



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.

SECURING A FORTUNE

Network Security in a Small Business Environment

A GCNT GIAC Practical

Author: Marc A. Mesmer
November - December 2001

Contents

- [Basic Assumptions](#)
- [Network Design](#)
- [The Network](#)
- [Hardware Considerations](#)
- [Software Considerations](#)
- [Steps For Securing](#)
- [Applying Security Settings](#)
- [Locking Down The DMZ](#)
- [Database Security](#)
- [Securing The Keys](#)
- [IPSec And The Network](#)
- [The Unimplemented Enhancement](#)
- [Physical Security](#)
- [Continued Threats](#)
- [Conclusion](#)
- [References](#)
- [Addendum](#)

Basic Assumptions

GIAC Enterprises is an e-business that deal in the online sale of fortune cookie sayings. This document outlines the requirements of the internal network used by the organizations employees. It includes the following departments:

- Research and Development (R & D)
 - Sales and Marketing (S & M)
 - Finance and Human Resources (F & HR)
- additionally, this enterprise will include:
- Business Administration (BA)
 - Information Technology / Mission Information Systems (IT / MIS)

This enterprise, in this instance, is assumed to be a single-office company. Each department listed above currently consists of about ten members each making in all about a fifty people, but with an eye toward expansion in the future. Expansion will occur in phases over the next few years both locally and into regional offices. The new office building for which the current design is to serve, will be connected by a high-speed LAN and presents to the world a domain name of giac.com, a web server at www.giac.com (any

resemblance to any business real or fictitious, is purely coincidental) and a registered Class B IP network address.

This document presupposes a familiarity with the Windows 2000 operating environment and Active Directory. Only cursory attention will be given to concepts such as the promotion of a server to a domain controller (DC) or Replication of the Global Catalog. This document will focus on securing systems and will only touch on the basics of Active Directory where it has a direct security implication.

The design of this environment is in ongoing, iterative process. The document that you are now reading bears little resemblance to the designs which were first considered, and it is assumed that the requirements for the environment, as well as the perceptions of that environment by those participating will change over time. The goals of this design are to provide a flexible structure into which growth and changes can be facilitated while maintaining the requirements of a secure environment. Of course, not all changes will be possible without compromising the goals of security, and not all abilities of the flexible design, just because they are possible, should be put into practice.

Certain aspects of the design cannot be specified except once the design is put into practice. Delegation over Printers and Shares and other administrative tasks shall be determined more in a few days of practice than months of theorizing could possibly provide. The basic tenant to be followed shall be to delegate as little as possible, and then only when there is a compelling reason to do so. In such a case, an existing group or a group created for the purpose shall be granted the right under consideration for the most limited scope possible. In this way it is hoped that that practice will discover what theory cannot, but principle will prevent indiscriminate exercise of arbitrary authority.

The design and implementation of a secure computing environment cannot occur in a vacuum. It will require the efforts of several individuals dedicated to the tasks of installing and then maintaining and monitoring the systems. It is primarily the IT or MIS department which shall implement this plan, and as their duties are integrally related to its implementation, they shall be here enumerated:

1. Director of IT services Responsible for establishing service level agreements between other departments, scheduling and allocation of resources, overseeing of budgetary issues, etc. This position shall serve as the coordinator of the others on the team.
2. Chief Security Officer: Monitoring of intrusion detection systems (IDS), maintaining of current virus definitions, scanning and archiving of logs, etc. Monitors security lists for up to date security information and evaluates the latest security patches, software releases and hotfixes.
3. Mail / Web / Database Server Administrator: Maintaining servers established for these purposes. Interacts primarily with the designated users from other departments who co-ordinate the use of these services for their departments.
4. Active Directory Administrator: Maintains User and Computer base as represented in active directory and monitors replication and other AD functions throughout the enterprise.
5. Hardware/Software Tech: Installation of hardware and software.
6. Backup Technician: Performs regularly scheduled maintenance and backups as well as restores files from backups as needed. Services log file archiving.

These functions are intended to be representational, rather than absolute. The duties of two may be consolidated into a single individual at times, or the workload may expand to require two to fulfill the duties of a single position.


Network Design:


The network design chosen for this project is one which is closely modeled on the "Contained Domain

Model Diagram" topology suggested on page 3 and in Appendix A in the NSA guide "Microsoft Windows 2000 Network Architecture Guide". This is a network design for a single domain which does not extend across any networks not controlled by the organization. [1] For a small organization, such a design, with specific modifications to suit the specific demands at hand, will admirably fit the bill.

Much of the work that goes into securing this basic network, however, can be applied as the basis for larger, more complex projects, and if properly planned for, can provide a sound basis for the expansion of this model into a larger enterprise should the growth of the company warrant.

As the network can in this case be began "from scratch", with any existing computers upgraded before being brought on-line, the goal is a Windows 2000 (Win2K hereafter) Network operating in Native mode, thereby enjoying maximum benefit of the security features of Win2K and avoiding most pitfalls of the "backwards compatibility" and "mixed mode" of operation. All Servers and Domain Controllers shall be new or reformatted machines and so be essentially "new". Although Native mode is the goal, the switch to native mode is "one-way" and therefore the last step in the implementation. One should first determine that all systems and applications are fully Windows 2000 compatible before taking this step. [2] This is believed to be a fully realizable goal.





What is not in this design?

1. Dial in network access. Providing this service is a convenience to remote users, but possibly one of the largest security risks to a network. It should be avoided unless there is a compelling reason to include it. A dial-in access point provides a back-door into an otherwise secure network. Also, access provided to an authorized user is providing access to a computer of uncertain safety. A computer with more than one network interface may provide a bridge between otherwise secure networks whose administrators might not wish them to be so connected. One may also be faced with the loss of a certified laptop with remote access privilege to a secure network, entirely compromising security... It may be that some proprietary information needs to be provided to the traveling sales force. This could be placed in an authenticated section of the public web server. Then, in case of a compromise, only the published data is compromised, rather than the entire network with all of its data.
2. External web access (even via proxy server). Should this service be necessary, it may be prudent to provide a separate connection, for instance, via a dial-up ISP on a dedicated computer with an attached printer and *no floppy drives* to allow for internet research. This should help in a threefold manner:
 1. Time spent on the internet will be kept to a minimum and directed to specific purposes, and
 2. such internet access will not interact with the intranet in any way and
 3. resources spent maintaining and monitoring a proxy server can be better directed to other pursuits.

This suggestion, draconian as it may appear, is presented in the interest of network security.

3. Network Address Translation (NAT): Often this is used to mask the internal network from the external. However, this design shall be making extensive use of IPSec Authentication Headers (AH) which do not work with NAT. NAT will change the very header information which is "checksummed" for verification by IPSec, hence the two protocols will not interoperate.

This should not be much of a difficulty, as there are only three machines in a DMZ exposed to the internet, and traffic between the rest of the intranet and the internet is blocked, so there shall be no need of NAT in this design.

4. Dynamic Host Configuration Protocol (DHCP): A part of the original design for reasons of convenience in network administration, it became more and more evident that there were reasons for its omission. First, DHCP should not be used with servers and services that require statically configured IP addresses. Secondly, DHCP servers should not be run on Domain Controllers, as this can leak sensitive DNS information. Finally, the design began to link specific IP subnets to physical network segments to facilitate application of IPSec filter rules applied by Organizational Units whose computers are assigned addresses in a given subnet and are connected on a given segment. Then, it became desirable to assign a static IP to each machine to aid in its identification and location. DHCP will not cope efficiently with this type of rule, as it defeats the point of DHCP. Also, DHCP will more willingly grant a connecting computer an IP from an available IP pool and it is more difficult to tell

which IPs are used by legitimate computers. At the cost of more administrative overhead, one can tell with manually configured IP addresses (without DHCP) whether an IP in use is authorized, and to what workstation. One cannot tell if that is indeed the machine using that IP, but conflicts will more certainly arise when two machines attempt the use of the same IP. Scanning logs for out-of-bounds (unassigned) IPs can then detect the presence of unauthorized machines on the network.

The Network:

The Network consists of the following sections:

The internet: Over this there is no control and all connections should be considered suspect/potentially hostile.

External Router: Cisco 2600; The First line of defense: Here initial packet filtering can occur. All traffic directed toward the intranet should be blocked. Attempts to gain access to the intranet should be logged. DMZ traffic should only be the particular allowed ports desired for DNS, WEB and MAIL on the particular hosts.

PIX Firewall: First real line of defense. Requests for the publicly available services are routed to the DMZ. Performs "Sanity checking" to prevent spoofing attacks. prevents unauthorized access from the intranet to the internet or DMZ and visa-versa.

Basic Firewall Access Rules			
From	To	Allow	Deny
Internet	DMZ Webserver	http(s) 80, 443	All Others
DMZ Webserver	Internet	http(s) ACK Responses	All Others
Internet	DMZ Mail Server	Mail 25	All Others
DMZ Mail Server	Internet	Mail 25	All Others
Internet	DMZ DNS	DNS Requests	All Others
DMZ DNS	Internet	DNS Responses	All Others
Intranet	DMZ DNS	DNS Queries from DNS Server	All Others
DMZ Web	Intranet	Web Server Mirroring from Web Server	All Others
DMZ Mail Server	Intranet	Mail Exchanging with Exchange Server	All Others
Internet	Intranet	None	All

Network Switches: Cisco Managed Switches: Entire network based on Gigabit Ethernet

DMZ: (De-Militarized Zone): This leg of the network contains those servers visible to the internet at large. These are the only servers which are publicly advertised in public DNS records.

- **External DNS Server:** This should contain only records of the public servers, the DNS Server itself, the public Web server, and the mail forwarding server. The internal DNS server will be unable to communicate with the external net directly, but will be configured to resolve queries through this server.

- Public Web server (www.giac.com): Most internal users will publish to an intranet mirror site, which will be periodically copied here by the Web Server Administrator. Access to this and other DMZ servers should be kept strictly to a minimum.
- Mail Forwarder: Public mail point, A store and forward mail server, which is the only known point of contact for public email, the point of sending public email, and also sends email notifications as required by web applications.

The Intranet:

BackBone: Gigabit Ethernet The main network LAN, to which all other segments connect. Certain servers will connect direct to this directly:

- DC2/DNS Secondary Domain Controller
- DC3/DNS Tertiary Domain Controller (Present for Fault Tolerance and Redundancy)

Server Segment 1

- DC1 Primary Domain Controller
- Cert 2 Issuing Certificate Server
- Tape/Backup device
- Admin Workstation
- Log Server

Server Segment 2

- File/Print Server
- DB Database Server
- WEB Web Server
- R&D Server
- Mail Server

IS Segment

- IS Workstations(x10)

HR Segment

- HR Workstations(x10)

R&D Segment

- R&D Workstations(x10)

Sales/Marketing Segment

- Sales Workstations(x10)

Admin Segment

- Admin Workstations(x10)

In addition, there is a single, stand alone Root Certificate Server, which is offline, and not connected to the

network, which is kept in a locked room, with a removable, removed hard drive, locked separately.

