



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 (Security 401)"
at <http://www.giac.org/registration/gsec>

Connecting a classified network to the Internet. A case study.

Introduction

A classified (computer) network is one that stores information that is sensitive and should not be made available to the general public. This may be;

- Military secrets
- Financial records
- Personnel records
- Trade secrets

Generally only the first type of information, military secrets, will make classification mandatory, but lower classifications (NATO RESTRICTED or equivalent) may be imposed by civil law.

The purpose of this document is to point out some common elements from the guidelines published to regulate computer security and suggest administrative action and technical solutions to build a network that may be connected to the Internet, and still obtain/retain a classification up to and including NATO RESTRICTED.

The author is not aware of any standard that will allow a system classified at NATO CONFIDENTIAL or higher to be connected to a public network (Internet). Note that the standards mentioned (BS7799, ITSEC, CommonCriteria a.s.o.) only tell what to achieve, not how.

In Europe certification schemes are coordinated by the European co-operation for Accreditation and will be quite similar. The US have adopted the Common Criteria initiative (www.commoncriteria.org), but not to ISO17799, so there might be some differences here to be aware of. Nevertheless it should be noted that there is differences from country to country when it comes to accreditation and the reader is advised to contact the accrediting agency, or a private IS security company for guidance before committing to a any strategy.

For the rest of this document the internal, classified network will be referred to as the R-net. (R=Restricted)

What's changed?

For years the standard way to protect R-networks from intrusion was to simply not connect them to any external nets. The rationale here was that the hacker can't compromise what he or she can't physically connect to. This holds true today; network devices such as hubs, routers and switches merely extend the physical cable. Interrupt the stream of electrons (or radiowaves /optic signals, and data traffic is stopped instantly!

With the widespread use of the Internet, the situation has changed. Even organizations that have information that must remain protected, find that they need an Internet-connection in their day to day tasks, or that their employees demand it. A lot of them

have simply connected their R-net to the Internet just separated by a network device called a "firewall", and pray that they will not be compromised.

In most countries the fact that organizations, both governmental and private, store information that may be highly sensitive on networks that might be accessible from the public Internet is unnerving. Governments in the western world including the United States and most of the European Union have now passed laws that require all classified networks to be protected. In Europe usually a government agency is required to check and approve a network before it may store classified information. These agencies also publish guidelines to how a network can conform to the standards of a given classification. In the US the National Security Agency Infosec Service Center (<http://www.nsa.gov/isso/index.html>) publish guides on network security, but it is unknown to the author if they inspect and certify computer networks.

Where to start.

The procedures in this document might be helpful in these situations:

- The network will be connected to the Internet and is storing classified data.
- The network is presently connected to the Internet, and you're going to store classified data.
- You are presently both connected to the Internet and storing classified data and praying your firewall will protect you (bad idea...).

Although this document assume that the reason for protecting the internal network is to obtain/retain an official classification, more and more organizations realize that they store sensitive information that they are morally and maybe legally required to protect.

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

The first thing to get in place is backing from management. To do this, an official **information security policy (ISP)** identifying what assets need protecting and why, should be written, and signed by the CEO. This is also the first document an accrediting body will look at. It is not easy to find templates for security policies, but the research project "What do I put in a security policy" by *William Farnsworth* (<http://www.sans.org/infosecFAQ/policy/policy.htm>) on the GSEC pages of SANS is a good starting point. You might find that the provided sample Security Policy has all you need.

Now you are ready to plan how you will meet the demands set by the policy (or the approving body) on your network. A good roadmap to this process is the British Standard 7799 (BS7799 or ISO/IEC 17799). It comes in two parts;




Figure 1

1. The standard code of practice, issued with permission of Gamma Secure Systems Ltd.
 2. Information Security Management System (ISMS) standard specification.
 It is part 2, the management standard, which is interesting at this point. The process is described in the graphic Figure 1 on the previous page.

