



SANS Institute

Information Security Reading Room

The Changing Face of Distributed Denial of Service Mitigation

Justin Stephen

Copyright SANS Institute 2021. Author Retains Full Rights.

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

The Changing Face of Distributed Denial of Service Mitigation

On February 7, 2000, Internet mainstay Yahoo.com experienced a lengthy service outage. What became clear over the following hours was that the site had been victimized by a distributed denial of service (hereafter abbreviated to “DDoS”) attack from hundreds of geographically dispersed Internet-connected machines sending millions of request for service packets. Over the next few days, additional attacks were launched against six other major Web sites, among them some of the favorite sons of the then-burgeoning e-commerce revolution. The ultimate victims were Yahoo.com, Amazon.com, Buy.com (attacked a mere hour after their initial public stock offering), ZDNet.com, E-Trade.com, eBay.com, and CNN.com. According to the Yankee Group, estimated costs of the attack totaled \$1.2 billion cumulative and the attack on Amazon alone cost between \$200,000 and \$300,000 per hourⁱ. Losses of customer goodwill, corporate reputation and public trust may have been even greater.

Mainstream media coverage of these attacks was very heavy because of the sheer scale and likely because one of their own, CNN, was among the victims. Although the first recorded DDoS attacks had occurred years earlier, these February 2000 incidents marked a public coming out party for this type of cyberattack. Almost more disconcerting than the attacks themselves was the revelation of the identity of the perpetrator. A 15-year-old Canadian teenager, who went by the alias “Mafiaboy”, had researched and downloaded several hacker tools, such as AMDEX, Trank, and Slice 3, and launched the attacks using a DDoS tool called Tribe Flood Network (a.k.a. Trinoo). By some estimations, the only reason he was ultimately caught was because he bragged about his exploits in Internet chat rooms.

Major DDoS attacks still make the news. In January, Microsoft became the victim of such an attack. Microsoft’s main Web site and affiliated sites for MSN, online travel site Expedia.com, the auto sales site Carpoint, and the Microsoft email service Hotmail were unreachable for several hours. This past May, a DDoS was launched against the CERT Coordination Center, the government-appointed InfoSec watchdog that, for many, symbolizes security on the Web. By some estimates, losses from this attack total \$100,000.

“We get attacked every day,” said Richard D. Pethia, a CERT director. “This is just another attack. The lesson to be learned here is that no one is immune to these kinds of attacks. They cause operational problems, and it takes time to deal with them.”ⁱⁱ

Just last month, the Internet-connected world was rocked by the Code Red worm. Exploiting a buffer overflow vulnerability in Microsoft Internet Information Server, the worm was able to infect 359,000 machines worldwide in just 14 hours. Those machines hosting sites whose default language was English were defaced and all infected

machines served as a springboard for vicious propagation code that attempted to spread the worm to other machines. Part of the attack pattern (“phase 2”) of the original iteration of this worm was to launch a DDoS attack against whitehouse.gov. Fortunately for the White House IT staff, not only was the worm hard-coded to check to make sure that port 80 at whitehouse.gov was active before launching, the IP address to be attacked was hard-coded as well. Whitehouse.gov systems administrators simply turned off the DNS server at the target IP (192.137.240.91), rerouting all requests to the other server. Additionally, ISPs worked together to “black hole” packets sent to the target IP.

Mainstream media coverage of the Code Red worm has also been very heavy, most focusing on the rapid spread of the worm. Truth be told, however, we all dodged a very large bullet with this worm. Despite its impressive rate of propagation, minimal damage was done. Hopefully, Code Red will serve as a wake-up call. It should also serve as an extremely nefarious omen of bigger and nastier DDoS attacks to come. Instead of the traditional model of DDoS slave, or “zombie”, acquisition employed by Mafiaboy and others, wherein it can take weeks or months to crack into the slave machines needed for a large attack and plant the attack software, the Code Red worm built a slave army of 359,000 machines in just about 14 hours.

There still does not exist a tool or process that can fully protect a Web site from a DDoS attack. By many accounts, those seven Web sites victimized by Mafiaboy in February 2000 are only marginally better prepared to thwart such attacks today, well over a year later. The frequency of DDoS attacks continues to increase, going up 60 percent in the past 3 years. One-third of the respondents to the 2001 Computer Crime and Security Survey report having experienced denial of service attacks. It is also safe to say that the problem is under-reported. Many attacks go undetected at all and many organizations, fearing bad publicity and the consequent effect on their customers and stockholders, do not report those that are detected. Additionally, the attack tools available for launching these attacks are becoming more and more sophisticated and their schemes are getting increasingly complex. Security experts have identified more than seven primary DDoS tools and variants are appearing continuously. The painful reality is that any bored teenager can download most of these tools from the Web and launch his/her own DDoS in relatively short order.

