DDoS: threats and mitigation

Steve Mansfield-Devine

Steve Mansfield-Devine, editor, Network Security

It's likely that 2011 will be remembered by many as the year of Distributed Denial of Service (DDoS) attacks. There's nothing new about this kind of threat, but its use has increased and it has even achieved a kind of mainstream notoriety thanks to the antics of self-publicising groups like Anonymous. But how is the threat evolving? And what can you do about it?

According to a report from Prolexic, one of the first and largest companies offering DDoS mitigation services, attack traffic rose 66% over the course of a year (to Q3 2011).[1] Network-layer attacks accounted for 83%, the rest being application-layer attacks. The average duration was 1.4 days and the average bandwidth consumed was 1.5Gbps.

The size of attacks is getting bigger, too. In July 2011, Prolexic announced it had mitigated what it believed to be the largest packet-per-second DDoS attack ever seen in Asia. Consisting of SYN and ICMP floods, the attack deployed 176,000 bots (compared to the 5,000-10,000 bots more normally seen by Prolexic) to generate 25 million packets per second. According to the company, the majority of high-end border routers typically forward 70,000 packets per second. It mitigated the attack by distributing traffic among its Tier 1 carrier partners and scrubbing centres.

Ben Petro, senior VP network intelligence & availability at Verisign, traces the rise of DDoS attacks back another year. He says that for years there was little awareness of the problem but that, "2010 was a dramatic shift – not only in the size, scale and trajectory of DDoS, but also in its proliferation and the number of different types of organisation that were hit."

In the years 2006-2008, he adds, the average attack was somewhere around 40Mbps. "And then, all of a sudden, coming in 2010 we started to see 2Gbps, then 5Gbps, then 8Gbps and 15Gbps attacks, culminating in the largest that we've seen coming in at 84Gbps sustained for a week and a half. That is an incredible amount of traffic when you get down to packets or queries per second – you're over the hundreds of thousands of queries per second per location for Verisign, and we have 165 locations. So it's an enormous volume."

The rising popularity of DDoS attacks may be connected directly to their effectiveness. Paul Sop, CTO at Prolexic, says, "people get creative and they use and invent new ways and reasons to use DDoS. When an idea catches on in the industry there's a tipping point – people say, hey here's a type of attack we can launch and it's very likely we won't go to jail unless we're pretty dumb about it."

He adds: "It is like asymmetric warfare – the attackers have so much advantage in terms of the country they operate from, the size of the attack in terms of victim's infrastructure, and it's very difficult, because of the large number of these attacks, for law enforcement to prioritise and go after any specific attacker."

Potential targets

Online gambling firms have long been subject to DDoS attacks – that's where Prolexic got its start. And there are other industries that might be viewed as high-risk – financial organisations, for example, e-commerce operations or online gaming. And these do remain at the top of lists of industries affected. But it seems that, now, anyone can be a target. For example, in August 2011, Prolexic dealt with a two-day attack against Spafinder.com, an online resource for spa and wellness services and products.

"We never really expected to be the target of a DDoS attack," said Pete Ellis, chairman and CEO of Spafinder. "We had a DDoS mitigation solution in place from a hosting company just in case. Unfortunately, that solution couldn't stop the attack."

The firm called in Prolexic, which identified the most common sources of the DDoS as Kazakhstan, Belarus, Peru and the UAE, among others. "As we deployed our mitigation tools and real-time monitoring, the attack would trickle down to being almost non-existent, and then another wave of attacks with a different type of signature would start," says Neal Quinn, vice-president of operations at Prolexic. "The attack actually spanned over two days after we began mitigation because the attackers changed the signature every time they realised we were successfully blocking the attack."

In its Q2 2011 report, Kaspersky recorded that 25% of attacks were

Ben Petro, Verisign.

