



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

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

Interested in learning more?

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

Protecting Large Industrial Organisations from the new breed of Virus Attack.

C. Leigh Schouten

GSEC Option 1, Version 1.4b

6th November 2003

Introduction.

The aim of this paper is to help large industrial organisations prepare for the threat of the modern virus. Early viruses were simple and fairly harmless. They affected a few machines and were easy to clean. Examples include traditional viruses that came via email attachments or trojan programs. However in recent times we are seeing an increasing amount of more sinister "worm" type of viruses such as the "Blaster" and "Natchia" Worms. These worms can infect the world in a matter of hours choking networks and disrupting systems. This paper discusses a number of ways in which an organisation can prepare their systems to minimise the impact and recover quickly from a virus attack.

Today's viruses are now more intelligent and more penetrative than any previous ones. It is simply not enough to assume that having an up to date virus-checking tool will be effective. Businesses also need to have a management strategy covering all aspects of today's modern PC environment.

Virus Protection

A Virus is defined as a program or piece of code that is loaded onto your computer without the user's knowledge and runs against the computer users wishes. Viruses can also replicate themselves. All computer viruses are manmade. Viruses usually spread by being attached to an email message that fools the recipient into opening the attachment. The message might say "In response to your question " or "Microsoft patches required" When the user opens the attachment a malicious program contained in the attachment is run infecting the computer, often damaging data or programs on the computer.

A worm is a special type of virus that can replicate itself across networks, and use memory, but cannot attach itself to other programs. Worms take advantage of security flaws in programs. A worm can scan thousands of IP addresses looking for computers that are running a particular version of a program. For example, this worm (known as W32.Blaster, W32/Lovsan,

WORM_MSBLAST.A or Win32.Posa.Worm) scans the Internet and attempts to connect to systems that are RPC (remote procedure call) enabled

The first computer viruses appeared in the 1980's on the Apple II computer. However as more and more pc's were developed viruses were found on other types including Atari, Amiga and of course the IBM PC. The first IBM PC virus discovered was the Brain virus and this appeared in 1986. This was a boot sector virus that was run at boot up of the pc. Since their advent viruses have become increasingly clever and more complex.

Different methods have been employed by malicious code writers to deploy viruses. These include attaching to program files (EXE and COM file), Trojan horse's (virus hides inside another program and is deployed on activation), email attachments and more increasingly worms. Organisations can guarantee that they never get a virus by having no email and no network connectivity but this is impossible in today's society. If a computer is in use it is almost certain that at one time or other it will be subjected to some type of virus attack. This necessitates the use of virus protection in the form of anti-virus software.

The computer operating system software is now quite complex and as soon as companies like Microsoft (www.microsoft.com/security/virus) announce a security flaw virus creators will try to exploit it. As seen in CAIDA's Slammer Worm report this worm had the ability to double it's infection rate every 8.5 seconds. Within 30 minutes of it's release it achieved a scan rate of 55 million scans per second. Newer viruses spread at these ever-increasing aggressive rates. At the time of writing no detailed analysis was available for the more recent Nachia worm.

Virus protection software vendors such as McAfee (www.mcafee.com) and Sophos (www.sophos.com) to name a few can only ever hope to be reactive to a virus threat. With aggressive infection rate's businesses need to be able to react quickly. The faster a business can react and isolate a virus threat the less damage will be inflicted.

Virus software vendors are releasing pro-active products to identify vulnerabilities before a virus incident emerges. An example product is McAfee's Threatscan. Threatscan is a risk-assessment tool for performing scans against intranets, web servers, firewalls, and routers to identify virus related vulnerabilities on a network. The product is similar to Microsoft's Baseline Security Analyser in it's operation. The product searches the network to discover all the devices on the network. It can identify every machines operating system and it's relevant patch level in the organisation. It also finds machines vulnerable to specific virus infections and discovers all listening ports. Thus an organisation can perform a self check on how many of it's machines are actually fully protected to how many it thinks are protected (often thought about by how many machines it is patching). The threat scan product is not an antivirus product but rather a tool to work in conjunction with the traditional antivirus toolset.