Not considered part of the network, but essential to its operation, will be a testlab, of servers and workstations capable of emulating the various domain controllers, servers, and workstations in the production environment. Here not only will the various applications be tested for Win2K compliance to ensure the network can function in Native mode, but also to confirm various security settings will not interfere with the operation of the business. All changes should be well tested here against a variety of scenarios and configurations before being activated in the production environment. The test lab will not, however, be connected in any way to the inter/intra nets.

Authentication:

Internally, the machines shall authenticate to each other using the MS-Kerberos v5 Protocol. This is the preferred method when available. However, for authentication of the machines in the DMZ, which shall be configured as stand-alone servers and hence *not part of the Active Directory*, this method shall not be available. Therefore, in this situation, certificates shall be installed on these machines prior to deployment, to be used in the process of authentication. Also, for use in https:// protocol over the Web on the public web server, a public certificate from a third party such as VeriSign should be obtained.

Machines on the intranet shall communicate with those on the DMZ only in a very constrained manner. The internal server alone may initiate a connection. IPSec shall be employed between partnered servers; the Mail Servers shall communicate, the Web servers may communicate, and the DNS servers may communicate. All other combinations shall be prohibited.

Authentication shall also use IPSec Authentication Headers (AH) with Secure Hash Algorithm 1 (SHA1) or alternately Message Digest 5 (MD5) if for some reason SHA1 is not available (more on this below).

DNS:

DNS must be established before the first Active Directory Domain Controller is established on the network (using DCPROMO.EXE). Once this is done, the primary Windows domain is established and cannot be altered. Sub-domains can be added to the tree, and other domains can be added to the forest, but the primary domain should be established first. Therefore, one of the early tasks in creating the network is the establishment of the DNS server and defining the Primary Domains (in this case, the only domains, giac.com and win.giac.com).

Active Directory:

Active Directory shall be created first in a primary Domain controller, then in two more domain controllers, although one more would be sufficient functionally, the third is added for redundancy and fail-over. All three are functionally identical machines, but the first, by virtue of being the first to be promoted to domain controller (via running DCPROMO.EXE) becomes the Primary Domain Controller. As the design is for a single domain, all three may be Global Catalog Servers.

The design choice of a single domain was suggested by the Bishop of Occam, who originally suggested the KISS rule (Keep it Simple, Stupid!). This will ultimately simplify administration but could also raise some questions relative to differing security requirements among departments. While such requirements were considered, they were rejected in this case, which is the case of a smaller business. Only were a business to grow beyond a threshold of, say, ~10K users or departments with radically different security requirements would such distinctions begin to outweigh the advantages gained by the simplifications of a single domain

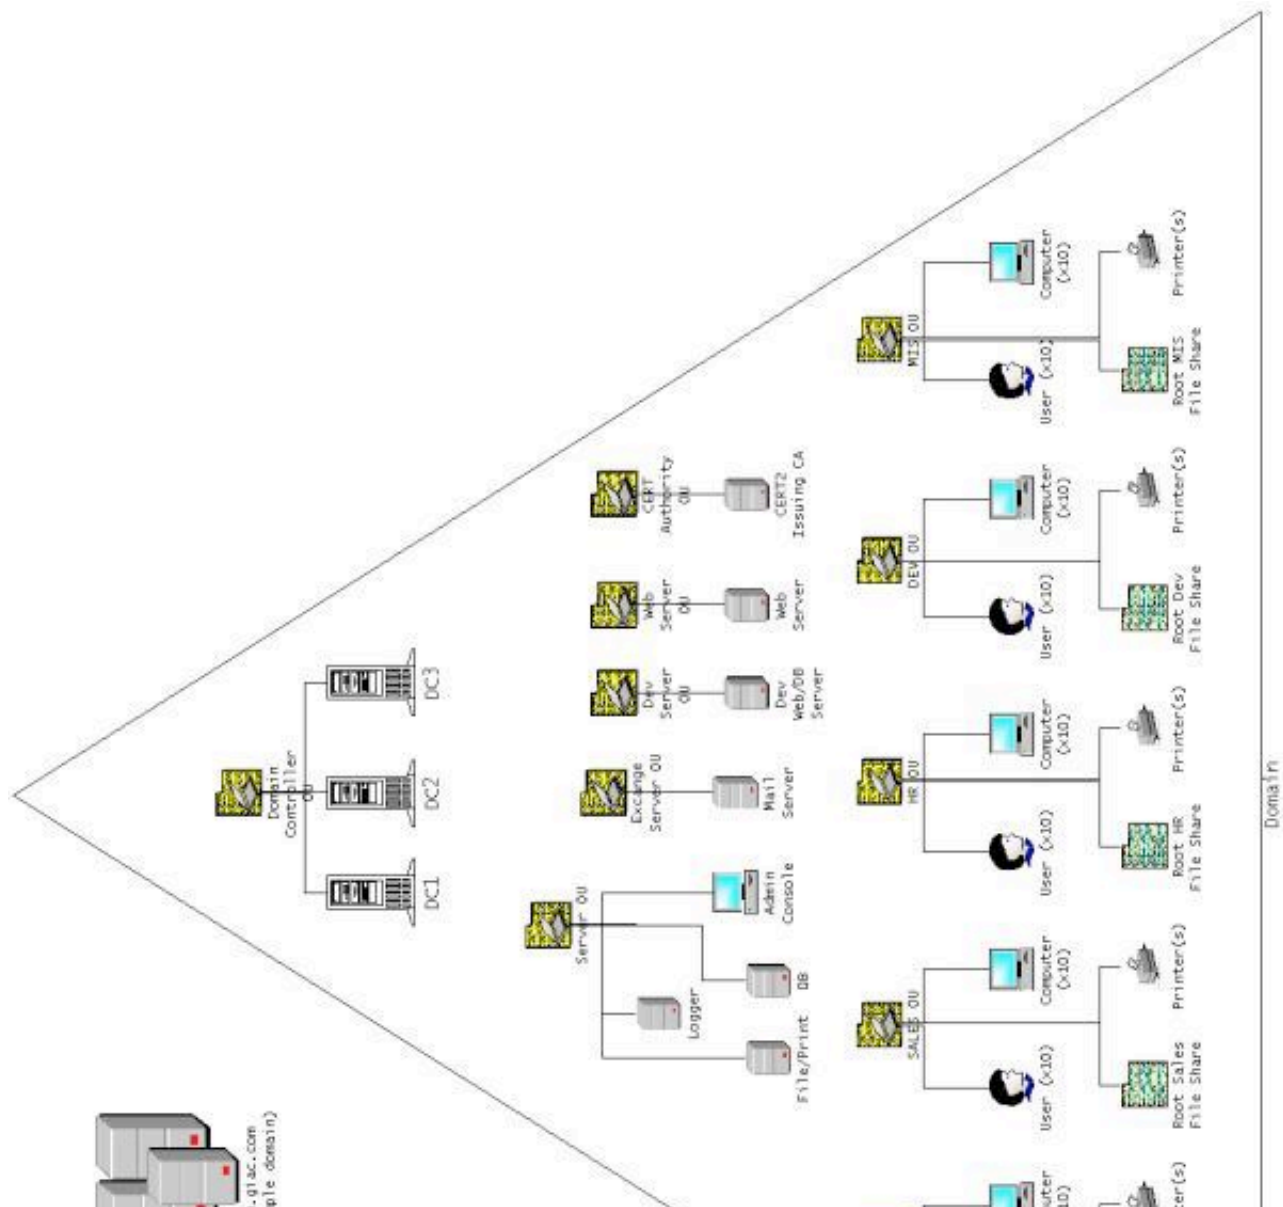
administration. Since this Active Directory design consists of a single domain, the AD hierarchy itself is simply a single domain, further subdivided into functional units. The advantages realized by a single domain are:


- Single view of the Directory
- Unified security policies and settings
- Simplified administration
- The ability to affect globally all users/computers by a single action.

The major disadvantage are:

- Specialized security measures are more difficult to enforce in a single domain.
- Bandwidth and Replication limitations

However, with the ability to fine tune administrative rights and controls through DACLS and Group Policy, the major disadvantage is far outweighed by the advantages gained in a streamlined administrative overhead of the simpler design. This has been an overall guiding principal of the design. Also, with a single site, replication issues are nonexistent, and even in a multi-site domain, these issues can be resolved without resorting to multiple domains to limit replication.





As the entire network is available via a high speed connection, it is defined as a single site at this time. In later phases, as the company expands, additional sites at regional offices may be added, but even then these will probably be included in the same domain. The model will then expand to be a Single-domain, Multi-site "Extended Domain Model".

Hardware Considerations:

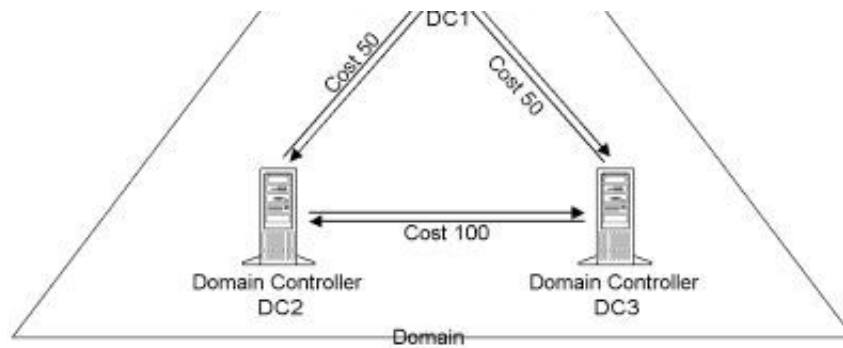
Domain Controllers (DC):

Each of the 3 DCs on the network shall be an exact hardware copy of the same configuration.

- Dual Processor Pentium 3 or better
- 2 independent SCSI 3 Ultra controllers
- 1st SCSI Controller:
 1. Mirrored System Drive (C:)
 2. Raid 0,5 DriveLogical Partitions:
 - Active Directory Drive (D:)
 - Application Partition (E:)
- 2nd SCSI Controller:
 1. Mirrored Swap Drive (F:)
 2. Raid 0,5 Log File DriveLogical Partitions:
 - one for AD Logs(G:),
 - one for system logs(H:)
- CD-ROM(Z:) Set to high letter drive to avoid conflicts during drive swaps.

