TECHNICAL NOTE 06/02

RESPONSE TO DISTRIBUTED DENIAL OF SERVICE (DDOS) ATTACKS

2002

This paper was previously published by the National Infrastructure Security Co-ordination Centre (NISCC) – a predecessor organisation to the Centre for the Protection of National Infrastructure (CPNI).

Hyperlinks in this document may refer to resources that no longer exist. Please see CPNI's website (www.cpni.gov.uk) for up-to-date information.

NISCC Technical Note 06/02:
Response to Distributed Denial of Service (DDoS) Attacks

Background

This NISCC technical note is intended to provide information to enable organisations in the UK's Critical National Infrastructure to respond to Distributed Denial of Service (DDoS) attacks. The note expands on the advice given in NISCC Technical Note 01/02, "Protecting your Computer Network". This technical note does not promote the use of networking technologies provided by any particular vendor, but some of the URLs referred to in this paper contain advice from particular vendors.

2. DDoS attacks first came to public attention in the first quarter of 2000 when a number of prominent US web sites were attacked [1]. In 2001 the first version of the Code Red worm was designed so that hosts compromised by the worm would perform a denial of service attack on the White House in the USA. More recently, the Slapper worm, which targeted Linux machines in September 2002, was designed to create a massive denial of service network of infected machines controllable by the virus writer [2]. DDoS activity in the UK has also been reported to UNIRAS recently, and on 21 October 2002 four to six of the 13 root Domain Name Servers (DNS) that manage global Internet traffic were attacked.

3. Given the wide availability of DDoS tools there is a significant risk that, if not properly protected, a large number of machines will be compromised and could be used to mount a DDoS attack. Too many machines have not been properly patched or protected, and are exploitable as 'zombies'. They are also vulnerable to the attentions of 'script-kiddies', whose aim is to infect as many machines as possible. The threat of DDoS attacks occurring, however, will then depend on the intentions of the attacker controlling the DDoS network.

4. The following URLs link to papers that provide alternative or more detailed coverage of particular aspects of DDoS, including analyses of attack tools and history:

- http://www.nipc.gov/ddos.pdf
- http://ciac.llnl.gov/ciac/documents/CIAC-2319_Distributed_Denial_of_Service.pdf
- http://www.cert.org/tech_tips/denial_of_service.html
- http://www.cert.org/archive/pdf/DoS_trends.pdf
- www.ietf.org/rfc/rfc2827.txt
- http://www.sans.org/ddos_roadmap.htm
- http://www.cert.dfn.de/dfn/berichte/db093/behringer-ddos.pdf
- http://www.cs.virginia.edu/~survive/ddos/
- http://staff.washington.edu/dittrich/misc/ddos/
- http://www.anml.iu.edu/ddos/
- http://www.cisco.com/warp/public/707/newsflash.html

What is DDoS?

5. DDoS is a type of denial of service attack performed by and synchronised between more than one attacking host. A denial of service attack is an attack which is designed to render the target system incapable of offering the service that is being targeted. In some cases the target system may crash and become unusable, in other cases the attack consumes resources on the target system. Attacks designed to crash a service or system are called "nukes"; resource consumption attacks are known as "Floods".

6. Common techniques for denial of service are as follows [3]:

- SYN floods
- ICMP floods (including "Smurf" attacks)
- UDP floods (including "Fraggle" attacks)
- Application level floods
- Nukes (malformed or specially crafted packets)

7. A **SYN flood** is a sequence of TCP session initiation packets, often from incorrect (or "spoofed") IP addresses. The result is that the target tries and fails to establish a number of TCP sessions, which consumes resources on the target. An **ICMP flood** is a sequence of ICMP echo request packets from spoofed IP addresses. Echo requests are usually answered by an echo reply packet if the target is operational, so such an attack will consume resources on the target. This attack will also consume resources on the spoofed source IP addresses as they will receive a number of ICMP echo replies. The same idea is used in a **UDP flood** where a sequence of UDP packets, often from spoofed IP addresses, are sent to UDP ports such as 7 (echo) or 13 (character generator). Smurf and Fraggle attacks exploit the fact that a source IP address sending an ICMP or UDP flood will be flooded with reply traffic, by flooding the broadcast address of a target network with ICMP or UDP packets. All operational hosts on the target network will respond to the spoofed source IP address.

