Global Information Assurance Certification Paper

Intrusion Detection in Depth

GCIA Practical Assignment

Version 3.0 (revised August 13, 2001)

Stephen Bonner

Assignment 1 -Describe the State of Intrusion Detection

Features of Enterprise -scale IDS systems.

This paper reviews the sta ndard approach to Enterprise -scale intrusion detection and highlights the particular challenges faced by large -scale deployments. A series of alternative approaches to detecting intrusions is presented, focusing on new developments within this field.

An Enterprise network contains at least several thousand network connected devices, and is used by many thousands of staff. For this size of installation, effective intrusion detection is vital.

The fast-paced change of IT equipment and configuration within many organisations of this size makes complete protection unworkable without an intolerable overhead on business functionality. There will always be new systems deployed, or modifications to existing systems, without complete security. As well as the risk from "zero day" vulnerabilities[1], once a patch exists for a known vulnerability it will take significant time to be deployed across the enterprise. During this time, the enterprise is exposed to risk and yet cannot disconnect without affecting the busines s. Effective intrusion detection allows company information security teams to manage this risk by identifying, responding to, and recovering from incidents.

The standard approach to providing intrusion detection in such an organisation is to deploy netwo rk-based IDS[2] at entry points and on key LANs and then host -based IDS[3] on critical servers. Given the scale of enterprises of this nature such a deployment is likely to include around 50 network sensors and a few hundred host - based sensors. These sensors w ill have to be monitored and managed by a series of global teams covering different time zones.

The major issues that such a deployment faces include:

Network sensors are blind to some threats.

1. LANs are normally provided via switches with few, if any, co llision domains containing multiple servers.
2. Network speeds are increasing to Gigabit volumes, faster than sensors can analyse.
3. Use of network and application -level encryption (VPN, SSL and IPSEC) is increasing.
4. Enterprises run their own proprietary applic ation protocols across the networks.

Deploying and maintaining intrusion detection has a high overhead.

---

[1] Cliff, Exploits: Zero Day Exploit
[2] Bace, p49.
[3] Amoroso, p55.

5. Difficulty in gaining buy-in from systems administration teams to deploy, maintain and respond to host-based IDS systems due to memory and processor overhead on systems
6. The difficulties of rolling out updated detection signatures to meet new threats and vulnerabilities.
7. The high level of alarms, both those that turn out to be false alarms and also the high level of questionable use of enterprise networks.

Risks are introduced into the environment by deploying intrusion detection.

8. Adding to the information security risks by providing a centralised monitoring point and an access/reporting method to every critical host.
9. Legal and regulatory issues of monitoring staff behaviour[4].

Other issues.

10. The detections tend to be signature-based, with a small number of known out-of-spec generic detections.
11. Host-based solutions are not available for all platforms.
12. Few effective tools exist to integrate output from intrusion detection systems into problem management tools to allow prioritised response to attacks.

Many of the issues around network sensors' blindness can be tackled using Hybrid intrusion detection systems. These are host-based sensors that, as well as reviewing operating system logs, also include a "shim" into the TCP/IP stack so that network traffic can be reviewed post-decryption and before being passed to the operating system. These tools generally also include a host-based firewall system to allow the automated blocking of detected bad packets. By reviewing the packets at the host after decryption, before passing the data to the host, the problems of faster network speed and encrypted traffic are avoided. Also as there is one sensor per host, protection is still provided irrespective of the collision domain the host is connected to.

However, these hybrid sensors are still blind to proprietary protocols and applications. They also require higher deployment time and effort, and higher maintenance costs, as there are more copies of the intrusion detection software to be managed. Also, some of the advantages of network-based intrusion detection, such as the ability to defend multiple machines with a single sensor, low impact on network performance and host CPU/memory and the ability to defend hosts with unusual or legacy operating systems are lost.

The overheads of deployment are an ongoing challenge that require the assistance of intrusion detection developers to improve deployment processes for new signature of the kind that are currently used for anti-virus protection and additional capabilities to tune signatures and alerting to reduce the workload on intrusion analysts while improving the ability to detect actual attacks amidst the chaff of false alarms and automated probing.

There are a number of new approaches that can complement the traditional partnership of network and host-based signature-based intrusion detection, allowing improved detection of higher-skilled attackers, without further adding to the burden of

[4] Northcutt, p30-32

time and skill on corporate information security teams. In this paper, integrity checking, honeypot alarms and vulnerability correlation are discussed.

Integrity Checking

The majority of successful intrusions result in the modification of key system or application files, with the exception of some worms (like CodeRed) that modify only system memory.[5]. By using an integrity checking system such as Tripwire [6], or ISS System Scanner Frozen Files [7] regular checks can be performed to calculate cryptographically-sound checksums of file contents. These can be compared with known good values held on read-only media. If a checksum has altered and no authorised change has been carried out on the system, then an intrusion has occurred.

This method of detecting intrusions has been so successful that attackers have begun to develop methods to circumvent its detective capabilities. As the file checker relies upon operating system information about the file, it can be fooled by kernel module alterations that provide false information [8]. Although booting from known good read-only media and mounting the drive will defeat these attempts to avoid detection, this is a luxury that can only be carried out on suspected compromised machines or on a very infrequent basis, due to the downtime required.

These databases of checksums are also useful in forensic work: by eliminating known good operating system and application files, attacker-modified data can be identified. The hundreds of megabytes of a standard Microsoft Office install can take days to sift through, to ensure that no Trojan code or sniffed passwords have been hidden within the many files and directories, but if a they can be quickly compared to a set of default install files, then the needle within the haystack of data can be easily found.

Honeypot Alarms

Most enterprises suffer from Armadillo-style information security: they have a good tough exterior that is well-defended, surrounding a soft chewy inside that can be easily attacked once the outer shell has been compromised. Such organisations suffer a high risk of insider attack, from disgruntled staff or a workstation that has had a Trojan successfully installed. Traditional intrusion detection is limited against this threat, as identifying staff that are attempting to exceed their authority can be a difficult challenge.

One approach is to configure services or even entire systems that do not contain real information, but only bait files. These are closely monitored in ways that would generate false alarms if used on the real data, but can identify staff browsing around for information they are not authorised to access. For example, a server can be added to the HR LAN, named like other HR servers, but with a file share containing a spreadsheet called salary.xls. Any attempts to read, modify or access this file can be flagged by the host-based intrusion detection system. The network sensors can be

[5] CERT, Section III Solutions.
[6] Tripwire
[7] ISS
[8] Seifried

tuned to look for the contents of this file. As there is no legitimate reason to access the file, the problem of false alarms is greatly reduced. Also, as the data has no value, staff can be lightly informed that they are doing wrong when they stumble across it, allowing them to cease probing without losing face. This also prevents them from accidentally accessing data, finding that nobody noticed and then moving on to steadily larger and larger policy deviance until a significant loss occurs.

It is important to note that deploying such systems on externally facing systems is ill-advised. Not only are they likely to attract increased attention from external attackers, as honeypot devices will respond as if vulnerable to probes, but also if the press are informed of the weakness they will find it hard to understand that sensitive information was exposed intentionally, even if the data was fabricated, and might report it as if it was a real weakness.

Vulnerability Correlation

Another approach to improving enterprise scale intrusion detection is the integration of vulnerability data to reduce false alarms and overhead on analysts. These systems reduce priority of alerts if no weakness exists that matches the attack, and increase the priority is a known weakness is being targeted. Data is collected by host- and network-based vulnerability scanning tools about the operating systems, services and patch levels of the enterprise and held centrally for comparison with IDS data.

Many organisations carry out a similar process when tuning the performance of their intrusion detection systems. If they do not run IIS but only have Apache, they might manually disable the IIS signatures to allow their sensor to work at higher network loads.

This approach carries some risk, as systems are often run without the knowledge or approval of the security team, and approved components might include vulnerable systems without the knowledge of the defensive team. For example, various Cisco products come bundled with IIS [9] exposing the organisations to CodeRed risk. It only takes a few minutes to install a copy of Vmware and add a surprising operating system to an environment.

Even the collection of vulnerability information can complicate intrusion detection systems. Unless a team are willing to filter all traffic from scanning hosts and introduce the risk that traffic spoofed from those addresses is ignored, then an intensive manual process must be carried out to remove alarms caused by network scanning.

Many organisations still need to know about attacks against systems they don't run, in order to catch the internal audit department testing their abilities. Attackers often have a series of attacks they can use and by downplaying the early attacks, the time available to block their later attempts might be reduced. No signature-based IDS can identify all attacks, and if all known ones are filtered as the organisation is patched, then the chance to alert before the unknown attacks are launched has been lost.