The second and third domain controllers are identical and present for redundant fault tolerance. The replication cost for the channel between them is set to double that of the path between either and DC1, so that, all other factors being equal, information generally shall take the path from DC1 to DC2 or DC3 and from DC2 or DC3 back to DC1, or from the "periphery" towards the "center" and back "out".





Other Servers:

- Dual Processor Pentium 3 or better
 - 1 SCSI 3 Ultra controller
 1. Mirrored System Drive (C:)
 2. Mirrored Swap Drive (D:)
 3. Raid 0,5 Drive
- Logical Partitions:
- one for Applications(E:)
 - one for system logs (F:)

Log server:

The log server is a multi-processor machine which requires a large disk space for receiving logs from all machines on the network, and to which is attached a CD burner tower. Logs are transmitted to this machine where they are committed to write-once media and archived for reference in case of need. Meanwhile, scripts can process the files for irregularities and evidence of suspicious activity, generating administrative alerts as necessary.

Admin Workstation:

A System for orchestrating the control of Domain controllers, system backups, Logging, IDS monitoring, and other system administration tasks. Attached to the same leg of the network as the primary domain controller so that it is accessible even if the rest of the network is blocked.

Workstations:

- Single Processor
- Single Hard Drive, 2 Partitions for system and apps
- Ghosted system install for each department

Network Interfaces:

All Network Interface Cards must support 168 bit encryption in order to offload operations from the CPU.

Software Considerations

Software changes even faster than hardware, and patches and security updates are released almost daily. The basic configuration for the network is as follows:

- Windows 2000
Advanced Server for Servers
Professional for Workstations
- Service Pack 2
This is *essential* for advanced encryption options.
- Cumulative security updates to date.
These change from day to day and must be monitored. Rather than enumerate every one, a pointer to a good list of the critical patches and updates for Windows 2000 at <http://www.activewin.com/win2000/patches.shtml> should suffice.
- Anti-Virus Software
Norton Anti Virus is the current choice du jour, but some good commercial grade product is essential. This may not detect every virus out there, but every one that it does is hours of work and much data saved that would otherwise have been wasted and lost. Updates should be obtained frequently and distributed to the intranet from a central point of distribution. Hosts on the DMZ will require direct installs of updates to keep interactions with the intranet to an absolute minimum.
- Intrusion Detection System
There are both Host based and Network based systems, and hybrid systems that combine features of both. A host based system will check for changes to registries, critical files and logs, basically monitoring a given host machine. Network based systems monitor packets through the network for evidence of malicious intent. The best systems use both methods in a hybrid solution. This area has been changing and evolving quite rapidly, and the recommendation here is not for any one specific system, but only that one be in place *and monitored*. A hybrid solution seems optimal. For more details, see: <http://www.securityfocus.com/infocus/1520>

Steps for Securing:

All systems shall be converted to NTFS5 (Win2K FS) > `convert volume /FS:NTFS [/V]` (i.e.: `convert C: /FS:NTFS /V` [the /V gives verbose output]) Note: Upon converting the filesystem, the Everyone Group has Full Control of all Files on the volume! Since the Everyone Group includes all users, including Null sessions, random hackers, and just about anyone else, this should probably be changed before much else gets done. This should be altered both as a local setting and as a Domain Setting in Group Policy. (See below) A cursory examination of most of the default security templates provided by Microsoft shows that this group still persists through most of them. The NSA is on the right track in ruthlessly eliminating it from all file systems in their secure templates. The group itself still exists, but should have all privileges revoked. This will be accomplished when applying Security Settings in the section below.

Security Settings:

Security Settings are applied by Group Policy in four places, of which three are recommended, and the fourth is not. They are the Local, Site, Domain, and Organizational Unit. (We are assuming a homogeneous WinNT network, and so ignoring legacy NT4 Policy.):

1. At the local level: Account and Local policies can and should be applied here through the MMC snap-in "Group Policy" as soon as practicable. Policy at this level is determined by the type of machine; Domain Controller, Server, or Workstation. This Policy will be overridden by Domain and OU policy, but a strong policy at this level will fill in any gaps left in those policies.
2. Linking GPO's to a Site [3]

Linking a GPO to a site is a way to assign a Group Policy to more than one domain. Any given site may contain computers from one or more domains. If a site contains users and computers from more than one domain, the site GPO settings will apply to all users and computers in that site, regardless of the domain in which the user or computer resides. However, linking GPOs to sites introduces a number of considerations, such as:

Site GPO permissions - With a GPO linked to a site, anyone with read and write permissions to that GPO could make changes to the GPO. Because the GPO is linked to a site, its policies would propagate to the entire site, possibly affecting computers in multiple domains.

Network traffic patterns - By default, the GPO for a site is created in the root domain of the forest. GPO implementation uses some amount of network bandwidth. Placing a GPO in the domain root could have a negative effect on inter-domain traffic and perhaps GPO refresh.

User rights troubleshooting - Because of the flexibility in Active Directory to layer GPOs combined with GPO inheritance, troubleshooting user rights problems can become a difficult issue. Also, since Active Directory object permissions are not inherited among domains, receiving a domain or OU GPO from a site may be an unusual concept for some administrators.

To reduce unnecessary complexity and avoid misconfiguration, it is recommended that GPOs generally not be linked to sites.

3. Domain Level GPO's This is the primary application of Group Policy. One should strive for as thorough an application at this level as possible, as users can log on to the domain and bypass the OU GPOs. Any security traits that can be abstracted should be implemented at this level. Unfortunately, there is much that probably will have to be left for more specific groupings.
4. Organizational Units At this level, GPOs functionality should be divided. One can choose to activate only a portion of a GPO, that applying to a User or Computer, and Computers can be divided in their OUs by function, Workstations in the common OUs with the Users, Servers in their own OUs, Domain Controllers or Special Purpose Servers such as Certificate Servers in their own OUs, where specific GPOs can be linked to the containers without contradiction.

Security settings are complex, and complexity is generally antithetical to security. One way to simplify the complex is to begin with a set of security settings from a trusted source (or a recommended source, if a trusted one is not available) scrutinize those, making modifications to suit one's own environment, and apply those. It cannot be stressed enough that any such changes should be applied first in a test environment, and a small controlled set at a time before releasing them in a production environment.

The settings considered here began from the NSA's recommendations for domain controllers, servers and workstations available from <http://nsa.gov> ([4] and [5]). Most of the screen shots show a modified domain controller, although the listings indicate settings considered for servers and workstations, as well. Note: The NSA has added some settings to the default MS templates. In order to use these, one must

1. copy their extended sciregvl.inf to %SystemRoot%\inf directory
2. run: "regsvr32 scecli.dll" to register the changes

This will register the additional security template choices:

```
; Added for NSA security templates
AutoAdmin = Allow Automatic Administrator Logon
DisableAutoplay = Disable Media Autoplay
DACD = CD-ROM Drives
DAALL = All Drives
```

This is the actual text from the modified `sceregl1.inf` file supplied by the NSA. It adds the possibilities to enable/disable automatic administrator logon and enable/disable media autoplay as noted below. One may use the NSA provided templates without this process, in which case, the additional choices will not be available.

Importing Security Templates:

Security Template Files should be imported. One may copy them into the default directory `%SystemRoot%\Security\Templates\` or create a separate directory for templates. In the latter case, it should receive the same protections as the default directory. Template files should be copied before modification so that they can be restored to their original state. The originals can be used as the basis for different files for different purposes (i.e., web server, mail server, etc.). There are both command line and GUI tools for investigating and applying security templates, each of which have advantages. When first importing a template, and after modification before applying it, the file should be validated. The command line tool `secedit.exe` will do so with the syntax: `secedit.exe /validate filename` Note: A Domain controller template will fail this test on a workstation or on a server before it has been promoted (before `dcpromo.exe` has been run) due to the lack of a `SYSVOL` share. This is expected, and does not reflect an error in the file. If the file is a valid template, it may be used to

1. analyze the current system; this will compare the current settings to those called for in the template and report any changes that would be made were the template applied, or
2. applied to the system; in this case, the changes are actually applied to the system. Warning: file permissions, registry and other changes may be made, which, if improperly set, may render the system inaccessible or unbootable. Approach this task with caution.

Either of these two tasks can be done with the `secedit.exe` command line tool or through the "Security Configuration and Analysis" MMC Snap-In. The primary advantage of the command line tool is the `/areas` switch that allows the application of a section or sections of the template at a time (SECURITYPOLICY, GROUP_MGMT, USER_RIGHTS, REGKEYS, FILESTORE or SERVICES). With the GUI approach, applying the template is an all-or-nothing affair.

Security Templates can also be imported and configured with the "Security Templates" MMC Snap-In. This allows a Tree-view to drill down to any desired setting in the template to adjust its setting, and the root template may be saved or "Save-As" to a new file as desired.

Last but not least, once a template has been configured, it may be associated with a Group Policy Object, and thence linked to any container in Active Directory which allows GPOs to be linked: Domains, Sites, and Organizational Units.

What follows is a rather tedious, yet important, consideration of some of the more important settings available through Security Templates and suggested values (Those marked with an asterisk differ from the NSA recommended values)...

Account Policies:

Account policies should be set on each machine in the domain. To insure that the policy is enforced

consistently across the domain, the template should be applied to the Domain via the domain's Group Policy Object.

Password Policy	
*Enforce password history	12 passwords remembered
Maximum password age	90 days
Minimum password age	1 days
Minimum password length	12 characters
Passwords must meet complexity requirements	Enabled
Store password using reversible encryption	Disabled

A reasonably large pool of remembered passwords is essential, as many users will cycle familiar passwords as soon as possible. Too soon an interval for password changes becomes an annoyance, and too long a change becomes too great a risk. The minimum age prevents changes too frequently, which can cause users to forget the changes, as well as stressing the system with unnecessary replication of information. Password lengths could even be longer, but every character longer than eight characters increases permutation combinatorics. Better complexity requirements can only help security, so long as users are cautioned against sharing or recording of passwords.

A one-way hash as an internal storage of a password is arguably more secure than the reversible variety, which should only be used when absolutely required by interoperability constraints, and even then, those constraints should be closely examined first and avoided if at all possible before taking such measures. Note: The premise at this point is that logins shall be by password authentication. It is hoped eventually that smart-card PKI can be implemented at some point in the future, but this would involve the deployment of an additional CA and other hardware, more technical support and user training, and does not appear to be an attainable goal for the immediate future. This course should not, however, be abandoned, but pressed in stages and as such, security settings are included below to incorporate its gradual adoption.

