



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>

High Level Network Monitoring for Security.

by Robert Lowe
17 May, 2004.

Abstract

This paper examines the high level monitoring of networks. It assumes a large network with limited control over internal clients and/or may be forced to provide a liberal (and possibly insecure) set of applications/access. Examples of such organizations are ISPs and universities.

Dunn¹ examines some security applications for Cisco NetFlow data, but this paper is pitched at a higher level, by not focusing on a specific protocol and employs a different approach. This approach is to first examine the network borne threats to organizations, with emphasis placed on the network signatures typical of these threats. From here, requirements for the types of information needed to detect these threats can be defined. The aim of this paper is to make it possible to evaluate the session level network traffic data formats such as Cisco NetFlow and Argus to determine if sufficient detail is included in those formats in order to facilitate initial detection of such incidents.

Bejtlich² makes four levels of distinction when examining data collected from a network: full content, alert, session and statistical data. Due to privacy restrictions and/or the volume of data generated, it is often not practical to collect full content data from large networks. Alert data is flagged by Intrusion Detection Systems based on known signatures or a set of heuristics, which may allow certain threats to go undetected. Therefore, this paper shall only examine the use of session data ("Session data represents conversations or flows between parties") and statistical data ("Statistical data represents broad trends in network activity") which, for this paper, is considered "High Level".

The use of Intrusion Detection Systems and Firewalls is out of scope of this paper, as is finer grained network monitoring (e.g. sniffers). Some threats are better detected using these types of tools. A high level view of network traffic has the strength of being a good initial indicator to deeper problems which can be further collaborated and analysed using more specialised security tools.

¹ Dunn, Jana. URL: <http://www.sans.org/rr/papers/10/778.pdf>

² Bejtlich, Richard. SysAdmin April 2004 - Volume 13 - Number 4 (2004): 24 -28.

ABSTRACT.....	1
THREATS	2
(DISTRIBUTED) DENIAL OF SERVICE ATTACKS.....	3
ICMP ping (echo request/reply) based attacks	3
TCP based attacks	4
UDP based attacks.....	4
IP based attacks.....	5
MALICIOUS SOFTWARE	6
SPAMMING.....	7
Collateral spamming.....	7
TARGETED ATTACK.....	7
Reconnaissance.....	8
Scanning and OS Fingerprinting.....	8
"Browsing".....	8
Digital Snooping.....	8
Execution.....	9
Brute force attack.....	9
Exploitation of vulnerable software services	9
Masquerading.....	9
Exploitation of a trap door or back door	9
Social Engineering.....	10
Spoofing.....	10
Result	10
EQUIPMENT/SOFTWARE MALFUNCTION AND PHYSICAL THREATS.	10
COMPUTER AND NETWORK MISUSE	11
Bot Net.....	11
Unauthorised Web Sites.....	11
Intermediate Points for Further Attacks.....	11
Proxy for Anonymous Internet Access	12
Unauthorised FTP Server.....	12
Unauthorised Mail Server.....	12
Peer to Peer File sharing.....	13
Remote Control software.....	13
REQUIREMENTS FOR NETWORK SECURITY MONITORING	14
REQUIREMENTS SPECIFIC TO THREATS/COMPUTER MISUSE	14
GENERAL REQUIREMENTS.....	15
Definition of Exceptions.....	15
Determining a Baseline and Detecting Trends.....	16
Meaningful Summarisation.....	16
CONCLUSION	17
REFERENCES.....	18

Threats

A non-exhaustive list of computer security threats is available online³. From this list, the Network borne threats have been selected and have been regrouped for

³ URL: <http://www.caci.com/business/ia/threats.html>

the purposes of this paper (e.g. virus, Trojan horse and mass mailers have been combined). Several threats are omitted from this listing, including Denial of Service (DoS) and Distributed Denial of Service Attacks which shall be covered separately. I have chosen to list the forms of computer misuse in a separate section and several additions (e.g. Peer to Peer File sharing). Also, “Tunnelling” is covered by targeted exploitation and worms and “Masquerade” is included in the “Execution” section.

(Distributed) Denial of Service Attacks

A Denial of Service (DoS) Attack is an attack on a system or networks availability. The goal of the attack is to make that system or network unusable or to cause complete failure, requiring manual intervention to recover. Distributed Denial of Service Attacks (DDoS) are DoS attacks which originate from many independent sources.