[9] Cisco

Summary

Large-scale intrusion detection systems are increasingly required to properly defend a large-scale organisation from computer misuse. However, reliance on solely network - and host-based systems can cause significant and wide -ranging problems. Alternative approaches, including integrity checking, honeypot alarms and vulnerability correlation, should also be considered to provide proper d efence in depth.

References

Amoroso, Edward. Intrusion Detection - An Introduction to Internet Surveillance, Correlation, Trace Back, Traps and Response . Sparta: Intrusion.Net Books, 1999.

Bace, Rebecca. Intrusion Detection . Indianapolis: Macmillan Techn ical Publishing, 2000.

CERT. "CERT® Advisory CA -2001-19".Original release date: July 19, 2001 Last revised: August 23, 2001. http://www.cert.org/advisories/CA -2001-19.html (Sep 1, 2001)

CISCO. "Cisco Security Advisory: "Code Red" Worm - Customer Impact". Revision 2.2. For Public Release July 20, 2001 Last Update August 11, 2001. http://www.cisco.com/warp/public /707/cisco-code-red-worm-pub.shtml (Sep 1, 2001)

Cliff, A, "Intrusion Detection Systems Terminology, Part One: A - H". July 3, 2001. http://www.securityfocus.com/frames/?focus=ids&content=/focus/ids/articles/idsterms.html (Sep 1, 2001)

ISS, "Security Assessment: System Scanner" Jan 4, 2000. http://www.iss.net/securing_e -business/security_products/security_assessment/system_scanner/ (Sep 1, 2001)

Northcutt, Stephen. Network Intrusion Detection An Analyst's Handbook . Indianapolis: New Riders Publishing, 1999

Seifried, Kurt. "Creating and Preve nting Backdoors in UNIX Systems". June 28, 2000. http://www.securityportal.com/closet/closet20000628.html (Sep 1, 2001)

Tripwire. "Tripwire for Servers for Security & Network Manage ment". July 17, 2001. http://www.tripwire.com/products/servers/ (Sep 1, 2001)

.

Assignment 2 - Network Detects

Detect 1 - Smurf attacks.

```
'Smurf' event detected by the RealSecure engine at 'sensorXXXX'.
Details:
      Source Address: 24.93.35.210
      Source MAC Address: 00:D0:XX:XX:XX:XX
      Destination Address: MY.NET.239.24
      Destination MAC Address: 00:90:XX:XX:XX:XX
      Time: Tue Sep 11 00:08:17 BST 2001
      Protocol: ICMP (1)
      ICMP Type: Echo Reply
      ICMP Code: None
      Priority: low
      Actions mask: 0xa44
```

1. Source of Trace.
A company network.

2. Detect was generated by:
RealSecure

3. Probability the source address was spoofed:
Low, ICMP floods use unwitting amplification hosts to send the attack to the target, the packets used to initiate the attack has the target source spoofed as the source of those packets so that the replies go to the target.

4. Description of attack:

Increasing numbers of internal desktop machines became the target of Smurf attacks as detected by RealSecure. Investigations of the source of t he ICMP packets traced them to University, Cable Modem and Dial -up machines. The classic targets of compromise for use in DDoS attacks. However router logs indicated that the floods stopped at 100 packets every time and therefore were unlikely to affect sy stem performance. Further investigation using Snoop identified that ICMP request packets were being sent by the internal desktop machines and forensic examination discovered that they were running Napster.

Before Napster initiates a download it pings the hosts that offer that file to identify the fastest downloads. As RealSecure was stateless in reporting ICMP floods, the flurry of replies from multiple machines around the world triggered this signature.

It is an interesting fact that the same hosts that provide the ripe targets for DDoS zombies are exactly the same sources of the majority of Napster servers.

5. Attack mechanism:

High volumes of ICMP replies from many hundreds of machines globally arrive at the target machine. This high volume of traffi c floods networks and machine resources.

6. Correlations:

http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-1999-0513
http://www.cert.org/advisories/CA-1998-01.html

7. Evidence of active targeting:

Yes, only existing machines on the internal protected network were attacked.

8. Severity:

(Criticality + Lethality) - (System Countermeasures + Network Countermeasu res) = Severity

(3 + 3) - (3 + 1) = 2

Criticality = 3, desktop machines are not critical to the organisation
Lethality = 3, Floods can delay a machine but ICMP based ones are unlikely to cause a crash
System Countermeasures = 3, Desktop hosts are mo stly up to date with system patches
Network Countermeasures = 1, At this time ICMP replies were allowed through the firewall without requiring a corresponding request. This allowed the confusion to occur over the stimulus that caused these echo replies. T his has now been upgraded to stateful ICMP inspection so that any Smurf attacks will be blocked..

9. Defensive recommendation:

Upgrade firewalls to allow only stateful replies to ICMP. Bill Burns (http://people.netscape.com/shadow/work/inspect/fw1_icmp.html) has the details of how to do this for Firewall-1. Ban Napster from internal networks, as it can be used to distribute copyright infringing material.

10. Multiple choice te st question:

If you see traffic from 24.93.35.210 a working hypothesis for the reason for the traffic should be ?

A) A compromised DDoS zombie is attacking from that address.
B) A trusted business partner is trying to exchange information using EDI.
C) The addres s is RFC1918 Private address space and must be spoofed.
D) This a cable modem home user on the RoadRunner network.

Answer D), This address resolves to

ServiceCo LLC - Road Runner (NET-ROAD-RUNNER-3-A)

13241 Woodland Park Road
Herndon, VA 20171
US

Netname: ROAD-RUNNER-3-A
Netblock: 24.92.160.0 - 24.95.255.255
Maintainer: SCRR

Detect 2 - DDoS Backscatter

```
snoop -v -x 0

ETHER:  ----- Ether Header -----
ETHER:
ETHER:  Packet 124 arrived at 17:26:3.82
ETHER:  Packet size = 60 bytes
ETHER:  Destination = 8:0:XX:XX:XX:XX, XXX
ETHER:  Source      = 8:0:XX:XX:XX:XX, XXX
ETHER:  Ethertype = 0800  IP
ETHER:
IP:   ----- IP Header -----
IP:
IP:   Version = 4
IP:   Header length = 20 bytes
IP:   Type of service = 0x00
IP:         xxx. .... = 0  prece dence
IP:         ...0 .... = normal delay
IP:         .... 0... = normal throughput
IP:         .... .0.. = normal reliability
IP:   Total length = 40 bytes
IP:   Identification = 12774
IP:   Flags = 0x0
IP:         .0.. .... = may fragment
IP:         ..0. .... = last fragment
IP:   Fragment offset = 0 bytes
IP:   Time to live = 241 seconds/hops
IP:   Protocol = 6  TCP
IP:   Header checksum = 4d33
IP:   Source address = 194.204.128.94, 194.204.128.94
IP:   Destination address = MY.NET.107.91, MY.NET.107  .91
IP:   No options
IP:
TCP:  ----- TCP Header -----
TCP:
TCP:  Source port = 40728
TCP:  Destination port = 34587
TCP:  Sequence number = 0
TCP:  Acknowledgement number = 3400998452
TCP:  Data offset = 20 bytes
TCP:  Flags = 0x14
TCP:        ..0. .... =  No urgent pointer
TCP:        ...1 .... = Acknowledgement
TCP:        .... 0... = No push
TCP:        .... .1.. = Reset
TCP:        .... ..0. = No Syn
TCP:        .... ...0 = No Fin
TCP:  Window = 0
TCP:  Checksum = 0x55fa
```

```
TCP:  Urgent pointer = 0
TCP:  No options
TCP:


          0: 0800 209b e138 0800 8713 d7c4 0800 4500    .. .8........E.
         16: 0028 31e6 0000 f106 4d33 c2cc 805e XXXX    . 1.....M3...^.0
         32: 6b5b 9f18 871b 0000 0000 cab7 1e34 5014    k[...........4P.
         48: 0000  55fa 0000 1955 a012 2798              ..U....U..'.
```

1. Source of Trace.

A company network

2. Detect was generated by:

A review of firewall logs and then Solaris snoop. Snoop is a tool that is broadly equivalent to TCPDump and is shipped with Solaris by default.

3. Probability the source address was spoofed:

Very low, TTL of packets broadly matched the reverse traceroute. Although the source address of the packets that caused this response were definitely spoofed.

4. Description of attack:

The network is receiving many thousands of packets with ACK and RST set, however no corresponding SYNs have been sent. The possibility that it was a scan of some sort was discounted once it became clear that no useful data was being returned, the scan was so slow that it would take thousands of years to complete and that the source of the ACK/RST traffic was a router.