8. **Application level floods** will depend on the application being flooded. The commonest application level flood is multiple requests for web pages directed against a web server. Similarly mail servers can be flooded with email and/or email with large attachments. In the case of email, the email sender and recipient addresses are usually spoofed.

9. As indicated above, **nukes** are designed to crash remote systems. Nukes can be extremely varied depending on how the IP packets are malformed or crafted. Common examples include "Land" (where the source and destination IP addresses are the same), "Christmas Tree" (where the FIN, URG and PUSH TCP flags in the packet are set) and "Teardrop" (where IP fragments overlap when reassembled). Nukes can also occur at the application layer: witness the recent exploits for Microsoft Windows Server Message Block (SMB) traffic [4].

10. A typical architecture of a DDoS attack network consists of three layers:

- A client computer that is operated by the attacker
- A number of handlers (also known as masters) which are controlled by the client, and
- A number of agents (also called zombies or daemons) which are controlled by the handlers and which perform the denial of service attack.

11. The client scans hosts for a particular set of exploitable vulnerabilities. Details of those that are vulnerable are recorded. Those vulnerable hosts are then compromised by the client. The handler software is then installed automatically on the vulnerable hosts. The handlers then perform further automated scans for further vulnerable systems to compromise, which then become agents. The agents then perform attacks controlled by the handlers, which are in turn directed by the client.

12. DDoS tools often hide themselves on the compromised systems so that system administrators and users will not be able to detect that those tools are present. They may do this by using plausible file or process names, but many DDoS tools include altered operating system commands that aim to make the DDOS tool and its processes and network activity invisible to the system user. Programs of this kind are known as "rootkits", and a number exist for UNIX systems and, to a lesser extent, for Windows systems.

13. The following are common DDoS attack tools [5]:

- Mstream
- Shaft
- Stacheldraht
- Tribe Flood Network
- Tribe Flood Network 2000
- Trinoo
- Trinity
- WinTrinoo

14. As mentioned in the background, worms (such as Slapper) are being increasingly used in order to compromise hosts and install DDoS tools.

15. DDoS attacks can be very difficult for the victim organisation to handle. It is common for DDoS agents to spoof the IP address of the source computer so that the agent is not flooded with reply traffic. Spoofing the source IP also obscures the origin of the attack and can make tracing the attack difficult. When a target system is attacked from several sources at once in a DDoS attack, it can be very difficult to trace back and block traffic from a number of spoofed IP addresses.

Responding to DDoS attacks

16. There are four different types of measures that can be taken to respond to DDoS attacks, the first of which an organisation can implement today, the last of which are still research proposals for network infrastructure vendors to implement:

- Immediate remedial measures
- Longer term remedial measures
- Preventing your organisation from becoming a DDoS handler or agent
- Research proposals designed to minimise DDoS attacks

17. The remainder of this paper considers these four types of measures in detail.

Immediate Remedial Measures

18. There are four measures that need to be taken to handle a DDoS attack:

- Detection of the attack
- Identifying the attack
- Mitigating the attack
- Tracing the attack

Detection of the Attack

19. There are three common indicators that a system is under a DDoS attack:

- Complaints from system users that a service is unavailable
- Abnormally high CPU (Central Processor Unit) usage on the target system and on hosts en route, including routers
- Security alerts from network management software or from firewalls, routers or network monitoring software

20. An abnormally high CPU load is a good indicator of DDoS attack when the system is connected to untrusted networks that only require basic system utilities. Security alerts on the other hand require the relevant devices to be configured to detect that inbound network traffic is unusually high and to have the capability of alerting the network administrator, either individually or via a network management workstation.