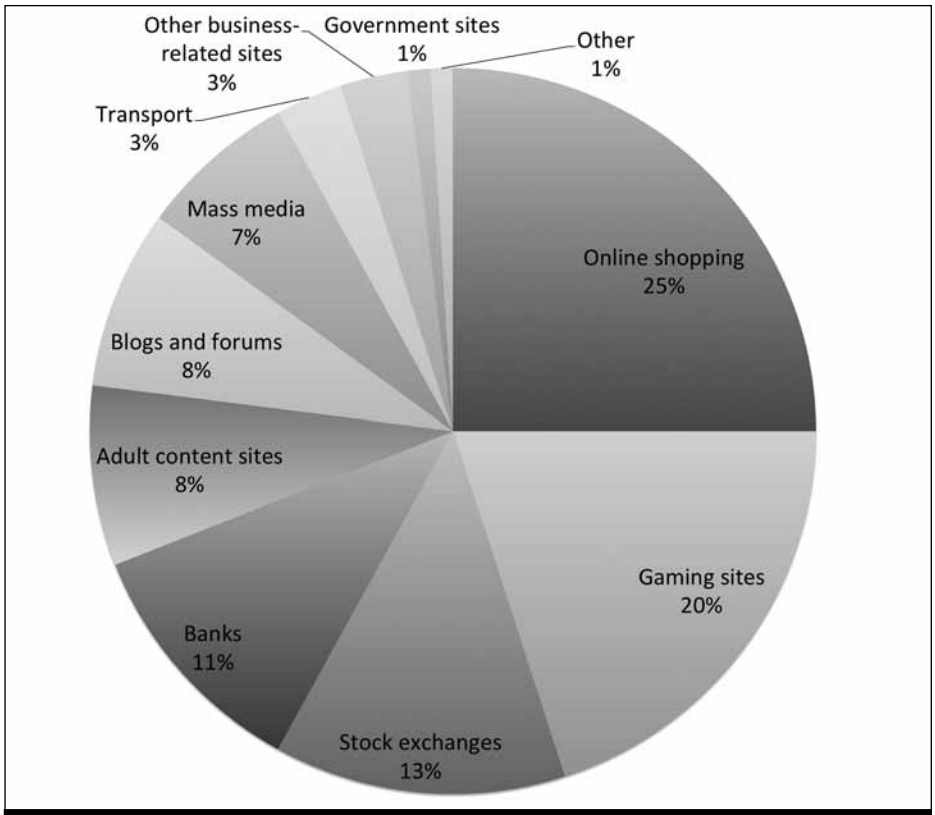

Breakdown of attacked sites by areas of activity, Q2 2011. Source: Kaspersky.

against online shopping sites. Gaming sites were the next most popular (20%) followed by stock exchanges (13%) and banks (11%). The firm also reported that the longest attack duration it had seen was just over 60 days.

Of course, for many firms even an hour is too long. "For example, in the financial services community, every single query is important," says Petro. "Some queries are worth a million dollars, so it's a big deal." But he's also seen the range of victims widen. Verisign carried out research and, even after removing the obvious targets from the figures, such as financial services and e-commerce, the firm found that 70% of the surveyed companies – many in unlikely verticals such as hotels and hospitality – had been attacked. And 60% of those had seen up to six attacks a year.

Simon Woodhead, Simwood.

Attack motivation

So why are the attackers doing this? Extortion has long been one of the key motivations: a firm is warned that, unless it pays up, its website or Internet connectivity will go down, usually at a critical time for its business. With online gambling companies, for example, this might coincide with a major sporting fixture. But the range of motivations seems to be increasing too.

"We still see extortion used on a daily basis," says Quinn. "That one's definitely alive and well." He also sees companies being attacked as part of stock market manipulation. "We've seen evidence of major information outlets being attacked as a way of limiting the ability of companies to propagate information that investors need. You couple that with other things, such as automated stock trading algorithms causing things to dump very rapidly, and it's not that hard to see a scenario where somebody uses this in combination with the knowledge that that's going to happen. And it would be very difficult for someone to find who the direct beneficiary was in such an environment."

Sop adds: "One of the earliest cases we saw of that was someone sending out fake PR Newswire press releases about a business that was 'going bankrupt', and then they DDoSed the website."

Some instances might be baffling at first, with the motivation only becoming clear over time. "For example, a cigar company came to us and said we need DoS protection," says Verisign's Petro. "Not a very big company either and it's kind of an expensive service. And about two weeks later, another cigar company. It turns out they're DoSing each other."

Occasionally the motivations are more trivial, he adds, and this demonstrates just how easy and common denial of service attacks have become. "In the University of California system, if grades aren't ready at a particular time it's an automatic 'A' or pass. So we found that students DoSed the system so that when you went to get your transcript, they just auto-A'd you."