5. Attack mechanism:

Spoofed SYNs are sent to the target router to fill state tables and provoke a denial of service. A router is often the target for this as they tend to have less logging enabled than end host devices and by disabling ISP border routers even end users with DHCP addresses can be effectively disconnected from the network with no recourse within their power.

The attack has a second order feature in that if your network is selected as the spoofed source of the traffic then reputational damage can ensue if the attacked system administrators do not realise that your network isn't the source of this attack.

In extreme cases the backscatter from the attack can flood the spoofed sourced network, however in many cases the attacker chooses a wide range of spoofed sources so that load is distributed across many networks without significant issue for any one of them.

6. Correlations:

http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-1999-0116

http://www.cert.org/advisories/CA -1996-21.html

7. Evidence of active targeting:

Low, as the entire range of addresses was hit it is likely that the spoofed source would have been many networks and that it was only chance that our network was involved.

Recent papers highlight that this activity is very common,
http://www.cs.ucsd.edu/~savage/papers/UsenixSec01.pdf

8. Severity:

(Criticality + Lethality) - (System Countermeasures + Network Countermeasures) = Severity

(5 + 2) - (3 + 4) = 0

Criticality = 5 Every host within the address range had packets launched at it, including all critical internal servers
Lethality = 2 Alth ough from a technical prospective the attack would score a 0 as low rates of ACK,RST do nothing to a host, as a PR attack it can be lethal.
System Countermeasures = 3 Hosts run a mix of patch levels although some have host based IDS.
Network Countermeasur es = 4 A correctly configured firewall stopped these packets and raised alerts from the logs so that a response could be made. Even if the packets had penetrated the firewall internal routers would have dropped them as they did not originate from the appl ication proxy servers within the DMZ.

9. Defensive recommendation:

Maintain a well configured perimeter firewall and review the logs to identify such behaviour to eliminate the technical risk. The PR and political risk can be mitigated by having a good w orking relationship with the abuse department of your ISPs so that they are likely to call you rather than assume you are launching an attack if such a thing is reported to them. ISPs can also filter the traffic within their network and reduce the bandwidt h consumed. However for the low volume of traffic it is likely that the overhead of the additional filter on the ISP router would outweigh the performance benefit of dropping this traffic.

10. Multiple choice test question:

If you see a port scan with AC K/RST set across your network range, the best response is to

A) Complain to the source address that they are scanning you and request that they take administrative action against their staff launching scans at your network.
B) Shun that address by reconfiguring the firewall and pass the address to other organisations so that they can add it to their watch lists but do not inform the

attacker organisation in case you alert them to the fact you have detected their attack.

C) Check for SYNs originating from your own m achines, and if there aren't any report that you are the spoofed source SYN flood to your ISP and the source of the ACK/RSTs.

D) Disconnect your network as it is a sign of a total compromise of your critical servers.

Answer: C)

Detect 3 - Welcome all guests !

```
Event: Guest user login
Event Time: Monday, September 17, 2001 10:48:18
Event Priority: 1
SystemAgent = DomainController
Destination = Default
```

1. Source of Trace.

A company network.

2. Detect was generated by:

RealSecure

3. Probability the source address was spoofed:

0, a full login occurred and files were modified making the chances of a correctly spoofed three-way handshake very low.

4. Description of attack:

The default posture of the majority of NT security administrators is to rename and disable the guest account. The guest account is a special built-in account in NT that can be used to access some resources if not controlled carefully. Within the network a large number of alerts that the Guest account was being used were reported. Investigation highlighted that these access attempts where in fact by a new member of staff (Paul Guesten), with NT username GuestenP.

5. Attack mechanism:

By using default accounts an attacker can gain privileges that are normally reserved for registered users.
http://cve.mitre.org/cgi-bin/cvename.cgi?name=CAN-1999-0546

6. Correlations:

Although many references can be found to diabling the Guest account, e.g.

http://www.cert.org/tech_tips/win_configuration_guidelines.html#V . No evidence could be found of other teams experiencing this problem with Guest accounts.

7. Evidence of active targeting:

Yes, critical internal systems were connected, authenticated and accessed.

8. Severity:
(Criticality + Lethality) - (System Countermeasures + Network Countermeasures) = Severity

(5 + 4) - (3 + 1) = 5

Criticality = 5, key servers were accessed
Lethality = 4, if someone can re-enable and gain access to a Guest account they are using a new exploit that may be used to gain access to other disabled accounts.
System Countermeasures = 3, Host-based IDS raised the alarm although access was achieved.
Network Countermeasures = 1, No internal filters stopped this user from being able to access the resources using the Guest account.

9. Defensive recommendation:

If staff are unwilling to change their surname then ensure that IDS pattern matching is correct. This kind of false positive is annoying but the complementary false negative that as the Guest account has been renamed its use will not trigger the signature is more significant. Guest and other default accounts should have difficult passwords set, be disabled and if possible re-named to provide multiple layers of obscuration to attackers.

10. Multiple choice test question:

Which of these regular expressions is the right way to match a guest login from the NT logs ?

a) GUEST
b) @(GUEST|guest)
c) @(GUEST|guest)$
d) S-1-5-21-[0-9]*-[0-9]*-[0-9]*-501

Answer d) - NT accounts should be tracked by SID, not name to avoid the risk of renamed accounts avoiding detection.

Detect 4 - ISS RealSecure Kills

```
'RealSecure_Kill' event detected by the RealSecure engine at
'SensorXXXX'.
Details:
      Source Address: 203.34.60.45
      Source Port: 4471
      Source MAC Address: 08: 00:XX:XX:XX:XX
      Destination Address: MY.NET.254.4
```

```
      Destination Port: HTTP (80)
      Destination MAC Address: 08:00:XX:XX:XX:XX
      Protocol: TCP (6)
      Priority: high
      Actions mask: 0xa44
      Event Specific Information:
            CUSTID: 1401902

'RealSecure_Kill' event dete cted by the RealSecure engine at
'SensorXXXX'.
Details:
      Source Address: 202.56.229.83
      Source Port: 2045
      Source MAC Address: 08:00:XX:XX:XX:XX
      Destination Address: MY.NET.254.4
      Destination Port: HTTP (80)
      Destination MAC Address: 08:00:XX:XX:XX:XX
      Protocol: TCP (6)
      Priority: high
      Actions mask: 0xa44
      Event Specific Information:
            CUSTID: 6514002

'RealSecure_Kill' event detected by the RealSecure engine at
'SensorXXXX'.
Details:
      Source Address: 62.172.160.50
      Source Port: 13816
      Source MAC Address : 08:00:XX:XX:XX:XX
      Destination Address: MY.NET.254.4
      Destination Port: HTTP (80)
      Destination MAC Address: 08:00:XX:XX:XX:XX
      Protocol: TCP (6)
      Priority: high
      Actions mask: 0xa44
      Event Specific Information:
            CUSTID: 5339601
```

1. Source of Trace.

A company network.

2. Detect was generated by:

RealSecure

3. Probability the source address was spoofed:

Very high, RealSecure Kills are generated by the network sensor to reset traffic it is monitoring that does not meet policy. However, as the RS -KILL originates very near the source of the traffic it can be hard to identify by TTL cross referencing.

4. Description of attack:

MY.NET.254.4 is a public facing web -server containing many documents, images and scripts. Reserve Bank of Australia, Nestle Indi a ltd And Abbey National plc are running RealSecure to defend their networks. They are detecting a SYN flood

launched from their network to our web server and are sending automatic response generated RS-Kills to try to stop the connection. This is an error because SYN flood nearly always has a spoofed source so by sending RSTs to the spoofed source they are adding to the network load without any hope of resolving the problem and this is not a SYN flood as the ACKs were sent by our web server, however the large number of objects on each page causes a large number of SYNs within a short timeframe, triggering the alarm.

5. Attack mechanism:

A RST is spoofed into the packet stream to tear down the connection.

6. Correlations:
Northcutt, p168 mentions this attack and RealSecure users discuss the same issue http://archives.neohapsis.com/archives/iss/2000-q3/0229.html

7. Evidence of active targeting:

Yes, one of our key web servers is attacked with this spurious RSTs.

8. Severity:

(Criticality + Lethality) - (System Countermeasures + Network Countermeasures) = Severity

(5 + 4) - (1 + 3) = 5

Criticality = 5, critical servers were attacked.
Lethality = 4, although it will not kill the box, it will stop all traffic from their site to our servers costing us lost business from these users.
System Countermeasures = 1, The system happily accepts and handles RSTs and will tear down connections based on the contents.
Network Countermeasures = 3, although the firewall allowed the RST in, it was detected as spurious and enough information gathered to allow a response to the sender.

