Global Information Assurance Certification Paper

Protecting the Desktop Against Blended Threats with Symantec Client Security

GIAC Security Essentials Certification (GSEC)
Ver 1.4b Option 1

Daniel Brown Feb 4th 2004

Table of Contents

Abstract

The desktop threat is changing. Over the last decade we have witnessed the evolution of malicious code from boot sector and file based viruses through to polymorphic viruses and self replicating network based worms. More recently we have seen hybrid attacks, consisting of viruses with the ability to execute like a worm and also scan for compromised hosts to exploit known vulnerabilities.[1] As this threat profile evolves, so do the tools required to protect an organisations desktops and servers from these attacks.

The aim of this paper is to examine the functionality of Symantec Client Security, paying particular attention to features that constitute a layered security approach, taking the traditional desktop antivirus model and adding client firewall and intrusion detection capabilities. My current work schedule involves advising and implementing antivirus solutions for corporate organisations and I have leveraged my knowledge in this area to assist in writing this paper. I have included a summary of the W32.Blaster.Worm as a blended threat example to help justify the need for 'security in depth' at the desktop.

Introducing blended threats

Blended threats are attacks that combine the characteristics of viruses, worms, trojan horses and mobile malicious code with software vulnerabilities.[1] These types of attacks have been increasing in terms of frequency and complexity. By utilising multiple vectors of attack, blended threats can spread rapidly and cause widespread damage [1]. As these types of attacks evolve and become more sophisticated, the level of knowledge required to create these exploits is becoming less, due to the the increasing availability of automated hacking tools, and advice on hacking software and techniques being widely published on the Internet.

Examples of propagation vectors used by blended threats

Traditional virus and worms
Self-replicating mass mailers (using SMTP engine)
Use of network shares
Port scanning for known compromised hosts with vulnerabilities
Remote access trojans

Traditional antivirus protection

As blended threats evolve and become more complex and sophisticated, so too are the measures an organisation will need to put in place to successfully protect against these threats. Traditionally an organisation might have deployed an antivirus solution and manage daily or weekly definition updates to reduce the risk of infection from a possible virus outbreak. With this approach, perhaps using a management console, an administrator can define and enforce an antivirus policy (including a definition update policy) and rest assured that as long as new definitions are being deployed in a timely manner, the risk of exposure to a new attack is being managed effectively.

Antivirus protection is effective for viruses and trojan horses whereby a signature can be written to detect the presence of malicious code and be deployed to update an incremental database. Antivirus software can also include a measure of heuristics to pro-actively detect a possible virus based on certain file behaviour and characteristics. Heuristic technology used in this way can create false positive results which need to be managed accordingly.

1 Symantec Internet Security Threat Report Trends for January 1, 2003 – June 30, 2003

Where antivirus protection becomes less effective is when the threat is undetected by the antivirus software installed (most malicious code writers will test their creations on commercial antivirus software before release). This situation can yield different results according to the severity of the infection i.e virus/worm or blended threat and the measure of protection installed on the infected machine in the first place. If an antivirus product fails to detect this attack, the computer could be susceptible to the payload of the malicious code in question. If the attack is a blended threat, the implications for the organisation could be disastrous. However, if the target computer had a client firewall with intrusion detection capabilities in addition to antivirus there is a possibility the attack could be contained if not prevented entirely.

As mentioned previously, traditional desktop antivirus updates comprise of a process of deploying signatures in a timely manner. The process works well when automated. The antivirus vendor will release a daily or weekly update which can be retrieved by a managed antivirus server (and tested if necessary) and then deployed across all the desktops in the organisation. As long as the desktop definitions are kept up to date the desktops will be protected from all known viruses.

The disadvantages of this approach are the reliance on vendor posted definitions. If the definitions are corrupt or internal testing reveals a false positive with a trusted application the organisation relies on, a window of exposure is introduced into the equation. Managing a virus outbreak effectively relies on the administrators ability to deploy successful definitions. With this process in doubt, the window of exposure and cost of attack cannot be defined or contained. This scenario would also apply to a completely new threat that has yet to be diagnosed and therefore has no signature associated with it.

Introducing Symantec Client Security

Symantec Client Security builds on the traditional managed antivirus model and adds client firewall with intrusion detection capabilities. With Symantec Client Security an administrator can manage client firewall policies built with specific groups of clients in mind, and enforce these policies using a management console to ensure that managed clients connecting to the organisation do not pose a security risk.

Symantec Client Security is based on the following components;

Symantec Antivirus Corporate Edition Client (SAV CE Client)
Symantec Antivirus Corporate Edition Server (SAV CE Server)
Symantec System Center (SSC)
Symantec Client Security Administrator (SCFA)
Symantec Client Firewall (Symantec Client Firewall)
Symantec Packager
Symantec LiveUpdate Admin Utility

Symantec Client Security clients include installations of both the Symantec Antivirus Corporate Edition Client AND the Symantec Client Firewall. These two applications install separately (or in sequence using the Symantec Packager). Once installed they both use the same client processes, and are managed using the Symantec Antivirus Corporate Edition Server (which is a Symantec Client Security server when managing Symantec Client Security clients).

Symantec Client Security integrated client components

The Symantec Client Security client features antivirus and client firewall with intrusion detection. The antivirus and client firewall components are integrated i.e if the client firewall detects an incoming file that it perceives to be infected with a virus it can execute a scan of that file using the antivirus component and proceed to block access. Symantec Client Security also includes a set of intrusion detection signatures that have been optimised for a windows client; these signatures can be enabled or disabled as part of the policy. The Symantec Client Security client also supports trusted and restricted zones. Zones allow the administrator to identify computers that will be excluded from the client firewall policy (trusted) and computers that are completely blocked from access (restricted).

Implementing a Symantec Client Security architecture

Symantec Client Security is managed by running the Symantec System Center console. An administrator needs to create at least one server group (a container of managed Symantec Client Security servers/clients). This arrangement of server groups together represent the Symantec Client Security domain, which is a logical domain not dependant on any existing networking infrastructure. A Symantec Client Security server group must contain at least one computer with the SAVCE server component installed. This server will be configured as the primary server for the group. Further SAVCE servers can be added to the group. These servers are called secondary servers. Primary and secondary servers with client responsibilities are called parent servers. There is no limitation to the number of SAVCE servers that can reside in a group or the number of server groups that can be created to represent the Symantec Client Security logical domain. A Symantec Client Security server can support up to ten thousand clients per server.