The Internet is now classed as critical infrastructure by many governments and many businesses have a high dependence on it for daily business functions. Therefore, the ability to quickly detect and react to network and system outages is a concern of many organisations. Given that many DoS and DDoS attacks generate a large amount of network traffic, high level network monitoring can play a crucial role in the early identification of such attacks.

Cole⁴ and attrition.org⁵ both provide a comprehensive, but dated list of DoS attacks. These lists (as well as other sources, which shall be referenced where appropriate) are used as the basis to summarise the major forms of DoS and DDoS attacks. However, specific attacks (against a particular platform) are many and varied, so this paper shall focus of more general, network flooding DoS attacks and a small cross section of specific, single packet attacks. This paper groups the attacks into protocol types.

This section could well be covered as part of the “Targeted attack” section (as that is what a DoS attack generally is). However, a critical distinction is often drawn between the ability to execute code and a DoS when assessing the severity of a vulnerability. Because of this distinction, it is worth dedicating a separate section to the examination of DoS attacks.

ICMP ping (echo request/reply) based attacks

Name	Description	Signature
Ping of Death	Large ICMP echo requests are sent to the target. Vulnerable operating systems (e.g. Windows) will crash, but there patches are available.	One or more large, incoming ICMP echo request packets, greater than 65Kb.
SSPing	Unexpected Fragmentation options	One or more large,

⁴ Cole, p. 182-227

⁵ URL: <http://www.attrition.org/security/denial/>

	are set in ICMP packets which, on receipt, may cause older versions of Windows and Macintosh to crash.	fragmented, incoming ICMP echo request packets.
Smurf	ICMP packets with a forged source address (of the target) sent to a broadcast address on an intermediate network (the “amplifier”). All machines on the (amplifier’s) network respond back to the (forged) address – flooding the target with responses. CERT Advisory CA-1998-01 ⁶ , examines this DoS in detail.	On the amplifiers network: one or more incoming ICMP packets directed towards broadcast addresses. And many (virtually simultaneous) outgoing ICMP echo replies to a single host. On the target network (may be the same as the amplifier): many incoming echo replies, from a single network, directed to a single host in a short space of time.

TCP based attacks

Name	Description	Signature
Bubonic	A DoS attack tool which randomly sets TCP flags in order to crash vulnerable Windows systems.	One or more incoming packets with a single source address (however, this may be varied each time the tool is run), with illegal or illogical flag combinations.
LandExploit	A TCP SYN packet with identical (spoofed) target and source IP addresses and port numbers, which may crash some routers.	One or more incoming TCP SYN packets with the same source and destination IP addresses and port numbers.
SYN Flood	By sending many TCP SYN packets to initiate a 3-way handshake, but never fully completing it after the ACK returned by the target.	Many incoming TCP SYN packets, without a following ACK packet.

UDP based attacks

Name	Description	Signature
DoS attacks using the	Certain DNS queries may return significantly more data than sent in	On the amplifier’s network: DNS queries.

⁶ URL: <http://www.cert.org/...>

Domain Name System (DNS) ⁷	the query itself. By sending many such DNS queries, potentially to several servers – the amplifier(s), using a spoofed (target) source, many (potentially large) responses can cause a DoS on the target.	One the target network: many DNS responses (from potentially many targets) for which no query was sent.
Fraggle	When an attacker sends a large number of UDP echo (ping) traffic to an IP broadcast addresses (on the amplifier's network), all of it having spoofed source address of a target network. It is the UDP variant of a smurf attack.	On the amplifier's network: one or more incoming UDP packets directed towards broadcast addresses. And many (virtually simultaneous) outgoing UDP echo replies to a single host. On the target network (may be the same as the amplifier): many incoming echo replies directed to a single host in a short space of time.

IP based attacks

This section lists layer 3 based attacks which don't use a higher layer protocol.

Name	Description	Signature
Checkpoint Firewall-1 Vulnerability	A vulnerability exists where by sending packets from a large number of IP addresses to the internal interface of the firewall.	Many packets each with different IP address destined for the internal interface of a Checkpoint Firewall-1 system.
Jolt2	Various vulnerable IP implementations (e.g. Windows 2000, Cisco 4500, Checkpoint Firewall-1) will consume large amounts of resources when processing large numbers of identical, fragmented IP packets.	A large number of identical, incoming IP packet fragments.