9. Defensive recommendation:

Do not enable automated response on signatures that are highly likely to contain a spoofed source. Tune SYN flood detectors to only alert on high floods. ISS recommends disabling the "Flag RS Kills" within the RS console. This is not advised as it reduces the chance of falsely attacked systems from being able to alert the sender of a poorly configured IDS. Any attacker who is capable of collecting RST information is more than able to drop RSTs and continue to attack.

10. Multiple choice test question:

If setting up an automated response which two of these signatures would make the most sense to send a RST to ?

A) SYN Flood
B) BackOrifice connection
C) Email containing a Virus
D) SMURF

Answer B & C. By sending a RST on a Trojan remote control session access can be disrupted and the resetting of an SMTP connection will eventually result in the mail containing the virus being returned to sender.

Detect 5 - CodeRed bottom feeders

Attacker51 MY.NET.254.1
Attacker51 MY.NET.254.2
Attacker51 MY.NET.254.3
Attacker51 MY.NET.254.4
Attacker51 MY.NET.254.5
Attacker51 MY.NET.254.6
Attacker51 MY.NET.254.7
Attacker51 MY.NET.254.8
Attacker51 MY.NET.254.9
Attacker51 MY.NET.254.10
Attacker51 MY.NET.254.11
Attacker51 MY.NET.254.12
Attacker51 MY.NET.254.13
Attacker51 MY.NET.254.14
Attacker51 MY.NET.254.15
Attacker51 MY.NET.254.16
Attacker51 MY.NET.254.17
Attacker51 MY.NET.254.18
Attacker51 MY.NET.254.19
Attacker51 MY.NET.254.20
Attacker51 MY.NET.254.21
Attacker51 MY.NET.254.22
Attacker51 MY.NET.254.23
Attacker51 MY.NET.254.24
Attacker51 MY.NET.254.25
Attacker51 MY.NET.254.26
Attacker51 MY.NET.254.27
Attacker51 MY.NET.254.28
Attacker51 MY.NET.254.29
Attacker51 MY.NET.254.30
Attacker51 MY.NET.254.31
Attacker51 MY.NET.254.32
Attacker51 MY.NET.254.33
Attacker51 MY.NET.254.34
Attacker51 MY.NET.254.35
Attacker51 MY.NET.254.36
Attacker51 MY.NET.254.37
Attacker51 MY.NET.254.38
Attacker51 MY.NET.254.39

Attacker51 MY.NET.254.40
Attacker51 MY.NET.254.41
Attacker51 MY.NET.254.42
Attacker51 MY.NET.254.43
Attacker51 MY.NET.254.44
Attacker51 MY.NET.254.45
Attacker51 MY.NET.254.46
Attacker51 MY.NET.254.47
Attacker51 MY.NET.254.48
Attacker51 MY.NET.254.49
Attacker51 MY.NET.254.50
Attacker51 MY.NET.254.51
Attacker51 MY.NET.254.52
Attacker51 MY.NET.254.53
Attacker51 MY.NET.254.54
Attacker51 MY.NET.254.55
Attacker51 MY.NET.254.56
Attacker51 MY.NET.254.57
Attacker51 MY.NET.254.58
Attacker51 MY.NET.254.59
Attacker51 MY.NET.254.60
Attacker51 MY.NET.254.61
Attacker51 MY.NET.254.62
Attacker51 MY.NET.254.63
Attacker51 MY.NET.254.64
Attacker51 MY.NET.254.65
Attacker51 MY.NET.254.66
Attacker51 MY.NET.254.67
Attacker51 MY.NET.254.68
Attacker51 MY.NET.254 .69
Attacker51 MY.NET.254.70
Attacker51 MY.NET.254.71
Attacker51 MY.NET.254.72
Attacker51 MY.NET.254.73
Attacker51 MY.NET.254.74
Attacker51 MY.NET.254.75
Attacker51 MY.NET.254.76
Attacker51 MY.NET.254.77
Attacker51 MY.NET.254.78
Attacker51 MY.NET .254.79
Attacker51 MY.NET.254.80
Attacker51 MY.NET.254.81
Attacker51 MY.NET.254.82
Attacker51 MY.NET.254.83
Attacker51 MY.NET.254.84
Attacker51 MY.NET.254.85
Attacker51 MY.NET.254.86
Attacker51 MY.NET.254.87
Attacker51 MY.NET.254.88
Attacker51 MY.NET.254.89
Attacker51 MY.NET.254.90

Attacker51 MY.NET.254.91
Attacker51 MY.NET.254.92
Attacker51 MY.NET.254.93
Attacker51 MY.NET.254.94
Attacker51 MY.NET.254.95
Attacker51 MY.NET.254.96
Attacker51 MY.NET.254.97
Attacker51 MY.NET.254.98
Attacker5 1 MY.NET.254.99
Attacker51 MY.NET.254.100
Attacker51 MY.NET.254.101
Attacker51 MY.NET.254.102
Attacker51 MY.NET.254.103
Attacker51 MY.NET.254.104
Attacker51 MY.NET.254.105
Attacker51 MY.NET.254.106
Attacker51 MY.NET.254.107
Attacker51 MY.NET.254.1 08
Attacker51 MY.NET.254.109
Attacker51 MY.NET.254.110
Attacker51 MY.NET.254.111
Attacker51 MY.NET.254.112
Attacker51 MY.NET.254.113
Attacker51 MY.NET.254.114
Attacker51 MY.NET.254.115
Attacker51 MY.NET.254.116
Attacker51 MY.NET.254.117
Attacker5 1 MY.NET.254.118
Attacker51 MY.NET.254.119
Attacker51 MY.NET.254.120
Attacker51 MY.NET.254.121
Attacker51 MY.NET.254.122
Attacker51 MY.NET.254.123
Attacker51 MY.NET.254.124
Attacker51 MY.NET.254.125
Attacker51 MY.NET.254.126
Attacker51 MY.NET.254. 127
Attacker51 MY.NET.254.128
Attacker51 MY.NET.254.129
Attacker51 MY.NET.254.130
Attacker51 MY.NET.254.131
Attacker51 MY.NET.254.132
Attacker51 MY.NET.254.133
Attacker51 MY.NET.254.134
Attacker51 MY.NET.254.135
Attacker51 MY.NET.254.136
Attacker 51 MY.NET.254.137
Attacker51 MY.NET.254.138
Attacker51 MY.NET.254.139
Attacker51 MY.NET.254.140
Attacker51 MY.NET.254.141

Attacker51 MY.NET.254.142
Attacker51 MY.NET.254.143
Attacker51 MY.NET.254.144
Attacker51 MY.NET.254.145
Attacker51 MY.NET.254 .146
Attacker51 MY.NET.254.147
Attacker51 MY.NET.254.148
Attacker51 MY.NET.254.149
Attacker51 MY.NET.254.150
Attacker51 MY.NET.254.151
Attacker51 MY.NET.254.152
Attacker51 MY.NET.254.153
Attacker51 MY.NET.254.154
Attacker51 MY.NET.254.155
Attacke r51 MY.NET.254.156
Attacker51 MY.NET.254.157
Attacker51 MY.NET.254.158
Attacker51 MY.NET.254.159
Attacker51 MY.NET.254.160
Attacker51 MY.NET.254.161
Attacker51 MY.NET.254.162
Attacker51 MY.NET.254.163
Attacker51 MY.NET.254.164
Attacker51 MY.NET.25 4.165
Attacker51 MY.NET.254.166
Attacker51 MY.NET.254.167
Attacker51 MY.NET.254.168
Attacker51 MY.NET.254.169
Attacker51 MY.NET.254.170
Attacker51 MY.NET.254.171
Attacker51 MY.NET.254.172
Attacker51 MY.NET.254.173
Attacker51 MY.NET.254.174
Attack er51 MY.NET.254.175
Attacker51 MY.NET.254.176
Attacker51 MY.NET.254.177
Attacker51 MY.NET.254.178
Attacker51 MY.NET.254.179
Attacker51 MY.NET.254.180
Attacker51 MY.NET.254.181
Attacker51 MY.NET.254.182
Attacker51 MY.NET.254.183
Attacker51 MY.NET.2 54.184
Attacker51 MY.NET.254.185
Attacker51 MY.NET.254.186
Attacker51 MY.NET.254.187
Attacker51 MY.NET.254.188
Attacker51 MY.NET.254.189
Attacker51 MY.NET.254.190
Attacker51 MY.NET.254.191
Attacker51 MY.NET.254.192