- Step one, the Information Security Policy, is not discussed in detail here.
- To define the scope of the ISMS you should describe the following:
 - Conceptual aspects of security.
 - a) The extent and purpose of the network.
 - b) What classification you want to obtain.
 - c) How traffic between the classified and unclassified network should be handled.
 - d) If commercial hardware and software (COTS) or tailor made solutions should be used .
 - e) A drawing of the network as you plan to implement it.
- Risk analysis and Risk management.
 Since the assets and their value are defined in the ISP you now need to evaluate the

What is risk? Here is a definition:
Risk is the combination of a threat exploiting some vulnerability that could cause harm to some asset.

risk of loss to each of these values. Arrange the results in a matrix. The highest valued asset with the most risk of loss will be your first priority. You then proceed all the way to the opposite end of the matrix (least valued asset with little risk of loss). This is now your prioritized list of objects to protect.

With your risk analysis in mind you need to describe how you want to safeguard your assets (risk management). You may choose one or a combination of the following:

1. Reduce the threat.
2. Reduce the vulnerability.
3. Reduce the value of the asset.

This may be difficult to understand so let's do an example;

You decide that your customer database is a valuable asset, so you regulate access to it by setting a password. You have now reduced the threat of someone viewing, altering or deleting data. You may then install a system patch to stop unauthorized users accessing the database through some exploit. You have reduced the vulnerability of the database system. Finally you encrypt the tables in the database, so to make it unreadable to an intruder. You have reduced the value of the asset (at least for the thief...).

- **Control Objectives and Controls**

Here you describe, in detail, how you plan to implement the ISMS.

- a) A list of all the documents that is needed to regulate computer security (contingency plan, access policy, and computer security policy.)
- b) The hardware (secure placement, cabling, marking and electronic radiation protection).
- c) Handling of data storage devices (floppy's, tapes, harddisks and so on).
- d) Handling of print devices
- e) Data Communication (e.g. encryption; yes/no, what kind)
- f) System configuration control and maintenance.

Note that this part of the plan is what will differ the most based on regional laws.

Often the requirement specification is a direct copy of some official document. Save yourself a lot of work and check this first!

It is beyond the scope of this document to go into detail on the technical solutions.

For some pointers, please see the sample at the end of the project.

- Finally you should produce a document called the Statement of Applicability. The document identifies all your controls, justifies how they will mediate the risk, and how they map to the BS7799 or whatever standard you are using. If you deviate from the standard you must be able to explain why your implementation is better.

The classification process.

OK. With all your paperwork in order you should contact whatever body you want to classify your network (either government agency or high level management). If you don't seek official classification you should also present a cost analysis at this point. Management might care about this aspect, a government agency will not!

Note that few or none of the countermeasures in your plan are in place, but not to

worry. When you first present your plan, an official body more often than not have some additions or changes you need to include. Finally, when you get acceptance of your policy and plan then its time to start thinking about implementation.

Doing the work

The first part of implementation is a lot of writing. You will need to write procedures for just about every aspect of your IT operation; backup, restore, user management, file and print access, logging, mail and web access, and so on. You also need a full detailed inventory of your computer network configuration. The procedures are referenced in the applicable policy.

When this is done you should turn your attention to your physical network. You will need to evaluate if you are going for a strong perimeter defense, hardening the inside or both. Of course a strong perimeter defense AND a hard inside is ideal, but building this to perfection might be a daunting task...

Further you need to decide how you should handle traffic going to and from the public network to your protected network (or how to avoid it). Usually your greatest concern is to stop data from your protected network being sent to the public network thus compromising the confidentiality of your data, but having unwanted data from the public network getting in is not good either, basically threatening both integrity and availability.

You need to do some degree of auditing. At a minimum you must log all file transfer from your classified network to your unclassified network. It will not stop sensitive data from being compromised, but you will have accountability! You should also configure your e-mail system so that it will search for some kind of verification that an outbound message is unclassified (like having the sender write "unclassified" in the subject field...) before sending it out. Again accountability, not prevention is the objective here.