⁷ URL: <http://www.auscert.org.au/80>

Malicious Software

Malicious software (also known as malicious logic or “malware”) covers a variety of computer software designed to perform malicious activity. Viruses, Trojan Horses and Worms are all types of malware. This paper will not focus on the definition of each of these as this is a contentious issue and not contribute to purpose of this paper. Instead, this section will attempt to cover the propagation and infection of all malware, with distinction of specific malware types where relevant.

Currently, the major vectors for the spread of malware are email, previous malware infections, poorly secured Windows file shares, P2P (Peer to Peer applications) and software vulnerabilities (exploited by Worms). Message Labs, an email filtering provider lists the three “most active [email borne] viruses for all time”⁸ are W32/MyDoom.A, W32/Sobig.F and W32/Netsky.B. These viruses propagated by sending email with attachments which appeared as legitimate applications or documents, however contained malicious code. During such wide spread propagation email traffic will increase, but generally it is not possible to separate malicious email messages from legitimate ones, without content inspection (e.g. Virus scanners, mail filtering or IDS). However, recent malware contains it's own SMTP engine in order to send further messages which may show manifest as unexpected email traffic – see also the section entitled “Computer and Network Misuse: Unauthorised Mail Server” for more details.

Infection via P2P applications occurs when the malware detects P2P software is running on the system and copies itself onto directories it believes are shared across P2P networks. The malware often renames the copy of itself to entice potential victims to download and open it. P2P is covered in more detail in the section entitled “Computer and Network Misuse: Peer to Peer File sharing”.

Similar to the above propagation method described for P2P, malware will scan for available unsecured Windows file shares and when found, will copy itself to these shares for other potential victims. This activity will most likely generate increased and/or abnormal SMB (Simple Message Block) traffic when scanning for insecure shares.

Once infected, malware will often listen on a well defined TCP port (commonly know as “backdoors”) in order to receive later instructions. The facilities provided by these backdoors range in sophistication, depending on the malware. Certain types of malware will attempt to use the backdoor of previously released malware to infect a system. Malware which offers this facility of remote control may also be used to build up bot nets (see also the section entitled “Computer and Network Misuse: Bot Net”). Manual (human initiated) or automatic (malware initiated) scanning for, or connection attempts to, the listening backdoors would be characterised by TCP connection attempts to the backdoor port.

⁸ URL: <http://www.messagelabs.com/viruseye/threats/default.asp?toptenduration=all>

Once a worm (self propagating malware) infects a computer system via a vulnerable service, it proceeds to find other systems to infect, this may be done using ICMP echo requests to determine if a host is answering at an IP address or direct attempts to infect systems (without first checking their existence). This activity will generate large amounts of ICMP echo requests or TCP connection attempts to a vulnerable port.

Spamming

Spamming is the process of sending commercial email to a recipient who has not solicited it. Spammers (the senders of these messages) may attempt to target a large number of addresses in a single domain. This can cause inconvenience to the recipient(s) or occasionally Denial of Service conditions by consuming storage on mail servers and/or generating excessive network traffic.

Often, spammers do not use their own systems to send spam as this practice is generally contrary to the acceptable usage policies of ISPs (and the law in some jurisdictions). The use of systems as unauthorized mail relays is covered in the section entitled "Computer and Network Misuse: Unauthorised Mail Server".

Large amounts of spam will cause an increase in the amount of incoming email traffic, however if the recipient is valid, this traffic will be impossible to separate from usual email without content inspection. There are several solutions, such as filtering software/services and black hole lists which are available to reduce the impact of spam.

Collateral spamming

Collateral spamming⁹ is an informal definition (see [JANET reference]) given to the act of a spammer forging the "From:" address of their spam to appear to come from a legitimate (and innocent) organisation. This causes the "bounced" (undeliverable) messages and complaint emails to sent back to the innocent organisation.

Like incoming spam, these messages are not easily distinguishable from legitimate messages – indeed undeliverable messages are required normal behaviour. However, the undeliverable notifications will come in soon after the original spam is sent, so a sharp increase in email traffic is expected, however detection depends on the email volume considered normal by an organisation.

Targeted attack

During a targeted attack, there are several (potentially optional and non-distinct) stages: Reconnaissance, Execution and Result.