Simon Woodhead at Simwood, a firm that specialises in services for the Voice over IP (VoIP) industry, has even witnessed fancy dress shops being DDoSed in the run-up to Halloween – but without any attempt at extortion. In fact, the motivation for the attacks remains unclear, and this is true for a surprising number of attacks. In early 2011, the Darkshell botnet, originating from China, seemed to be particularly focused on food-processing firms, for no readily apparent reason.[2]

The volume of activity is unknown – but everyone agrees it's high. Simwood operates a darknet – machines using IP addresses that have never been issued or used and which, therefore, should get no traffic. But it does – a lot of traffic. "Some of this traffic might be the result of a misconfiguration or an error," says Woodhead, "but we see more traffic than you would predict from that. You see general network scans, malware looking for exploits, and frankly it's shocking. In the last 24 hours we saw 247,000 events from 25,000 IP addresses. We see a massive proportion – it varies, but it's something around 87% – targeting Windows

exploits, so it's malware looking for compromisable Windows hosts. Port 445 represents about 75% of it, typically, and 139 is the next one."

That malware is looking to create new bots, many of which will be used for DDoS attacks. And they will take advantage of the victims' resources.

"Instead of it being my home computer, which five years ago might have had maybe a meg out to the Internet, now we're talking about corporate IT infrastructure that's compromised," says Petro. And much of this malware is going uncaught. "You're finding very intricate code that slips right through the most sophisticated security products that are out there."

Types of attack

Denial of service attacks generally fall into two categories – Layer 3 (network) 'floods', which usually attempt to overwhelm the bandwidth available to the victim; and Layer 7 (application) attacks that exploit the limitations of a specific application, such as a web server.

The SYN flood is typical of a network-layer attack, and still among the most popular. According to Prolexic's Q3 2011 report, this method accounted for 24% of all attacks. The attacker initiates a TCP handshake by sending a SYN packet. The victim responds with a SYN/ACK as normal but then the attacker simply doesn't complete the handshake by sending an ACK. By itself, this doesn't have a significant effect: even low-cost, domestic routers offer SYN flood protection by aggressively timing out such incomplete connections and dropping them. Nevertheless, some resources are tied up for a brief time, and by sending enough SYN packets, opening additional connections, a simple router can suffer resource exhaustion fairly quickly. Higher performance routers are harder to overwhelm, but where a DDoS succeeds is in using connection attempts from thousands of machines in a botnet. The volume of SYN packets alone may be enough to clog the victim's bandwidth. ICMP (ping) and UDP floods work in similar ways and, by Prolexic's counting, comprised 22% and 19% of recent attacks.

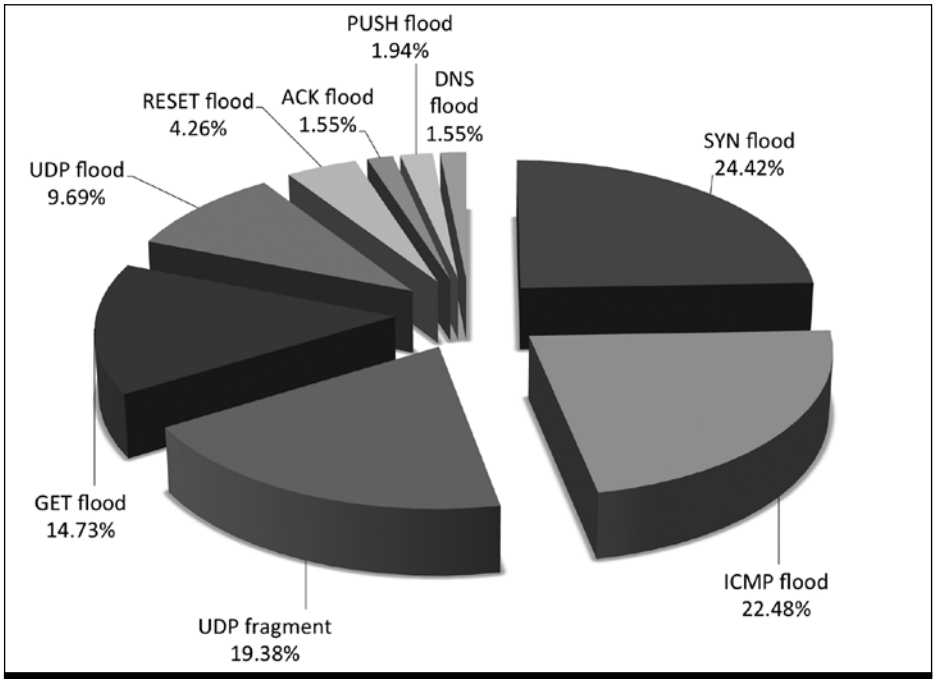

Types of DDoS attack in Q3 2011. Source: Prolexic.