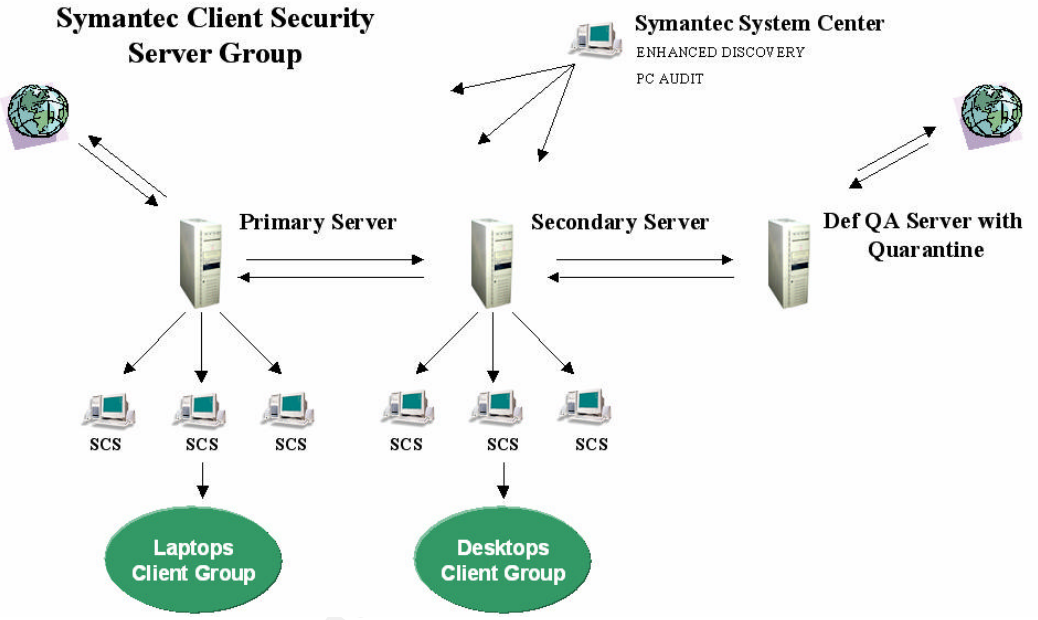

Figure 1 Managing a Symantec Client Security Server Group [2]

Symantec Client Security policies

The Symantec Client Security antivirus policy is defined (around 5k in size) using the Symantec System Center for the primary server of the group. The primary server automatically updates all other servers and clients in the group when this policy file is changed or updated. The Symantec Client Firewall policy can be created on a local client firewall installed computer and imported into the Symantec Client Firewall Administrator to allow for administrator only options to be added. The Symantec Client Firewall Administrator program can be run from a drop down menu within the Symantec System Center or as a standalone application. Both Symantec Client Security policies are enforced and maintained from within the Symantec System Center.

In addition, client groups can be created (within the server group) to own an exclusive Symantec Client Security policy. This flexibility allows for specific sub policies to be defined and associated with a client group independent of the group policy created for the primary server.

2 Diagram based on principles taken from Symantec Client Security Administratrator's Guide

The primary server includes the Alert Management System (AMS2) component to manage virus alerts for the server group. If there is a virus event or policy violation on a remote client computer, the event will trigger an alert on the client, which is forwarded to the primary server resulting in a configurable alert action e.g SNMP trap, SMTP message etc whereby the Symantec Client Security administrator would be notified.

The Symantec System Center utilises the Microsoft Management Console (MMC) framework and can be installed on any Microsoft Windows NT4, 2000 or XP computer. The Symantec System Center is used by the administrator to manage all aspects of the Symantec Client Security policy ; including client installation, policy rollout, changes to policy, definition update configuration, configuration of alerts and quarantine management. The Symantec System Center can also be used to view virus history, client firewall alerts and intrusion attempts.

In the event of an outbreak, a quarantine server can be configured to receive virus submissions from any managed client and automatically forward the infected file(s) to Symantec. Symantec can respond with a special definition (uncertified) that can be tested on a definition Q/A server and then deployed to the managed environment using the Virus Definition Transport Method (VDTM). VDTM is one of the antivirus definition update processes available for configuration using the Symantec System Center.

In addition to, or as an alternative, an internal LiveUpdate Server can be used to act as a central repository for definition updates. This computer can be scheduled to contact Symantec's FTP/HTTP servers to check for new updates. Symantec Client Security managed clients and servers can be configured to pull new definition updates from this server, based on a schedule configured using the Symantec System Center. All Symantec Client Security antivirus definition update mechanisms are reviewed in this paper.

Installing Symantec Client Security servers

The Symantec Client Security server component can be installed using the following methods:

From software CD's.
Using the Symantec Packager
Using the 'AV Server Rollout' option within the Symantec System Center

Please refer to Symantec documentation for server system requirements [3]

Installing Symantec Client Security clients

After installing and rebooting the server, Symantec Client Security clients can be deployed. Symantec Client Security clients can be installed using the following methods:

Pushed directly from the Symantec System Center (NT/2000/XP)
Web based install
From a Symantec Packager MSI package
Login script
Using 3rd party deployment tools
From software CD's

Please refer to Symantec documentation for client system requirements [3]

Once the administrator has deployed the Symantec Client Security clients, the Symantec System Center can be used to define and enforce the Symantec Client Security policies for the organisation. The Symantec Client Security policies include antivirus and client firewall with intrusion detection.

3 Symantec Client Security Installation Guide

The Symantec Client Security hierarchy consists of the server groups containing servers and their connected clients. Server groups also contain client groups. Client groups are logical groupings of clients. Clients can be placed into client groups to inherit a Symantec Client Security policy which is exclusive to that client group. In the example used in this paper, I have created two client groups, one called 'laptops' and one called 'desktops'. Symantec Client Security clients that are remote access users and typically leave the office may require a more restrictive policy and can be placed in the 'laptops' group. Symantec Client Security clients that remain connected to the Local Area Network (LAN) will benefit from other security measures in place e.g perimeter security in the form of a firewall, these clients will be placed in the 'desktops' group with fewer restrictions applied.

Symantec Client Security client and server communication

When the Symantec Client Security client services load, the SAV CE service will attempt to contact its parent server. The client will connect to a listening parent server port using the User Datagram Protocol (UDP). The client and parent perform a UDP handshake and the client is processed and added as a registry key on the parent server. The parent server will update this registry value each time a client successfully contacts or 'checks in' with it. During this process of communication the parent server will check the Symantec Client Security client's policy time and date and will copy down a new policy if required. When a SAV CE secondary server's service is loaded, it too will attempt to contact the primary server for the group and receive updates in the same way.

Installing the Symantec System Center

The Symantec System Center is the management console for Symantec Client Security. The Symantec System Center snaps in to the Microsoft Management Console (MMC) and allows the administrator to manage all aspects of Symantec Client Security. The Symantec System Center can be installed on any Windows NT/2000/XP workstation or server and can be installed on as many different computers as desired. Each computer running the Symantec System Center will be running the Symantec System Center Discovery Service. When the Symantec System Center is loaded it performs a network discovery. The network discovery uses a random selection of unreserved UDP ports to find Symantec Client Security primary servers. The network discovery can be tailored to discover all networks or localised for a local subnet. Once a primary server is located its registry values are read (including all connected clients and second ary servers in the server group). This process is repeated until all primary servers are located within the scope of the network discovery performed (see Figure 2).