Account Lockout Policy	
*Account lockout duration	20 minutes
Account lockout threshold	3 invalid logon attempts
Reset account lockout counter after	20 minutes

Generally, these values should be long enough to discourage dictionary and guessing attacks, but not so long as to completely lock out a hapless user. Many authorities suggest 15 minutes, and for that reason, many hackers will guess that interval. Fewer than three invalid attempts would be silly, as it is all too easy to mis-type a password or not notice the CAPS-LOCK is on. Too many attempts gives too much ammunition to a would be intruder against what might already be an easy to guess password. (Spouse's Middle Name, Your Cat, or Your Birthday Backwards)

Auditing Policy:
Path: Local Policies - Audit Policy

Audit account logon events	Success, Failure
Audit account management	Success, Failure
Audit directory service access	Failure (domain controllers)
Audit logon events	Success, Failure
Audit object access	Failure
Audit policy change	Success, Failure
Audit privilege use	Failure
Audit system events	Success, Failure

Audits can produce huge amounts of data very rapidly. It is important to tailor the audit to the information that is necessary to detect the behaviors one wishes to detect and/or deter. Activating an audit policy after the fact is locking the barn door after the cow is gone. Audit policies are also useless unless attention is paid to them. Much of this can be automated with scripts, using regular expressions and pattern matches which trigger some action on a positive result.

perl is the scripting tool par excellence for this sort of task, (a Windows compatible version is available from [ActiveState](http://activestate.com) at <http://activestate.com>) but any language with regular expression support such as VBScript can be made to do duty. One can find scripts ready made, and also construct them to perform any number of repetitive tasks one finds ones self repeating.

It is also important to remove logs from the machine on which they are generated as soon as practicable, and to store them on a write-once medium. This way they can be compared to the logs on a suspect machine for signs of tampering. A tool such as "diff" is quite useful for this purpose.

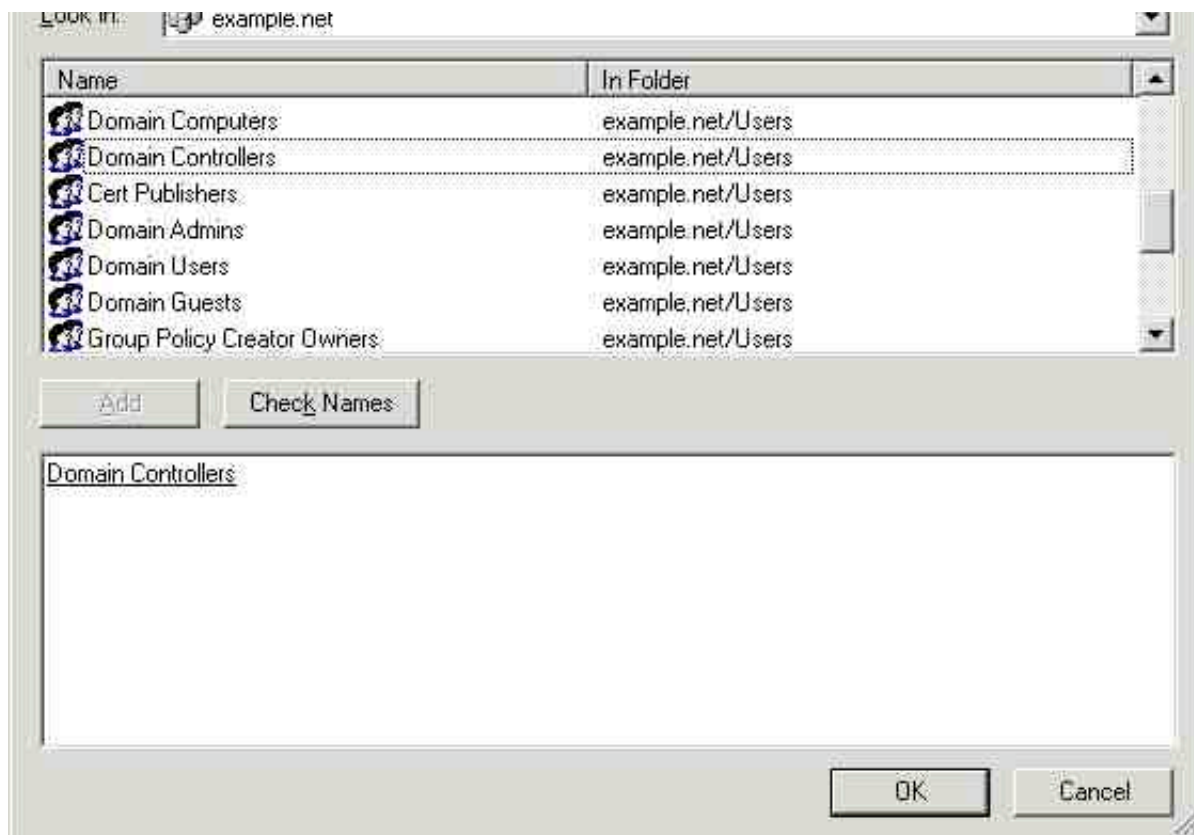
User Rights Assignments		
Path: Local Policies - User Rights Assignment		
Security Option:	Workstation	Server
Additional restrictions for anonymous connections	No access without explicit anonymous permissions	
Note: see Knowledge Base article Q246261 http://support.microsoft.com/support/kb/articles/Q246/2/61.asp		
Allow Automatic Administrator Logon	Disabled	Disabled
Allow Server Operators to schedule tasks (domain controllers only)	Not defined	Disabled
Allow system to be shut down without having to log on	Disabled	Disabled
Audit the access of global system objects	Enabled	Enabled
Audit the use of backup and restore privilege	Enabled	Enabled
Clear virtual memory pagefile when system shuts down	Enabled	Enabled

Digitally sign client communication (when possible)	Enabled	Enabled
Digitally sign server communication (when possible)	Enabled	Enabled
Disable CTRL+ALT+DEL requirement for logon	Disabled	Disabled
Disable Media Autoplay	All Drives	All Drives
Do Not Display last user name in logon screen	Enabled	Enabled
Lan Manager authentication level	Send NTLMv2 response only\refuse LM & NTLM	
Prevent users from installing printer drivers	Enabled	Enabled
Prompt users to change password before expiration	*7 days	
Recovery Console: Allow automatic administrative logon	Disabled	Disabled
Restrict CD_ROM access to locally logged on user only	Enabled	Enabled
Restrict floppy access to locally logged on user only	Enabled	Enabled
Secure channel: Digitally encrypt or sign secure channel data (always)	Disabled	Disabled
Secure channel: Digitally encrypt secure channel data (when possible)	Enabled	Enabled
Secure channel: Digitally sign secure channel data (when possible)	Enabled	Enabled
*Secure channel: Require strong (Windows 2000 or later) session key	Enabled	Enabled
Send unencrypted password to connect to third party SMB servers	Disabled	Disabled
Shut down system immediately if unable to log security audits	Enabled	Enabled
Smart card removal behavior	Lock Workstation	
Strengthen default permissions of global system objects (e.g., Symbolic Links)	Enabled	Enabled
Unsigned driver installation behavior	Warn but allow installation	
Unsigned non-driver installation behavior	Warn but allow	

Restricted Groups (Domain):

Path: Security Templates - Template - Restricted Groups Right Click - Add Group





Here is an excellent aspect of Group Policy Security. Restricted Groups contains a list of groups for which the list of members should remain constant. When this policy is implemented (and each time it is refreshed), any member of the group that is not listed in the policy is removed from the group. Any member of the list (user or group) who is not a member of the group is added to the group. This guarantees that even if a group is modified, it will not remain so for a period longer than the Group Policy refresh interval.

Exchange Server Settings

Exchange Server requires special settings to function, and many of these will be outlined here. (This is one reason for placing the exchange server in its own OU.)

Permission	Description
Log on as a service	The account which is to run Microsoft Exchange will require this User right.
Manage auditing and security log	Exchange Enterprise Servers group requires this right on the Domain Controller.
Registry Key: \\MACHINE\\SYSTEM\\CurrentControlSet\\Control\\SecurePipeServers\\winreg	Exchange Domain Servers group Full Control on Exchange servers and Domain Controllers

Note: Converting Exchange Server from an earlier version to Exchange 2000 is a complex process and goes far beyond the scope of this document. There is far more than just security consideration to take into account. Should one be considering such a conversion, there are a couple of excellent guides to serve as

starting points:

- "20 tips for Exchange 2000 Migration" Windows 2000 Magazine, October 2001 P.94 [6]
- Building Enterprise Active Directory Services Notes from the Field [7]

Applying Security Settings:

Having enumerated these settings in some detail, the actual implementing of these policies should be discussed. Although a simple process, it is fraught with peril, and care should be exercised.

First, backup the system. Changes can be made which can affect the registry and other parts of the system, and the only way to guarantee that those changes can be reversed is to have a back-up of the state before the changes were made.

Second, validate the templates: the command line utility `secedit.exe` can be used to do just that, with the following syntax:

```
secedit /validate filename
```

Example:

```
"> secedit /validate C:\WINNT\security\templates\securedc.inf"
```

Once one has a valid template to work with, there are two ways with which to work it. From the command line, one can use `secedit.exe` to analyze or configure system security, or one may graphically do so using the mmc snap in: "Security Templates".

As a command line example to analyze a security template:

```
> secedit /analyze /DB secedit12-09-01.db /CFG
```

Exceptions to Group Policy:

1. Exceptions for someone:

One may set a particular user (often the creator of a policy or an administrator) so that a Group Policy is not applied to that person, even though they belong to a domain or OU to which that policy would normally apply.

2. Another group altogether:

A preferred method would be to place exceptional individuals into a separate group to which a different security policy would be applied, instead. This would be more of an alternative rule, rather than an exception.

Locking Down the DMZ

The servers on the DMZ are configured as Stand-Alone servers. They are independent of Active Directory and must be configured individually. Although tools exist to allow remote login to servers, their use is highly discouraged, as they will become additional targets for intrusion. These servers should be managed by console login only, if at all possible. The exception is the updating of web material from another server, where the contents an internal site is mirrored to the external site. But most routine administrative tasks should be performed on DMZ machines on the console.

What follows is a list of suggestions for securing machines exposed to the internet. It is by no means exhaustive. Entire books have been written on the subject. These are general guiding principles and a few specific important points, only. Much of this was covered in the SANS course on "Securing IIS" [8] .

Securing other web available services follows a similar pattern, only the service changes.