⁹ URL: <http://www.ja.net/mail/junk/collateral.html>

Reconnaissance

Scanning and OS Fingerprinting

Two distinctions will be drawn for the purposes of this paper: network scanning and host scanning. Network scanning is the scanning of particular network range in order to determine the number of hosts on that network and their IP addresses. Host scanning is the process of targeting one particular IP address in an attempt to determine which operating system (OS fingerprinting) and services that host is running. It is possible to combine host and network scans.

A network scan will generally cause a high frequency of ICMP, TCP or (less frequently) UDP packets from a single IP address to a range of internal network addresses, some of which may not exist or be intended to be contacted directly.

A host scan attempts to contact a range or all ports of a system in order to determine which are open. Of the ports that are open, connections may be initiated to obtain information about the host OS. More advanced OS fingerprinting may be done by sending out of band or unspecified data and the OS deduced by examining such responses. Host scans will typically generate large amounts and a variety of network traffic to a particular host.

There is also the concept of “low and slow” scanning. This is when the data sent to a site is minimised and spaced out over long periods of time, in order to avoid detection. For large amounts of network traffic, it may be very difficult to correlate this scanning activity to determine the nature of the reconnaissance.

“Browsing”

An attacker may use online information resources, such as search engines (e.g. Google), a corporate web site or network information service (e.g. whois and DNS) to gain information about an organisation and its network. This type of activity (if it does involve direct connections with an organisation) will most likely be indeterminable from other legitimate activity. For example, requesting a web page from a company web server or performing a DNS query on the organisations name server.

Digital Snooping

Digital Snooping is the unauthorised capture of network traffic. This information may contain confidential data which is valuable to an attacker (e.g. using that data to launch further attacks). A properly used sniffer, in itself does not cause any network traffic and can be extremely difficult to detect. If the sniffer is located on systems outside external to the targeted organisation, it will be almost impossible to detect by examining an organisation’s network traffic alone.

In incidents where the attacker has installed a snooping device within the organisations network, it may cause some network traffic when attempting to send that information back to the attacker. There are a number of methods employed to do this: by sending captured details to a specific email address,

posting to a designated web site, via ftp, a terminal connection. It may be difficult to detect this type email, web or ftp traffic. However, a long, unexplained terminal connection may be an indicator of such activity. There is also the possibility of the data being stored on the system and retrieved later by the attacker. This may be done by setting up an ftp, web server or some type of terminal service on the system.

Execution

Brute force attack

A brute force attack consists of sending many login attempts to a public service which requires authentication. Examples of such services include: web applications (e.g. internet banking, auction or payment sites, internet based email accounts), windows shares, ssh/ftp/telnet servers and some VPN solutions (which require a username and password to authenticate). This type of attack will create an increase in the amount of traffic. If the users of these services are finite and known, then it may be possible to examine network traffic for exceptional connections (e.g. those originating from an unexpected net block).

Local attacks (using physical access to the console) and offline attacks (obtaining and attempting to crack the password databases) can be attempted with little or no network activity.

Exploitation of vulnerable software services

No software is perfect and will have vulnerabilities. Attackers attempt to exploit such vulnerabilities in order to make the system execute normally restricted commands, obtain or change protected files or cause the service to fail (causing a Denial of Service – covered in section <Blah>).

The execution of most successful targeted attacks is done with a minimum of network traffic and across allowed services. Such attacks are not easily detected in large networks as they may be confused for legitimate traffic at the high level and only content inspection (performed by an IDS or specialised firewall) may detect malicious intent.

Masquerading

A masquerading attack occurs when the attacker pretends to be a valid user or service. The very nature of this threat will mean that a successful attack will appear normal and may blend in with legitimate usage. There is a chance that this attack may be detected by examining network connections from unexpected sources.

Exploitation of a trap door or back door

Some software has the ability for a user to gain access or increase their privileges without being subjected to the normal authentication mechanisms. This attack is often committed by an “insider” or someone associated with the target organisation or the developers of their software that has intimate knowledge of the systems employed.

Also, once a successful attack has been performed, the attacker may attempt to install a backdoor or create their own user accounts in order to enable faster and easier access in the future.

This type of unauthorised access may not exhibit the properties of expected traffic. The network signature of such an attack may be administration traffic from an external (or unexpected) source or traffic with an anomalous destination (e.g. connections to a seldom used port).

Social Engineering

Social Engineering is the process of eliciting trust from a person, which is not deserved. The classic example is when a user receives a call from an attacker, masquerading as a support officer, who asks for the user's password in order to perform some maintenance or trouble shooting.