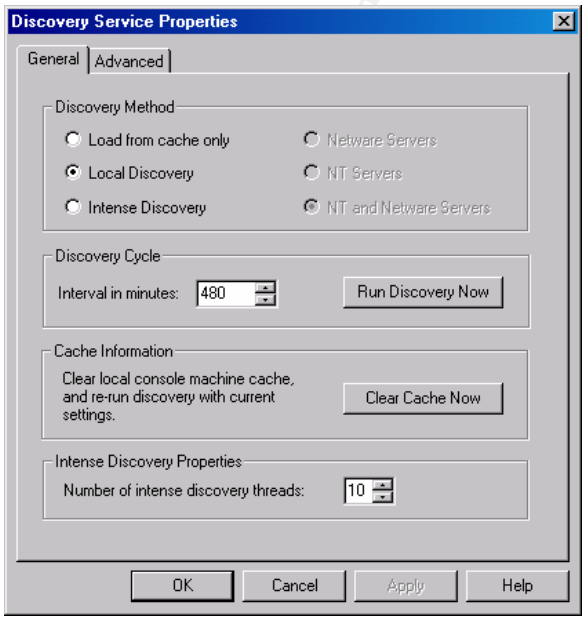

Figure 2 Using Symantec System Center to discover primary servers

Once the network discovery is complete the console view is populated with all properties of the Symantec Client Security logical domain. Equipped with the server group password the administrator can proceed to unlock a server group and perform management operations. The network discovery will find all clients that have 'checked in' with their parent server and all secondary servers that have 'checked in' with their parent servers respectively. The Symantec System Center supports three different views to assist with Symantec Client Security management – default console view, client firewall view and antivirus view. Each view displays information about the Symantec Client Security clients and servers according to the selected level in the system hierarchy.

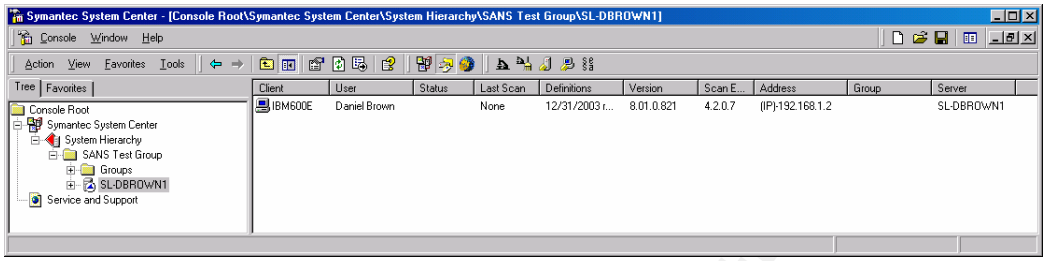

Figure 3 Using the antivirus view from within the Symantec System Center

Using the Symantec System Center for troubleshooting

Because client and server communication uses UDP there may be occasions when clients that fail to connect with their parent server will not appear in the console view. The Symantec System Center includes functionality to help troubleshoot this scenario with the 'find computer' and 'pc network audit' features. Computers can be located that have failed to communicate with their parent servers or are running an alternative antivirus product. Unmanaged or misconfigured computers can constitute a potential point of entry for malicious code, so this feature can prove invaluable for finding unprotected computers and eliminating the risk of exposure they present.

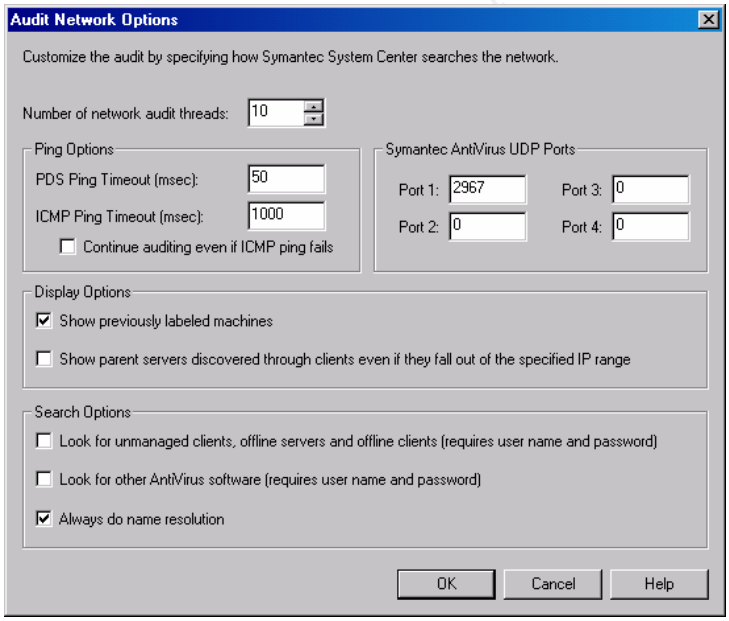

Figure 4 Using the Symantec System Center to configure Audit Network Options

Configuring antivirus policies

Symantec Client Security antivirus policie s can be defined and rolled out using the Symantec Sys tem Center. All aspects of the ant ivirus policy can be configured in this way. Policy settings can be managed for groups of servers and clients, client groups or just clients, depending on which level you are at in the system hierarchy. Antivirus policies comprise all scanning options, inc luding manual, scheduled and real-time scanning.

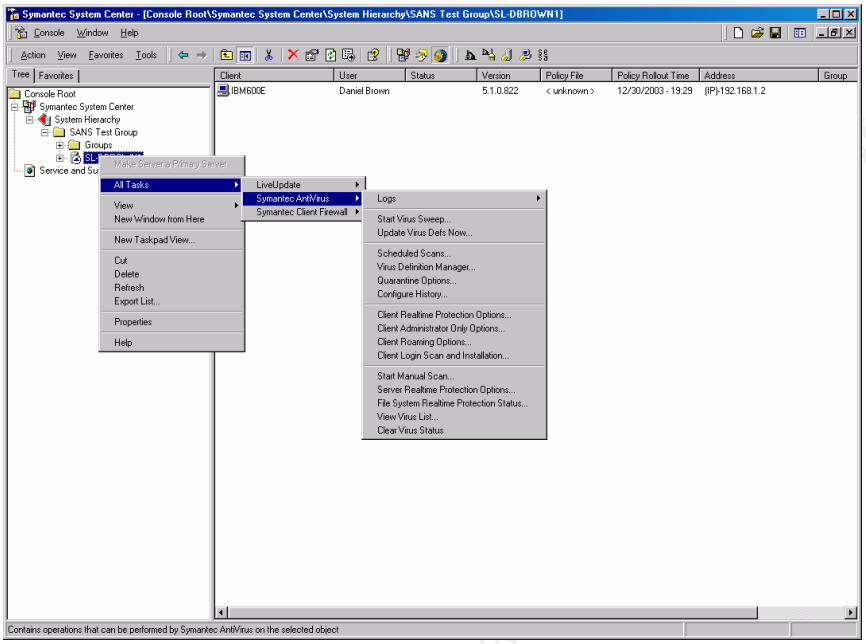

Figure 5 Using the Symantec System Center to configure antivirus settings

How antivirus policies are rolled out to client and servers.