Relying on anti-virus desktop protection is not enough. Ensure that antivirus scanning and file filtering is enabled on the organisation's mail gateway. Banning all file attachments except the common *.doc, *.pdf, *.zip etc is essential to prevent the simpler viruses coming through. Consider implementing secondary controls to stop malicious code spreading throughout the corporate network. These controls can be implemented both by policy as well as physically. Examples of secondary controls include: -

- Disallow Internet Mail i.e. Hotmail. Web based email pose threats by automatically opening and executing file attachments that can contain viruses.
- Browser Configuration (Netscape and Internet Explorer). Tighten the default security settings. This can include disabling ActiveX and Java controls. Ensure that the Web browser has the most up to date patches etc.
- Establish Mail Gateway content filtering. Steps include Applying filters against new viruses as they occur and filtering against older viruses should new variant's occur. Generic code or code that normally should not be in an email should be filtered as well.
- Refine operating system configuration i.e. disable unnecessary applications, change file associations to associate script files with a text editor.

Risk Assessment

The assessment of threat and management of risk to critical information assets within a business should be a continual process. The focus must be on protection of critical information assets.

There are a number of well-accepted best-practice techniques and frameworks for IS risk-management. One example is the OCTAVE methodology, developed by Carnegie Mellon in the US. Operation Critical Threat Asset Vulnerability Evaluation (OCTAVE) is a framework for managing and identification of security risks. It defines the evaluation method that allows identification of assets and threats to those assets. It allows management to understand what information is at risk, and then an appropriate Business Continuity Plan (BCP) and Disaster Recovery Plan (DRP) can be formed.

Techniques such as OCTAVE (amongst others) have been applied successfully at many large corporations in the Hi-Tech, Financial Services and Telecommunications sectors globally

The OCTAVE method differs from other risk assessment technology tools in that it focuses on organisational risk and practices. During the OCTAVE process people from both the business and IT department work together to assess the security requirements of the organisation. The OCTAVE method is totally driven by the business people with the analysis team managing the evaluation process and making any security decisions.

The analysis team as part of the process identify the information related assets and base their activity on those that are judged to be the most important to the organisation. They then focus on the interaction between those assets.

This method uses a three phase approach.

1. **Build Asset-Based threat profile.**- This uses internal staff to conduct an assessment and identify important information assets. It also identifies threats to those assets and what are the security requirements.
2. **Identify Infrastructure Vulnerabilities** – This is an evaluation of the information systems infrastructure. This uses the data gathered in phase 1 and examines the weaknesses.
3. **Develop Security Strategy and Plans** – All the risks are identified in this part of the exercise. The information gathered in the first two phases is evaluated and analysed. Risks are evaluated and prioritised based on the effect of the business impact to the organisation. A protection plan and mitigation strategy is created and developed for the higher impact risks.

The OCTAVE method has a defined beginning and end. Once a business has run the process for the first time this will set a “baseline”. The output being a defined set of action plans. When a period of time or a major event such as new infrastructure is introduced the process may need to be ran again. The OCTAVE method is not the only method available however it does look at risk assessment from an organisational perspective rather than a purely IT perspective.

A risk assessment is an essential tool in determining likelihood and consequences of a virus attack. It can be used for other purposes including disaster recovery and legacy replacement strategy etc. Once the risk assessment is completed the business will have an understanding of what systems are the most vulnerable, categorised systems by importance and have a plan of action to mitigate the risk.

Incident Response Team

When a virus incident is in progress it is too late to be defining roles and responsibilities for people involved in the management team. The business should have a predefined “Incident Response Team” (IRT). This means that individuals need to be in already assigned roles so that they know exactly what their roles and responsibilities are during the incident.

Some Defined roles include :-

- **Response Team Leader** – Makes decisions based on information supplied by the rest of the team
- **Process Support Coordinator** – Ensures established processes are followed and realigns the team if they digress.
- **Technical Advisor** – Supplies technical advice as required.

- **Administrator** – Records meeting minutes and all other administrative tasks. Ensures group is not distracted by these tasks.
- **Safety Coordinator** – Makes sure team follows safety guidelines and procedures. Tasks include time management etc.

A “war room” needs to be defined and set up. This work needs to be carried out in advance of any incident. This room needs to be large enough to contain all interested parties. The room must also have phone, fax and network connections.

The contact details for all named individual's need to be made accessible. If possible these individuals need to be contactable 24 hours a day 7 days a week. If individuals are unavailable a list of substitute staff should also be available.