Attacker51 MY.NET.254.193
Attacker51 MY.NET.254.194
Attacker51 MY.NET.254.195
Attacker51 MY.NET.254.196
Attacker51 MY.NET.254.197
Attacker51 MY.NET.254.198
Attacker51 MY.NET.254.199
Attacker51 MY.NET.254.200
Attacker51 MY.NET.254.201
Attacker51 MY.NET.254.202
Attacker51 MY.NET. 254.203
Attacker51 MY.NET.254.204
Attacker51 MY.NET.254.205
Attacker51 MY.NET.254.206
Attacker51 MY.NET.254.207
Attacker51 MY.NET.254.208
Attacker51 MY.NET.254.209
Attacker51 MY.NET.254.210
Attacker51 MY.NET.254.211
Attacker51 MY.NET.254.212
Attacker51 MY.NET.254.213
Attacker51 MY.NET.254.214
Attacker51 MY.NET.254.215
Attacker51 MY.NET.254.216
Attacker51 MY.NET.254.217
Attacker51 MY.NET.254.218
Attacker51 MY.NET.254.219
Attacker51 MY.NET.254.220
Attacker51 MY.NET.254.221
Attacker51 MY.NET .254.222
Attacker51 MY.NET.254.223
Attacker51 MY.NET.254.224
Attacker51 MY.NET.254.225
Attacker51 MY.NET.254.226
Attacker51 MY.NET.254.227
Attacker51 MY.NET.254.228
Attacker51 MY.NET.254.229
Attacker51 MY.NET.254.230
Attacker51 MY.NET.254.231
Attacker51 MY.NET.254.232
Attacker51 MY.NET.254.233
Attacker51 MY.NET.254.234
Attacker51 MY.NET.254.235
Attacker51 MY.NET.254.236
Attacker51 MY.NET.254.237
Attacker51 MY.NET.254.238
Attacker51 MY.NET.254.239
Attacker51 MY.NET.254.240
Attacker51 MY.NET.254.241
Attacker51 MY.NET.254.242
Attacker51 MY.NET.254.243

Attacker51 MY.NET.254.244
Attacker51 MY.NET.254.245
Attacker51 MY.NET.254.246
Attacker51 MY.NET.254.247
Attacker51 MY.NET.254.248
Attacker51 MY.NET.254.249
Attacker51 MY.NET.254.250
Attacker51 MY.NET.254.251
Attacker51 MY.NET.254.252
Attacker51 MY.NET.254.253
Attacker51 MY.NET.254.254

1. Source of Trace.

A company network

2. Detect was generated by:

Firewall logs of attempts to exploit .ida holes.

3. Probability the source add ress was spoofed:

Low, the TCP handshake was completed to allow HTTP requests to be sent.

4. Description of attack:

Code Red uses the buffer overflow within Microsoft Index server .ida files. Much attention was focused on the worm that exploited this vu lnerability but manual attempts can also be made to scan networks..

5. Attack mechanism:

This process has been well documented at http://xforce.iss.net/alerts/advise89.php , http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/bulletin/MS01-033.asp and http://www.eeye.com/html/Research/Advisories/AL20010717.html The interesting feature of this attack is that although the probe is lifted directly from the worm code the probing is not random but sequential.

6. Correlations:
http://www.incidents.org/react/code_red.php contains valuable information about Code Red and the sterling work done to notify administrators and prevent an Internet meltdown.

7. Evidence of active targeting:

Yes, unlike the majority of Code Red activity which targets random IP addresses, this attacker is using the noise of those scans to run his own tool against our addresses.

8. Severity:

(Criticality + Lethality) - (System Countermeasures + Network Count ermeasures) = Severity

(4 + 4) - (4 + 4) = 0

Criticality = 4, Key web servers and critical networks were scanned.
Lethality = 4, Code Red allows the running of arbitrary code with administrator rights.
System Countermeasures = 4, IIS is disabled on all non-required servers and is patched on those that it is required.
Network Countermeasures = 4, IDS and firewall logs allowed the user carrying out the scan to be identified as an anomaly within the Code Red random probes.

9. Defensive recommendation:

Although protected against an automated worm, be careful to review data that matches the worm pattern to ensure it is not being used to provide cover for other attacks.

10. Multiple choice test question:

The sign of the automated code red worm hitting your network is

A) .ida buffer overflow attempts sequentially against your network
B) .ida buffer overflow attempts from a single server against IIS servers randomly on your network
C) Unicode attempts from multiple servers against IIS servers randomly on your network.
D) .ida buffer overflow attempts from multiple servers against random addresses on your network

Answer D), worms in general do not have targeting information and therefore strike at random rather than targeting vulnerable servers and Code Red uses a .i da buffer overflow.

Assignment 3 - "Analyse This" Scenario

List of chosen files

Alerts
alert.010901.gz
alert.010902.gz
alert.010903.gz
alert.010904.gz
alert.010905.gz

Scans
scans.010901.gz
scans.010902.gz
scans.010903.gz
scans.010904.gz
scans.010905.gz

OOS
oos_Sep.1.2001.gz
oos_Sep.2.2001.gz
oos_Sep.3.2001.gz
oos_Sep.4.2001.gz
oos_Sep.5.2001.gz

Overview of analysis

Describe process

The process used to analyse the data was to take the chosen days files and concatenate them together. Ad hoc Sed and Awk scripts were used to extract key information of source, target, date and time and type of attack. The header rows were removed.

Analysis to remove noise and identify interesting features. A first cut identified the following events that are not worth y of further consideration. This assumes that this kind of traffic is authorised for this network.

20453 MISC traceroute	If this kind of traffic externally is not required then UDP and ICMP should be blocked at the perimeter.
15458 CS WEBSERVER - external web traffic	Access to the web server via web is expected
14853 INFO MSN IM Chat data	Instant Messenger is a discussion tool that is useful, although offers little protection to the confidentiality or integrity of the data sent
8603 INFO napster login	Napster is a tool to share music
1812 INFO Napster Client Data	As Napster Login
1782 INFO Inbound GNUTella Connect accept	GNUTella is a freeware clone of Napster

1296	ICMP traceroute	As per MISC traceroute
1214	INFO Possible IRC Access	IRC is a discus sion and file sharing tool.
520	INFO Outbound GNUTella Connect accept	As per Inbound GNUTella
376	INFO FTP anonymous FTP	All connections are inwards but it would be wise to check the top internal hosts to ensure that the data they are sharing via this me thod is appropriate. 15 MY.NET.130.122 32 MY.NET.253.105 35 MY.NET.100.165 36 MY.NET.70.148 53 MY.NET.99.85
113	CS WEBSERVER - external ftp traffic	If MY.NET.100.165 is an authorised FTP server then this is fine.
59	WEB-MISC count.cgi access	Given the hi gh number of these alerts from a small number of hosts it is likely that this is a false alarm and that count.cgi is enabled to allow a webhit counter to work.
43	WEB-CGI scriptalias access	Same as count.cgi
38	WEB-FRONTPAGE fpcount.exe access	Same as co unt.cgi
37	WEB-FRONTPAGE _vti_rpc access	Same as count.cgi
33	INFO napster upload request	As Napster Login
30	WEB-IIS _vti_inf access	Same as count.cgi
26	WEB-CGI csh access	Same as count.cgi
25	WEB-FRONTPAGE shtml.dll	Same as count.cgi
19	X11 outgoi ng	X should not be used over the wide area network as it offers no confidentiality, SSH should be used to tunnel X for this purpose.
19	INFO Outbound GNUTella Connect request	As per Inbound GNUTella
17	WEB-CGI redirect access	Same as count.cgi
17	INFO - Web Cmd completed	Same as count.cgi
16	WEB-FRONTPAGE fourdots request	Same as count.cgi
11	INFO Inbound GNUTella Connect request	As per Inbound GNUTella
10	WEB-IIS encoding access	Same as count.cgi
7	INFO napster new user login	As Napster Login
5	MISC Source Port 20 to <1024	Legitimate FTP traffic from academic software site
4	WEB-CGI upload.pl access	Same as count.cgi
4	WEB-CGI files.pl access	Same as count.cgi
3	WEB-CGI ksh access	Same as count.cgi
3	WEB-CGI archie access	Same as count.cgi
2	WEB-FRONTPAGE author.exe access	Same as count.cgi
2	WEB-CGI w3-msql access	Same as count.cgi
2	WEB-CGI finger access	Same as count.cgi
2	WEB-CGI calendar access	Same as count.cgi
2	External FTP to HelpDesk MY.NET.53.29	Access to FTP site
1	TELNET access	Access from a University.
1	SMTP chameleon overflow	False alarm, mail traffic from a mail server at a travel company is to be expected.
1	External FTP to HelpDesk MY.NET.83.197	Access to FTP site
1	External FTP to HelpDesk MY.NET.70.50	Access to FTP si te

Top alerts were .ida buffer overflow exploit and WEB -MISC cmd exec which accounted for 76% of the data so this was pulled from the core data and reviewed, no attacks from the inside outward,
http://cve.mitre.org/cgi-bin/cvename.cgi?name=CAN-2001-0500 .ida buffer, more commonly known as CodeRed and CMD exploits.