Whenever a policy change is made in the Symantec System Center, t he change is written down to clients and servers in the form of the GRC.DAT file. This file is around 5k in size, stored on the parent server and contains all the policy data. When a client or server computer 'checks in' with its parent server it will comp are GRC.DAT file properties, and if the contents of this file on the computer 'checking in' is older than that of the parent server, a new GRC.DAT file is copied down. In the same way, whenever the administrator makes a policy change from the Symantec Syst em Center, the server GRC.DAT file is copied down to the Symantec Client Security client or server computer in real - time.

Configuring virus scans

An important part of the antivirus policy is the configuration of virus scans. Symantec Client Security supports manual scans, scheduled scans and real -time scanning. All three scan types share similar configuration propertie s. The real-time protection option monitors the file system continuously and will use system resources accordingly. Figur e 6 shows the real-time protection configuration screen.

Regardless of the scan type, Symantec Client Security performs one of two actions on discovering a virus, both for macro and non-macro viruses. The primary action might be to clean. If this is not possible a secondary action takes place i.e. the vi rus can be deleted, quarantined locally or just logged. Policy options can be locked down completely to ensure the integrity of the desired settings once deployed.

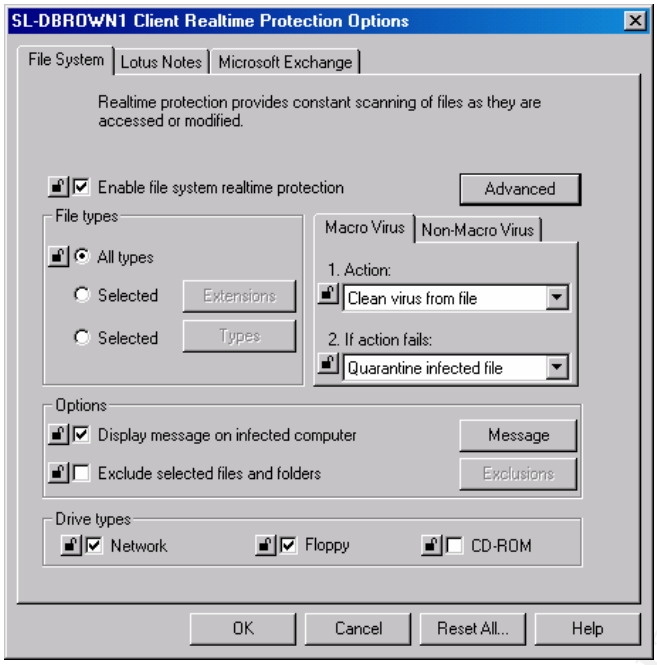

Figure 6 Using the Symantec System Center to configure real-time protection

Configuring antivirus definition updates

Updating antivirus signatures/definitions is a crucial part of providing effective protection for Symantec Client Security managed clients and servers. Symantec Client Security supports a numb er of different mechanisms to achieve this goal. The update methods described below and illustrated in Figure 4 are not mutually exclusive and can be combined together to provide redundancy and failover in a virus outbreak situation.

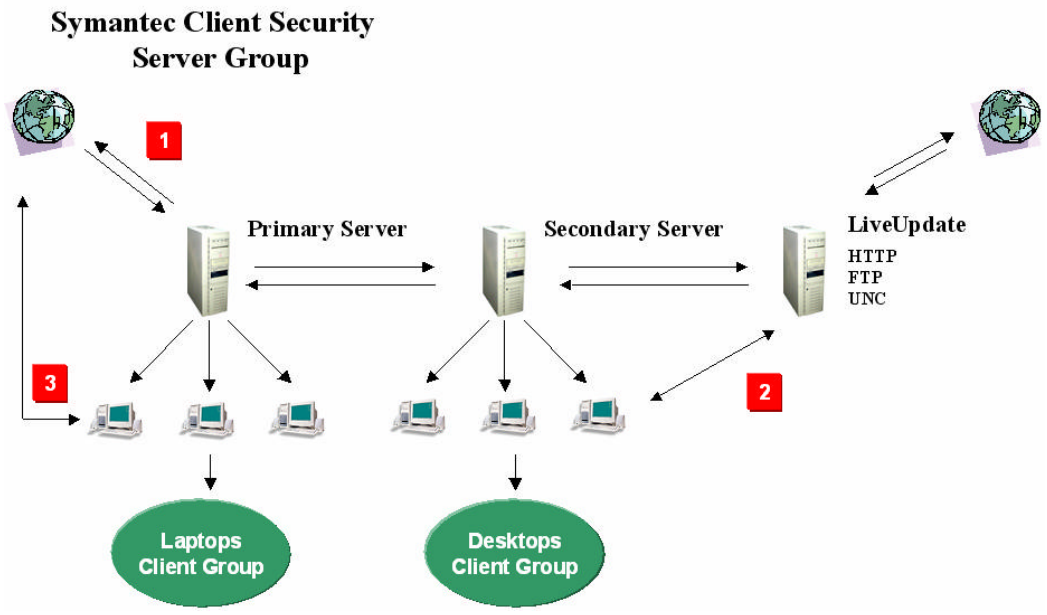

Figure 7 Symantec Client Security supports multiple definition update methods [2]

2 Diagram based on principles taken from Symantec Client Security Administratrator's Guide

Method 1 using Virus Definition Transport Method

Symantec Client Security supports the 'Virus Definition Transport Method' (VDTM). VDTM allows the administrator to allocate a Symantec Client Security primary server to retrieve definitions from a Symantec LiveUpdate server on the Internet. If a Symantec Client Security primary server has this responsibility it is called a master primary server. The master primary server will retrieve definitions manually or on an automated schedule. Once these definitions have been downloaded they are automatically copied down to all clients connected to the primary server and all secondary servers in the group. The secondary servers will then proceed to copy definitions down to their clients respectively; VDTM definitions are based on increments of data and are operating system specific. If the definition update process is kept frequent, these definitions can be kept as small as 15k daily or 80k weekly.