The incident response team needs to look after resource planning. However if a threat can be resolved in 24 hours this is not required. In a large organisation such a short incident time however is highly unlikely. Steps need to be taken to ensure the availability of a skilled workforce for the entire duration of the emergency. People need to be rostered to ensure that resources will be available throughout the day and night. People will tend to try to work till exhaustion. Thus there needs to be numerous people to fill one role within the Incident Management Team; each individual must have a handover at predefined intervals.

The team needs to maintain an independent network connection to the outside world. Assuming the virus will affect the network an external ISP connection will be required. Such connections can be made by modems or standby DSL services can be established. This connection is essential for software updates, communications etc. The team needs to be able to communicate with the organisation. This can be done by use of faxes to predefined groups or communicating via PSTN and mobile phones.

The IRT needs to record all activities that have been undertaken as well as all decisions that have been made and implemented. These notes will be essential and useful during debriefing sessions after the crisis has passed.

Patch Management

It has become increasingly apparent that a Virus protection strategy is not enough. Virus management and patch management need to work in tandem. At the time of writing the MSBlast (Blaster) and Nachia Internet worms had hit the internet. They were two of the fastest spreading viruses in history. The Nachia worm exploited the same RPC (Remote Procedure Call) vulnerability as the Blaster worm. Yet both of these worms caused untold havoc in many organisations. Just as organisations were recovering from Blaster they were hit by Nachia. The following timeline shows a general outline of the history of the recent worm virus major events.

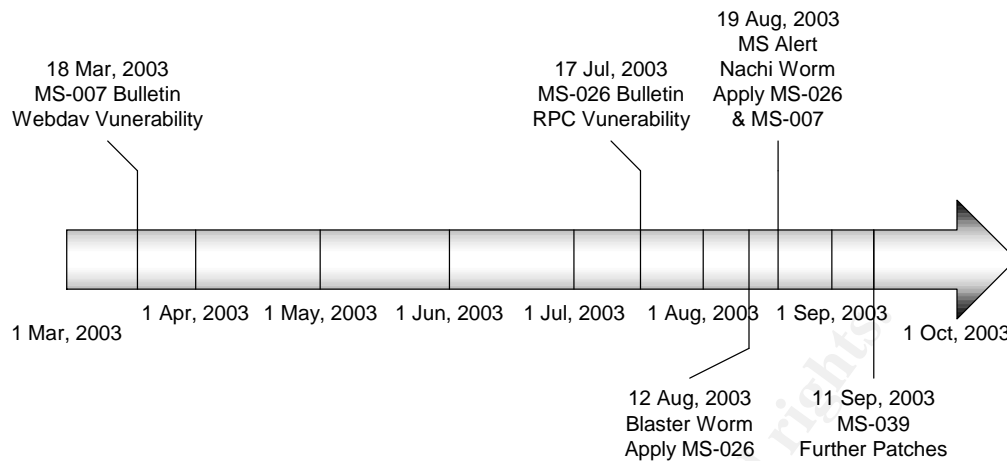


Figure 3. Event Timeline

As seen above an organisation with an effective patch management strategy could have avoided this particular infection or at least minimised the impact.

Patches can be classified by their level of criticality. Microsoft has outlined a severity rating in a patch management document (www.microsoft.com/security/whitepapers/patch_management.asp) which can be applied across all platforms including UNIX, Apple etc. The ratings are :-

Rating	Description	Time frame
Critical	A vulnerability whose exploitation that can allow a virus / worm to spread without any user intervention	5 days
Important	A vulnerability whose exploitation could lead to the compromise of user data or compromise system resources	10 days
Moderate	Vulnerability that can largely be guarded against by configuration or exploitation is difficult	30 days
Low	Vulnerability that is extremely difficult to exploit or impact is quite minimal	Normal attrition

Table 1. Vulnerability rating and their descriptions.

For vulnerabilities that fall within these ratings there should be an agreed time frame for patch completion. A suggested time frame is shown in the table. An agreed timeframe must be reached to assure that the patch management process is not drawn out and never completed.

Patches however cannot be deployed without a defined approval process within a large organisation. A balance between patch severity and business impact needs to be considered. A simple process is outlined below:-

1. Categorize computer systems by criticality.
2. Develop the patch deployment method.
3. Plan the patch rollout so that records are obtained and a back out strategy is developed. (Important should the deployment go wrong).
4. Obtain change management approval from the entire organisation.
5. Execute the rollout.
6. Perform ongoing maintenance to catch systems that have been missed.