The most common application-layer attacks are against web servers. They may involve sending specially crafted HTTP requests, or the attacks may consist of nothing more than getting a botnet to flood the server with more requests than it can handle.

"I know of a customer who came under a DDoS attack that consisted of a couple of hundred simultaneous clients downloading a certain PDF file from the server," says Amichai Shulman, CTO at Imperva. "There's no vulnerability there, it's just a matter of sizing. When the application was deployed, the assumption was made that there would be 10 or 20 people who would be interested in that resource and would be downloading it at any given time. And the attack consisted of an order of magnitude more clients. In terms of network traffic, it was negligible."

This is something that needs to be carefully considered when configuring systems. And it can be tested. Rapid7, producer of the famous Metasploit package, provides penetration testing services of which DDoS simulations form a part. As Marcus Carey, security researcher and community manager for the firm, puts it: "Companies like banks want to be DDoSed, and we do it to them."

He says he's seen a significant increase in the use of 'slow' HTTP attacks. This is where apparently legitimate GET requests are sent to the server, but omitting the final carriage return. The server has a connection tied up while it's waiting. It's possible to set the server to time-out such connections quickly, but if thousands of such requests are being made, this may still take down the server. "It doesn't take a botnet anymore," says Carey, "it takes a couple of computers. This is what all the kids are using."

Slow POST requests work in a similar manner. "It doesn't require a lot of bandwidth, either," adds Carey, "because you're sending a very simple request, very small packets – so you can go through TOR and proxies."

Carey's impression of the frequency of such attacks is backed up by Kaspersky, which has a very different view from

Amichai Shulman, Imperva.

Marcus Carey, Rapid7.

Prolexic as to which are the most common attacks. Kaspersky said that HTTP floods accounted for 88.9% of DDoS attacks in Q2 2011. This may have been connected, to some extent, with the activities of Anonymous, LulzSec and other hacktivists. Or the discrepancy in the figures may simply reflect the two firms' different customer bases. However, most observers agree that there is a distinct increase in the use of application-layer attacks.

These attacks often target specific parts of an application – login pages or authentication servers may be easier to clog up than general web pages. Or the attackers may go after other protocols. German hacker group, The Hacker's Choice, released a free tool – THC-SSL-DOS – that attacks using SSL, as this protocol tends to be resource-hungry (although not so much on HTTPS port 443 as many firms use an SSL accelerator).[3] Attacks against DNS servers can also be very effective.

It's actually comparatively rare for DDoS attacks, even those targeting the application layer, to exploit specific vulnerabilities in the victim's system. In August 2011, however, there was a certain degree of panic about a bug in the Apache web server, known about since January 2007, that could lead to memory exhaustion.[4] An exploit posted on the Internet by 'Kingcope' used GET requests with specially crafted Byte-range headers.[5] By mid-September, Apache had released an unscheduled security update that fixed the problem.