Method 2 using Internal LiveUpdate

Symantec Client Security supports the deployment of an internal LiveUpdate server. The LiveUpdate server acts as a repository and publishes the definitions for download using the following transports: HTTP/FTP/UNC. As part of their Symantec Client Security policy, clients will attempt to pull their definitions from the internal LiveUpdate server on a scheduled basis. This client LiveUpdate schedule has a number of configuration options available including 'missed event' management and 'randomisation' to stagger client updates and assist in managing network bandwidth. LiveUpdate definitions are currently available as weekly increments and are between 150k/300k in size. Symantec Client Security can support multiple internal LiveUpdate servers as part of the configuration for redundancy purposes.

Method 3 using Symantec LiveUpdate

Symantec Client Security supports definition updating using Symantec's HTTP/FTP LiveUpdate servers. If the Symantec Client Security policy includes this option, the client will attempt to connect to the Internet to retrieve its definitions. The same 'missed event' and 'randomisation' options apply and the Symantec LiveUpdate server can also be added to an existing internal LiveUpdate server configuration. This would introduce a measure of failover as if the internal LiveUpdate server fails, as long as the client has an Internet connection, definitions can be downloaded successfully from Symantec LiveUpdate servers.

All of the above methods can be used together without duplication of resources. All Symantec antivirus definitions are processed in a specific directory on the client, if one method has been successful, the other methods will not execute.

Symantec will test antivirus definitions before publishing them for download, but it is considered good practice to test newly downloaded definitions in house prior to rolling out or publishing them for client access. It can be useful to test definitions on standard desktop builds to detect any issues before users experience them. If the worst happens and a defective definition is deployed, VDTM includes a feature that allows the administrator to 'rollback' the defective definition and deploy a different one.

Updating Symantec Client Security intrusion detection signatures

The IDS signatures can be updated for Symantec Client Security using LiveUpdate, these signatures are usually updated quarterly.

Symantec Client Firewall policies

Symantec Client Firewall policies can be configured for each individual build or configuration that exists within the organisation. Policies can be created with specific network enabled application rules in mind or system wide rules set to allow or deny certain protocols and ports to communicate. Within the Symantec System Center, these policies can be deployed to a whole server group or to individual client groups. In the example used in this paper the organisation has a number of laptop users that work remotely. These users do not benefit from perimeter security in place on the LAN, so a more restrictive policy can be created for these users in line with the company firewall policy. By placing the laptop users in a separate client group, the Symantec System Center can be used to enforce the laptop client firewall policy just to those users. The example also includes another client group called 'desktops' allowing the administrator to deploy a client firewall policy which is less restrictive, to reflect security measures that are already in place on the LAN.

Creating the Symantec Client Firewall policy

Symantec Client Firewall policies can be built on the client computer, and then imported into the Symantec Client Security Administrator so administrative functions can be applied. To first build a firewall policy, the administrator needs to select a client computer the new firewall policy will represent and install the client firewall application to that computer.

Internet Access Control Application Scan

Once Symantec Client Firewall is running, an Internet Access Control (IAC) Application Scan can be performed. The application scan (see Figure 8) will search local drives for executables, and try to match them with potential rules (pRules). The application scan will discover all network-enabled applications that Symantec have profiled for inclusion into Symantec Client Security and create potential rules (pRules) for these applications. pRules are application rules, which are not written to the registry until they are activated, thus saving registry space on the client. When the application scan is finished, lists of network-enabled applications are listed and the administrator can include, exclude or modify these rules as desired.

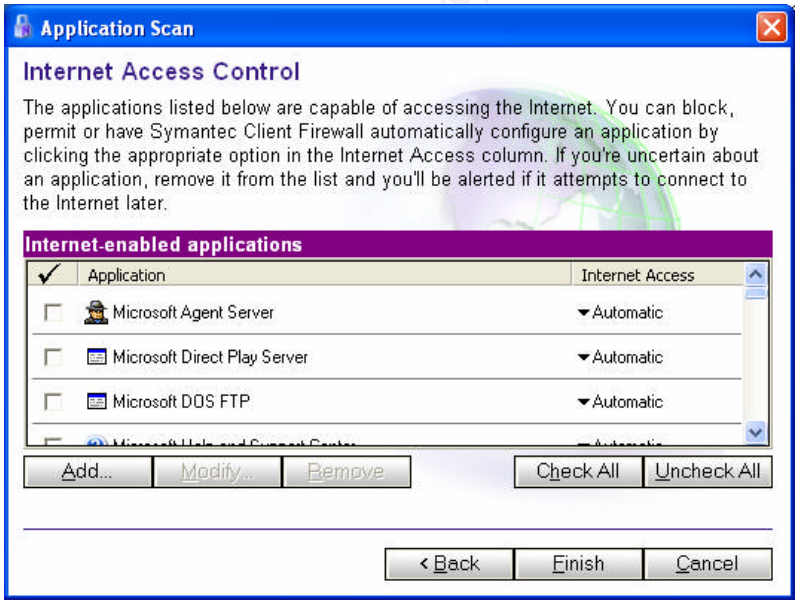

Figure 8 Internet Access Control Application Scan

Once the Symantec Client Firewall policy has been created, it can be imported into the Symantec Client Firewall Administrator for further manipulation before being rolled out to groups of clients using the Symantec System Center.

Loading Symantec Client Firewall policies into the Symantec Client Firewall Administrator can be achieved in different ways. The administrator can import the policy from the active client (that is a computer with the Symantec Client Firewall and Symantec Client Firewall Administrator installed). This method will import the current configuration settings. Another method would be to open a previously saved policy file.

Symantec Client Firewall policy files

Symantec Client Firewall policy files are saved in one of two formats:

.xml	The .xml format saves all configuration data with the exception of Intrusion Detection System (IDS) exclusion settings. For this reason the .xml is smaller than the .cfp format.
.cfp	The .cfp format saves all configuration data into a compressed policy file including all the IDS exclusion data. For this reason the .cfp file is larger than the .xml but is necessary if the intended Symantec Client Firewall policy contains IDS settings.

Using FIO.EXE to manually import/export client firewall policy files.

Symantec supports a file import/export utility for reading and writing Symantec Client Firewall policy files on the local computer. This command has to be executed from the local Symantec Client Firewall directory and uses the following command syntax:

Fio.exe i |o [path] package filename

The I parameter will import a specified policy and the o will export current configuration settings and create a new .xml file

Symantec Client Firewall rules

Symantec Client Firewall rules are broken down into three categories:

System Wide Rules	System wide rules are based on protocols and services i .e. TCP, UDP, ICMP, DNS, NETBIOS.
Application Rules	Application rules can be added according to the configurati on on the target client. Application rules can be added on demand. An application rule can be tailored to allow safe execution of only trusted business applications.
Trojan Rules	Trojan horse rules examine network communication for characteristics of known malicious programs.
pRules	pRules are application rules that are not written to the registry until they are processed for the first time. pRules can include match criteria to ensure the executing application is genuine. The match criteria can include an encrypted digest value from the trusted application to ensure validation and integrity.