One way to implement a patch management strategy would be to establish an update server within the organisation. Microsoft provide the Software Update Service (SUS Server) to implement a management strategy for Microsoft platforms. (www.microsoft.com/windows2000/windowsupdate/sus/) This package is free and was developed by Microsoft to give Network Administrators a way to automate patch deployment while maintaining control of the process. The advantage of bringing the server within the perimeter is that patches can be “staged” to allow administrators to fully test patches before they are deployed to the fleet. This system works in a similar way to the Windows updater server but allows the administrators to only deploy patches that they need to. By implementing the SUS server any tested and approved updates could be implemented and distributed quickly. By having each non-critical machine on the network continually polling the server any patch released could be very quickly deployed. However there will always be other critical machines that cannot be patched automatically. These machines will always require attention by support staff.

Other vendors provide methods for patch deployment. Two examples are IBM's Tivoli Configuration Manager (www-3.ibm.com/software/tivoli/products/config-mgr/) as well as Unicenter's TNG Software Delivery Option (www3.ca.com/Solutions/Product.asp?ID=234).

Organisational Structure

A traditional large industrial organisation of many PC's might have a network topology that relies upon a perimeter defence arrangement.

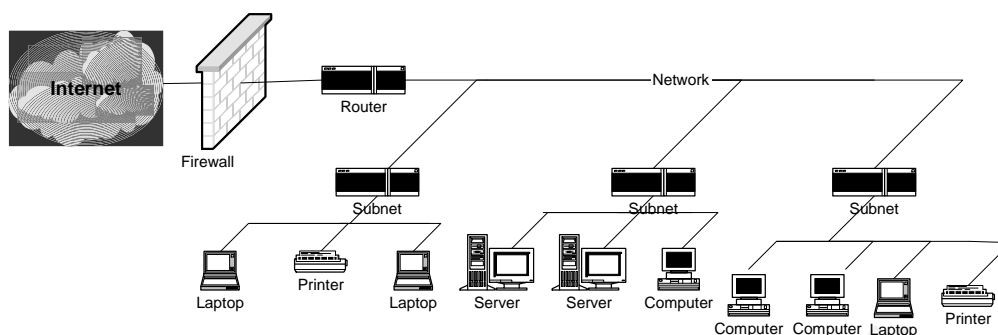


Figure 1. Traditional Corporate Network

There is only one single point of entry into and out of the corporate network. Once a worm subverts the main firewall the entire network is as vulnerable as if it was on the Internet itself. The ever-increasing number of mobile users present an increased risk of virus infection.

Large organisations are often spread across widely dispersed areas so you may need to consider a network segregation plan. The aim being to separate large parts of the Wide Area Network (WAN) into smaller manageable LAN segments. Each of these segments would be able to be isolated quickly from the WAN if they become infected. This step alone may contain a major outbreak to a defined area.

Knowledge of the network and its logical organisation is essential. A business must know what systems are attached to its network and the breakdown of systems on each subnet. PC's that perform critical functions to the business cannot necessarily have their network turned off without having other adverse effects on the business e.g. messaging from one pc to another.

Where possible computers should be categorised by the function they perform and how vital they are to the business. Critical computers should be placed onto totally separate networks (such as VLAN's) to standard office pc's.

Knowledge of the computer roles e.g. office, production, system, automation etc for each subnet is essential if software push tools such as SDO (Software Delivery Option) are to be utilised for rapid response to pc patching.

One option is to have each subnet fire walled off from each other.

