maintain.

Corporate security policies are a good place to start when developing application-level security policies. To do that, corporate security policies must be translated into rules that apply to the specific information the application is providing access to. The development of an application security policy is best accomplished by asking the business sponsor (the person or organization that asked for the application in the first place) and the data owners (the person or organization responsible for maintaining the information the applications will use) about existing policies. Data owners may have a written policy in place dictating the rules of access for the information they manage. More often, it is "just known" or implied by the way the information is currently protected and current users are managed. It should also be noted that the needs and desires of the business sponsor occasionally conflict with those of the data owner(s) and corporate level security policies. Identifying these conflicts and getting them resolved is a natural side effect of this process.

The application security policy identifies what information is required to implement the application, what types of access are anticipated (read, delete, add, modify), and who should have what types of access and under what conditions. The policy should be as simple and specific as possible so that it is easy to understand and implement, yet cover all the conditions anticipated for the application. An example policy for an application that provides access to employee records might be as simple as: Employees may view only their own personnel records and edit the home address and home telephone number in their records. Managers have employee access to their own records and may also view the employee records of the employees who directly work for them. Members of the personnel group will have employee access to their own records, may add new employee records, and may edit all employee records except their own. No user shall have the ability to remove personnel records. All modifications to and creations of employee records will be recorded in an audit log accessible only to the personnel manager and the security officer. Unauthorized users and authorized users who are not employees shall not have any access to any records maintained by this application. No employee information of any kind will be permitted on a network unless it is encrypted using a technique currently approved for corporate confidential information. The application will comply with corporate systems security policy for identification and authentication. The host systems shall be located, maintained and operated in compliance with corporate policies and procedures for systems containing personal and restricted information.

SECURITY ARCHITECTURE

Once an organization has identified its security needs by developing application-level security policies, it can start selecting a variety of devices and system capabilities to implement solutions. These technologies provide features that generally fall into one or more of the following categories:

- **Authentication:** Identifying users before granting them an identity
- **Identification:** The means to securely identify networked users
- **Authorization:** Limiting access based on some attributes of the user
- **Secrecy/privacy:** Preventing eavesdropping of traffic between the user and the application
- **Audit:** Keeping a record of all accesses to identify possible misuse

A security architecture is a combination of devices and features that implement and enforce the security policy. Often, individual network components such as routers and switching hubs have useful security-related side effects. Network operating systems such as Novell NetWare and Microsoft NT contain features that can be effectively used to implement security policies. Distributed security infrastructures such as Kerberos and public key infrastructures can be used to supplant operating system features and use common means to enforce policy across a network of dissimilar systems.

> Perhaps the most important step in protecting your data is to know what exactly you are sharing and the policies that apply to how and with whom that information can be shared.

A good security architecture supports the concept of "defense in depth." Like a secure facility, access to sensitive information should require the user to pass through more doors, and sensitive information should be used in secure facilities. A company certainly doesn't want to put its personnel database on its public web server that is accessible from the Internet, but it may want to give its payroll vendor some limited form of access to the network so that it can cut checks every payday. Of course, having to authenticate multiple times in order to get to an application is a usability problem, and can even decrease the overall security of information by encouraging users to write down their user IDs and passwords and keep them in a handy — and not secure — location. Distributed security infrastructures and network operating systems avoid this problem by identifying users with a non-forgeable credential that can safely be transmitted over the network and verified by the application or device.

Of course, a good security architecture should be cost-effective. Placing a firewall in front of every network server or requiring the use of a highly secure system as a server is probably overkill in many commercial situations. The expense of protecting the information needs to be balanced against the value of sharing it. Sometimes just a good backup plan is enough to protect the system.

PROTECTING THE NETWORK

Network design should keep public Internet access points separate. This includes not only the web server, but also any other publicly advertised service such as email servers. This permits other access points into and out of the network from the Internet to remain relatively anonymous, and therefore less likely to be attacked by someone with a specific goal to break into the system. Not only should Internet servers be protected by a firewall from random access, the internal network should also be protected from Internet servers by a second firewall in case someone figures out how to subvert them.

For internal access from a corporate network to the Internet, devices and software packages are available that can hide your internal addresses from external users — these include the network address translation capability available on many firewalls and proxy servers. Both of these reduce the amount of information an attacker can gather about an internal network and therefore, make it harder for him/her to identify weaknesses that may be present.

Dedicated lines offer communications between sites and business partners that are reasonably secure (at least in the United States), but these can be protected by using encryption devices if the data is more sensitive. Also, if the connection is for an extranet partner whose users are not supposed to have general access to the network, an access controller and/or firewall is an appropriate way to protect your network and restrict the systems and applications they can connect to. Virtual private networks (VPNs) avoid the costs of a dedicated line by creating a tunnel over the network and using encryption to protect the information going through it.