In this paper we will review the traditional best practices and tools for DDoS mitigation, discuss the inherent weaknesses of these best practices, review the developing legal issues and trends that may soon be forcing change on how DDoS attacks are combated, and look at the new generation of tools becoming available for mitigating these attacks.

Traditional Defenses

Many of the basic practices that can help prevent or mitigate DDoS attacks should be included in any defense-in-depth enterprise security plan, even one not overly concerned with this particular risk. Among these are:

- Timely application of patches and system updates, especially to potentially exposed machines. For example, update and maintain a current build of BIND on DNS servers.
- Deployment of only strictly necessary network services.
- Intrusion detection systems.
- Firewalls.
- Anti-virus software.
- Good password policies.
- Use of Tripwire or other similar tools to detect changes in configuration information or other important files.
- Paying heed to “Top 20” vulnerability lists provided by the information security community and evaluating these risks against one’s environment.
- Establishment and maintenance of regular backup schedules and policies.
- As a network is only as secure as its weakest link, protection of mobile and remote machines with personal firewall/intrusion detection software.

Other best practices that can be employed at the user organization level that will help mitigate the risk of denial of service attacks include:

- Carefully architect the DNS server network, distributing DNS servers around the edge of the network and consider establishment of back-up relationships with other parties. Poor DNS server network design was a crucial factor in the January DDoS attack on Microsoft mentioned above. Additionally, safeguard information about the architecture and thus vulnerabilities of DNS networks.
- Address filtering, also known as “egress filtering”, of packets leaving the enterprise. This can ensure that packets leaving carry source addresses within the ranges of those sites. It can also ensure that no traffic from unroutable addresses (see RFC 1918) leave those sites.
- Turn off ICMP echo and chargen services unless there is a specific need for these services. This will prevent “smurf attacks” and similar vulnerabilities.
- Patches are available to help prevent TCP SYN flood attacks. Test and install them.
- Establish baselines for normal activity. This will help enable administrators to determine if there is a problem.
- User organizations should check their systems regularly to determine whether they have malicious software installed. There are a number of tools, many of them free, to assist in this effort. Some examples:
 - National Infrastructure Protection Center (NIPC)’s “find_ddos” tool is able to detect several old and more current DDoS tools including mstream, TFN2000 client and daemon, Trin00 daemon and master, TFN daemon and client, stacheldraht master, client and daemon and TFN-rush client.
 - RID, from David Brumley at Stanford University, is able to detect Trin00, TFN, and stacheldraht agents.
 - Zombie Zapper from Bindview Inc. works against Trinoo, TFN, stacheldraht, Troj_Trinoo (the trinoo agent ported to Windows), and Shaft.

- Invest in hot spares, machines that can be placed into action quickly in the event that a similar machine is disabled.
- Invest in more bandwidth to lower your vulnerability to flooding attacks. Invest in redundant load-balancing networks and servers. If there are multiple versions of the same Web site operating on different network segments, rogue packets can be distributed evenly amongst them making it more unlikely that any given server will crumble under the weight of an attack.
- Education and communication throughout the community can be extremely helpful. When organizations fail to share information about attacks, this helps give the hacker community an even greater advantage. Systems administrators should participate in industry-wide early warning systems. Information about attacks should be disseminated to vendors and response teams so that it can be applied to the defenses of others.

There are certainly things that network and hosting providers do now that can assist in the DDoS mitigation efforts. NSPs can utilize ingress filtering, similar to egress filtering but on a larger scale, to help combat IP spoofing. They can, and often do, respond to information from their customers and from other NSPs to combat malicious packets. The “black holing” of packets destined for whitehouse.gov during the recent Code Red attacks are but one example of this. Lastly, they can perform traffic and load monitoring that can provide early warning of some attacks.

Weaknesses of Traditional Defenses

There are certainly drawbacks to the practices described above. Ultimately, these drawbacks can be summed up as onerous levels of effort and the ultimate inability of user organizations to truly determine their own fate when it comes to DDoS attacks.