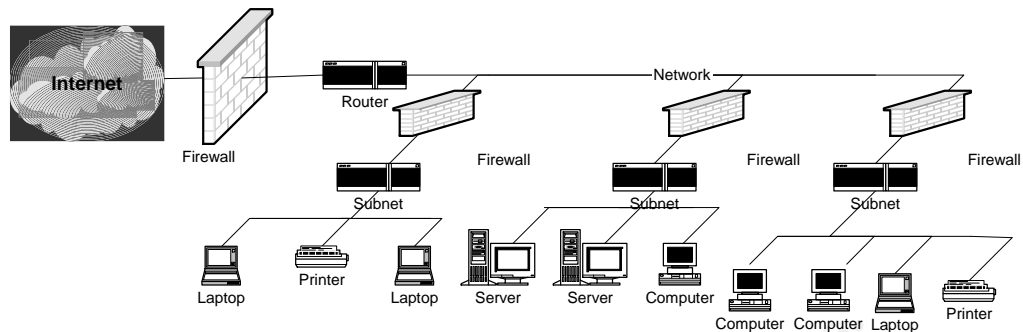


Figure 2. Firewalled Subnets

If malicious code does enter a particular subnet, isolating each subnet will keep the problem contained and isolated. Thus this would only affect a small percentage of the corporate pc's. However in a large organisation the task of managing many firewalls and maintaining the firewall rule set would become quite complex and time intensive.

Potentially, each computer could have it's own firewall however this makes network management difficult if not impossible unless the rule sets are centrally managed.

If it is not possible to isolate network segments areas from which an attack is more likely should be identified. For example, in a major organisation LAN there are many users with laptop computers. These computers can pose a high risk due to their mobility and promiscuous roaming nature. High risk areas could be placed on their own fire walled subnet.

A fundamental security control method is to remove unnecessary services from system configurations and so close any unnecessary open ports. This applies to all operating systems including Windows and Unix variants. This provides a basic security improvement without the additional complexity of a firewall. However if a port needs to be re-enabled each computer needs to be reconfigured.

Intrusion Detection Systems

Consider installing an Intrusion Detection System (IDS). Intrusion detection is the process of inspecting data for malicious activity. An intrusion detection system is designed to work in conjunction with a Firewall. While the firewall keeps traffic out and only permits certain inbound traffic. An IDS constantly monitors the IT system security and detects potential intrusions inside the network. Intrusion detection systems can assist the network administrators with notification when malicious or suspicious activity occurs.

There are 2 basic types of IDS:
 Host-based systems – HIDS
 Network-based systems – NIDS

1 . Host-based IDS

A host based intrusion detection system can monitor the file system of a computer for events such as file permission changes, privilege elevation etc. It's primary task is to audit the system logs. Certain personal firewall software act as host based IDS's. They can look at all network traffic flowing into and out of the machine. Examples include Internet Security Systems Black Ice Firewall (http://www.iss.net/products_services/blackice.php) and Tripwire (www.tripwiresecurity.com). Tripwire works in both Unix and Windows environments.

The HIDS looks at the data that flows into the system logs and can be configured in multiple ways to report on any activity. Examples could include multiple logon attempts or starting of unauthorised computer services. HIDS can also detect changes in system files. This is done by check summing system files against a checksum database. However the checksum database itself must be protected against attack. One way of doing this is to store the database on an external CD ROM.

2. Network-based IDS

A network based IDS analyses network packets. Network sensors are deployed at key network choke points. A NIDS constantly analyses all network traffic in its segment. Any NIDS needs to have sufficient processing power to handle the high data throughput volumes on any modern network. Suspicious data can be missed if the IDS cannot handle the throughput of a typical corporate network. A NIDS can notify the administrator the instant an attack is noticed.

There are two forms of an NIDS : Pattern Matching and Anomaly based. A pattern matching system contains prior information about specific attacks and network vulnerabilities. All inbound and outbound traffic is compared packet by packet against the signature database. If any pattern is matched the network administrator is immediately notified. To be effective though the pattern database must be kept up to date otherwise a new attack will not be recognised. An anomaly based NIDS learns a profile for normal network traffic. Once the profile is constructed any traffic that is outside this profile will be considered suspicious. Example of NIDS's include Cisco's Secure Intrusion Detection (www.cisco.com/warp/public/cc/pd/sqsw/sqidsz/), Enterasys's Dragon (www.enterasys.com/products/ids) or the open source Snort (www.snort.org)

Both the host and network based IDS's can work in conjunction to minimise the risk of a hostile attack. They can be set up in a reactive mode to respond to certain types of attack. However this can have the unwanted side effect of shutting down services when they detect a false positives. If the IDS is not tuned correctly it can generate large amounts of alerts causing large amounts of disruption. The IT department may ignore these warnings rendering the system ineffective. To be effective an IDS needs to be constantly maintained and tuned.