These stand alone machines should all be striped-down, armored boxes. As they are exposed to the internet, one may expect frequent probes and intrusion attempts. The firewall should first block all ports except those necessary for the services exposed:

Server	Service	Port
DNS	dns	53 udp,tcp
Web	http	80 tcp
Web	https	443 udp,tcp
Mail	smtp	25

All unnecessary traffic should be blocked at the firewall, and access attempts logged for intrusion analysis. No unnecessary service or software should be installed on any of these machines, and that which is installed by default should be uninstalled (if possible) or disabled.

- For the DNS Server, this means DNS only
- For the Web Server, this means IIS Common Files, MMC Snap in, and the WWW Server only.
- For the Mail Server, this means the IIS Common Files, MMC Snap In, and SMTP Server only.

Application and content data should be installed on volumes separate from the system partition. The properties tabs for the mail and web servers allow one to reset the locations for the mail and web roots to other directories. This will help prevent system crashes due to overloaded mail servers or some forms of DOS attacks.

Set Registry setting

CAL_MACHINE\System\CurrentControlSet\Services\Tcpip\Parameters\SynAttackProtect = 2 to reduce the effects of SYN Flooding.

Remove unused tools and executables. Cmd.exe, Format.exe, Cscript.exe, etc. may be a useful tools, but they can wreak havoc in the hands of a hacker. Thanks to Jason Fossen for the invaluable suggestion of creating an Admin Tools CD containing these and other useful but dangerous tools, which can be moved from machine to machine so that they are available to the administrator yet can be kept of the box and out of the hands of intruders.

Disable Printers and File Sharing

Secure the Metabase: (IIS for Mail and Web servers)

The IIS metabase contains sensitive information and must be protected. Both the metabase its backup files should be protected: %SystemRoot%\System32\Inetserve\MetaBase.bin (by default) and %SystemRoot%\System32\Inetserve\MetaBack*.MDO .

The files and registry keys which determine the files' locations

(HKEY_LOCAL_MACHINE\Software\Microsoft\INetMgr\Parameters\MetadataFile) should be set to full control for Administrators and System *only* . Audit all Failed attempts at access.

Since this is not a clustered server, the deletion of %SystemRoot%\System32\Inetserve\Iiisync.exe is an

additional step to help protect the metabase.

Authentication Issues

Since these machines are stand-alone, the domain's Kerberos authentication is out of bounds here. Therefore, each machine have public certificate chains installed for the machines which they are expected to trust. For the most part, this would be the internal mail server for the DMZ mail server, the internal web-server for the DMZ web-server, and the internal DNS server for the DMZ DNS server. In each case the trust is one way, in that the internal server is trusted by the DMZ machine, not visa versa. The internal machine must initiate the connection.

Connections between these machines should be on a second network interface from that presenting the public internet interface. This second interface should be configured to use IPSec with both AH (Authentication) and ESP (Encryption) so that the connection is both validated and secured from eavesdropping.

Rules on the DMZ hosts and Firewall shall be set to permit these connections from these designated hosts on their specific ports and to deny all others. The internet interface shall accept connections without authentication or encryption on the designated service ports only.

Web Server Specific

Microsoft publishes an IIS tool called `IISlockd.exe` which will assist in the configuration of a secure web-server. Microsoft also publishes `hotfix.exe`, the authoritative check for Security hotfixes for IIS. available at <http://www.microsoft.com/technet/security/tools.asp>.

In the MMC for IIS, right click on an IIS server - Properties choose the ISAPI Filters tab and remove any unneeded ISAPI Filters. `FPEXEDLL.DLL` is a good example.

`URLSCAN.DLL` is available from microsoft.com to search for URL patterns, and reject requests based on additional information in the URLs.

Sample and help files that should be removed where possible. However, the Error Message pages are included as a subtree of the help file path, so the entire help file tree cannot be removed indiscriminately.

Do not install Web Management scripts (and if they are there and you didn't ask for them, delete them: `\inetpub\AdminScripts`).

Disable WebDAV:

This appears to be a security hole waiting to happen, and should be avoided. Since the content management model in operation at this site is to create content on the intranet and then mirror it to the internet site, there should be no interactive manipulation of content on this site. Set No Access NTFS permission for Everyone on the `%SystemRoot%\System32\Inetsrv\httpext.dll` file to turn off WebDAV.

Note: CERT2, The online Certificate Server for generating certificates for the Web server, Mail server, DNS, and other machines with which secure communication is necessary without Kerberos, uses IIS to issue certificates. This system must be secured as a web server in much the same way as all other web-servers on the network. These steps need not be repeated for every web server, as they are identical, but should not be neglected in any case.

Database Security

Database security is more mature than many other aspects of security and offers an additional layer of security for sensitive data. Three particular areas will benefit from this in the design currently under consideration:

1. **Research and Development:** The latest and greatest developments can be stored in a database where they can be sliced, diced, integrated, analyzed and reported upon in a manner that is unparalleled by other systems.
2. **Finance:** For rows and columns of figures in multi-dimensional tables, a relational database may be just the place to keep this information.
3. **Human Resources:** Personal information beyond the basics in Active Directory will need a home. For ease of data entry and query and reporting, as well as a centralized repository for backup and recovery, a third party SQL Server application may be just the right tool.

Each of these areas (and others) contains sensitive information that should be protected from unauthorized access. One might extend Active Directory to contain additional properties on users such as Salary for employees, then render that property readable by only the HR OU, but there are reasons for avoiding this course. First, this information is widely replicated, and may prove difficult to control. Second, it could lead to the "all eggs in one basket" problem, where critical systems all depend on a single database. Overloading the user authentication system as a Human Resource tracking system as well may not be the best course.

There are many mature software packages developed for SQL Server which can be used for Financial and Human Resource needs. Databases can then be managed by well known methods. Database files can be restricted by NTFS permissions, User Groups and SQL Groups can be created and granted appropriate access permissions on database objects; tables, stored procedures, etc.

R&D often possesses the expertise to modify or create specialized database applications to perform tasks tailored to the business in a manner that is well suited to the demands of the enterprise. Along with the existing Active Directory and NT Security, a single sign-on access control can be instituted which is transparent to the user, but will apply granularity of permissions based on that users Group and OU memberships.

Management of SQL Server permissions can be done through scripts or through SQL Server Enterprise Manager, an MMC Snap-in, through the "Security" settings on a given database. Selecting the user (or better, a group) desired, one can grant Select, Insert, Update, Delete, and Execute permissions with the click of a checkbox.

Although the granting of permissions is simple, they should be granted in the context of a policy document which outlines the permissions to be granted to which groups (to what purpose), and also tested on a test server for efficacy before being used in a production environment.

Securing The Keys

One of the Key areas of Windows 2000 and Active Directory (pun intended) is the LSA secrets, and obscure corner where service account passwords and (apparently) portions of the decryption keys for the Encrypting File System and other vital security measures are protected. Along with some encryption keys, Users' and administrators' passwords, these should be protected by the strongest encryption and protection available. By default, this is not really the case.

There is a utility called `SYSKEY.EXE` to apply a special RC4 system key to further protect these "secrets". This step is, however, irreversible, so it is advisable to create an Emergency Repair Disk (with updated Security Information) before proceeding.

`SYSKEY` presents three options for the computer/operator as to how the System key which it creates will be stored. Each has its drawbacks and its advantages. The key may be:

1. Stored on the hard drive of the machine in a "complex obscuring function" (In other words, it is disassembled and the pieces are hidden). This option shall be selected only for the off-line Certificate Authority (CA1 in the Network Diagram). Its hard drive is hopefully stored in a vault and little is gained by adding further security constraints in this case.
2. Stored on a Floppy Disk. This option requires the floppy disk be present in the machine to boot. The floppy may be left in the floppy bay for unattended booting, but then, the floppy containing sensitive security information is also unattended. This option keeps the information off the hard drives and backup tapes, and in that respect is better. However, damage to the floppy will render the computer unbootable. One must then balance the propagation of floppies against the security of the system. This option is recommended for the DMZ servers only. In this case, two floppies should be made for each machine, one backup locked away and one left in the machine for unattended rebooting. These machines need to be constantly on line and the level of information they contain is not critical to the enterprise.
3. Password Startup - The password is hashed with MD5 to create a key. The password can be up to 128 characters long and should be close to that, using as random a choice of characters as possible. The password must be entered for the machine to boot. Two lists of boot passwords can be maintained in separate, secure (non-electronic) environment, with logged access. (One on site, the other in a safe deposit box.) About three trusted individuals should have access to the list, to be able to access it to boot machines should the occasion arise. The machines will not be able to boot unattended, but they will be considerably more secure than with either of the two other options.

perl code for generating random password lists
<pre>usage: perl randlist.pl 128 12 for a list of 12 passwords each 128 characters long. Each will contain Upper & lower case, numeric and non-alphanumeric characters.</pre> <pre>#!/usr/bin/perl use strict; usage() if scalar(@ARGV) <1 scalar(@argv) > 2; my \$limit = shift; #input length of passwords \$limit = 4 if \$limit < 4; #must be at least 4 or loop will be infinite... my \$listlen = 1; #one password is the default if(@ARGV) { \$listlen = shift; #but more if you like... } #not the best rand seed, but something needs to be seeded... srand(time ^ (\$\$ + (\$\$ <<15))); for (1..\$listlen) { my \$string = "";</pre>

```

while (($string !~ /[A-Z]/)
      || ($string !~ /^[A-Za-z0-9]/)
      || ($string !~ /[a-z]/)
      || ($string !~ /[0-9]/)) {
    $string = "";
    for (1..$limit) {
        $string .= sprintf("%c",rand(94) + 33);
    }
    print $string, "\n";
}

sub usage {
    print <<EOP;
prints 1 or [listlen] random strings of characters [length >= 4] characters.
usage: $0 length [listlen]

EOP
    exit(0);
}

```

IPSec and the Network

The network is divided into physical segments, logical divisions by department or function, subnet groups, and organizational units. By convenient convergence in design, most of these divisions, although representing different levels of abstraction, coincide. For instance, all workstations for the Sales and Marketing department are located on a single physical network segment; they will all be addressed under the same subnet mask of the Class B net; and they are assigned to the same organizational unit. Each of the other departments, and the server groups are similarly grouped.

One advantage of this design is that one may establish rules for access based on sub-net rather than on a per-machine basis. As long as a machine is in the proper sub-net, the rule will apply.

One may apply policy to the Organizational Unit and have it enforced essentially for the physical network segment/sub-net group. One may now consider any of these terms to be loosely equivalent (in this design setting only).