Keeping up to date with, researching, testing and implementing every applicable software patch and system update is a time consuming process, as are tuning, monitoring and updating firewalls and intrusion detection systems. Additionally, firewalls and IDSs were designed to detect discrete attacks against individual hosts or Web servers, not to detect and counter attacks against the network. As such, they do not provide the ability to monitor and characterize floods of abnormal network traffic in real time. By the time the attack hits, customer traffic has already been affected.

Egress filtering and use of tools to find DDoS malware on our own systems has far more benefit to our peers than to our own networks. Yes, if everyone took these steps, the incidences of DDoS attacks could be lowered dramatically. However, “depending on the other guy” is hardly a strategy likely to bear much fruit when the attacks can come from any Internet-connected system anywhere in the world.

The type of rate-limiting or service denial that works well against such ICMP-based attacks as “smurf attacks” is almost useless against TCP traffic. Current filtering capabilities on most routers is too coarse-grained, inflexible, and slow to effectively handle the work we need them to do. Firewalls, on the other hand, possess

sophisticated filtering capabilities, but their performance is limited because they need to perform additional analysis not specific to denial of service protection.

Investing in hot spares and additional bandwidth could prove to be more of a financial burden than some user organizations are willing or capable to bear. Additionally, throwing more and more bandwidth and load-balancing at the problem, while possibly effective in the short haul, is akin to getting into an arms race with attackers. They will eventually catch up with a target network's capacity. There are always plenty of slave machines to throw into the attack. University machines, for example, have always been a common target due to lack of money dedicated to protecting them, transient students who run and tinker with machines and then leave, and lack of accountability. Machines from Stanford, U.C. Santa Barbara, the University of Washington, Oregon State, and James Madison just to name a few were used by Mafiaboy in his attacks of February 2000. Thankfully, the publicity from these and other attacks have prompted many educational institutions to reevaluate their security practices, or lack thereof. However, as educational facilities get more on the information security ball, the proliferation of high-speed, always-on DSL and cable modem connections to lightly or unprotected private machines continues to explode. In short, plenty of potential DDoS slave candidates are available and will continue to be for the foreseeable future.

Just the act of detecting a DDoS, or detecting it soon enough to limit damage, is often extremely difficult. Existing monitoring tools and other active defense methodologies are often unable to differentiate between an attack and heavy but legitimate "flash crowd" traffic. Most lack the performance and level of detail needed to handle the attack. Often, the only symptoms of an attack are sluggish performance and/or router and server failures. Many victims haven't even realized they were under attack until they have rebooted their routers numerous times to no avail. Alas, the most effective means of fighting and tracing a DDoS is the telephone conference call, with operators at the various victim sites and NSPs collaborating to trace the sources of the attack and block them off. While the results of this type of collaboration can be extremely effective, the precious time lost in the process can be expensive hours for the victims and their customers.

To date, most DDoS countermeasures have focused on the enterprises themselves, at the router and firewall level. Some denial of service attacks exploit bugs which can be patched and fixed. However, many DDoS attacks simply consume server resources by exploiting legitimate features. No matter how well an enterprise protects itself, large attacks consume bandwidth upstream of the ultimate target, making downstream filtering of any kind mostly ineffective.

Legal Issues Related to DDoS Mitigation

Who is most capable of taking effective precautions to prevent or reduce the frequency of successful DDoS attacks? Who will be subject to legal liability in the event of damage from such attacks? These two questions are closely linked because the

entities most capable of taking effective precautions are most likely those that the courts will impose tort liability on if they do not take these precautions.

Tort liability is generally understood to have two purposes; to deter bad acts and to compensate those harmed by them. Counter this with the purposes of criminal law; to deter bad acts and punish those who commit them. DDoS attacks are already illegal under the federal Computer Fraud and Abuse Act and numerous state statutes. However, these criminal laws alone do not deter many would-be attackers. Many attacks originate from overseas, from countries a million miles away from the long arm of U.S. law. Many cybercriminals believe they can do as they wish and not get caught. More often than not, they aren't wrong. The inherent difficulties of successful criminal investigation of cybercrime combined with the complexities of international law make successful apprehension and prosecution of cybercrime unlikely. As such, little real deterrent value is realized.

The impacted parties of DDoS attacks are the NSPs and hosting providers, the operators of portals, commercial Web sites, and services delivered over the Internet (the typical targets), and those individuals who do business with these sites and services. For example, to say that only the Web site and its operators are adversely effected if E-Trade.com suffers a major outage from a DDoS isn't accurate at all. E-Trade's customers are also victims and they might realize legitimate and significant financial losses from the outage.

