



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

**Remote Access--Protecting the Internal Network
Or
How to Allow “Them” in While Keeping “Them” out**

January 16, 2001

Joe Quinby

© SANS Institute 2000 - 2005, Author retains full rights.

1. <u>The Problem</u>	3
2. <u>Separating “Them” from “Us”</u>	3
2.1. <u>Cover yourself with an MOU</u>	4
2.2. <u>Define the User Population</u>	4
2.2.1. <u>A Question of Identity</u>	6
2.2.2. <u>Individuals or Groups ACLs</u>	6
2.2.3. <u>Centralized authentication</u>	7
2.3. <u>Create a secure infrastructure</u>	7
3. <u>Controlling the flow</u>	8
3.1. <u>DMZs</u>	8
3.2. <u>Extranets</u>	9
3.3. <u>Reverse Proxy</u>	9
3.4. <u>VPNs</u>	9
4. <u>Connectivity</u>	9
4.1. <u>Dedicated Lines</u>	10
4.2. <u>Dial-up modems</u>	10
4.3. <u>Satellite Modems</u>	10
4.4. <u>Cable Modems</u>	11
5. <u>In Summary</u>	11
References	11

© SANS Institute 2000 - 2005, Author retains full rights.

Remote Access--Protecting the Internal Network Or How to Allow “Them” in While Keeping “Them” out

1. The Problem

Over the years, the computing environments at many corporations have evolved by way of mutation, rather than via a formal program of designed development. Mergers between corporations and the lack of clear-cut policy have resulted in amorphous systems of networks tied together with numerous configurations of software and security, implemented as the need has arisen. Although we have come to recognize the need for standards within our computing environments, many Information Technology (IT) professionals cannot seem to agree on the most basic tenet of access control. Some corporations put great dependence on their firewall, and feel that within the internal network there should be no restrictions--that all company associates should have unfettered access to anything. Others have an attitude of “This data is Company Confidential and we mustn’t share it with anyone, not even our own people.” Still others have already invited too many outsiders past their firewall, and now face the need for access control within their internal network because the outsiders” are already roaming around inside. Although all these arguments may have merit, corporations need a firm policy stating what is and is not acceptable, else the dam bursts and all access control is lost.

In this paper, I am taking the stance that outsiders are just that-- outsiders, and should not be allowed into internal networks. While a company may have some control over employees, they do not have the same control over outsiders. They cannot perform the same screening of outsiders as with their own employees, they cannot apply the same disciplinary action to outsiders as they can with an employee, an outsider has no loyalties to the company as an employee does. This paper will explore some methods of remote access and (hopefully) will provide the reader with some ideas on how to implement them. It is not highly technical, but may be useful to explain to management the need to provide well-defined remote access policies and controls to protect a corporate internal network.

2. Separating “Them” from “Us”

Here is a simple analogy: how many people outside your family have a key to your house? Perhaps a trusted neighbor, whom you have known for years, has a key in order to put the paper and the mail on the table and water the plants when you’re away. There exists the expectation that this neighbor will not go “snooping” into your medicine cabinets or dresser drawers. You trust this neighbor because, over the years, you’ve had the opportunity to screen them and they have passed your standards for integrity. You have very likely also made a reciprocal agreement, an unwritten “memo of understanding”, where you in turn have a key to their house for the same purpose.

2.1. Cover yourself with an MOU

I am not saying that outsiders do not have a valid need for information contained on corporate networks. There are legitimate needs, and just as in the analogy above, there is a mutual benefit by sharing information. Unfortunately, corporations don't have the advantage of living next door to all their outside partners and hence, cannot gain first-hand knowledge of their integrity. It is therefore necessary to draw up a more formal Memorandum of Understanding (MOU) between the corporation and the outside parties wherein both parties legally acknowledge the rules of behavior. It is then the corporation's responsibility to take steps to ensure the outsiders stay out of the medicine cabinet if they are not included by name on the prescription. A Conditions and Limitations agreement, stating that both the rules of use and the consequences of behavior not consistent with those rules, is the main part of the MOU. Rather than state what they can and cannot do, it usually states what outsiders are allowed to do and that any other use is forbidden. MOUs usually state who is responsible for what (i.e., who maintains the system, software, client hardware/software) and lists points of contact for getting assistance. Finally, officers from both the corporation and from the outside organization need to sign the formal document.