Vulnerability Scan

In a large organisation pc's are normally supplied by a 3rd party such as IBM. If machines are supplied with a preloaded set of software and configuration, (a Standard Operating Environment or SOE), ensure you understand the details of what is being supplied. The higher turnover of machines in a large organisations can introduce variation in models, components and software over time. There is the risk of "SOE drift". It is good practice to carry out random audits of new machines and assess them for security deficiencies. The freely available Microsoft Baseline Security Analyser or the CIS (Centre for Internet Security) auditing tool are useful for this task.

These tools can be used to determine both patch and service pack levels of supplied machines. It is also capable of determining if unneeded services are running, open shares, IIS vulnerabilities, Internet Explore security zones etc. MS Baseline security analyser is a tool that is capable of checking basic security of every MS system on the network and establishing a security baseline for the organisation.

After implementing all the security changes, shutting down unnecessary ports and services for every machine on the network there is only one thing left to do. Test it ! Performing a vulnerability scan provides several benefits.

- It acts as a check on all the work that has been done to secure the network.
- Acts as a check to confirm correct security decisions have been made.
- Establishes a network security baseline.
- Discovers previously unknown assets.

A vulnerability scanner checks to see if systems are active, which ports are open, Versions of services that are running and if the service has a vulnerability.

There are many tools that are available to do a vulnerability scan. Many are commercially available and expensive. However the freely available Nessus (www.nessus.org) is an excellent tool and is consistently one of the best. Other tools available include, MacAnalysis (www.securemac.com/macanalysis.php) (For Apple Macintosh) and RETINA (www.eey.com/html/products/Retina/index.html)

If it is not possible to do a vulnerability scan on the whole network then vulnerability scans should be at least performed on business critical machines. Before performing any vulnerability scan though, all relevant people in the organisation should be informed. Unleashing Nessus or any other vulnerability scanner on the corporate network could send protection systems into reactive mode to an attack. This could alone cause more major disruption to an organisation than if a real attack was underway.

Strengthened Corporate Network

Using the steps outlined in the previous sections an example of a strengthened corporate network is shown in figure 3. below.