By comparing the data of these two events over time a clear pattern emerges

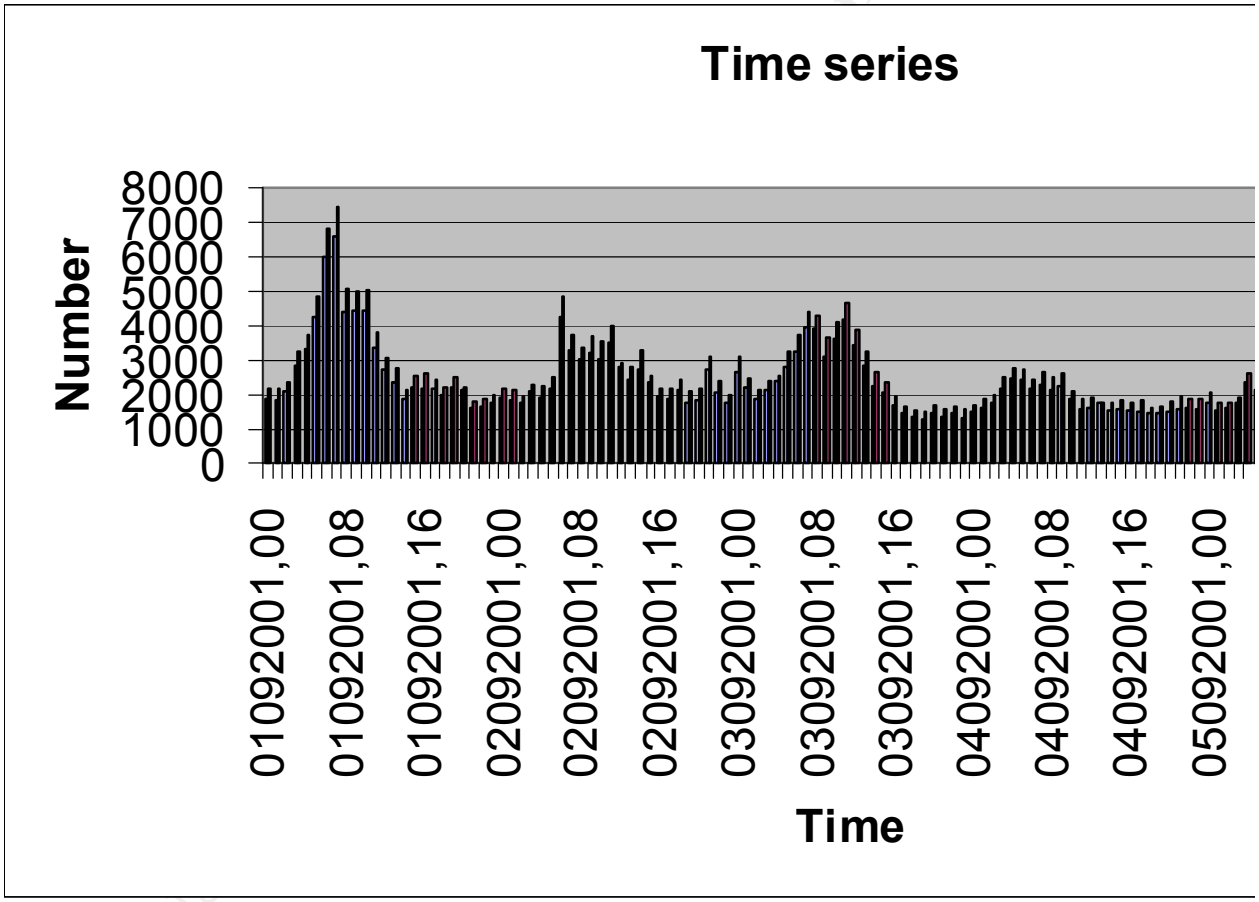

If increased in size it can be seen that the peaks and dips match on both these attacks, and than a regular peak occurs at around 8am on each of the days, however the size of this peak decreases over time.

It is assumed that this daily increase as due to addition al machines being turned on during the day and being able to scan but stopping their scans as they were powered off in the evening. The decrease in the size of the peaks can be explained by the steady fixing and containment of the infected servers. Sadly, it is likely that many of the servers that have gone quiet are because they have drawn attention to themselves as vulnerable boxes and have been broken into, backdoored and silenced to be used for DDoS zombies.

Top 10 talkers from Scans, Alerts and OOS

Top ten internal talkers in Scans

Number	Src Address
70655	MY.NET.160.114
31409	MY.NET.218.78
27329	MY.NET.218.50
21385	MY.NET.202.102
17837	MY.NET.234.198
15110	MY.NET.233.202
13195	MY.NET.201.42
12955	MY.NET.212.150
8047	MY.NET.205.230
7979	MY.NET.204.190

Top ten external talkers in Scans

Number	Src Address
15469	212.199.28.76
14869	216.162.3.20
6446	217.128.232.163
5949	205.188.246.121
5547	64.37.156.9
5226	210.95.106.2
4711	130.89.229.75
2602	129.2.144.201
2458	130.161.37.101
2200	217.11.167.47

Top ten internal talkers Alerts

Number	Src Address
16091	MY.NET.14.1
14701	MY.NET.16.5
5302	MY.NET.226.18
4690	MY.NET.30.2
3591	MY.NET.98.190
3431	MY.NET.208.82
2401	MY.NET.226.118
2189	MY.NET.235.14
1991	MY.NET.235.106
1953	MY.NET.218.50

Top ten external talkers Alerts

Number	Src Address
21934	211.90.176.59
11358	211.90.164.34
9813	211.90.88.43
8904	61.153.17.244

7468 200.250.65.1
6976 211.96.99.59
6677 217.57.15.133
6662 61.153.17.24
6290 200.26.105.130
6013 130.206.73.191

Top Ten talkers in OOS
Number Src Address
71 151.38.11.166
58 198.186.202.147
20 128.46.156.155
13 212.194.4.183
11 151.38.84.194
6 24.147.31.25
5 193.137.96.74
4 24.39.170.205
4 24.28.134.6
4 213.23.38.230

Top alert types from al erts
Number Type
305468 WEB-MISC Attempt to execute cmd
268112 IDS552/web-iis_IIS ISAPI Overflow ida nosize
32311 ICMP Destination Unreachable (Communication Administratively Prohibited)
20678 MISC Large UDP Packet
20453 MISC traceroute
19590 MISC s ource port 53 to <1024
15458 CS WEBSERVER - external web traffic
14853 INFO MSN IM Chat data
13086 Portscan
12258 WEB-MISC prefix-get //
10805 ICMP Echo Request Nmap or HPING2
8603 INFO napster login
6097 Possible trojan server activity
5315 Watchlist 000220 IL-ISDNNET-990517
4706 ICMP Destination Unreachable (Network Unreachable)
4319 High port 65535 tcp - possible Red Worm - traffic
3598 ICMP Destination Unreachable (Host Unreachable)
3380 Tiny Fragments - Possible Hostile Activity
1970 Null scan!
1947 spp_http_decode: IIS Unicode attack detected
1812 INFO Napster Client Data
1782 INFO Inbound GNUTella Connect accept
16433 Other