An important aspect to the MOU is the fact that it is a legal document, a binding contract, if you will. Without it, a corporation has almost no recourse against the outside party who performs mischief. A much-abbreviated version of an MOU might be a legal notice or warning banner, offering users the opportunity to either accept the rules of use, or log off. Particularly, it should prohibit unauthorized use, and notify the user that their actions may be monitored, and that they should have no expectation of privacy. This applies to employees as well as outside users. It is next to impossible for a court to rule in favor of the company if they do not have such a banner.^{1,1} Most Federal organizations require them.

2.2. Define the User Population

Let's look at the typical composition of a corporation's "family and friends". We find

Employees (on- and off-site)	w/constant interaction
Temporary Staffers	Off site contractors/businesses
On-site contractors	with varying needs for
On- and off-site regulatory	interaction
agency personnel	Student interns
On-site visitors	Supply Vendors
Satellite offices of the	Schools & Universities
company	Worldwide Public
Parent companies	
Off-site contractors/partners	

Of these, which are "family" and which are "friends"? For the most part, family should be given access to just about everything on the inside. Note, I said

“just about everything”. There may be some information that, by law or some other agreement, must not be shared with anyone outside a particular programmatic area. For example, not all employees have the need for Research and Development information, nor do student interns, temps, corporate spies, etc. It is to be made available only to those working in the program. A corporation must, then, segregate their user population into some type of groupings, with each group having some common need for access to information.² A few (but certainly not all) sample groups might be:

- Managers (perhaps even several levels)
- Employees (all locations? Or by geographical or administrative organization?)
- Temporary Staffers
- Project-Specific (may include partners, contractors, universities, other sites, etc., but all with specified access requirements)
- Interns
- Regulatory agencies
- Vendors
- Contractors
- Worldwide Public

Many of these groups might be broken into various sub-groups. Project-Specific groups often vary in size and number. These are generally short-term assignments, generated at a departmental level, for a lifespan of the project.

Which of these are considered “family” and therefore to be given access to the internal networks? And if not “family”, how should the information be provided to the “friends”? Certainly managers and employees would be “family”—after all, they live here in the house. Yet, temporary staffers and student interns live here too. How should they be treated? The rest (usually) come to visit and are either invited into the house (escorted by the host), or perhaps they might remain out on the front porch where drinks and sandwiches are brought out to them.

These “visits” are electronic, and therefore cannot be escorted in the traditional manner. Either we must provide them a cordoned-off walkway through the house to the information they need, or we must bring the information out onto the front porch. What about remote offices of the corporation, or even home-based telecommuters? They have a need for an extended stay and require access to real information, not just to a webserver published for the benefit of Internet users. They effectively have a need to become extended family. We need to invite them in, and give them the same run of the house that the typical employee gets. This is offered by one of three ways;

² Most worthwhile operating systems today offer methods to put users into groups and provide file access controls (commonly called Access Control Lists, or ACLs) that allow only certain groups access to those files. Later in the paper, more discussion is given to this subject. It is considered a fundamental requirement that the internal network be capable of segregating users from data and allowing information to be accessed only by certain groups who have a valid need for that information.

1. We identify each remote *individual* and provide them an account (with Access Control Lists (ACLs)), or
2. We identify the *source IP addresses* and protocols and allow data to pass to and from this location only (again, with ACLs), or
3. We allow virtually anything to pass (not recommended).

Option 1 is the most recommended, for several reasons. By identifying an individual, you can then place them into the particular ACL groups, you can control and audit an individual's whereabouts on the network, and, if necessary, you can attribute a particular action to an individual (sometimes handy in the courtroom!) Of course, this depends on the users not sharing userIDs and passwords with each other; else, there is no point. Option 2 might be preferable in a few cases where it is known that all users at the remote location have a common need-to-know for information, and you are confident that physical access to their computer systems ensures that only they can get on the network. Because this does not allow for tracing actions to individuals, it may come back to bite you if disciplinary action is ever necessary. A combination of Options 1 and 2 (along with other conditions discussed later) provides much greater control. Option 3 is equivalent to burying your head in the sand--anyone can get anything from anywhere.

2.2.1. A Question of Identity

The rest of this paper will assume that individuals must be identified. The term "our" network indicates the internal corporate network that we are trying to protect. For the remote user, we have a choice of identifying them ourselves, or accepting the identification scheme of the network at the visiting site. I caution against the latter for this reason: if someone can illegally become a member of the other site's network, that person can then access our corporate network's data unchallenged. Many (if not most) remote visitor networks have connectivity to other sites around the country (their "friends"), which have no need for access to our corporate information. Option 2 (identifying the source IP address) can help to thwart this, but not by itself. Ideally, at every entry point to our network, we should identify and authenticate (userID and password) each individual requesting access.