Once a social engineering attack has been successful, the network signature of the use of the information obtained would be minimal. The only indicator would be the external use or attempted use of data (e.g. username and passwords) from unexpected sources.

Spoofing

A spoofing attack is when an attacker provides false identification data (generally IP addresses) in an attempt to gain increased system access to a target because of a trust relationship that exists between the target and (forged) source.

A properly configured, packet filtering firewall will be most effective in addressing this threat as well as IDS systems which may detect suspicious precursors or slight irregularities in network traffic. Typically, minimal unexpected network traffic will be generated in such an attack, making it difficult to detect.

Result

Generally, once an attack has been successful and unauthorised access has been gained, then an attacker will look to use this access for some purpose. The possible purposes of the attacker will generally fall into one of the categories outline in the "Computer Misuse" section.

Instead of, or in addition to the results covered in that section and attacker may also: attempt to obtain confidential data from an organisation, install a kit or place back doors in the system to facilitate later access and/or patch the system in order to prevent the system being compromised by another attacker. All of this activity may produce a large variance in the amount of network traffic generated.

Equipment/Software malfunction and physical threats.

All equipment and software is susceptible to failure and while some steps may be taken to increase system availability, it can never be perfect.

A sharp decline in network traffic can sometimes be a first indication that a network component has malfunctioned. It may also provide a valuable source of information in determining the nature and extent of that failure.

Computer and Network Misuse

The following sections describe the common unauthorised functions performed by systems. This may be the result of a threat being successful (e.g. targeted attack, worm) or from internally (“trusted”) sources.

Bot Net

A (ro)bot is software agent, which performs a set of tasks (generally DoS attacks or key logging) and is remote controlled. Large groupings of Bots are called BotNets and are centrally controlled and are effectively used to carry out DDoS attacks.

Control of bots have commonly be performed by configuring the bot to listen on specific Internet Relay Chat (IRC) channels and wait for commands. However, recent activity points to the use of P2P (Peer to Peer) technology used more commonly in file sharing as a means of controlling botnets. A recent example of this is Phatbot¹⁰.

The detection of bot nets may be aided by the monitoring of IRC traffic, which defaults to TCP port 6667. However it is possible that bots are configured to connect to servers on non-default ports or use P2P communication mechanisms. Therefore this traffic may be better detected by content aware systems such as IDS.

Unauthorised Web Sites

Malicious attackers or ignorant local users may decide to commission unauthorised web servers on systems within an organisations network. This may be for distribution of inappropriate or illegal material or simply for convenience.

Once again, the web site may be at port 80 or 443 (the default HTTP and HTTPS ports, respectively) or an arbitrary port chosen by the web server owner and content inspection should be employed for a more robust detection solution.

Intermediate Points for Further Attacks.

A common use for a compromised machine is to use it to compromise other machines. By doing this, it becomes harder to trace the malicious activity back to the actual attacker.

The requirements defined by the targeted attack threat, specifically the subsections entitled “Scanning and OS Fingerprinting” and “Exploitation”; describe the signature behaviour of an organisations host being used for further

¹⁰ URL: <http://www.lurhq.com/phatbot.html>

attacks. However the key difference is that an internal network will be the source of these attacks.

Proxy for Anonymous Internet Access

A proxy is an intermediate point which makes a connection to a server at the request of a client. The most common proxy is a web proxy (also known as a web cache) however, most other protocols can be proxied. Proxies are often legitimately used by an organisation to cache popular web pages in order to save on internet traffic costs and funnel outgoing web connections from a single host so that security can be tightened around this single point. However, an organisation should not generally allow its proxies to be used by external users – if this is the case; this is referred to as an “open proxy”. Open proxies exist because of deliberate setup (potentially by the administrator or a miscreant) or because of an administrator’s mis-configuration.

Most proxies run on a well known (TCP) ports (namely 3128, 8080, and 8088), however, as with almost all server software this may be varied. Once again, content needs to be examined to determine if specific traffic is proxy traffic.

Unauthorised FTP Server

FTP (File Transfer Protocol) servers are used specifically for the distribution of files. FTP servers are often set up by miscreants for the distribution of copyright material (“warez”) or maybe by a legitimate user for various (seldom valid) reasons. Or alternatively a legitimate server may be incorrectly configured to allow unprivileged users to upload unauthorised files.