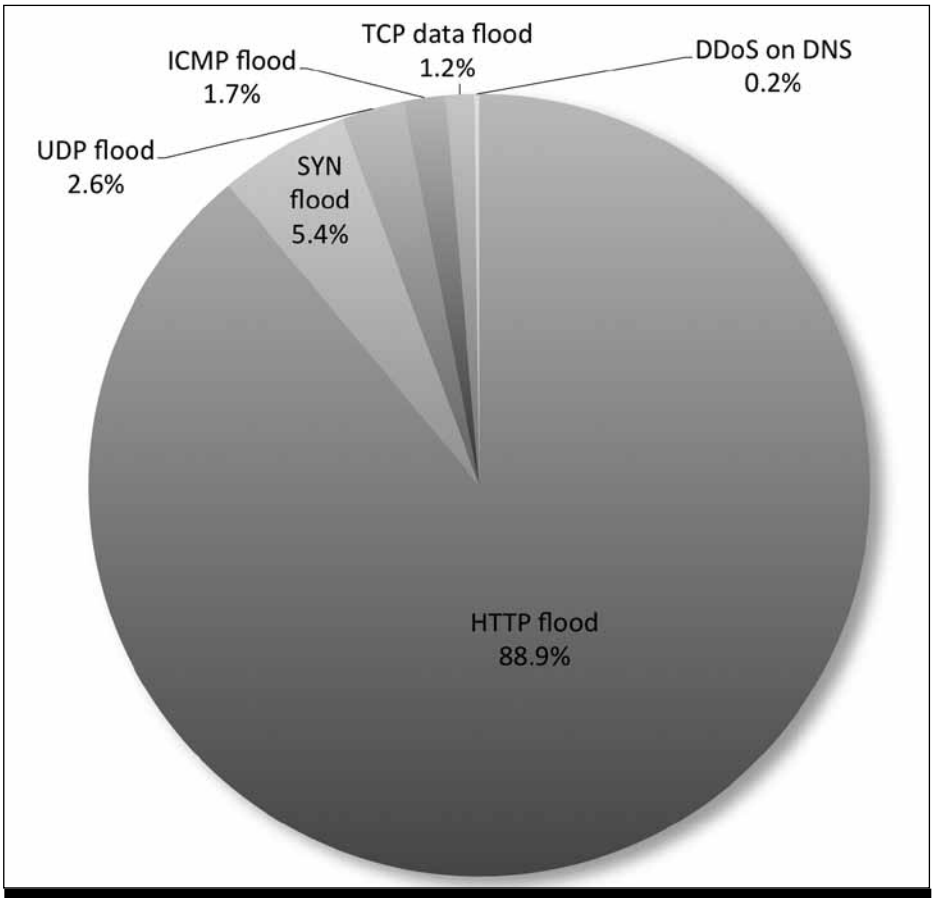

Types of DDoS attacks, Q2 2011. Source: Kaspersky.

Who are the attackers?

Hacktivism is what has driven DDoS into the mainstream press. The Internet Crime Center (IC3) highlighted the availability of hacktivist tools, such as the Low Orbit Ion Cannon (LOIC) as a possible reason for the rise in DDoS attacks.[6] Yet for all the headlines, and for all that Anonymous, LulzSec and their ilk may have contributed to the number of attacks, many still see hacktivism as something of a side issue when assessing the real threat of DDoS.

"It's mostly a nuisance," says Shulman. "Most organisations are not targets for hacktivists. But most organisations can be blackmailed."

Yet the hacktivism trend does highlight the ease with which even people with relatively limited technical skills and resources are able to mount DDoS attacks. On underground forums, the crude but effective Aldi Bot code is available for as little as 5.[7] Careful Googling will find bot sourcecode for free – not necessarily the most sophisticated or effective but good enough to take down an unprepared firm.

"We're seeing DDoS as a growing problem among our customers, probably as a result of the commoditisation of botnets," says Shulman. It's becoming cheaper for attackers to get their hands on larger botnets."

Cyber-criminals remain the primary users of DDoS. And they are becoming more creative. "For example," says Sop, "some of the attacks are used to distract. We were protecting a subsidiary, a UK bank in Russia, and that bank was attacked to divide the IT resources – fraud was happening, bank accounts were being emptied at the same time as large-scale disablement of ATMs, point of sale and money transfer functions."

DDoS attacks are also being deployed as part of stock pump-and-dump scams. A DDoS often has the effect of depressing the stock price of the victim. The attacker can then do lots of small transactions involving that stock, which is very difficult to track or prosecute.

And then there are state actors. According to Sop: "We see attacks come from certain Asian countries and, because

of these countries' use of censorship and Internet controls, there's absolutely no way that these attacks would be let out of the countries in the volumes that we have seen without some kind of state complicity."

According to its 'Anatomy of a Botnet' report, Arbor Networks says there has been a rise in politically motivated DDoS attacks around the world.[8] In 2007, the taking down of key systems in Estonia by attackers operating out of Russia is among the most notorious.[9] But Arbor has also noted attacks targeted at Iran, South Korea, Malaysia, China and the US. Kaspersky researchers also noted that botnet code originating from China lacked any attempt at stealth.[10] It's not clear if this is simply because they don't feel the necessity to be cautious because they are safe behind the Great Firewall and, perhaps, the Government's protection. And at the end of 2010, Harvard's Berkman Center for Internet & Society issued a report that showed how DDoS attacks are being used to silence civil liberties groups, human rights activists and independent media sites.[11]

Whoever is doing it, many of them are very skilled. "They understand the limits of load balancers, firewalls, web app firewalls, DDoS mitigation devices and, in general, the limits of the organisation that is mounting the defence," says Prolexic's Quinn. "It's very easy for them to see who is mounting the defence as well. People notice very early in an attack that we're involved, and we see things ramp up in a way that's really interesting."