2.2.2. Individuals or Groups ACLs

Once we've identified individuals upon entry to the network, we should not let them roam, but instead need to introduce them to the individual server(s) to which they need access. Rather than forcing hundreds of host systems to each add hundreds of users, we should create a handful of approved groups, such as those mentioned earlier. At the entry point to the network, the user is identified with a particular group. Individual host administrators can set their systems to allow access to information by a particular group rather than by named individuals. When a new employee or an approved outsider is granted an account, their names are placed into the appropriate groups, thus giving them access to every system which that group has access to. The individual host systems need not change their access

control list at all whenever an individual user's access changes. At a single corporate top-level account-granting server (perhaps tied to the Human Resources department), the individual's new account is placed into a group (or groups) once, and equally as important, removed from the group(s) once. (How many times has an employee left the company, yet still has a valid account because nobody removed it from all the servers?) If data owners have no restrictions as to who can access the data, they need do nothing, and the default will be that all authenticated users can see it. If the data owner has data that only a certain group of users can see, then it is up to that owner to set permissions to that data for only the group(s) that is(are) authorized to see it. Only certain information, commonly needed by all users would then be open to "insider" viewing. Information, which must have access control, can be controlled by using the ACL functions of the operating system access controls. At many corporations, even diverse operating systems (OS) such as Windows NT and UNIX, can have their system access methods synchronized so that one OS recognizes the other's groups, effectively providing one set of global groupings at the corporation. To ensure such access controls, a corporation must have a standard requiring operating systems that provide for ACLs and group management. If an operating system (OS) does not offer such access controls, they must either utilize an add-on product that will offer them, or work only as a standalone system. NT, UNIX, and any other allowed OS/add-on must be able to recognize the group management structure present in the corporation's domain.

2.2.3. Centralized authentication

A corporate-wide authentication scheme, such as the use of a Kerberos³ server, is highly recommended, and each host computer with data to protect should use the corporate authentication scheme. This reduces the need for individual system administrators to maintain access control lists for every application that may be stored on their servers. Companies should front-end individual host computers and applications with some type of identification and authorization scheme. This, in conjunction with file and directory ACLs, is known as protection of data at the source. We've already pointed out that not all inside users have a need to see absolutely everything on the network, so by authenticating at the source, we can fix the problem for both insiders and outsiders-who've-become-insiders. Kerberos offers an effective method of verifying users, and needs only a small bit of code added to an application to allow Kerberos to screen users before granting access to an application. (Please note: Kerberos is useful on your internal network, but is not suggested as the authentication method for remote users, as it uses cleartext passwords. For remote users, it is better to use either completely encrypted transmissions, or one-time passwords, such as SecurID.⁴)

2.3. Create a secure infrastructure

Within the internal corporate network itself, there is a danger of allowing information to be seen by malicious insiders, via "sniffing" the network. This can be done simply by any computer that can set its network interface card (NIC) to

“promiscuous mode. This allows it to look at all data that passes through a LAN. By design, Ethernet has a characteristic of broadcasting information to every machine on the LAN segment. A router can limit that broadcast to only those machines on the segment, but by utilizing switches, you can stop the broadcast nature of Ethernet and send data to only the machine it is destined. This reduces the desirability of sniffing as an attack, as the hacker would have to put a sniffer machine on every wire.

Depending upon the sensitivity of the data, another method to foil the insider is to encrypt every bit of traffic that goes across the network. IPSEC, designed for Ipv6, has been implemented in Ipv4 and offers encryption and assurance that the traffic has not been modified. This is usually performed by the operating system, and tends to slow down the throughput of the system.

A faster method of implementing IPSEC is via a NIC card that performs the encryption on the card itself, relieving the computer’s CPU of the intensive processing job. Encrypting NICs can be bought today for about \$20.00 additional per card.

Along the same lines, encrypting of stored data on disk can protect the information even after the disk is removed from the computer. This is performed either by the OS, or by a third party application.

And, of course, separate the inside network from the outside by either an air gap (no connectivity to the outside) or a system of firewalls, which leads us to:

3. Controlling the flow

I say a “system of firewalls”, rather than “firewall” because, depending upon the size of your company, a single machine becomes a chokepoint for all traffic entering or leaving the network. It also becomes a single point of failure should it break down. Better to have several machines that can leverage the load between them. Perhaps, one set can manage web access, others may manage FTP or other protocols, another for email, etc. A Network Address Translator (NAT) can hide internal IP addresses from the outside. These are all forms of “proxies”, and together, can provide an effective firewall system.