The X.X. refers to the base IP Address

Segment	Department	Network	OU
0	Backbone	X.X.0.0 to X.X.0.255	Domain Controllers
1	DMZ	X.X.1.0 to X.X.1.255	N/A (Not in domain)
2	Research And Development	X.X.2.0 to X.X.2.255	DEV OU
3	Sales And Marketing	X.X.3.0 to X.X.3.255	SALES OU
4	Finance and Human Resources	X.X.4.0 to X.X.4.255	HR OU
5	Administration	X.X.5.0 to X.X.5.255	ADMIN OU
6	Mission Information Systems	X.X.6.0 to X.X.6.255	MIS OU

Note: The analogy does not extend quite evenly to servers, whose OU divisions are more functional, and do not reflect the physical segment placement with such accuracy as the workstation segment map to departmental segments.

All systems on the network should be set to authenticate at all times (IPSec AH). This could be the single most powerful deterrent to an intruder should he penetrate the outer defenses.

Any Server which handles sensitive information, of which it would be considered tactless, at least, to disclose, should also be configured to use encryption in its communications. This would be, at the least, the File and Database and R&D servers. One might suppose that the Intranet Web server would publish public information and would not require Encryption, only Authentication. Should specific web info require stronger protections, https could be used in that specific case.

When one considers that network traffic is essentially like sending postcards, and readable by any system through whose hands the information passes, one should strongly consider the widespread use of encryption throughout the enterprise. As it is, we have identified over half the enterprise which requires such protection, and with a little judicious inquiry, one could probably discover grounds for protecting the data of the rest.

On the downside of encrypting all network traffic is the performance hit the machines will take. This is not much of an issue on the individual workstation, particularly with much of the work off-loaded to the network interface card, but it shall quite increase the demands upon the server, even with hardware off-loading. But such is the price of data safety. It would be senseless to go to extremes to protect one's data while stored on a disk only to broadcast the information the moment any legitimate user should attempt to read it. One must bear in mind that access to a network should not automatically grant permission to access all information that passes over the network. In practice this is often the case.

Another problem with encrypted traffic is that it makes network intrusion detection much more difficult. One must depend more upon host based IDS systems and other means of detection, as the packets in transit offer little evidence of their nature.

Note: "A few types of IP traffic cannot be secured by the design of IPSec transport filters in Windows 2000, including:

- Broadcast addresses - usually ending with .255 - with appropriate subnet masks.
- Multicast-addresses from 224.0.0.0 through 239.255.255.255.
- RSVP-IP protocol type 46. This is to allow RSVP to signal Quality of Service (QOS) requests for application traffic that may then be IPSec protected.
- Kerberos-UDP source or dest port 88. Kerberos is itself a secure protocol, which the IPSec's IKE negotiation service uses for authentication of other computers in a domain.
- IKE-UDP dest port 500. This is required to allow IKE to negotiate parameters for IPSec security." [9]

However, Kerberos and RSVP traffic between domain controllers can be secured by setting a registry key: `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\IPSEC\NoDefaultExempt = 1`
This will enable the behavior, but one must also define IPSec filters as well. (see Microsoft Knowledge Base: Q254728)

With all that said, we will proceed to consider the configuration of IPSec for 1) workstations, 2) servers, and 3) DMZ hosts. The settings in the first two cases shall be distributed through the appropriate Group Policy Object linked to the Organizational Unit representing the department to which the machine is assigned.

1. Workstations:

Workstations should be configured to use authentication in their communications. Usually it will be the workstation that will initiate the interaction, and should only communicate with and authenticated

server. There will be very few times that a peer-to-peer communication between workstations shall be necessary which is unmediated by a server, and should such an occasion arise, the eventuality can then be dealt with. The basic rule is that communication is allowed from the subnet to the servers, if authenticated.

2. Servers:

Most servers contain sensitive data, with the exception of the Web server, which may contain some secure information, but mostly publicly viewable information (public to the intranet, that is).

File servers, the database server, and other servers may contain sensitive information that should be protected in transit. In order to do so without auditing every file, a blanket policy can be implemented on the OU for the servers to insist upon encryption for all server requests.

3. DMZ Hosts:

Unlike the preceding two examples, these hosts cannot be configured through Group Policy and must be configured individually. The basic settings and rules will be similar, however.

Before deploying these servers to the DMZ, they should have a Certificate distributed to them for internal use, and the Web server a certificate for the Web server from a public source such as Verisign for use with SSL/https authentication.

The DMZ hosts will be configured through the MMC snap-in for "IP Security Policies on Local Machine" as they will not be connected to Active Directory. These hosts will be set to use certificate authentication rather than Kerberos as the primary method, again as they are not part of an Active Directory domain.

IPSec Policy

Only one IPSec Policy may be in force at any one time, but a policy may define multiple rules.

IPSec offers the choice of Authentication Headers (AH) and Encapsulating Security Payload (ESP). The former provides authentication of the packets origination only, while the latter provides encryption of the packets contents as well. The choice the combination of these two, as well as the methods used to effect these goals, are effected by the IPSec Policy in effect at the time an interaction is negotiated.

A Policy shall be established at the domain level, with the "Default Response" rule enabled. This shall be followed up by specific policies in specific OUs that require more specific rules, to override the domain policy.

From the MMC, Add the IP Security Policies Snap In

- There is a choice of Local Computer or Active Directory
- One should choose Active Directory when specifying for the Network
- One should choose Local Computer for the DMZ machines

There are three default policies already available, although none of these should be used, as they use the same GUID numbers which could cause problems. Rather, one should use these as guides when creating policies tailored to one's own needs.

To create a policy, Right click on the IP Security Policies Icon in the MMC. Select "Create IP Security Policy" which will launch a wizard to guide the creation process.

Begin wizard:

Name and Description:

Give your Policy a name and a description which indicate its purpose, i.e., "Default Domain Policy"; "Use AH with SHA1 for default communication if no other policy is set."

-> Next

Choose Activate the default response rule Windows 2000 default (Kerberos V5 protocol)

[For intranet machines generally, exception, when a machine will connect to a DMZ machine, Kerberos shall be 1st, with Certificate the second choice. This is covered elsewhere in this document]

[For DMZ machines, Certificate is the only option to choose, no Kerberos choice should be available, as these machines will not be part of an Active Directory domain.]

-> Next Edit properties: This will present the Properties editor for the policy when the wizard exits.

-> Finish wizard.

Edit Default Rule:

Check Session key Perfect Forward Secrecy

Remove "None" entries (When using ESP, leave only these when using only AH)

Remove DES Entries (in all cases)

In other words, require "3DES" for encryption with ESP

Double click each security method, and Click Custom and Settings for each, and choose SHA1 for the Data integrity algorithm as first choice.

The second choice on the list should be identical to the first, except that it uses MD5 as an alternative. Usually, the first choice will be used unless for some reason this is unavailable.

When using AH, choose the first checkbox "Data and address integrity without encryption (AH):
choose SHA1 if this is the first choice, MD5 if this is the second, alternate method.

When using ESP (most of the policies except for the default domain policy and that for the Web Server) check the second box, "Data integrity and encryption (ESP)"

Again, the first choice is SHA1, second MD5.

The encryption algorithm is always 3DES, never DES.

Check "Generate a new key every 100000 kbytes" (100MB max) or less, if performance will tolerate, but never more.

-> Ok. (Closes dialog)

-> Ok. (Closes dialog)

On the second Edit Rule Properties tab, "Authentication Methods", one should see Kerberos already in place. For machines that will be connecting to the DMZ.

One should have a local certificate installed for the machine with which communication will take place.

Selecting "Add" on this tab will prompt for the certificate store and this machine's store should be indicated.

This should be moved to the second choice, which will be failed over to as the second choice, should Kerberos, the first choice fail. Except in the case of the machine that validates the certificate in question, communication should fail outside of the intranet.

Third Tab, "Connection Type" should indicate "All network connections"

-> Ok (close dialog)

General Tab

Name and Description can be edited, if desired.

Check for policy changes every 90 minutes (1/2 the default) should be fine.

Key Exchange Settings

-> Advanced... Button

check Master key Perfect Forward Secrecy to generate automatic key changes during session.

Authenticate and generate a new key after every "30" minutes and 0 sessions (will change to 1 session and gray out when Master key PFS is checked).

-> "Methods" Button:

Remove DES methods, leave only 3DES choices.

Prefer SHA1 for Integrity Algorithm, 3DES for Encryption, and Medium(2) for Diffie-Hellman Group

The Diffie-Hellman Group is important, as it determines the key-size which is generated. The other choices are the standard SHA1 + 3DES that are generally preferred.

-> Ok.

Remove DES choices

Leave MD5 + 3DES + Medium(2) alternative as second choice

-> Ok

-> Close.

Add new rules...

Begin Wizard...

This rule does not specify a tunnel... (as it specifies a transport.)

-> Next

All network connections

-> Next

Windows 2000 default (Kerberos v5 protocol)

-> Next IP Filter List

-> All IP Traffic Add

--> Next

Source My IP Address

Destination Any IP Address

Protocol type: Any

Properties: Edit

-> Finish

Insure that the Filter is Mirrored.

Description should be descriptive:

"This filters all IP traffic from this IP to any IP Address on any Protocol and is mirrored."

Select Radio for new Filter...

Filter Action Tab

Require Security...

The preceding has set up a Policy that will require Encryption for all connections to any machine in the container to which the policy is associated. Encryption for both Phase 1 negotiation and Phase 2 encryption shall be 3DES with SHA1 the preferred authentication method.

Since the protocol has been mirrored, this policy will be activated in both directions when this policy is active on one of the machines.

The following table shows the various policies that need to be in place for the network. They should be associated with the Domain and OUs that contain the affected machines:

Domain (Default w/no filters)				
Source	Destination	Protocol	Mirrored	Action
My IP	Any	Any	Yes	AH with SHA1 then MD5
Server OU				
Source	Destination	Protocol	Mirrored	Action
My IP	Any	Any	Yes	AH + ESP 3DES / SH1
Exchange OU				
Source	Destination	Protocol	Mirrored	Action
My IP	Any	Any	Yes	AH + ESP 3DES / SH1
Web Server				
Source	Destination	Protocol	Mirrored	Action
Any IP	My IP	http/https	Yes	AH with SHA1 only

Exchange OU				
Source	Destination	Protocol	Mirrored	Action
My IP	Any	Any	Yes	AH + ESP 3DES
Certificate Authority OU				
Source	Destination	Protocol	Mirrored	Action
My IP	Any	Any	Yes	Block
ADMIN Workstation(by IP)	CERT2 (by IP)	Any	Yes	AH + ESP 3DES
In this case, access is blocked, except from the admin workstation. all access must be by console or from the Admin station except to the certificate web server. This must be modified to allow connection by other systems as needed, but only as needed for as long as needed.				
Each Department OU				
Source	Destination	Protocol	Mirrored	Action
OtherDepartmentSubnet	ThisDepartmentSubnet	Any	Yes	Drop
This filter needs to be repeated for each of the five subnets, although four will be drop, the fifth will be authenticate only. In this case, IPSec is used to packet filter packets from one department to another, so that each department can function much as an independent subnet. Departments can still access the servers and communicate via email and other shared applications on servers to other departments, but direct peering to other departments is blocked. When connecting to the servers, the server rules to use encryption and authentication shall be enforced, and the communications will be free from eavesdropping.				