FTP ports default to 21 for control and 20 for data (both TCP and UDP). However, again, it is possible for a user to modify these default ports to any of their choosing for that server. The use of legitimate ftp servers to host unauthorised material is difficult to detect using this method an abnormal increase in data transfer to and from the server may be an indicator, but logging of the ftp server is the best way to determine if this type of activity is taking place.

Unauthorised Mail Server

Email is transmitted using SMTP (Simple Mail Transfer Protocol). Unauthorised mail servers may be commission by miscreants or an existing incorrectly configured mail server may be used to send spam. Spam has already been defined in the section entitled “Spamming”. The use of an organisation’s resources to send spend is undesirable. Apart from tangible traffic costs, spam originating from an organisation can be very damaging for its reputation and can lead to systems being added to black lists which prevent them from connecting to some mail servers. An “open relay” is a specific type of proxy for internet access (the sending of electronic mail), similarly it is brought about by a mis-configuration of an organisations mail server or the deployment of a unauthorised mail server (either by an attacker or malicious insider).

Such a mail relay, like other servers, can be configured to accept connections on various ports, however TCP 25 is the default for SMTP (Simple Mail Transfer Protocol). That host will have to connect to other mail servers in the process of sending mail and this will almost certainly be done on the default destination port of TCP/25.

Peer to Peer File sharing

Peer to Peer (P2P) traditionally referred to a communications architecture which was facilitated by identical end points (as opposed to Client/Server architectures) is now often used as a term to describe Peer to Peer file sharing networks. Napster was probably the most (in)famous P2P file sharing application, attracting a great deal of publicity (and litigation). Currently, an ever increasing number of P2P file sharing applications are available, these include: Gnutella, KaZaA Media Desktop, Morpheus, WinMX, eDonkey and Bit Torrent, just to name a few.

In addition to facilitating copyright infringement such file sharing adds another vector by which malicious code may propagate. An example of this was Beagle.H¹¹ and MyDoom.B¹² both of which attempted to propagate by copying themselves to shared directory as a file with an enticing name.

Given the large, increasing numbers of P2P applications, each with their own protocol, there is no one set of ports which can be used to track P2P traffic. Lists¹³ are available, but are not always updated. Another challenge for network monitoring of P2P applications is that "Certain P2P clients already use port 80 (usually reserved for Web browsing) when they detect the presence of a firewall blocking other ports."¹⁴ However, a typical song in MP3 format will be around 2-5 Megabytes, therefore regular usage of this software will cause quite a large amount of traffic to be generated.

Remote Control software

There are many remote administration and remote control software packages available. Some of these are written with malicious intent, such as Back Orifice¹⁵, Subseven¹⁶ or Netbus¹⁷. Some may be put to legitimate use for remote administration or trouble shooting, such as Microsoft's Terminal Services (Windows 2000/NT) or Remote Desktop (Windows Server 2003) products, Virtual Network Computer (VNC) and PCAnywhere.

Like P2P software, the wide variety of Remote Control software and the ability to run it various ports will mean that traffic may have a wide variety of signatures. The size of the traffic can vary.

¹¹ URL: <http://securityresponse.symantec.com/avcenter/venc/data/w32.beagle.h@mm.html>

¹² URL: <http://securityresponse.symantec.com/avcenter/venc/data/w32.mydoom.b@mm.html>

¹³ URL: http://www.outpostfirewall.com/guide/rules/preset_rules/p2p.htm

¹⁴ URL: <http://www.epic.org/privacy/student/p2pletter.html>

¹⁵ URL: <http://www.symantec.com/avcenter/warn/backorifice.html>

¹⁶ URL: <http://www.f-secure.com/v-descs/subseven.shtml>

¹⁷ URL: <http://www.windowsecurity.com/pages/article.asp?id=453>

Requirements for Network Security Monitoring

The purpose of this section is to isolate the data required to identify the threats and misuse discussed previously and to examine more general requirements which may assist in adapting to future threats and computer misuse.

Requirements specific to threats/computer misuse

The following table summarises the network monitoring requirements which may need to be satisfied in order to detect the threats and computer misuse outlined previously in this document. If a particular threat or computer misuse can not be readily be detected via from network traffic it has been omitted from this table.