If a pRule is triggered and the associated application is run, the rule dat a is saved to the registry as an application rule and treated as an application rule from t hat point onwards.

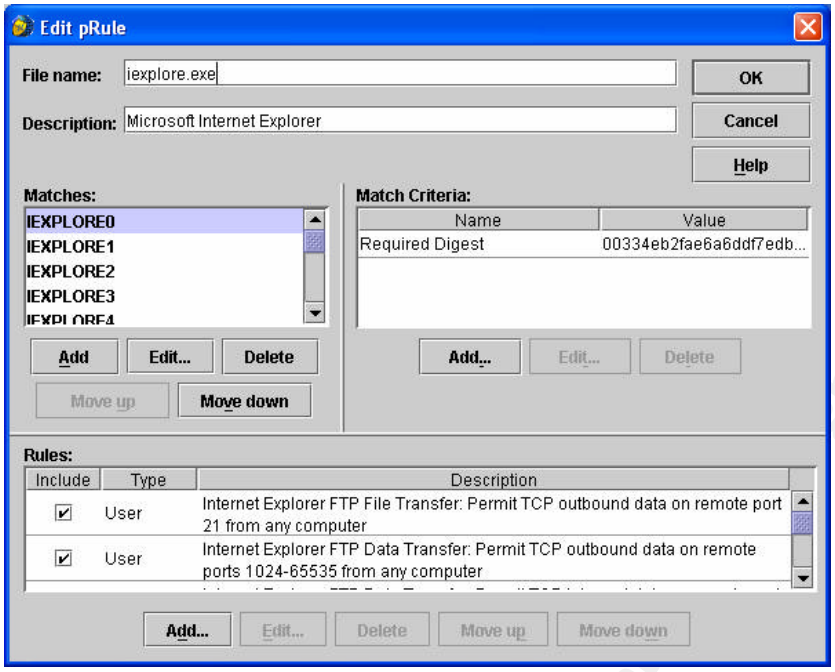

Figure 9 Using Symantec Client Firewall Administrator to edit a pRule

Symantec Client Firewall access permissions

Symantec Client Firewall supports three types of user access permissions:

Admin	This user has the ability to add, modify and delete all rule sets and perform administration functions.
Normal	This user can has the ability to add, modify and delete user rule sets, view firewall data and set privacy control settings.
Restricted	This user has no ability to change any rules or settings.

Rule processing order

Admin	System Wide Rules
Admin	Application Rules
User	System Wide Rules
User	Application Rules
Admin	Trojan Horse Rules
User	Trojan Horse Rules

Rules are processed from top to bottom based on the category. System wide rules are always processed first, application rules next, followed by trojan horse rules. The rules are also processed by user type as shown above. Admin rules always take priority over user rules except for trojan horse rules.

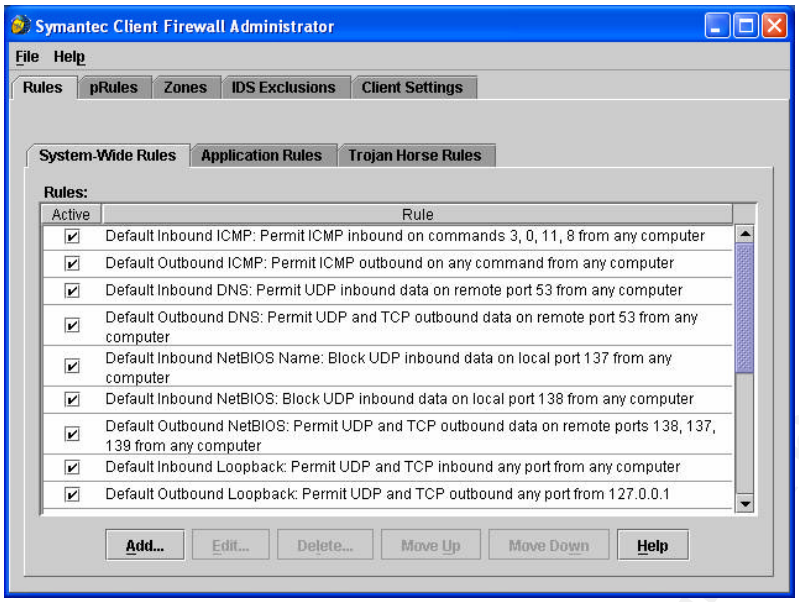

Figure 10 Using the Symantec Client Firewall Administrator to view System Wide Rules

Intrusion detection

The intrusion detection feature allows Symantec Client Securit y to detect inbound port scans based on a list of attack signatures appropriate for a windows client. These network based attack signatures are updated periodically. When Symantec Client Security detects incoming activity it suspects to be a port scan it can be configured to automatically block the originating IP address (using a fe ature called AutoBlock). Symantec Client Security intrusion detection signatures can be included or excluded as part of the policy.

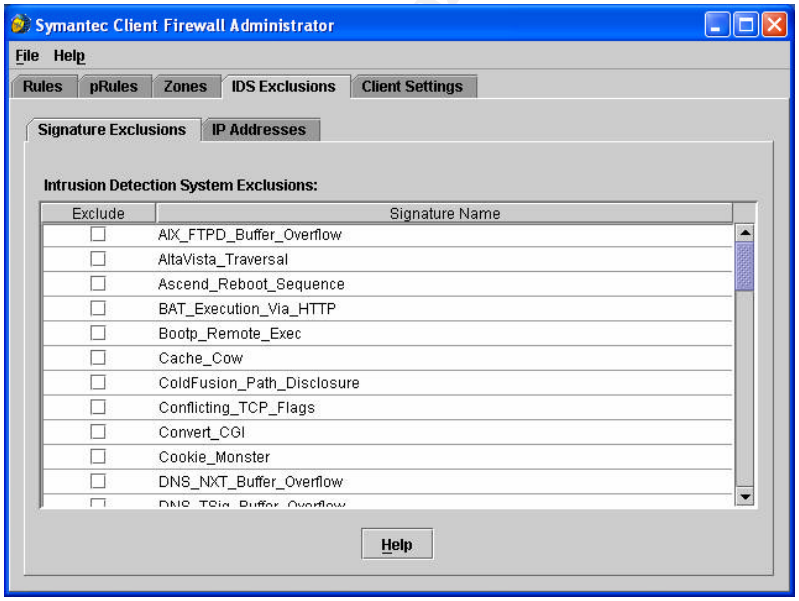

Figure 11 Using Symantec Client Firewall Administrator to exclude IDS signatures

Internet zone control

Symantec Client Security also allows the administrator to define lists of computers that are exempt from the firewall rules:

Trusted : Computers placed on this list are regarded as trusted and will be exempt from all firewall rules and allowed complete access to the computer.

Restricted : Computers on this list are regarded as untrusted and will be blocked completely and are allowed no access to the computer.