In considering tort liability, numerous courts have adopted an economic analysis known as "best cost avoider." This analysis says that legal liability should rest on the "best cost avoider" for the harm, or the actor who is in the best position to know the risk and take precautions against it. Being in the "best position" means the actor situated such that it can develop the information about the risk and implement the precautions most cheaply.

Best cost avoider analysis makes it pretty clear that those farthest downstream, the customers of the affected sites, are powerless to mitigate the negative ramifications of a DDoS attack. The intruders themselves, while most certainly culpable, are judgment-proof from tort liability because they have no money and aren't usually caught anyway. Right now, it certainly appears that the Web site operators themselves, while certainly able to take some precautions, are not readily available to protect themselves against many of these attacks because it isn't clear where the messages are originating from and the attacks themselves keep changing. Again, no matter how well an enterprise protects itself, large attacks consume bandwidth upstream of the ultimate target, making downstream filtering of any kind mostly ineffective. This leaves those entities at the top of the stream, the NSPs.

Courts generally set the legally required standard of care, the practices necessary to avoid tort liability for negligence. If wholesale prevention technologies were to become available to them, courts would have reason to place the liability on these network intermediaries because it will give them the incentive to implement the most effective and efficient protective strategies.

NSPs might try to contractually force customers to take adequate precautions as a condition of providing service. However, “adequate” is a relative term as we have already established that what the target sites can actually do to protect themselves is limited. Additionally, this contract strategy might backfire, driving customers to NSPs with less onerous contracts.

NSPs might also try to force customers to take on the liability themselves by using contracts that disclaim liability on the part of the service provider. This type of strategy is not always legally effective. First, not all contracts are valid and enforceable. Two ways a contract might be ruled unenforceable are:

1. invalid information, where the courts find that no agreement was formed.
2. invalid content, where the courts find that public policy disavows such an agreement.

Some jurisdictions consider exculpation of one’s own negligence to be contrary to public policy. Second, contracts are not binding on third parties.

A Web site, having been sued by its customers for monetary losses resulting from a DDoS-spawned outage, might choose to initiate an indemnification lawsuit against its hosting or bandwidth service provider. Indemnification involves one party attempting to shift its liability to another, either via a contract or a lawsuit.

To go back to the E-Trade example, the customers of E-Trade who experienced financial losses resulting from an outage might end up suing E-Trade’s network service provider. The contract between the NSP and E-Trade would be worthless as a defense as the E-Trade customer and the NSP do not have contractual privity with one another.

Those entities at the top of the stream may want to consider the possible changing face of DDoS-related liability and implement well-considered risk management strategies. Aggressive use of preventative technologies is an important part of risk management.

These issues may ultimately get decided not in court, but in Congress. Victims of DDoS attacks might seek legislation to impose liability on NSPs. Victim sites might also seek legislation to immunize themselves against, or at least curtail, liability.

Of course, most of this depends on there actually being preventative technologies available that NSPs and hosting providers can actually use to mitigate DDoS attacks. Until recently, no such technology existed making the legal theory presented above a bit of a moot point. Not surprisingly, no reported court decisions have held e-businesses liable for DDoS attacks. This author is also unaware of any court decisions that have held network intermediaries responsible either. However, new technologies are becoming available that may change all that.

New Trends and Products

A lot has happened since the DDoS attacks of February 2000 and a lot will continue happening. High profile portals and NSPs have formed industry groups to consider the prevention of DDoS attacks. Some examples include the DDoS Working Group, the RFC-2267-plus Working Group and the IT Information Sharing and Analysis Center (IT-ISAC). Additionally, organizations dedicated to information security, such as the SANS Institute, have prepared and disseminated best practice lists for dealing with these attacks and recommendations for changes that will assist mitigation of these attacks in the future.

Router and switch vendors have taken steps to integrate IP address filtering directly into the operating system of these devices. This should dramatically reduce the performance hit previously associated with more manual address filtering at this level.

At the same time, promising technologies have been and are being developed that might forestall DDoS attacks on a network basis by preventing attack traffic from exiting service provider networks and entering a victim Web site.

Backed by huge investment from many venture capital firms and, in one case, by additional investment from Cisco Systems, three firms have recently unveiled new products designed to be utilized by network intermediaries to combat DDoS attacks. Arbor Networks of Waltham, MA has introduced its Peakflow DoS line of products, Mazu Networks of Cambridge, MA has introduced its TrafficMaster line of products, and Asta Networks of Seattle has recently released Vantage System. This section is not intended to provide a granular comparison and contrast of the three products nor does it recommend any particular product. The intention here is to provide a high-level overview of how these three products may very well change the face of DDoS mitigation.