Detecting an attack

However, knowing that you're under attack isn't always easy. At first, you could easily mistake the increase in traffic for just a good day for your website.

"If you've got a huge network, you could be seeing relatively big flood events all day, every day and not even notice them," says Woodhead. "Equally, you could have a small network that sees a small flood and it's completely disabled. Layer 7 or slow-type attacks, for example, are imperceptible to the customer until they get to a certain level. They may manifest themselves as a requirement for additional hosting capacity and go completely unnoticed."

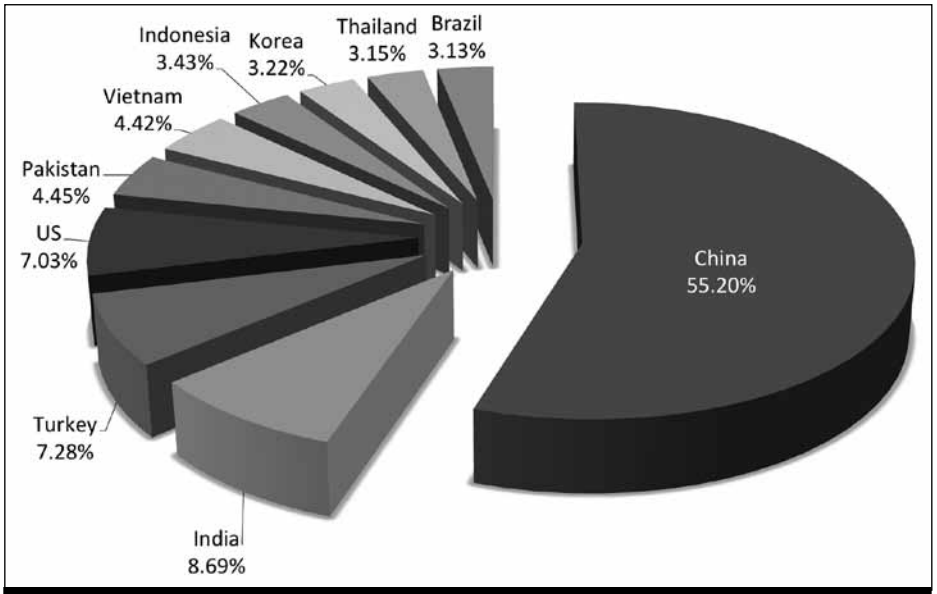

Country sources for DDoS attacks in Q3 2011. Source: Prolexic.

One of the difficulties is having the right expertise on tap. Automation can work to a degree, but isn't always perfect. "One interesting thing that people often overlook is the importance of analysing the attack to determine what the attacks vectors are in detail," says Sop. "If you just put the traffic through an automated device you can have a large number of false positives, and sometimes the appliance causes as many problems as the attack. So experience with these things is something that people often underestimate – how much variety there is in this ecosystem and how clever these people are. This is not like spam – there's a brain on the other end using very sophisticated system directed exclusively at you."

He adds that Prolexic uses a number of monitoring systems: "A lot of these are flow-based, which a lot of people use, so they take net flow feeds from routers." The firm also deploys analytical and correlation technology on its appliances. And it monitors the Internet. "We have an IP reputation database that we're running, currently tracking 10 million IPs," he says. "And we share this – we have data sharing agreements with some of the big credit card companies, at the national level with many different governments and law enforcement. We also work with our partners, and some of them have some deep monitoring on the Internet at large – one of our partners has monitoring feeds on over 50 backbones around the world, and they're able to see traffic globally. They also have correlation systems that allow them to identify command and control servers and such."

Quinn adds: "One of the other things we do is take both an inside and outside view of the customer's network – inside view via the appliance, outside view from Internet monitoring. Is the site functioning properly? Has the response time changed dramatically in recent time? Also we keep an ear to the ground; we monitor attack networks where possible to see if someone is a scheduled target, or is something ramping up."