21. In order to be able to alert abnormal inbound traffic at least one of the following conditions will need to be met:

- A rule set on the firewall is present that detects the DDoS traffic
- Access Control Lists (ACLs) set on the organisation's border routers to identify DDoS traffic
- Traffic flow monitoring is enabled on the organisation's border routers
- Traffic monitoring is used on the organisation's network
- Network based intrusion detection is used on the organisation's network

Identifying the Attack

22. There are eight pieces of information that will need to be identified in order to identify the attack:

- Target IP address(es)
- Protocol (eg TCP, UDP, ICMP)
- Target TCP or UDP ports/ICMP type
- Source IP addresses
- Ingress interface
- Packet type (eg SYN, SYN/ACK, ACK, FIN TCP packets)
- Application level information (eg buffer overflow)
- Traffic flow statistics

23. The IP address, protocol and port/service information will be apparent by using the same techniques that are preconditions for alerting, namely those identified in paragraph 19. It is important that system administrators have a good understanding of the normal traffic flows on their network to provide a baseline for comparison if a DDoS attack is suspected.

24. ACLs and router traffic flow monitoring can also be used to identify the ingress interface of the attack when the organisation has more than one border router. For large organisations and Internet Service Providers (ISPs) where traffic volumes make it impractical to perform traffic flow monitoring and application of ACLs results in an unacceptable degradation of service, alternative methods are needed. If there are a large number of sufficiently random spoofed addresses a technique called "backscatter" [6] can be used to identify the ingress interface. The technique works by statically routing all traffic from an unused IP address range to a particular interface of a designated router, and then by selecting a router (called the "blackhole router") to propagate a routing table update to all other routers in the network (but not outside the autonomous system) to filter out (null route) all traffic directed to the target IP address. The routers then respond to the DDoS traffic with "destination unreachable" ICMP messages. Some of those that are spoofed are very likely to be routed to the router interface selected to route the unused IP address range. When this happens the ICMP "destination unreachable" message will identify the ingress router interface of the DDoS attack.

25. In many cases the attack can be determined by the protocol and port/service information, especially if the attack is a flood. However, more detailed information is required if the denial of service attack is a nuke attack rather than a flood. In this case a network based intrusion detection system may provide more information if it issues an alert. However, even an intrusion detection system may not record the traffic to enable a complete analysis of the attack to be performed. In this case a traffic monitoring device that records IP packets (such as Ethereal running on a computer with a network card listening in promiscuous mode) would be needed to enable the packets to be analysed. As noted above, this may be impractical for large organisations and ISPs.

Mitigating the Attack

26. Once the attack has been identified as far the organisation's network boundary, the priority will be to mitigate the attack, particularly for business critical services. It is recommended that the origin of the attack is traced in order to identify the attacker and help prevent future attacks. This can be done by tracing the packets back through routers belonging to other organisations (most likely ISPs and backbone providers), so that an organisation upstream can block the attack near its sources, but this may take a considerable time.

27. Mitigation can take two forms:

- Blocking the attack
- Rendering the attack harmless

28. Often the simplest way to mitigate an attack is block the type of traffic received by creating an ACL on the router or adding a rule to the firewall rule set. Thus blocking inbound ICMP echo requests and UDP traffic to the echo and character generation ports (which is good security practice in any case) will stop simple ICMP and UDP floods. However, SYN floods cannot be easily stopped by blocking a type of traffic.

29. Blocking target IP addresses on the ingress router interface can be effective, provided that you can route legitimate users to other routers or router interfaces. Blocking source IP addresses is rarely possible as it is likely that they will have been spoofed, and it is very difficult to determine whether or not the source IP addressed have been spoofed. Blocking source IP addresses will only be effective if those addresses are consistently reused and are not used by legitimate customers. In general it is not trivial to determine if a source IP address has been spoofed (unless it is a reserved non-routable IP address). Tracing the packet back is often the only effective test [7].

30. Another technique that can be used for packets with spoofed source IP addresses on some routers and firewalls is to drop traffic that claims to originate from a class of IP addresses not served by a particular interface. This check is done by looking at the routing table. This idea is behind Unicast Reverse Path Forwarding (URPF), which is available on some routers.

31. In the cases where your organisation cannot block the DDoS attack, perhaps because target IP addresses cannot be blocked on the ingress interface because legitimate users cannot be rerouted, contacting organisations upstream is essential. The organisations that have routers near the sources of the attack can block the traffic once they have identified those sources.