Assigning IPSec Policy via Group Policy

IPSec Policies can be associated with Group Policy Objects (GPOs) which are then linked to containers (Sites, Domains, and Organizational Units) and they are applied in normal order of precedence, LSD-OU (Local Policy, Site, Domain, OU). Since only one IPSec policy can be in effect at a time, the last applied will be the one in effect.

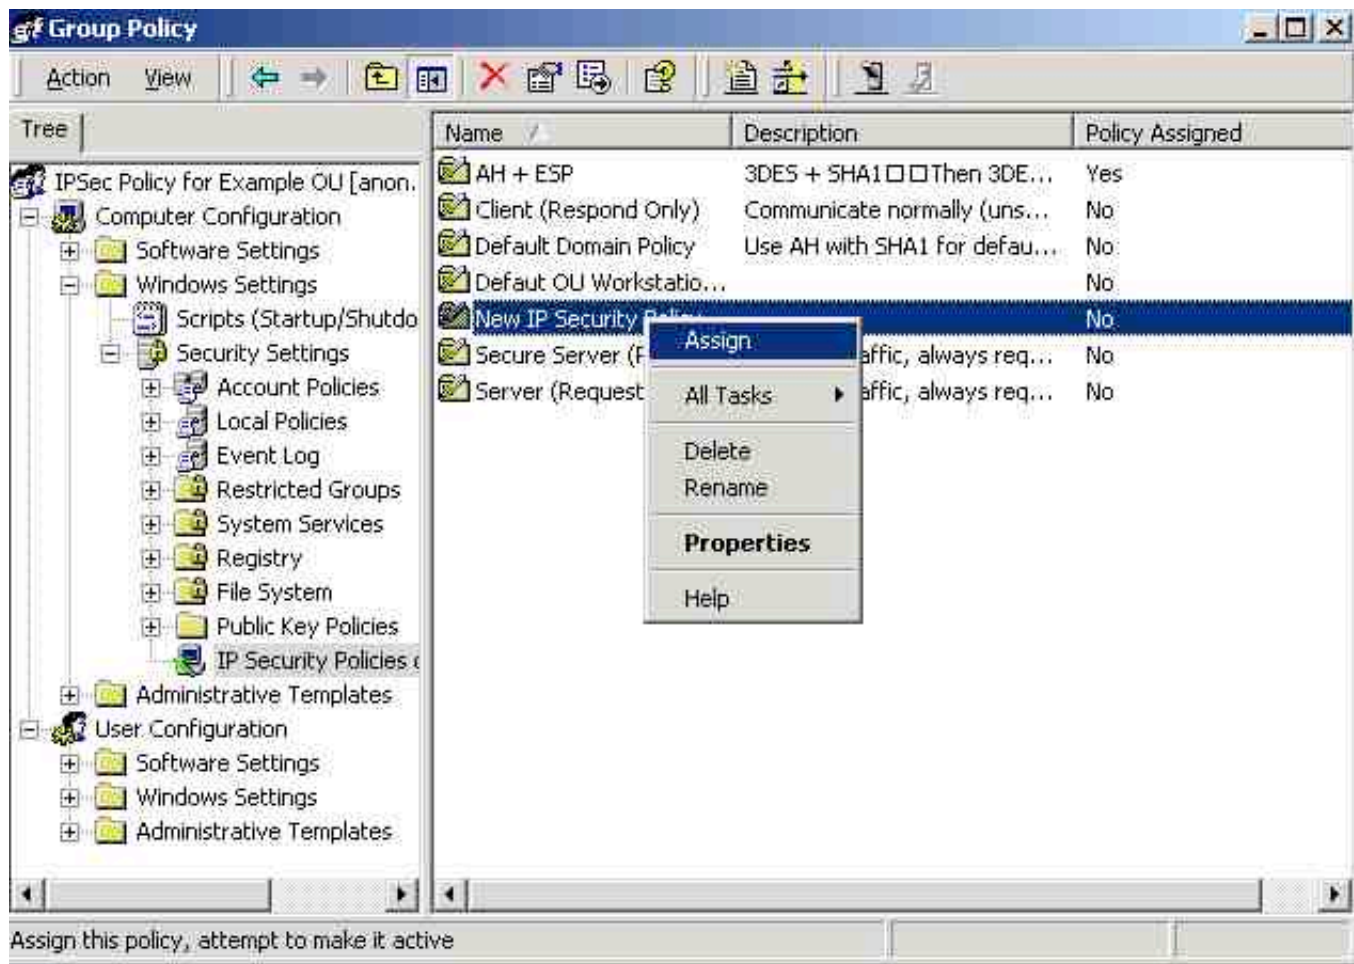
In this design, Local Policies shall only be in place on the stand-alone DMZ servers. For the intranet, the default policy is assigned to the domain, which contains only a default rule to authenticate (AH) both machines. This will essentially eliminate unauthorized intruders unless they gain access to an authorized machine. Beyond this, different OUs with more sensitive resources will apply policies listed above to require ESP in addition to AH.

Once policies have been defined in Active Directory, they are visible through a GPO by selecting a Group Policy Object. (Here we shall create specific GPOs for the explicit and sole purpose of assigning IPSec Policy.)

- Select Computer\Configuration->Windows Settings -> Security Settings ->IP Security Policies on Active Directory
- Select the desired policy, Right Click and choose "Assign".
- The column "Policy Assigned" will show "Yes" to reflect the choice.
- Contrawise, choosing "Un-assign" will cause the policy to be unassigned and "No" to appear in the same column.
- Since only one policy can be in effect at any time, assigning another policy will un-assign the current choice simultaneously.

It is important to bear in mind that even though a policy is assigned, it may not be in effect if a policy is assigned to a nested container (such as a contained OU). Mapping these relations and documenting changes is essential in maintaining the state of the network at all times.

Once created and assigned, the GPO and its links to containers should be controlled by limiting the permissions to alter them to either Administrators or a special IPSec Admins Group set up for the purpose. Since IPSec is one of the keystones of the security system, arbitrary changes to these settings should be avoided, and attempts at modification should be audited for success and failure.



Monitoring IPSec:

Two particularly useful tools for checking the current state of IPSec activity are: `IPSECMON.EXE` can monitor the state of IPSec activity on the local computer. It provides a GUI summary of the current state, refreshed at a user defined interval. `NETDIAG.EXE` a command line utility which can show both IPSec Phase 1 negotiation and Phase 2 packet statistics, in extremely detailed format.

`netdiag /test:ipsec /v` for Phase 1

`netdiag /test:ipsec /debug` for Phase 2

The Unimplemented Enhancement

It was originally desired to take advantage of the Encrypting File System to augment the security of the system, particularly because it is so well integrated with Active Directory and transparent to the users. However, scrutiny of the existing information leaves a little to be desired in the key area of Key Management.

The files are encrypted with DESX, a better choice than normal DES. But the key is stored *right along side the encrypted file*. True, the key is first encrypted, but this makes the private keys which encrypt those keys particularly subject to scrutiny.

It has been demonstrated time and again that the best security is provided by algorithms which are publicly presented to the review of peers. Only by such evaluation can the relative security of a scheme be honestly and fairly determined.

The best answer to the question "How are the Private Keys Protected?" that can believably be presented is "We don't know for sure." The party line appears to be, "Trust Us".

When determining the nature of the protection of Private and System keys on Windows 2000, one falls into the gray realm of undocumented corners, Obfuscation Algorithms and spin. Such a situation leaves the present author no choice but to recommend minimal use of PKI, particularly for EFS and file storage. This is unfortunate, as the outer manifestations of the system are well thought out and the features really desirable. Yet, if the seductive veneer does not conceal an underlying flaw, it is at least impossible to confirm at this point that it does not.

Physical Security:

- Even the most secure software will fail in an insecure environment. Physical access to hardware must be limited. Securing the desktop will probably present the greatest challenge, as they are the most accessible.
- The introduction of home/alien systems on the network must be strenuously avoided. Introduction of trojan horses and virii by this vector could be one of the greatest threats to an otherwise secure network. Likewise the introduction of user supplied media (floppies and zip disks).
- Workstations should be secured in their areas to make removal difficult and time consuming. A small amount of user training in proper secure habits will go a long way in preventing many casual mishaps.
- All servers shall be in locked server rooms where electronic access tokens monitor access times, both upon entrance and exit.
- Server rooms should be, if possible, away from exterior walls and secured from overhead/underneath access crawlspaces.
- All network access points, including routers and junctions shall likewise be in locked access closets with similarly logged access controls.
- Access logs must be monitored for them to be effective.
- Additional locked racks for servers, intrusion alarms, motion sensors, and cameras in the server room and work areas are also necessary.
- The root CA server should have its own, further secured room, removable disk drive stored in a separate vault.
- All equipment should be inventoried, marked, and tagged to aid both in identification and (hopefully) recovery of stolen property.

- Know your co-workers! Know your customers! Question those you don't recognize or don't look like they belong. Watching each other watching each other watching...
- Smart Cards with forced logout or Lock Workstation on removal are an excellent method of limiting access when users are absent. This has been slated as a future enhancement for the network. The security settings have been implemented so that when hardware is added, the system will respond appropriately. It is a major undertaking in hardware, software, infrastructure, and user training to contemplate this step, but one which for which the returns will not be small.

Continued Threats

In spite of the uncertainty, it appears that people are intent on engaging in ventures equivalent in today's world as perhaps to what gold prospecting was in California of the 1840's or sailing for the New World in centuries past. Without a doubt, in spite of the risk, they will continue to do so.

Continuing threats to the security of a system might be classified by many parameters. Here are a few categories that may be considered to some advantage:

- **Internal:** The "Disgruntled Employee" can potentially cause considerable damage to systems. The more access to systems a person has, the more the potential that access can be used for trouble. Very active Auditing and Monitoring will, if not prevent this, at least identify the source and hopefully mitigate the effects. Off-site backups should be entrusted to others than those who routinely service the same data. Dividing authority over any given area will help to prevent an "eggs in one basket" scenario.