Threat or Computer Misuse	Detection Requirement It shall be possible to detect...
Ping of Death DoS	ICMP echo requests larger than 65Kb, summarised by destination and number.
SSPing DoS	Fragmented, ICMP requests, summarised by destination and number.
Smurf DoS Amplifier	ICMP packets directed towards a broadcast address.
Smurf DoS Amplifier, Smurf DoS Target	ICMP echo replies, summarised by destination, time and number.
Bubonic DoS tool	TCP packets with illegal or illogical flag combinations.
Land Exploit	TCP SYN packets with identical source and destination addresses.
TCP SYN Flood DoS, SYN Scan	TCP traffic which consists of a SYN without a following ACK, summarised by destination IP address and number.
DNS DoS	DNS responses without initial requests.
Fraggle DoS Amplifier	UDP packets directed to broadcast addresses
Fraggle DoS Amplifier, Fraggle DoS Target	UDP packets, summarised by destination host, time and number.
CheckPoint FW-1 Vulnerability DoS	IP packets destined for the interface of a firewall.
Jolt2 DoS	Identical IP packet fragments.
Email borne malware infection, Unauthorised Mail server	SMTP traffic (destination TCP port 25) to hosts which are not authorised mail servers.
Remote control of malware	Connection attempts to known malware ports.
Worm propagation attempts, incoming network scan	ICMP ping or TCP connection attempts (to a known port) to a range of network hosts from a single source, summarised by source
Worm infections	A large amount of ICMP pings or TCP connection attempts (to a known port), summarised by source

	IP (infected system)
Network scan	A large amount of ICMP pings or TCP connection attempts, listing full details of each flow.
Host scan	Varied and numerous UDP and/or TCP packets destined for one IP, listing each flow.
Bot net control	Connections from hosts to IRC servers (TCP port 6667), summarised by source.
Unauthorised Web Sites	Incoming TCP connections to ports 80 or 433, for systems which are not authorised to host web servers.
Proxy for Anonymous Internet Access.	Incoming TCP connections to port 80, 3128, 8080, 8088 for systems which are not authorised to host proxy or web servers.
Unauthorised FTP Server	Incoming TCP or UDP connections to ports 20 or 21 for systems which are not authorised to host ftp servers.
Misuse of existing FTP Server	An unexpected increase in traffic to existing FTP servers.
Peer to Peer File Sharing	Inbound and outbound connection attempts to known Peer to Peer File Sharing protocol ports.
Remote Control software	Inbound connection attempts to known remote control software ports.

General requirements

In isolating the above specific requirements, it became obvious that many are the expression of more general requirements. These general requirements will be highlighted in this section.

Definition of Exceptions

Effective network monitoring requires the ability to define exceptions which are excluded from results. An example which occur frequently in the previous section is the ability to report on a particular type of traffic (e.g. HTTP or HTTPS) but exclude particular hosts from the result set (e.g. authorised web servers).

Another example might to obtain reports which do not contain a particular type of traffic in order to easily examine it for inconsistencies. For example, an analyst may wish to see the UDP traffic which does not have a source or destination port of 53 (DNS) or 123 (NTP).

Also, it may be useful to obtain high level statistical data without including certain known data sets. For example it may be useful to have a break down of traffic without the inclusion of FTP, HTTP, SMTP traffic in order to determine other broad trends in bandwidth usage.

Determining a Baseline and Detecting Trends

Many computer security incidents are first noticed when a security analyst notices something out of the ordinary. In order to notice something abnormal a security analyst must have an understanding of what is considered normal.

This is extremely difficult to automate reliably and completely, however, the next best thing is to provide the ability to quickly access similar historical data. Using this an analyst may be able to determine if particular behaviour is suspicious.

For example, an organisation may have automatically scheduled a complete vulnerability scan of their network on a quarterly basis. Therefore, an operator noticing this behaviour may compare it to past data to find that happens on a regular basis and this event may decrease in urgency.

Conversely, another example might be an operator notices an increase in FTP traffic to an organisation's authorised FTP server by a factor of twenty. Comparing this to historical data she sees that this has never occurred before and may then choose to further investigate the cause of the traffic, after some investigation she may see that the FTP server contains unauthorised files, whose popularity is the reason for the increased traffic. This may have otherwise gone unnoticed.

Meaningful Summarisation

The way in which data is summarised may add a great deal to it's usability. This will be highlighted with two of the above examples.