Ethernet and Token Ring, two of the best established network technologies, share a vulnerability in that in their basic form, every system has access to all the information that is sent across them. Routers and switching hubs can perform a security function by isolating traffic to a small group of systems. For instance, using either device to isolate the personnel department from the rest of the company can prevent eavesdroppers elsewhere in the company from sniffing out sensitive personnel information. Since switching hubs basically provide a "private" network between sender and receiver, these can really limit undesirable broadcasting of information while improving overall network performance.

External access to an internal network, be it through a dial-up connection or the Internet, should require authentication using a method that does not allow an eavesdropper to later impersonate an authorized user. Remember that if your company must support dial-in access from remote users in foreign countries (say, a sales force), not all of these countries offer as much protection for telephone conversations as does the United States. In addition, the use of cellular phones may expose user passwords and sensitive corporate data to eavesdroppers. The use of an authentication protocol such as SKIP and CHAP that protects the password from exposure, a one-time password mechanism, or the use of a non-forgeable credential such as public key authentication, should be required in order to gain access to internal systems. In addition, the use of a VPN, such as encrypted PPTP, should be used to protect sensitive information from prying eyes.

THE IMPORTANCE OF AUDITING AND MONITORING

The security technologies discussed thus far serve the role of fences. Their purpose is to separate people from information or access they are not supposed to have. Certain people have keys to the doors, and those keys need to be protected as well.

Unfortunately, fences and walls can be scaled, and any network security technology can be defeated given the desire, time and resources. Auditing and audit analysis are the equivalent of watchdogs and guards, and are perhaps the most under-utilized security measure today. Part of the problem is the lack of a good set of standards in this area. Another part is the lack of tools to help identify patterns of access that could be related to an attack.

Unfortunately, like a guard, regular system auditing is expensive. Still, a regular review of audit logs produced by firewalls, VPN devices, dial-up access controllers, and server operating systems and applications must be conducted in order to detect possible attacks in progress. With luck, you'll identify an attack before it is successful and be able to prevent any damage. If you are too late, audit logs could provide the means to identify the attacker and permit you to take action against him/her.

Alarms are a less expensive alternative to constant auditing. Many devices offer the capability to detect simple patterns of access (such as three failed logins in a row, a dozen rejected packets from the same source in a specified time period, and so forth) and send an alarm to a monitoring application. This can allow network operators the ability to detect attacks without the need for constantly reviewing every audit log produced by every device on the network.

SECURITY ADMINISTRATION POLICIES AND PROCEDURES

Besides the technical security measures such as those mentioned previously, a good security policy does not neglect provisions covering the areas of personnel, physical and procedural security. For instance, the policy might require that a security administrator's manual be maintained for each application and system to assist security personnel with overseeing their operation. The physical protection requirements for sensitive systems and information should also be covered — it makes no sense to spend a great deal of money on secure software and systems when a janitor can steal the entire system from inside the building after normal working hours. Such requirements should include requirements to support the backup and restoration of information and the physical and environmental protection of the backup storage media. Of special interest today are policies, procedures and recommendations for protecting laptop computers and the information they contain, since they are easily, and quite often stolen from employees when they are on the road. Finally, security policies often mandate periodic audits, usually by an external agency in large enterprises, to ensure that the application developers, system operators, and the security staff are following existing policies and procedures. Again, all of these requirements should be cost-justified, but simple measures such as requiring that a cable lock be purchased with each laptop and used by employees when traveling don't cost much and can prevent quite a few problems in the long term.

SUMMARY

Perhaps the most important step in protecting your data is to know what exactly you are sharing and the policies that apply to how and with whom that information can be shared. Once you have determined or established all the policies that apply to your information and application, it is a simple matter to develop specific functional requirements for your application.

A good security architecture, however, does not depend on a single security mechanism to protect your information. A "defense in depth" establishes multiple means of protection and limits who can accidentally or intentionally access your data in a way you did not intend. Firewalls and other external access control mechanisms permit you to control who can access your internal network — routers and switching hubs can be used to further limit who can see the data as it moves between server and client. VPNs and dedicated lines allow you to extend your network to remote sites and business partners with some control.

Any security infrastructure needs to be monitored and audited in order to help detect attacks, help ensure that the security policy is being correctly enforced, and that the various system components are working properly and being managed correctly. Measures such as establishing a good set of security operating procedures, physical security measures to prevent direct access to sensitive systems and information on media, and requirements to protect laptops and the information they contain are also necessary to protect valuable corporate information. Finally, periodic audits covering the security policies themselves and how they are implemented across the workgroup, data center or enterprise are valuable to ensure that they are up-to-date, relevant and are being followed. Combining these measures will help companies realize the benefits of networking information without exposing themselves to unnecessary risk and loss. ts

John Griffith is a professional services consultant for Gradient Technologies Inc.