Limitation of scope is a major tool in the securing of any set of resources. A user cannot compromise that to which they do not have access. Users should be only granted membership in groups with permissions to resources on an "as-needed" basis, with those permissions revoked upon completion of the tasks that made the access necessary. This notion is counter to the normal corporate strategy of accumulating access privileges as a status tokens. Access should be considered "on loan for the duration", rather than cumulative acquisitions.

"Circumvention of Security", usually in the name of efficiency, is a common practice. A physical analog of this is when one sees at construction site the safety guards on a table saw or other power tool taped back out of the way in the interest of boosting production. Where possible, this should be avoided by the enforcing of a rigorous policy with buy-in from the highest levels of management. Education must be encouraged at all levels of the organization as to the importance of the policies as well as the reasons for the same. If a policy seems to be random and arbitrary, it is more likely to be avoided or circumvented. Those who feel that they are participating in the protecting of their valuable assets will cooperate with a higher degree of willingness.

Unfortunately, anytime there is an easier or faster way to do something, someone will take the shortcut, even if policy or security is bypassed in the process. Users will share these "tricks and tips" amongst themselves like social virii to the detriment of the system. Not every act can be monitored, blocked or controlled. Herein may lie one of the system admins greatest challenges, social, rather than logical in nature.

- **External Threats:**

Some of these will come through the network. It is hoped that the firewall, intrusion detection, IPSec policies, authentication measures, etc. will prevent many of these. However, one should operate from

the assumption that at some point intrusion will occur. Then, Detection and Response take over where Prevention left off. Checking digital fingerprints of key files, analyzing logs and audit trails, preparing profiles of users' access patterns, etc. can all lead to indications that *something* is amiss. What to do when the excrement does indeed hit the propeller blade should be adequately detailed and documented. Determining what should be considered questionable or tainted data, when and how authorities should be notified, what information should be collected by whom, and much more, should be detailed in an incident response plan. This plan should delineate the duties and responsibilities of various principles, and goes far beyond the scope of this document to discuss further.

Social Engineering: Other threats may bypass the network proper and attempt to gain access through covert means, with the appearance of legitimacy. Education of employees is the greatest safeguard against this sort of ruse, but like that formerly discussed, one should assume that at some point, an attempt shall be successful. The results will normally be a breach of the network and then response will follow as outlined above. Recognizing the signs and symptoms of social engineering and sharing them with each other can be one of the most effective tools available.

- **Complexity:**

As already mentioned (to the point that it should ring like a mantra) complexity is the antithesis of security. Increasing the number of components in a system also increases the number of ways in which those components can interact. It is virtually impossible to enumerate, much less predict the effects of, all possible PC hardware and software configurations. Once a computer is hooked up to a network, the problem increased exponentially.

The solution, such as it is, is to maintain as modular a design as possible, and keep interactions to a well a defined an interface a possible. Simplifying interfaces can reduce some of the complexity, or at least reduce the multiplicative interaction.

One must be wary of the introduction of a new version of any component in the system, as it may have far-reaching implications in its interactions, which should be tested and evaluated as far as is practicable in a test situation before releasing the change in a production system.

- **Temporal:**

Today's state of the art is tomorrow's legacy curiosity. Any system can be rendered obsolete or at least vulnerable by a single Usenet posting. Changes can alter the constraints under which one believed the problem domain to be defined. New windows of exposure to hitherto unknown threats are constantly being opened.

The solution, such as it is, is eternal vigilance. Constant scanning of newsgroups and lists, manufacturers websites, and other such sources of information. One may never rest on one's laurels and be secure. Erosion at the very foundations of the superstructure one is attempting to secure is a given. This is both the challenge and the thrill.

- **Physical:**

With our nose to the monitor, watching the access patterns of the various users of the system and pattern matches from the weblogs, we failed to notice that someone has just made off with a truckload of merchandise from the loading dock. Even the most sophisticated system can be defeated by a determined user with physical access to the system.

Attack trees

Bruce Schneier presented the concept of Attack Trees [\[10\]](#) as a way to categorize and analyze methods of attack and their costs. This brief synopsis, of course, will not do the concept justice, but will hopefully offer a pointer towards this useful tool. It can be as simple as a set of index cards or as complex as an expert system or relational database application. The basic idea is to create a node in the tree for each attack, and assign it a cost (time, dollars). Attacks can be broken into sub-attacks, which are also classified.

One may then analyze an attack and judge the strength of an attacker to implement the attack: e.g.: a brute-force cryptanalytic attempt against one's 3DES keys could cost \$1 million for a machine (for example) to do the job which you know your competitor can't afford, 3 years (again, for example) which would not be timely, or the other option, bribing or hiring away an executive, (for much less, maybe) which would be more cost effective.

Studying one's opponents and the tree, one then can see what the real (rather than theoretical) threats might be or when the inflection point is reached when the influence of technological advance might shift the balance of probability from one form of attack to another. In the case above, one might be more concerned about the integrity of one's executives and their job satisfaction than their competitor's computing power.

Conclusion:

Security is a process, not an achievable goal. One cannot design and implement a system and then step back and say, with any degree of certainty, "This is secure", with the foreknowledge that this is and will continue to remain true. While following the design patterns laid out in this document and its referents, changes in the infoscape can alter any judgement at a moments notice. The discovery of a hitherto unknown or overlooked vulnerability can render an apparently impervious network open to every hacker on the planet in minutes. The best one can hope for is to reduce the "window of exposure". [\[11\]](#) to any know exploit while at the same time taking prudent steps to reduce one's exposure to unknown risks.

Unfortunately, as systems become more complex, they become more difficult to understand, and hence, much more difficult to "secure". It is well nigh impossible to evaluate the security of a closed, proprietary system as complex as Windows 2000, and the increasing number of exploits and security patches released indicates that even those with access to the source code are apparently unable to manage the diametrically opposed qualities of complexity and security. Part of the strategy of any business engaged in e-commerce should be to manage and mitigate the risk associated with the transaction of business online.

It is hoped that the design presented here has presented both a summarization of good practices, and some insights into design methodology that might assist in the approach of other similar problems. As emphasized above, the design process is iterative, and this document represents only the extent of the process as time and space allowed for the present cycle of iteration. It is not intended to represent a "finished" product, and in fact asserts that this process can never be finished...

References:

General Acknowledgement:

This entire work would not be possible without the notes and experience acquired from the Securing Windows training received from the SANS Institute as taught by Jason Fossen. While their contribution is acknowledged as indispensable, any errors or omissions are solely the responsibility of the present author.

References Cited in this Work:

[1] (p.3) Report Number: C4-051R-00
Microsoft Windows 2000 © Network Architecture Guide
Systems and Network Attack Center (SNAC)

Author:

Paul F. Bartock, Jr.

Paul L. Donahue

Daniel J. Duesterhaus

Julie M. Haney

Prentice S. Hayes

Trent H. Pitsenbarger

1Lt Robin G. Stephens, USAF

Neil L. Ziring

Updated: April 19, 2001

Version 1.0

<http://nsa2.www.conxion.com/win2k/guides/w2k-1.pdf>

[2] (p.19) Jason Fossen
Sans Institute Track 5 Windows 2000 Active Directory & Group Policy

[3] (p.5) Report Number: C4-007R-01
Guide to Securing Microsoft Windows 2000 © Group Policy
Network Security Evaluations and Tools Division
of the

Systems and Network Attack Center (SNAC)

Author:

Julie M. Haney

Updated: September 13, 2001

Version 1.1

<http://nsa2.www.conxion.com/win2k/guides/w2k-2.pdf>

[4] Report Number: C4-052R-00
Guide to Securing Microsoft
Windows 2000 © Group Policy:
Security Configuration Tool Set
Network Security Evaluations and Tools Division
of the

Systems and Network Attack Center (SNAC)

Author:

Julie M. Haney

Updated: May 17, 2001

Version 1.0

<http://nsa2.www.conxion.com/win2k/guides/w2k-3.pdf>

[5] Report Number: C4-053R-00
Group Policy Reference
Systems and Network Attack Center (SNAC)

Author:

David C. Rice

Updated: March 2, 2001

Version 1.0.8

<http://nsa2.www.conxion.com/win2k/guides/w2k-4.pdf>

[6] "20 tips for Exchange 2000 Migration" Windows 2000 Magazine, October 2001 P.94

[7] Building Enterprise Active Directory Services Notes from the Field

2000 Microsoft Press

ISBN 0-7356-0860-1

[8] Jason Fossen

Sans Institute Track 5 Windows 2000 Securing Internet Information Server 5.0

[9] Step by Step Guide to Internet Protocol Security (IPSec)

Posted: February 17, 2000

Section Entitled: Configuring an IPSec Filter List

<http://www.microsoft.com/windows2000/techinfo/planning/security/ipsecsteps.asp>

[10] Bruce Schneier Attack Trees Dr Dobbs Journal December 1999 "Attack trees provide a formal, methodical way of describing the security of systems, based on varying attacks. Bruce shows how you can use them to improve security by modeling attacks."

[11] Crypto-Gram Newsletter

September 15, 2000

by Bruce Schneier

Founder and CTO

Counterpane Internet Security, Inc.

schneier@counterpane.com

<http://www.counterpane.com/crypto-gram-0009.html>

This entire series is a highly recommended source of information.

Other URLs Referenced in this paper:

- <http://www.activewin.com/win2000/patches.shtml>
- <http://www.securityfocus.com/infocus/1520>
- <http://nsa.gov>
- <http://activestate.com/>
- <http://support.microsoft.com/support/kb/articles/Q246/2/61.asp>
- <http://www.microsoft.com/technet/security/tools.asp>

Addendum

As this was going to press, the following update came across the wire and seems an important enough of a software update to include this notice:

Source:

Security Alert Consensus

Number 128 (01.51)

Thursday, December 20, 2001

Network Computing and the SANS Institute

"*** {01.51.010} Win - MS01-058: Cumulative IE patch

Microsoft has released MS01-058 ("Cumulative IE patch"). This patch fixes all known security problems in Internet Explorer to date, including three new problems: the ability for a malicious Web site to execute arbitrary applications in IE 6; the ability to read files from the user's system; and a bug that could allow a Web site to trick the user into seeing a different file name in the download box.

FAQ and patch:

<http://www.microsoft.com/technet/security/bulletin/MS01-058.asp>

Source: Microsoft

<http://archives.neohapsis.com/archives/vendor/2001-q4/0053.html>"

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