Controlling also includes monitoring. Once you’ve set up a good infrastructure, how do you know if it’s working the way you intended it to? Regular auditing of activities is a necessary evil. Yes, it’s time consuming, but without it, you don’t know if you’ve been hacked, or if someone is doing something they’re not supposed to, or that a system is not responding properly to users’ requests. A good operating system will offer audit trails, some better than others. Use them. Review them. Add third-party intrusion detection tools—they can give you a vast amount of information about vulnerabilities in your system and offer methods to fix them. And, when someone does do something against the rules, have a pre-determined plan of how to deal with it.

3.1. DMZs

Often, an accepted practice is to create an area where outsiders can obtain and pass information to and from your company without actually letting them into

your internal network. This is called a DeMilitarized Zone, or DMZ. This is the “front yard”, where you can serve information on demand. It is a small network, external to and protected from your company network by various firewalls. Selected information can be sent out to the DMZ for privileged outsiders to pick up. A DMZ should be configured tightly so that no users (even employees on travel) can go from the outside to the internal network without following a controlled path you create for them. Because the machines inside the DMZ are relatively exposed to the world, you must take the stance that these are expendable, and you may have to rebuild them periodically as people attack them. Any information contained on these machines should also be expendable, in that, there should be no information available on them that you can’t afford to lose or be made public. There may be several different “need-to-know” groups with information being served in the DMZ. UserIDs and passwords can be used there as well, and encryption can ensure that a captured machine will not yield up its information to the hacker. These should be lean machines, with only the minimum of services available. They should be able to be rebuilt quickly. Most of the operating systems and services should be stored on read-only media to preclude an outsider from altering or exploiting the systems themselves.

3.2. Extranets

A cousin to DMZs is known as the “Extranet”. An Extranet is still be part of the internal network, but only a small subset of servers to which outsiders can gain access. It does not have the same connectivity to the internal network as a normal subnet, but data can be passed via a controlled “pipe”. Employees on the internal network can initiate communications to the extranet, but the reverse can be forbidden. This could prevent outsiders from becoming true insiders, and can also prevent outsiders from using the extranet servers as jumping off points to other sites. Each extranet server would still have the responsibility to identify and segregate users to only the information necessary.

3.3. Reverse Proxy

Most proxy machines are designed to allow the inside user restricted access to outside resources. The Reverse Proxy does just the opposite--it allows for the cordoned-off walkway through the house. Visitors can identify themselves to the proxy, and the machine provides them a pathway to get to only the servers they need. In fact, the data is gotten from the inside server by the proxy machine, and the proxy machine presents the data to the user. This is ideal for web-type access, but does not work well with other protocols, such as FTP or Telnet connections. The key is to limit access to only the internal systems (or even pages) necessary, and allow it from only known IP addresses of external systems we approve of.

3.4. VPNs

The concept of Virtual Private Networks is growing. This idea uses the Internet, or dial-up phone lines (or any connectivity method, for that matter) as a vehicle to “tunnel” an encrypted channel from the outside client’s desktop/network to the destination host network. This should be considered as a replacement for costly dedicated lines (i.e., T-1s) and inherently dangerous dial-ups. VPNs can

offer tremendous cost savings compared to several leased lines and hardware encryption pairs. The implementations vary. Most simply offer a connection to the network as a whole, or to an individual server/subnet, but not both. If a VPN solution is considered, a company must consider one with the capability of identifying and authenticating individual users (or the ability to route through an authentication server, such as Radius, TACACS+, Kerberos, SecurID, etc.) and the information must be encrypted while in transit over the public network (the Internet). In addition, the keys must be controllable from the central site. In the case of a traveling laptop, a stolen machine may still be considered an “authorized” connection. The VPN solution must therefore be able to disable lost or stolen keys immediately, and to be able to remotely supply a new key to a user on travel without the need to mail or fly it out to them.

4. Connectivity

This section discusses various methods of connecting one site to another. Although many companies consider these to be “safe” conduits, they are not, and should not be considered as secure infrastructure. By themselves, they provide neither access control, nor confidentiality—they simply bridge a remote computer or network to the internal network. Once users enter the network via a T-1 or any other physical line, they are instantly internal users with access to everything in the internal network. Most of these suffer from being able to be tapped at various switch points, so encryption of the data transiting these links should be considered. Various authentication schemes also need to be considered, as you don’t want cleartext passwords to be sent over these links. One-time passwords are valuable in this case, and can be used as the initial foot-in-the-door to create a connection to the network. After that connection is made, encryption can be applied to the link and then the user can make use of the standard internal network authentication methods.