At least from a high-level, all three products appear to work in very similar ways. Mostly intended for network intermediary customers (although Arbor offers a version of Peakflow DoS intended for enterprise customers), they provide a distributed managed availability solution that allows network engineers to counter denial of service attacks, hopefully before services to customers is degraded or halted. Since more technical architecture information was available on the Arbor and Asta products, the information described below is most applicable to them.

As traffic enters the service provider network, “sensors” or “collectors” analyze the traffic for anomalies without disrupting the flow of traffic. Asta and Arbor’s products do not look at each individual packet but look at traffic patterns to detect anomalies. Mazu’s product claims to look at every packet without slowing performance. Using proprietary algorithms and/or attack signature databases and also comparisons against network traffic baseline, the sensors monitor traffic patterns looking for known or new anomalies. If one is discovered, they communicate to a “coordinator” or “controller.” These central entities analyze the data from the various sensors to develop a network-wide analysis of probable denial of service activity and begin to trace the attack to its sources. At the same time, network engineers can be alerted via a variety of channels (e-mail, pagers,

SNMP, interfaces to trouble ticket software, etc.) and provided recommendations on best courses of action to filter and trace the attack. Both Mazu and Arbor claim that their products work with the hardware the NSP or hosting provider already owns. Asta's high-level documentation was not clear on this subject.

These products do not come cheap. As examples, Arbor's Peakflow DoS product starts at \$5,000 per monthⁱⁱⁱ. Mazu's TrafficMaster Inspector for DDoS has a list price of \$100,000 for a typical data center configuration^{iv}. However, since these new products may very well be the "preventative technologies" that NSPs need to embrace as part of a comprehensive risk management strategy, they may ultimately be dirt cheap compared to legal costs. Even if these new technologies do not prove to be 100% effective in stopping DDoS attacks (a safe bet), a firm that has shown willingness to address the problem by adopting state-of-the-art practices and technologies is far more likely to be able to defend itself in court against liability claims.

One thing is certain, DDoS attack tools continue to develop and evolve. Even if this new breed of products proves to be as effective as advertised and are put into ubiquitous use by service providers, the real backbone of any secure network system is constant monitoring. The systems administrators at the user organization level will always play a crucial role in preventing DDoS attacks. As Lance Hayden, manager of secure consulting services for Cisco Systems put it, "Proactive vulnerability assessment is the key."^v

ⁱ Unknown, "DDoS: It's Everyone's Problem.", p. 1.

ⁱⁱ Sullivan, p. 1.

ⁱⁱⁱ Unknown, "Arbor Networks Delivers Carrier-Class...", p. 1.

^{iv} Unknown, "Press Releases: Mazu Networks Unveils Packaged...", p. 1.

^v Fisher, p. 3.

Sources:

Ferguson, P. "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing." RFC 2827. May 2000. URL: <http://www.ietf.org/rfc/rfc2827.txt> (12 July 2001).

Fisher, Dennis. "Defenses still weak against DDoS attacks." 19 Jan 2001. URL: <http://www.zdnet.com/eweek/stories/general/0,11011,2676260,00.html> (30 July 2001)

Greene, Tim. "Forum warns of hidden DDoS legal liability." 2 Oct 2000. URL: http://www.nwfusion.com/archive/2000/108677_10-02-2000.html (12 July 2001).

Harrison, Ann. "Analysts: Mafiaboy Only Amateur Copycat." 24 Apr 2000. URL: http://www.computerworld.com/cwi/story/0,1199,NAV47_STO44528,00.html (2 Aug 2001).

Harrison, Ann. "Cyberassaults hit Buy.com, eBay, CNN and Amazon." 9 Feb 2000. URL: http://www.computerworld.com/cwi/story/0,1199,NAV47_STO43010,00.html (18 July 2001)

Jeon, DeokJo. "Understanding DDOS Attack, Tools and Free Anti-tools with Recommendation." 7 Apr 2001. URL: http://www.sans.org/infosecFAQ/threats/understanding_ddos.htm (12 July 2001).

Koh, Jay L. "Recent Developments and Emerging Defenses to D/DoS: The Microsoft Attacks and Distributed Network Security." 9 Feb 2001. URL: <http://www.sans.org/infosecFAQ/DNS/developments.htm> (12 July 2001).