You will probably need to decide on and buy software and hardware in this process. If possible, you should consult the list of software/hardware tested by the Common Criteria or ITSEC (<http://www.itsec.gov.uk/>). If you find what you need in the list chances for a smooth accrediting process is much better. And, you get software and hardware you can trust!

Some items, like backup software, are necessary in all networks. Others are specialized for securing data communication. At a minimum there must be some mechanism to control traffic access between the R-net and the Internet. For most this will be at the IP-address and protocol level, and implemented by a packet-filtering router (a "firewall"). Be aware that the firewall(s) are your most important network object and will be vigorously inspected and tested. Make sure it is evaluated to a level of E3, and configured to the standard of the classification body.

The firewall can be implemented on each host instead, but administrative overhead has made this a less popular option.

The router may be complemented with an application layer gateway, or proxy, which has access control on a higher level (application layer). But beware the no firewall (packet-filtering, statefull, proxy or other) is secure out of the box. For example; some well-known firewalls are default configured to allow all connections initialized from the inside. This is not desirable since if a Trojan is introduced in the R-net it may be able to open an outbound connection to its “master”.

Many implementations, like the one in the example, require an extra network segment between the R-net and the Internet. Let’s call it the controlled network (C-net). Its function is to be a controlled buffer zone that can prevent or delay an attack on the R-net, and to host services not able to run in the most secured environment. Such an implementation requires duplication of several systems, like the mail-server. Of course you will also need a second firewall.

Since logging is part of the requirements you might need some centralized logging functionality. You don’t want to run around checking each and every log on every network object in the middle of a hacker attack!!! There are COTS-products that can send both UNIX and NT events to a syslog server. Some kind of database is needed to systemize and store entries from the syslog.

Although not usually mentioned as a formal requirement the author strongly believe that the following should be found in all secure networks:

- Some form of user verification/authentication other than that of the host OS (Tacacs, Radius, and RSA...)
- An Intrusion Detection System (ISS, Shadow, Cisco Secure...)

Remember that your certification depends on the level of trust the classification body has in your network configuration and countermeasures. Going that extra mile can only give an impression that you are truly serious about security!

An example network

The network diagram below designed to give mail and web browsing capabilities, display a physical network layout that should pass most basic inspections. Just remember that you also need documentation and routines in place.


How it works:

NOTE: In the graphic the term modem is used for all types of interface (Frame Relay, dialup and so on.) to a WAN network

- Email:

When e-mail reach the external firewall it is forwarded to the anti-virus scanner. If it is clean the AV-scanner pass the message to the mail-server on the C-net. The mail-server in turn pushes a copy of the message through the internal firewall to the R-net mail-server. There is no session because the receiving mail-server is denied outbound connections.

IF the message need a reply, the recipient must to open a channel from his client to the external mail-server. In the example this is done with a Terminal Server. Through the TS the recipient can respond on the copy still available on the C-net mail-server. The TS-client is configured not to allow cut and paste from the internal network to the TS-session, thus preventing a user to add any attachments containing data from the R-net.



- Internet browsing is also performed through the Terminal server. There can be no transfer of internet-data into the R-net.
- The (not required) intrusion detection system here is set up with two agents, one on the C-net to detect intrusion from the internet that has passed the external firewall, and one on the R-net to spot any activity that may indicate that an insider is trying to defeat your countermeasures in order to compromise security.
- Any services offered to the Internet should be placed in a separate segment, DMZ, as is good practice in any network.
- The syslog-server is configured to accept messages from both C-net and R-net. For additional security, have a separate syslog-server on the R-net so that you can minimize traffic through the firewall from the R-net to the C-net. To get a central logging point, copy the syslog data (on a zip-disk, tape or other removable media) from C-net up to R-net.
- If data transfer from the R-net to the C-net is absolutely necessary, it would be possible to open up the internal firewall for FTP to a specific FTP-server on the C-net and give the user permission to *put* a file in a write only directory. The security officer should then manually screen the file before it is put in a folder where the

user can retrieve it from his Terminal Server session. This workaround may NOT be acceptable in all countries.