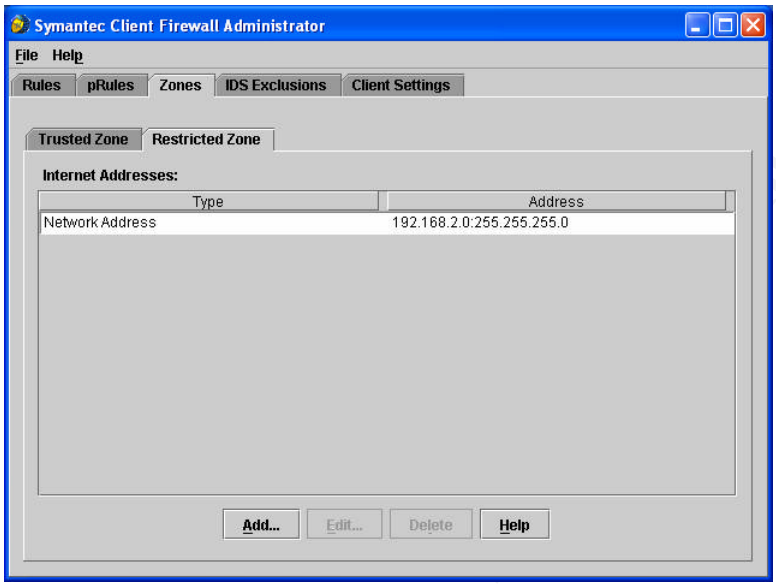

Figure 12 Modifying the restricted zone with Symantec Client Firewall Administrator

In the example above (Figure 12) computers residing on subnet 192.168.2.0 will be blocked from accessing this client computer.

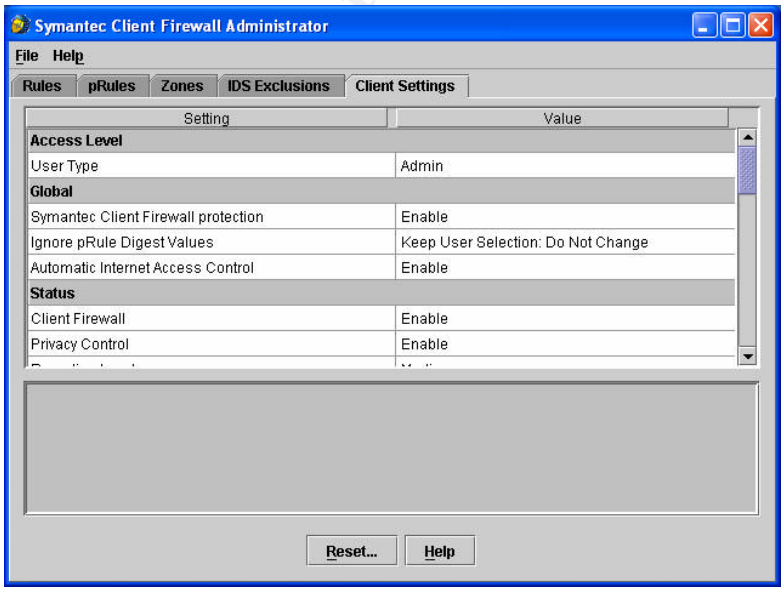

Figure 13 Using the Symantec Client Firewall Administrator to edit client settings.

Rolling out Symantec Client Firewall policies from the Symante c System Center

Once the Symantec Client Firewall policy has been finalised, it can be saved and rolled out to servers, clients and client groups from the Symantec System Center. On execution of a Symantec Client Firewall policy rollout a confirmation dia log is displayed for the administrator as shown in Figure 14. The Symantec Client Firewall policy file (.xml or .cfp) is copied down to each computer within scope of the rollout.

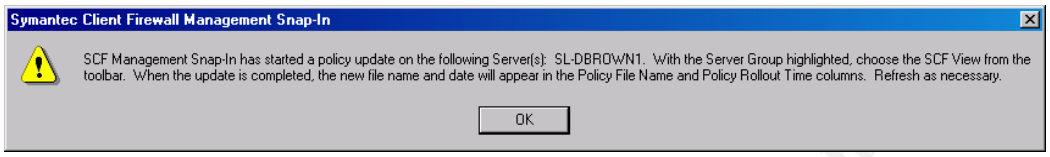

Figure 14 Using the Symantec System Center to rollout a Symantec Client F irewall policy file

Symantec Client Security logs and histories

All Symantec Client Security events are logged locally on the client and forwarded to the parent server. All events can be accessed through the Symantec System Center or exportable to a .CSV formatted file. For antivirus, logs are available for; virus histories, scan histories and configuration changes. Symantec Client Firewall logs are available for configuration changes, firewall violations, intrusion detection status and intrusion detectio n violations

Testing Symantec Client Security policies

Once Symantec Client Security policies have been defined, the administrator may want to perform a network port scan to ensure the polices behave in the correct manner. Running a port scan can verify that the correct level of protection is configured. Below is an example of the NMap configuration screen set up to scan the Symantec Client Security client (Figure 15). NMap supports a large number of scanning techniques such as: UDP, TCP connect (), TCP S YN (half open), ftp proxy (bounce attack), Reverse-ident, ICMP (pi ng sweep), FIN, ACK sweep, Xmas T ree, SYN sweep, IP Protoc ol, and Null scan.[4] These scanning techniques can be used to simulate an attack scenario.

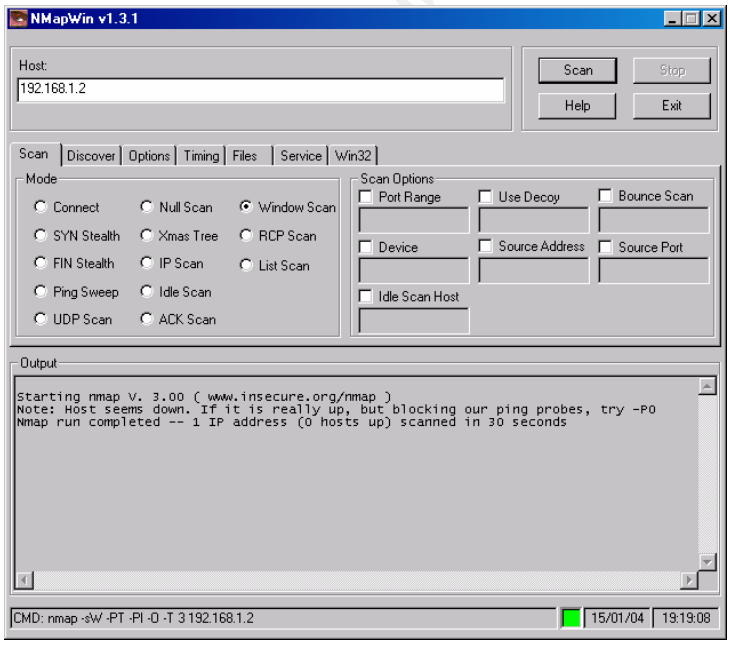

Figure 15 Using NMap to launch a port scan against a Symantec Client Security client

4 Nmap ("Network Mapper") is a free open source utility for network exploration or security

Symantec Client Firewall responds to the NMap port scan with a security alert according to the level of reporting set with the policy (Figure 16).