4.1. Dedicated Lines

Dedicated lines (T-1 [1.5Mbps], T-3 [4.3Mbps], OC-3 [155Mbps] up to OC-192 [10Gbps], Digital Subscriber Lines [DSL, 51Mbps], Frame Relay [similar to T-1 speeds] or FDDI [100 Mbps]) are simply an agreement with the phone (or cable) company to provide a certain level of performance to your corporation. It does not mean that the traffic follows only one path, it only means that you are guaranteed a particular speed and that they are responsible for providing alternate circuits to you at that speed should the one go down. This may be convenient, but they are costly, and provide no real security. Who knows what security requirements must be met at the far end of that link, if any? Moreover, although you’ve paid the phone company for a link between your company and a remote site, that link travels through the phone company’s network, and can be intercepted at any point by any individual with physical access to the phone company’s switching infrastructure. Sometimes, these links travel via microwave, so physical access is not necessarily a requirement.

4.2. Dial-up modems

Analog modems (300bps up to 56Kbps) or ISDN modems (28.8-128 Kbps) connect two computers. Though very inexpensive, allowing users to dial into their

desktop from the outside bypasses all the access controls mentioned previously in this paper. Modems should not be on individual's desktop machines! (Even dialing out bypasses the proxies!) A corporately controlled bank of dial-in modems can help, but they must be routed through some sort of authentication server (see the VPN process above) before placing the dial-in user onto the internal network. In general, this remote access method should be discouraged. VPNs provide a better method of controlling access from remote computers.

4.3. Satellite Modems

These provide higher speeds (~400Kbps) than analog or ISDN modems. The downside is that travelers cannot yet afford to have a satellite modem on their laptop. There is an inherent lag time between transmission by a satellite, and reception by a ground antenna, and until recently, only the downlink of data was available. Any uplink data still went through standard telephone lines, effectively taking 400Kbps down to little over 56Kbps. (Remember, any computer-to-computer communications uses two-way instructions.) Newer systems now offer two-way satellite communications, with the uplinks being transmitted at 128Kbps². Please note: the use of satellite modem uplinks is effectively broadcasting your data to the world!

4.4. Cable Modems

These can provide upload speeds of ~128Kbps and download speeds of slightly in excess of 2.5Mbps, however, an increase of subscribers on the network segment results in lower speeds. In practice, speeds of 1Mbps are considered optimum. Speed is also limited by the capabilities of the network path between the user and the server providing the information. And, just like dedicated lines, data is vulnerable to interception at any of the cable switching facilities. Even worse, in the privacy of someone's back yard, the cable can be tapped.

5. In Summary

A company can share information with those who have a valid need. The point of this paper is to first and foremost, CLEAN HOUSE! Take care of your internal network first to make sure you can control who gains access to what. Unless you are able to separate users from data and equipment on your internal network, all the guests you invite into your house are going to see all your dirt! Make certain you've identified every entry/exit point to your network, and apply a corporately-defined (not user-defined) access control method to every one of them. Then, regardless of the type of physical links you use, be certain that you can control the access at not only both ends of the link, but ensure the data is secure in transit.

References

¹ Lenzner, Terry. "World-class private eyes sharpen the focus on cyberspace: A CSI Interview with Investigative Group International (IGI)". Reprinted from the November 1998 issue (#188) of Computer Security Institute's monthly newsletter, Computer Security Alert. URL: <http://gocsi.com/world.htm>

² Grossman, Mark. "Protecting your Company from Hackers". 19 July 1999. URL: www.mgrossmanlaw.com/articles/1999/protecting_your_company_from_hac.htm

³ Garfinkel, Simson & Spaffor, Gene. "Practical UNIX and Internet Security". 2nd Edition. Sebastopol CA: O'Reilly & Associates, Inc. , 1996, 594-604.

⁴ Garfinkel, Simson & Spaffor, Gene. "Practical UNIX and Internet Security". 2nd Edition. Sebastopol CA: O'Reilly & Associates, Inc. , 1996, 252-254

⁵ DirecPC. "How does the Satellite Return system work?". 15 January 2001. URL:
<http://www.direcpc.com/consumer/owners/faqs/faqs.html#howSatwork>

© SANS Institute 2000 - 2005, Author retains full rights.