32. When an attack cannot be blocked quickly, it is sometimes possible to render the attack harmless. Some firewall and router products have defences to SYN floods, and a number of firewall products provide protection against malformed packets at the IP layer

and above. Firewalls that do this will generally be circuit level or application proxy firewalls, although stateful inspection firewalls may also provide protection [8]. Another technique to render an attack harmless is to limit the rate of traffic flow through the ingress router interface once the ingress interface has been identified. This technique is implemented by some routers.

Tracing the Attack

33. As mentioned in the previous section, tracing the attack can be essential if the attack cannot be mitigated by your own organisation. Tracing the attack can also be important if criminal prosecution or civil legal action is being considered against the attacker.

34. The available techniques used to trace an attack are at the time of writing the same for every organisation depending on the routers and firewalls deployed. These techniques can be used to identify the paths of the attack, router by router, organisation by organisation from target to attack source provided that the attack is ongoing or that sufficient information is logged and retained.

35. However, even if the DDoS agents are identified, those computers will have to be analysed to determine if the handlers can be identified. This could involve examining the DDoS tools' configuration files or monitoring for network communications to the DDoS tool. If a handler is identified, then there is a chance that the client computer can be identified and the attacker then identified and perhaps prosecuted or made subject to other legal action.

Long Term Remedial Measures

36. For any organisation that is a potential victim of a DDoS attack, the following mechanisms should be implemented:

- Deployment of firewalls at the organisational perimeter that provide stateful inspection, circuit level and application level proxies and have the capability to drop traffic that arrives at the wrong network interface and to defend against SYN floods
- Deployment of routers at the organisational perimeter that allow ACLs to be set, support traffic flow monitoring, flow rate limiting and URPF
- Deployment of network based intrusion detection systems on the internal network
- Deployment of network traffic monitoring devices on the internal network
- Deployment of backscatter network configuration
- Egress filtering

37. Egress filtering [9] is the practice of allowing only IP traffic that originates from your organisation's IP address range to leave (egress) your organisation's network. It is especially useful for organisations that use the Internet in order to avoid being used as a source of a DDoS attack using spoofed IP addresses, but it can also be used by ISPs. This can be implemented by firewall rules that allow only certain source IP addresses or by router ACLs.

38. For ISPs, the adoption of ingress filtering [10] is recommended. Ingress filtering is the practice of allowing only IP traffic that has a IP address source in domain which are known to be valid to enter (ingress) the ISP's network. This is best implemented by ACLs on each ingress router interface, as different IP ranges will need to be allowed on each interface depending on the ISPs peered with or on the domains routed through the interface.

39. The use of backscatter can also be beneficial for large organisations and ISPs. There is also a proposal called CenterTrack [11] to implement an overlay network of routers to isolate DDoS traffic: techniques identified in previous sections of this paper to identify and trace the DDoS attack may then be used.

40. While ingress and egress filtering are recommended, any type of filtering will have a performance impact on traffic flow. The use of filtering requires a trade-off between performance and security and needs to be judged against the risk of attack.

Preventing your organisation from becoming a DDOS handler or agent

41. In a DDoS attack, for every victim there are a number of compromised hosts. It is important to be a good neighbour on the internet for the security of your own and other organisations'. There may also be legal implications if your network is used in attacks against another organisation.

42. The following are some techniques that can be used to prevent compromises of hosts on your network:

- Application of patches and upgrades
- Maintenance of awareness of latest vulnerabilities and exploits
- Deployment of firewalls to enforce your organisational security policy
- Deployment of intrusion detection systems to indicate breaches of your organisational security policy

43. As indicated in NISCC Technical Note 01/02 [12], firewalls should be configured with a "default deny" rule: deny any service that is not explicitly allowed. Application proxy firewalls that also provide stateful inspection are generally the most secure. For advice on intrusion detection techniques, NISCC Technical Note 05/02 [13] should be consulted. Of course technology is no substitute for a good organisational security policy. Some suggestions on good security practice in a networking context are provided in NISCC Technical Note 01/02.