A bot net is discovered to be communicating with a particular (external) server on port 6667 and the security analyst is asked to report on any communicating hosts within an organisation. The primary concern is with the particular hosts (IP addresses) which have been communicating with this server, therefore, just summarising the unique IP addresses may be enough. However, if dynamic IP address allocation is activated, then also a listing of times for each of these IPs may be required.

Another example is an incoming network scan is performed across an organisation by an external host. If the security analyst wishes to report this activity to the controllers of the originating network then complete details, without summarisation should include (at a minimum) the: protocol, source and destination IPs, source and destination ports (if applicable) and the time, date and timezone of the traffic should be supplied if action is expected.

Conclusion

This paper attempts to distil all the major, current computer security threats and forms of security misuse. From here, it attempts to examine the reporting facilities necessary to examine (potentially) large amounts of data in order to determine if security incidents are occurring, rather than relying on other means, such as human reports and system/network failure.

Hopefully, gives a template against which to evaluate network traffic data formats and their associated reporting and analysis tools to determine if sufficient detail is included in those formats in order to facilitate initial detection of such incidents. Two such common formats are Cisco NetFlow¹⁸ and Argus¹⁹ (Audit Record Generation and Utilization System).

¹⁸ URL: <http://www.cisco.com/warp/public/732/Tech/nmp/netflow/>

¹⁹ URL: <http://www.qosient.com/argus/>

References

1. Dunn, Jenna. "Security Applications for Cisco NetFlow Data". 23 July 2001.
URL: <http://www.sans.org/rr/papers/10/778.pdf>
2. Bejtlich, Richard. "Integrating the Network Security Monitoring Model." Sys Admin Magazine. April 2004 - Volume 13 - Number 4 (2004): 24 -28.
3. CACI International Inc. "Computer Security Threats"
URL: <http://www.caci.com/business/ia/threats.html>
4. Cole, Eric. Hackers Beware: The Ultimate Guide to Network Security. New Riders, August 2001.
5. Martin, Brian. "Denial of Service Database". 1999.
URL: <http://www.attrition.org/security/denial/>
6. CERT/CC. "Smurf DoS". Date
URL: <http://www.cert.org/...>
7. AusCERT. "AL-1999.004 -- Denial of Service (DoS) attacks using the Domain Name System (DNS)". 13 August 1999.
URL: <http://www.auscert.org.au/80>
8. Message Labs. "MessageLabs Threats And Analysis: Threats - Overview". As at 20 April 2004.
URL:
<http://www.messagelabs.com/viruseye/threats/default.asp?toptenduration=all>
9. Tillotson, Rodney, JNT Association. "Collateral spam 13 Nov 2001". 13 November 2001.
URL: <http://www.ja.net/mail/junk/collateral.html>
10. LURHQ Threat Intelligence Group. "Phatbot Trojan Analysis - LURHQ". 15 March 2004.
URL: <http://www.lurhq.com/phantbot.html>
11. Magee, Maryl. "Symantec Security Response - W32.Beagle.H@mm." 3 March 2004.
URL:
<http://securityresponse.symantec.com/avcenter/venc/data/w32.beagle.h@mm.html>

12. Gettis, Scott. "Symantec Security Response - W32.Mydoom.B@mm" 4 February 2004.
URL: <http://securityresponse.symantec.com/avcenter/venc/data/w32.mydoom.b@mm.html>
13. Cox, Stephen. "Outpost Firewall Guide - P2P"
URL: http://www.outpostfirewall.com/guide/rules/preset_rules/p2p.htm
14. Electronic Privacy Information Center. "EPIC Letter on P2P Monitoring to Colleges and Universities". 2 November 2002.
URL: <http://www.epic.org/privacy/student/p2pletter.html>
15. Symantec Corporation. "Information on Back Orifice and NetBus"
URL: <http://www.symantec.com/avcenter/warn/backorifice.html>
16. Podrezov, Alexey. "F-Secure Computer Virus Information Pages: SubSeven".
URL: <http://www.f-secure.com/v-descs/subseven.shtml>
17. Henderson, William. "The Netbus Trojan". 18 July 2002.
URL: <http://www.windowsecurity.com/pages/article.asp?id=453>
18. Cisco Systems. "Cisco IOS Software NetFlow - Cisco Systems". 22 January 2004.
URL: <http://www.cisco.com/warp/public/732/Tech/nmp/netflow/>
19. QoSient. "Argus – Home".
URL: <http://www.qosient.com/argus/>

© SANS Institute 2004, Author retains full rights.