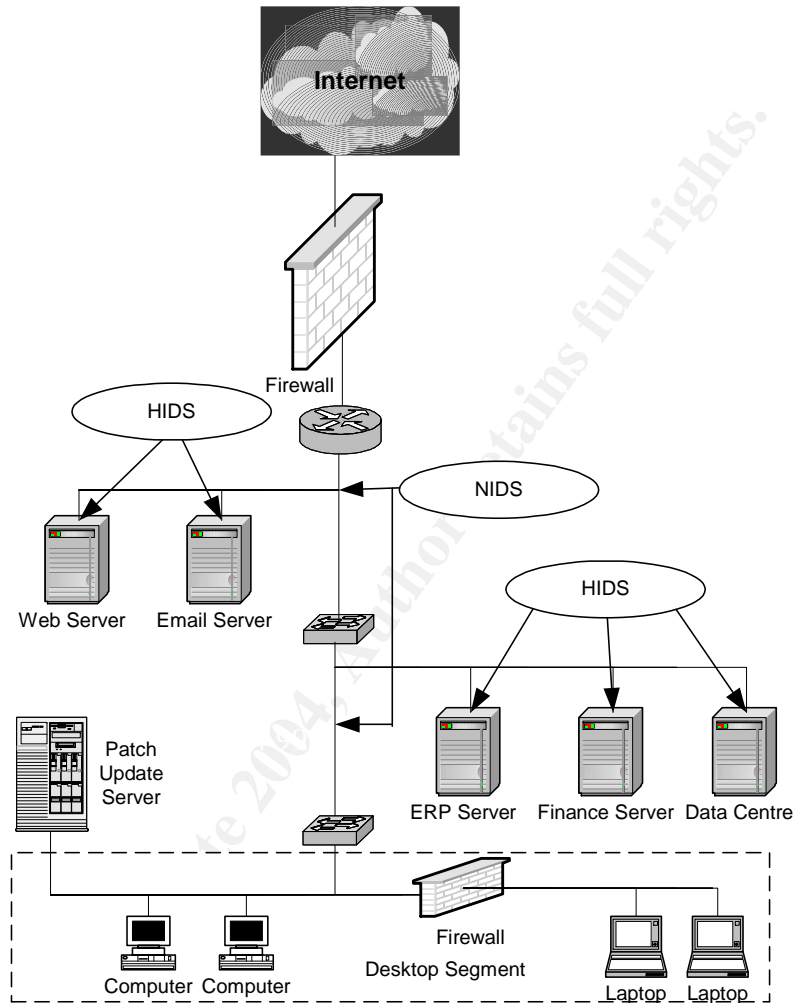


Figure 3. Strengthened Corporate Network

This network has controls in place for intrusion detection. Each business critical machine as determined during the risk assessment process has an host-based IDS installed on them. During this risk process mitigation plans will have been developed should key machines be attacked. The organisation will have a Incident Response Team with named individuals against specific roles. A network IDS is deployed at key points inside the network to guard against network attacks. All laptops are placed behind a firewall to contain possible infections originating from them.

To keep security levels of machines up to date a patch server has been deployed inside the network. This approach will allow for rapid patch deployment. The server approach will work for the majority of desktop

machines. However the more complex servers will still need to be attended to on a needs basis. Each of the desktop machines will have had unnecessary services shutdown as well as increased security settings in the browsers. If shutting down network ports manually is not achievable key desktop could have software firewalls installed with a centrally managed tool set.

Conclusion

This paper has endeavoured to help large organisations prepare for the next wave of virus attack. It has suggested multi faceted defence through the implementation of some basic policies, processes and controls. It has also suggested ways of preparing to for the next attack by forming an Incident Response team and defining roles for people. However there is no way to predict where the next “super virus” will come from. Implementing any or all of the steps above will help to minimise organisations vulnerability and enable it to respond and recover more quickly. However immunity from attack can never be guaranteed.

An organisation that has done it's homework and has well developed policies and procedures can react quickly to any incident and minimise any potential impact from the next virus or worm attack. One key factor is that the organisation needs to know the composition of it's computer fleet. It is impossible to protect any environment if the business does not know what is in it. However as shown from previous attacks any response needs to be swift and decisive due to the increasing aggressiveness of today's modern viruses. It must also be realised that the risk is not only to the “Microsoft” machines. Unix and other operating systems have had their own share of attacks over the years. These computers must also be covered as part of the organisations overall defence.

However the biggest risk of all is a lack of awareness and apathy to the risk of virus attack. Management needs to be constantly reminded of the risk otherwise or it will be soon forgotten until the next attack. If this happens the results can be catastrophic.

References

- Moore, David. Paxon, Vern. Savage, Stefan. Shannon Colleen. Staniford, Stuart. Weaver, Nicholas. The Spread of the Sapphire/Slammer Worm, Jan 2002
URL: <http://www.caida.org/outreach/papers/2003/sapphire/sapphire.html> (23 Sep. 2003)
- Tippet, Peter. Building “Synergistic” AV, May 2002.
URL: <http://infosecuritymag.techtarget.com/2002/may/synergisticav.shtml> (18 Sep. 2003).
- Alberts, Christopher. Dorofee, Audrey. Stevens, James. Woody, Carol. Introduction to the OCTAVE Approach, August 2003.
URL: http://www.cert.org/octave/approach_intro.pdf (2 Nov. 2003)

Dworakowski, Wojciech. Why is a firewall alone not enough? What are IDSes and why are they worth having? Aug 2002.

URL:

http://www.windowsecurity.com/articles/Why_is_a_firewall_alone_not_enough_What_are_IDSes_and_why_are_they_worth_having.html (30 Sep. 2003).

Radhakrishnan, Ramesh. A Patch Management Strategy for the Solaris Operating Environment, January 2003.

URL: <http://www.sun.com/solutions/blueprints/0103/817-1115.pdf> (12 Oct. 2003).

Houghton, Ken. Vulnerabilities and Vulnerability Scanning, July 2003.

URL: <http://www.sans.org/rr> (24 Sep. 2003).

Smith, Danny. Forming and Incident Response Team, Jan 1995.

URL: <http://www.auscert.org.au/render.html?it=2252&cid=1920> (7 Oct. 2003).

Trickle, Ian. Data Integrity – the unknown threat, Dec 2002.

URL: <http://www.itsecurity.com/papers/tripwire1.htm> (15 Sep. 2003)

© SANS Institute 2004, Author retains full rights.