Figure 16 Symantec Client Firewall responds to the port scan with a security alert

Symantec Client Security in action – W32.Blaster.Worm

In July 2003 W32 Blaster.Worm was released into the wild. This worm exploited the Windows Distributed Component Object Model (DCOM) Remote Procedure Call (RPC) interface affecting Windows 2000 and Windows XP systems that did not have the appropriate Microsoft security patch applied.[5] W32.Blaster.Worm exploited a hole in windows security and on infecting a host, it executed without user intervention and used different methods of propagation to find its targets. Compared to the Code Red and W32.SQLExp worms, W32.Blaster was quite rudimentary, and was surprising to some people that it was so effective.[6]

W32.Blaster.Worm infection lifecycle [7]

1. An infected W32.Blaster.Worm host will first perform a port scan using TCP Port 135. This port scan will include random IP addresses.
2. On locating a vulnerable host, W32.Blaster.Worm sends a sequence of specially crafted packets to cause a buffer overflow condition courtesy of the DCOM RPC flaw.
3. W32.Blaster.Worm then proceeds to create a remote shell on UDP port 4444 to receive TFTP commands from the infecting host.
4. Using TFTP to listen on UDP Port 69 the infecting host will download the W32.Blaster.Worm executable locally to the %windir%\system32 directory where the malicious code can execute and start the whole cycle again.
5. W32.Blaster.Worm also attempts a Distributed Denial Of Service (DDOS) attack by attempting to send a number of packets to windowsupdate.com on specific dates.

Using Symantec Client Security to protect against the W32.Blaster.Worm

An unpatched computer running with Port 135 open equipped with only traditional antivirus protection installed would have been susceptible to the W32.Blaster.Worm attack. Once compromised, this computer could infect other similarly configured computers. Antivirus signatures became available for download from antivirus vendors a number of hours after the first infections were reported, leaving a 'window of exposure' before response and cleanup operations could begin.

5 Microsoft Knowledge Base Article – 824146

6 Defense in Depth: A Comparison Study of Three Worm Families and Their Propagation

7 Symantec Security Response W32.Blaster.Worm virus writeup

Symantec Client Security had a default system wide rule enabled to block inbound communication on TCP and UDP port 135 (Default Block EPMAP) as part of its standard configuration, thus blocking inbound communication from an infected host computer. [8] In addition to an antivirus definition, an IDS signature was available to download for computers that used TCP and UDP Port 135 for other applications, and therefore could not incorporate this rule as part of their policy.
The W32.Blaster.Worm was interesting with respect to the period of time that had elapsed between publication of the vulnerability and release of exploit code (around 30 days). Compare this to the W32.SQLExp.Worm that emerged a number of months before and took 6 months from initial vulnerability awareness to the arrival of exploit code. [9] This trend in reduced time to exploitation for new vulnerabilities means that early warning information services combined with patch management systems could prove useful for protection against future malicious code attacks

Conclusion

This paper has examined the use of Symantec Client Security to provide layered protection for connected desktops in an organisation. Whilst proving an effective model in protecting against blended threats, it must not be regarded as a standalone measure. To be truly affective against hybrid attacks it must form part of an overall 'security in depth' model encompassing all aspects of security. Blended threats exploit new software vulnerabilities, so vulnerability management software could assist in identifying systems requiring new patches. Policy compliance tools can be use to ensure strategic servers are in line with best practice policies according to platform and version.

According to a recent study conducted by Symantec Security Response, eleven out of sixteen vulnerabilities were exploited in a widespread fashion less than sixty five days after the publication of the associated vulnerability. [10] It seems a realistic argument that an organisation will see a better return on investment by investing in layered client security and 'security in depth' prior to a blended threat attack than if they are trying to clean up and restore computers after an outbreak.

8 Symantec Knowledge Base Article How to configure Symantec Client Firewall to block W32.Blaster.Worm

9 The Spread of the Sapphire/Slammer Worm by Moore , Paxson, Savage, Shannon , Staniford, and Weaver

10 Vulnerability Versus Exploitation Latency by Sean Hittel, Jesse Gough, Bartek Kostanecki, Jensenne Roculan

References

1 Symantec Internet Security Threat Report Trends for January 1, 2003 – June 30th, 2003 URL: http://ses.symantec.com/PDF/malcodetrends10187539_SMCT_rp.pdf (Feb 4th 2004)

2 Symantec Client Security Administratr ator's Guide July 28th 2002 URL: ftp://ftp.symantec.com/public/english_us_canada/products/symantec_client_security/manuals/scsadmn.pdf (Feb 4th 2004)

3 Symantec Client Security Installati on Guide July 19th 2002 URL: ftp://ftp.symantec.com/public/english_us_canada/products/symantec_client_security/1.1/manuals/scs1.1_ig.pdf (Feb 4th 2004)

4 Nmap ("Network Mapper") is a free open source utility for network exploration or security auditing URL: http://www.insecure.org/nmap/data/nmap_manpage.html (Feb 4th 2004)

5 Microsoft Knowledge Base Article – 824146 September 10th 2003 URL:http://support.microsoft.com/default.aspx?kbid=824146 (Feb 4th 2004)

6 Defense in Depth: A Comparison Study of Three Worm Families and Their Propagation by Daniel Hanson, Bartek Kostanecki, Richard Jagodzinski, Jason V Miller November 27th 2003 URL:https://tms.symantec.com/members/AnalystReports/031127 -Analysis-ThreeWorms.pdf (Feb 4th 2004)

7 Symantec Security Response W32.Blaster.Worm virus writeup December 11 th 2003 URL :http://securityresponse.symantec.com/avcenter/venc/data/w32.blaster.worm.html (Feb 4th 2004)

8 Symantec Knowledge Base Article How to configure Symantec Client Firewall to block W32.Blaster.Worm URL: http://service1.symantec.com/S UPPORT/ent-security.nsf/docid/2003081213520248 (Feb 4th 2004)

9 The Spread of the Sapphire/Slammer Worm by Moore , Paxson, Savage, Shannon , Staniford, and Weaver URL: http://www.caida.org/outreach/papers/2003/sapphire/sapphire.html (Feb 4th 2004)

10 Vulnerability Versus Exploitati on Latency by Sean Hittel, Jesse Gough, Bartek Kostanecki, Jensenne Roculan May 3rd 2003 URL: https://tms.symantec.com/members/AnalystReports/030503 -Analysis-Vuln-vs-Exploit.pdf (Feb 4th 2004)