Virtual Private Networks

Although not strictly belonging in this discussion, I want to mention the issue of VPN in the context of classified networks. A VPN is an encrypted channel that uses a public network (usually the Internet) to transport non-public data between two LAN's. Since the sender and receiver don't exercise traffic control over the public network there is no guarantee privacy. That means others might pick up the transmissions and read the data. This can't be avoided, only mitigated by encrypting each datapacket to make its contents unreadable and ensure that the information is not altered.

There are several COTS solutions for VPN, but none is evaluated to a level acceptable for general classification. If you really need a VPN connection into your network, there are few references on what to choose and how to configure.

First and foremost: The chance of getting acceptance for a configuration that employ a VPN connection into the R-net is at best slim. The most viable alternative is to use a C-net/R-net configuration and terminate the VPN-connection in the external firewall, or in parallel with it. There are other actions you can take to better your chances of certification:

- Buy the VPN with the best encryption algorithm and bit length you can get your hands on.
- If there is a national manufacturer of a VPN product, or if a specific VPN has been chosen by a military or important governmental organization in your country, try to get that.
- Hardware crypto is usually more secure and faster than software crypto.
- Avoid VPN servers that depend on a well-known OS like Windows or UNIX if you can.
- If the number of VPN connections is small, go for a symmetric or private cryptosystem. Practice "perfect forward secrecy", which means that you change the keys frequently, giving an intruder a very narrow timeframe to access the system.

Summary

It's quite easy to obtain a given classification for a single computer, or for host on a LAN segment. That's because access to the data is easy to restrict and control. When a classified network is to be connected to the Internet, this triggers a whole new classification process, because new ways into the network exist.

As the risk of exposure, loss of data or loss of data integrity increases, Systems Admins around the world see the need of better protection for their networks. Connecting classified networks to the Internet bring new, often conflicting, aspects into the situation. By following some commonly accepted principles, the likelihood of succeeding in securing the network increase.

The BS7799 standard contain, among other things, a roadmap on how to plan, build and certify a internet-connected network that will contain classified data. The Common Criteria standard describe how to test hardware and software designs to a given security level, and national independent evaluation facilities carry out certification test on specific products.

If the ultimate goal is to get acceptance for a gateway onto the Internet, just following the common standards might not be satisfactory to your countries approving body. Additions and alterations to the generally accepted standard are quite common. Remember that a classification is not a guarantee that your network is safe from intrusion or tampering, it's just an expression of the level of trust the certifying body has in your countermeasures being able to mediate the identified risks.

Building a secure network is not just a matter of buying and configuring a firewall. The organization needs policies to regulate the access and use of the network, and to control the flow of data. The process of obtaining classification is just as much an organizational process as a technical one.

Bibliography

Commonwealth of Australia – "Gateway certification guide"

(URL) <http://www.dsd.gov.au/infosec/Gateway/>

Gamma Secure Systems – "How It Works, Part 2: The management standard"

(URL) <http://www.gammasl.co.uk/bs7799/works.html>

Checkpoint Software Tech. – "Virtual Private Network Security Components"

(URL) <http://cgi.us.checkpoint.com/rl/resourcelib.asp>

Forsvarets Overkommando/Sikkerhetsstaben – "Datasikkerhetsdirektivet"

(URL) <http://www.fo.mil.no/sikkerhetsstab/dsd/innhold.html>

Elizabeth D Zwicky, Simon Cooper – "Building Internet Firewalls, 2.ed"

O'Reilly & Associates 2000 (URL) www.oreilly.com