44. As mentioned in the previous section, the use of egress filtering on your organisation's perimeter firewalls will prevent your network being used as the source of a DDoS attack using spoofed IP addresses.

45. However good the organisational security policy and the technology deployed, it is important to check that your network is free from DDoS tools, rootkits and trojans. Specialist DDoS and Trojan scanners exist [14] and should be used with anti-virus products that have a Trojan detection capability.

Research proposals designed to minimise DDOS attacks

46. There are currently three research projects that have been proposed to attempt to solve the problem of identifying the source of spoofed IP addresses involved in a DDOS attack. These are as follows [15]:

- Pushback
- Traceback
- ICMP traceback

47. Pushback [16] is a proposal to rate limiting flows that works by successively asking each router along the traffic flow to rate limit traffic that belongs to the traffic flow. It requires special packets to be used between routers to request a rate limit from an upstream router and for a router to be able to respond to the request. Pushback drops traffic on the basis of the destination IP address(es) of the traffic.

48. Traceback involves additions to the IP packet header to include information that can be used to identify the path of the traffic. The exact details of what should be added have not been settled [17] but it is clear that details of the ingress path will be needed and that, due to the large volume of traffic involved, the overhead on routers can be reduced by sampling packets when countering DDOS attacks.

49. ICMP traceback [18] is similar in concept to traceback: the proposal is that routers send special ICMP packets to the victim for packets sampled with low probability. The special ICMP packets contain information about the last hop to the router that sends the ICMP packet. For a high volume DDOS attack the victim can reconstruct a path to the attacker.

Notes

[1] See, for example, http://www.businessweek.com/2000/00_08/b3669001.htm.

[2] See UNIRAS Briefing No. 219/02, http://www.niscc.gov.uk/niscc/docs/br-20020916-00320.html.

[3] See http://www.anml.iu.edu/ddos/types.html.

[4] For example, the SMBdie C source code available from the internet.

[5] See http://www.anml.iu.edu/ddos/tools.html.

[6] See http://www.secsup.org/Tracking for the original description of backscatter, and http://www.cert.dfn.de/dfn/berichte/db093/behringer-ddos.pdf for a graphical presentation of backscatter.

[7] In some cases it would be possible to test the spread and the randomness of the IP addresses, and tracing techniques such as backscatter would fail if none of the IP addresses were spoofed.

[8] The issue here is that the firewall should not forward malformed IP packets. A circuit or application level proxy does not forward any IP packets but establishes new proxy connections. Stateful inspection firewalls monitor the transport level state and can filter packets which do not reflect a valid state. Of course the firewall's TCP/IP implementation should also be hardened so as not to be vulnerable to nuking or flooding.

[9] See, for example, http://www.sans.org/y2k/egress.htm.

[10] See http://www.ietf.org/rfc/rfc2827.txt.

[11] See, for example, R. Stone "CenterTrack: An IP Overlay Network for Tracking DoS Tools" http://www.us.uu.net/gfx/projects/security/centertrack.pdf.

[12] See NISCC Technical Note 01/02 "Protecting your computer network - Guidance on securing LANs, WANs and internetworks", http://www.uniras.gov.uk/.

[13] See NISCC Technical Note 05/02 "Understanding Intrusion Detection Systems", http://www.uniras.gov.uk/.

[14] See, for example, David Dittrich's home page, http://staff.washington.edu/dittrich/misc/ddos/, for a list of DDOS detection and scanning tools.

[15] See R. Thomas "Tracking Spoofed IP Addresses" http://www.cymru.com/Documents/tracking-spoofed.html for an overview and references.

[16] See J. Ioannidis, S. Bellovin "Implementing Pushback: Router-Based Defense against DDOS Attacks" http://www.isoc.org/isoc/conferences/ndss/02/proceedings/papers/ioanni.pdf.

[17] See, for example, S. Savage, D. Wetherall, A. Karlin & T. Anderson "Practical Network Support for IP Traceback" http://cs.washington.edu/homes/savage/traceback.html.

[18] See, for example, S. Bellovin "ICMP Traceback Messages" http://www.ietf.org/proceedings/02mar/I-D/draft-ietf-itrace-01.txt