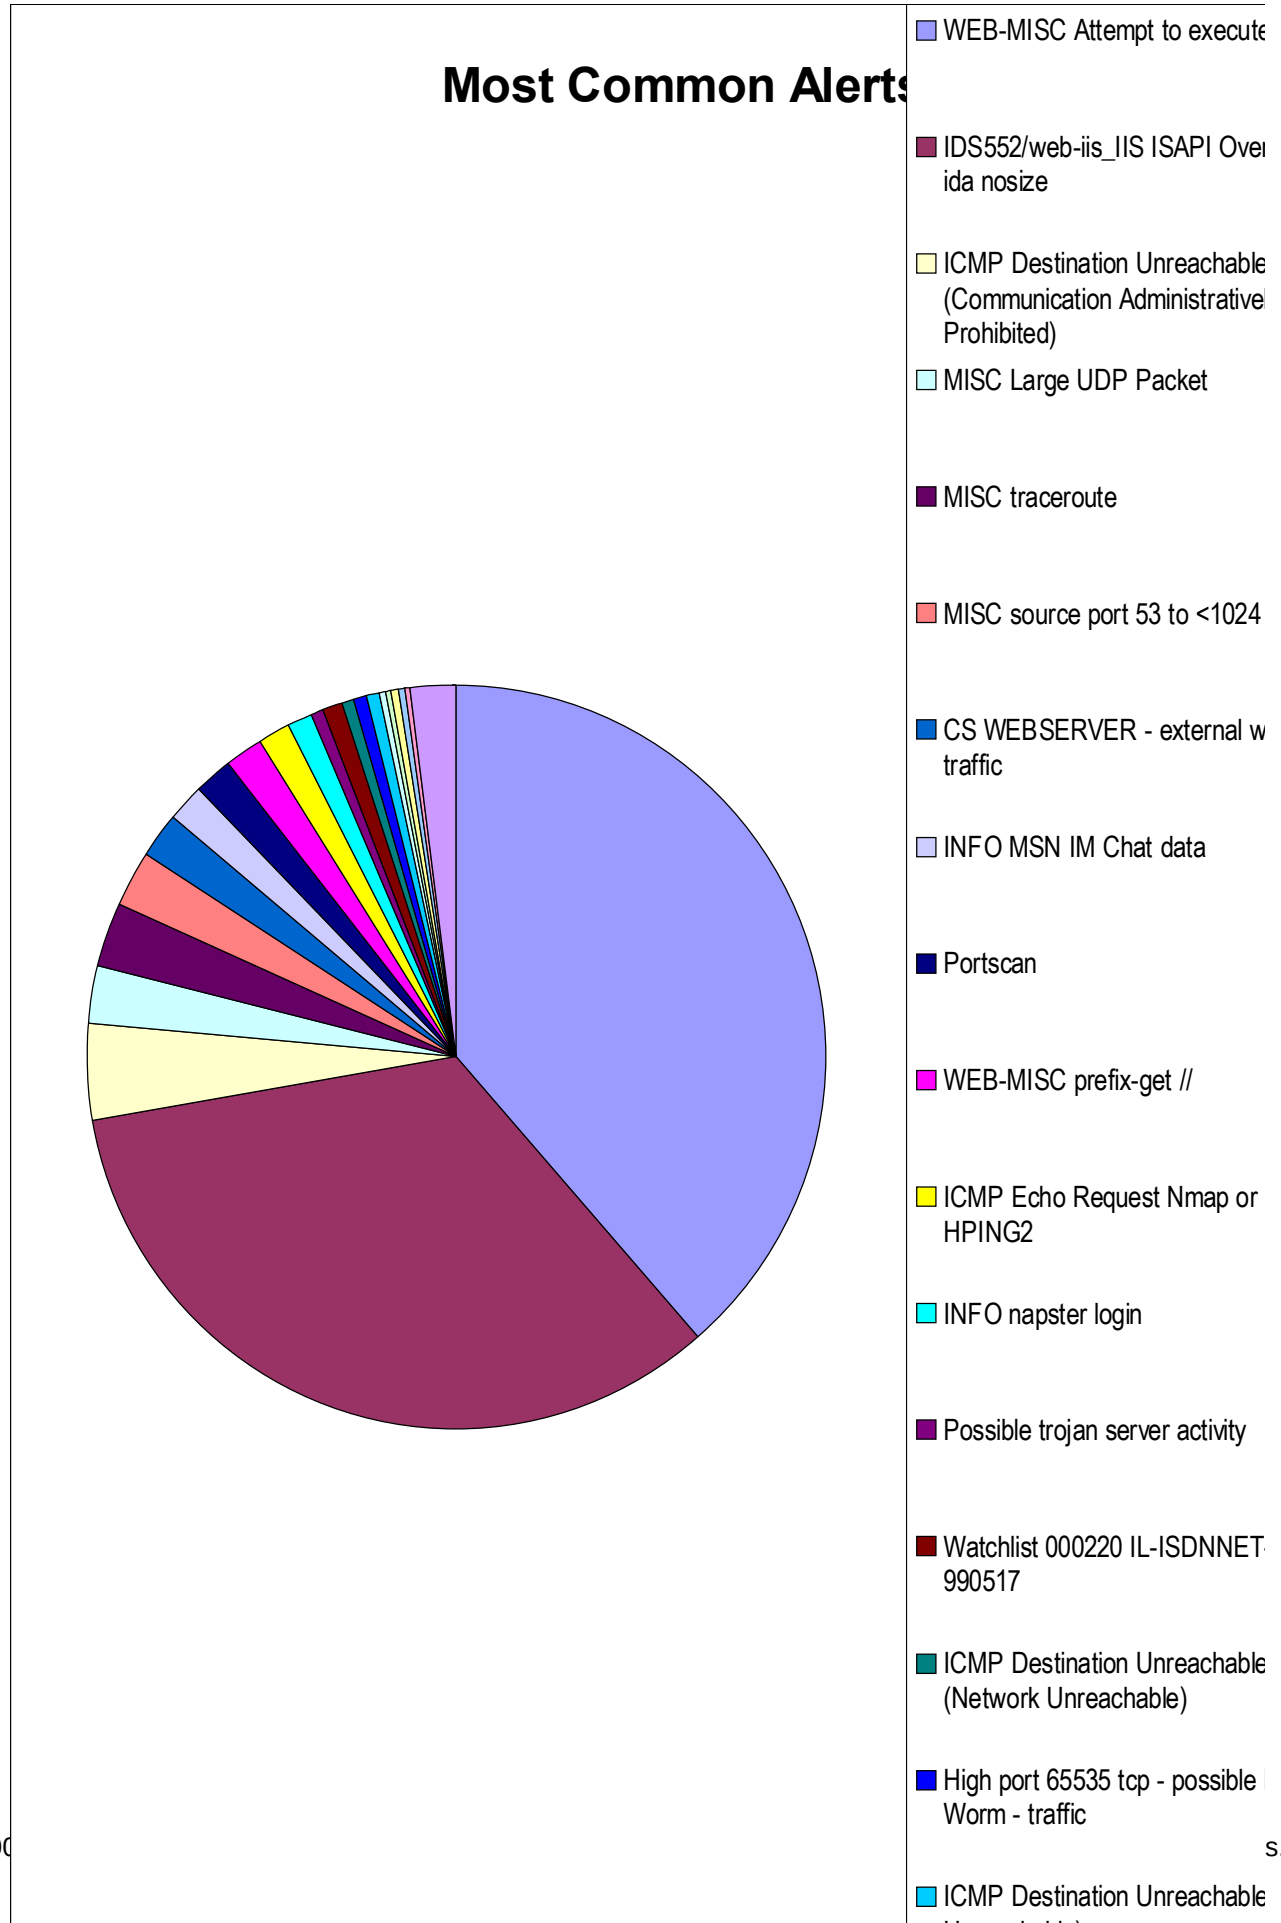

Lookup Information is spread through this write -up, but one network stands out as worthy of its own section.

With a large number of external scans and the majority of the CodeRed probes launched from this network it is definitely worth digging in further into this source to determine what they are up to.

inetnum: 211.90.0.0 - 211.91.255.255
netname: UNICOM
descr: China United Telecommunications Corporation
country: CN
admin-c: XL31-AP
tech-c: XL31-AP
mnt-by: MAINT-CNNIC-AP
changed: xiaqing@cnnic.net.cn 20000414
source: APNIC

person: XiaoMing Li
address: 6F Office Tower 3, Henderson Centre, Beijing China
country: CN
phone: +86-10-65181800-291
fax-no: +86-10-65181800-777
e-mail: lxmlxm@public3.bta.net.cn
nic-hdl: XL31-AP
mnt-by: MAINT-CNNIC-AP
changed: wangch@cnnic.net.cn 20000331
source: APNIC

Correlations to other practicals (209 +)

Link graph (plus analysis)