Lemon, Sumner. "Code Red worm exploits Windows NT flaw." 20 July 2001. URL: http://www.computerworld.com/cwi/story/0,1199,NAV47_STO62410,00.html (14 Aug 2001)

McAlearney, Shawna. "White House Dodges Massive DDoS." 26 July 2001. URL: <http://www.infosecuritymag.com/digest/2001/07-26-01.shtml> (30 July 2001).

Messmer, Ellen and Denise Pappalardo. "A year after meltdown: No silver bullet for DoS." 5 Feb 2001. URL: <http://nwfusion.com/news/2001/0205ddos.html> (2 Aug 2001).

Messmer, Ellen and Denise Pappalardo. "One year after DoS attacks, vulnerabilities remain." 8 Feb 2001. URL: <http://www.cnn.com/2001/TECH/internet/02/08/ddos.anniversary.idg/index.html> (2 Aug 2001).

Moore, David. "The Spread of the Code-red Worm (CRv2)." Aug 2001. URL: http://www.caida.org/analysis/security/code-red/coderedv2_analysis.xml (30 July 2001).

Niccolai, James. "Code Red Aug. 1 relaunch fizzles, for now." 1 Aug 2001. URL: http://www.computerworld.com/cwi/story/0,1199,NAV47_STO62701,00.html (14 Aug 2001).

Noack, David. "Expert: 'Mafiaboy' Begged for Hacking Tips." 20 Apr 2000. URL: http://www.apbnews.com/newscenter/internetcrime/2000/04/20/hacker0420_01.html (2 Aug 2001).

Pethia, Rich, Alan Paller and Gene Spafford. "Consensus Roadmap for Defeating Distributed Denial of Service Attacks." Version 1.10. 23 Feb 2000. URL: http://www.sans.org/ddos_roadmap.htm (14 July 2001).

Radcliff, Deborah. "University Computers Remain Hacker Havens." 12 Feb 2001. URL: http://www.computerworld.com/cwi/story/0,1199,NAV47_STO57605,00.html (18 July 2001).

Radin, Margaret Jane. "Distributed Denial of Service Attacks: Who Pays?" Date Unknown. URL: <http://www.mazunetworks.com/radin-es.html> (30 July 2001).

Sullivan, Brian. "CERT Web site hit by cyberattack." 23 May 2001. URL: http://www.computerworld.com/cwi/story/0,1199,NAV47_STO60799,00.html (18 July 2001).

Unknown. "Asta Networks Vantage System." 2001. URL: http://www.astanetworks.com/products/data_sheets/vantage.pdf (12 Aug 2001)

Unknown. "DDoS: It's Everyone's Problem." 2001. URL: <http://www.mazunetworks.com/ddos.html> (19 July 2001)

Unknown. "Denial of Service Attacks." 4 June 2001. URL: http://www.cert.org/tech_tips/denial_of_service.html (19 July 2001)

Unknown. "The Denial of Service Problem." 2001. URL: http://www.mazunetworks.com/wp_ddos.html (19 July 2001)

Unknown. "Peakflow DoS for Service Providers." Date unknown. Available on request from URL: <http://www.arbornetworks.com> (18 July 2001)

Unknown. "Peakflow DoS for the Enterprise." Date unknown. Available on request from URL: <http://www.arbornetworks.com> (18 July 2001)

Unknown. "Press Releases: Arbor Networks Delivers Carrier-Class Denial of Service Solution." 9 May 2001. URL: <http://www.arbornetworks.com/news2?cid=5&tid=9&nid=40> (30 July 2001)

Unknown. "Press Releases: Mazu Networks Unveils Packaged Products for Internet Data Centers that Slash Costs of DDoS and other 'Bandwidth Bandits.'" 25 June 2001. URL: <http://www.mazunetworks.com/n-pressreleases5.html> (19 July 2001)

Vijayan, Jaikumar. "Asta Launches DDOS Detection Software." 25 June 2001. URL: http://www.computerworld.com/cwi/story/0,1199,NAV47_STO61568,00.html (18 July 2001)

Weiss, Todd R. and Kim S. Nash. "Update: Microsoft Web sites hit by denial-of-service attack." 25 Jan 2001. URL:

http://www.computerworld.com/cwi/story/0,1199,NAV47_STO56873,00.html (18 July 2001).

© SANS Institute 2001, Author retains full rights