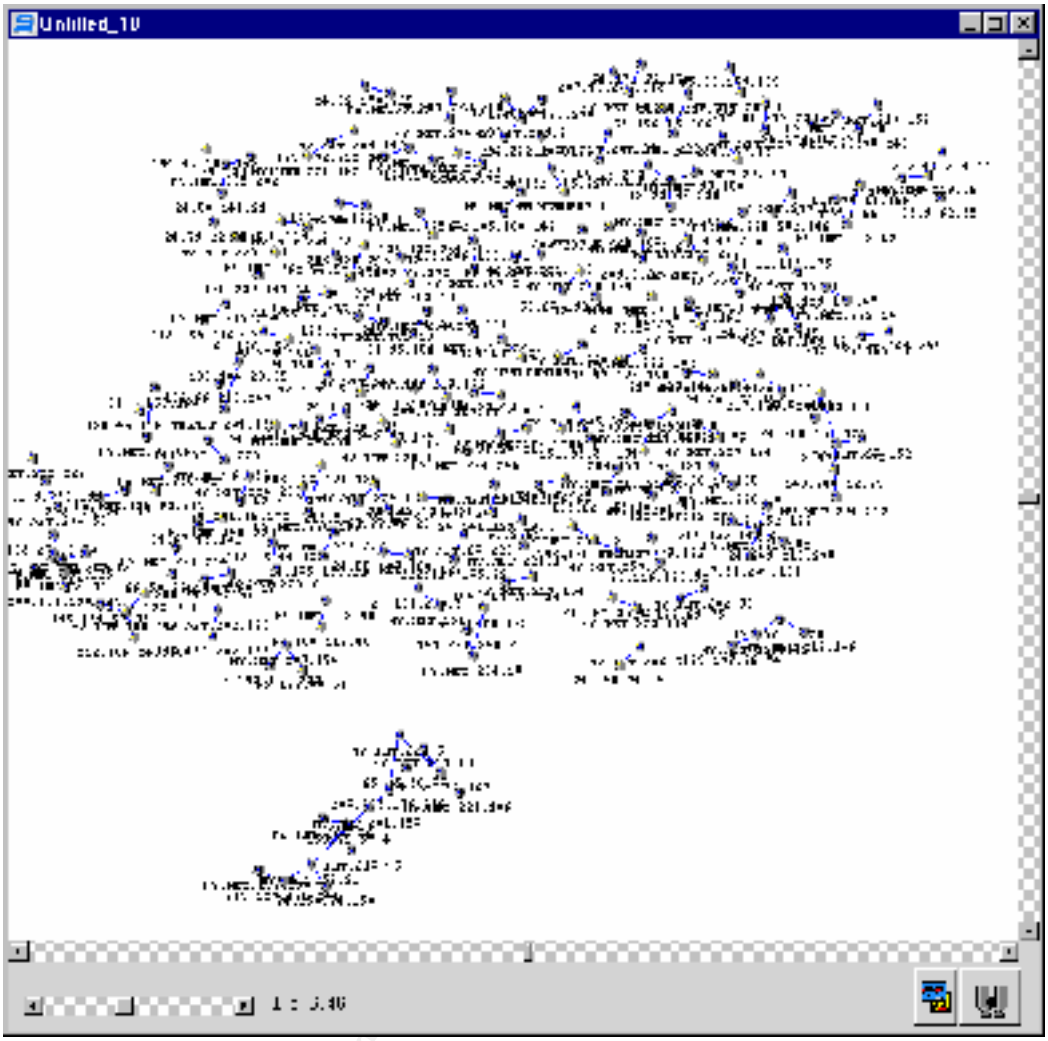

This link graph highlights that the majority of OOS packets are between pairs of hosts but that an interesting artefact can be seen at the bottom of the picture, focusing in provides this.

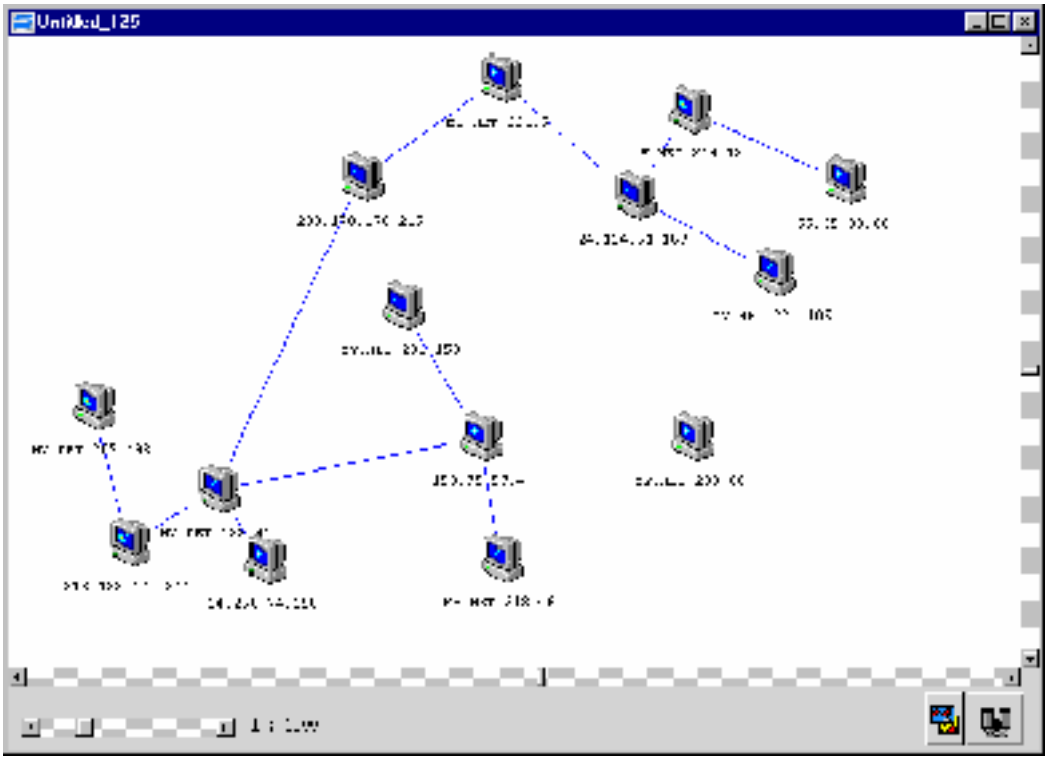

Investigations within this cluster of machines talking to each other highlights that of the 711 OOS packets captured over the 5 days, 167 were to or from port 6346. This port is used by Gnutella servers and clients. It appears that either something within the packet stream is false alarming or someone is hiding within the Gnutella bandwidth to carry out connections.

A review of the packets does not highlight issues beyond the use of the reserved bits within the TCP flags. Otherwise the fact that many hosts on the internal network and on the external network talk to each other is indicative of how popular the major servers for GNUTella content are.

Internal machines that indicate compromise.

In increasing order of volume of scans the following internal hosts have been scanning,

8047 MY.NET.205.230
12955 MY.NET.212.150
31409 MY.NET.218.78
70655 MY.NET.160.114

205.230 has been carrying out normal IM and IRC behaviour but was port scanned and then began a high volume of UDP scans on the ports 723,794,15,547,594,297,310 and 779 on the remote hosts 4.3.90.92 (Genuity DSL system), 24.130.66.57 (MediaOne RoadRunner), 64.219.197.245 (SWBell ppp dialup) and 4.67.139.203 (Genuity Satnet). Although it is possible that these are related to some kind of online game it seems more likely that this is some kind of Trojan trying to "phone home".

Certainly a machine worth checking. MY.NET.212.150 exhibits the same behaviour with UDP traffic to 4.3.90.92, 2 4.130.66.57 as above and also 24.200.31.233 Canadian Videotron CableModem.

MY.NET.218.78 has conducted a UDP scan of port 137 on some 24540 Internet hosts. This machine is either totally compromised or whoever is in charge is acting in a very unethical manner. This scan ran for 13 hours before it ended.

MY.NET.160.114 is conducting large volumes of UDP scans, however they all match the ports used by Half-Life http://www.incidents.org/detect/gamin g.php and are connecting to external systems on Cable Models and at Universities as you would expect for online gaming. Not a security issue.

MY.NET.98.190 and MY.NET.235.14 has launched a scan for SubSeven across a large section of the network. This smells of compromised boxes looking for already compromised machines to attack. Given this lame MO, it is likely that the only way this attacker could have broken into these machines was also via SubSeven, certainly worth checking the latest AV is installed on this box. This behaviour correlates with http://www.shmoo.com/mail/ids/apr00/msg00047.shtml

Defensive recommendations

Tune the IDS system.
A system that produces millions of alarms per week is not useful as it overloads the analyst team and leads to critical events being lost in the noise. The network should be segregated into various risk levels

Install Industry standard firewalls.
Many of the attacks detected and the machines compromised should never have been successfully broken into. Effective defences must be deployed and those devices that cannot deploy appropriate defences should be disconnected until they have done so, otherwise this network is a risk to itself and the rest of the Internet.

Use the tools in place.
Many of the attacks and scans launched from within the network continue for very significant periods of time. Consideration should be given to automatically blocking this behaviour if a manual response cannot be afforded. Scripts to automate this process for Firewalls are included in the documentation of the HoneyNet project.
http://www.enteract.com/~lspitz/intrusion.html