Mitigating the attack

So what about mitigation? First, you need a sense of realism. "There's no 100% way of protecting yourself against a DDoS attack," claims David Jacoby of the Global Research and Analysis Team at Kaspersky.

The next step is to make sure your systems are in good order. Arbor Networks found that many firms fell victim to relatively minor DDoS attacks because of poorly configured firewalls and IPS devices.[12] But that doesn't mean that tidying up your firewall rules will make you immune. There are DDoS mitigation devices available, and many organisations have equipment from the likes of Arbor or Cisco on their premises. The problem is that with many kinds of attack, by the time the attack traffic has reached your perimeter, it's too late to do anything about it.

David Jacoby, Kaspersky.

"The sad reality is that when someone has gathered enough horsepower to throw this kind of attack at your network infrastructure, nothing that you put on your end of the line will help you," says Shulman. "Basically the attack is about jamming the incoming network pipe into your organisation."

Having a lot of bandwidth will help. But the amount of bandwidth you need to deal with a DDoS attack can be expensive and will sit unused most of the time. In addition, most big firms use multiple network suppliers, for redundancy, which multiplies the problem. And it just sets a target for the attackers to hit. "If you've got 10Mbps out and you get an 11Mbps attack, your CPE [Customer Premises Equipment] is doing nothing," says Petro. "The pipe's full, let alone your state traffic."

Having your ISP provide a 'clean pipe' is a partial solution, but doesn't address the desire for redundancy. "This becomes difficult when you have larger enterprises with multiple connections to the Internet because they might need to deal with multiple providers with different levels of SLA," says Sop.

The solution works best when it's pushed further upstream. The Network Operations Centre (NOC) of a very large organisation might have the capacity to monitor all the traffic, analyse and filter it. But these resources are more likely to be found within your ISP. "They can do it in their core routers," says Jacoby. "A lot of big ISPs have contact with other ISPs – you need to co-ordinate this through different core routers around the world. They can use the null route – route it to 0.0.0.0 – so the attack can't get much beyond the routers, it might just go through two or three before the packet is dropped. Most of the traffic will, most likely, look the same even if it comes from a different IP. It will have the same TCP sequence number, the same TCP headers, the same structure of the packet. And if we can detect that kind of structure, we can either redirect or simply block every request with that structure."

Mitigation services

Increasingly, firms are turning to specialist service providers for their DDoS mitigation. It's becoming big business. For example, in September 2011, Tata Communications announced it was pushing out its DDoS protection services globally. This is a service that's rapidly moving downwards in terms of size of organisation. The biggest firms already have DDoS protection in place, which means the attackers are looking at smaller, easier targets.

Verisign's Petro gives the examples of stealing money from banks via hacking. This is now very unlikely with big banks – the security is just too good. But in 2010, $70m was still stolen from banks in the US – mainly from small, local banks and credit unions that don't have the resources or funds for the levels of security enjoyed by the big players.

"So we've seen the problem move from something that was million-dollar botnets focused at multi-million dollar websites for coercion, blackmail and other reasons," he says, "moving all the way down now to small and medium sized businesses."

He adds that Verisign was mainly selling to large corporates who were paying, perhaps, $500,000 a year. Now it's finding an increasing amount of business among smaller firms, with services costing around $3,000-$4,000 a month.

Verisign's service is effectively a DNS swing. It maintains servers ready and staged with the customer's IP. "We do this in two ways," explains Petro. "At the upper end of the market the customer may have CPE on prem, and their first reaction will be from the CPE. They'll swing their web server, so we'll effectively become their ISP during that timeframe. So DNS moves all web traffic to us and we send that to a major scrubbing centre – we have two, one in Amsterdam and one in the US." The European centre is important because certain types of traffic – online gambling, for instance, can't be routed via the US for legal reasons. "We scrub that traffic as your ISP and pass on the clean packets to your web server. And when the attack has subsided, at the customer's request we swing that back." This happens very quickly. "If we're set up and we're wired to the customer, then we don't have to worry about the TTL, we don't necessarily have to worry about propagation because we are that physical address now. With a proactive customer with CPE that we've already provisioned, it's instantaneous – it's when the customer presses the button."

Some services operate well upstream. Prolexic, for example, uses both DNS and BGP to attract traffic to its cloud. It often operates upstream of the ISP.

Simwood's Woodhead reckons that, globally, about 80% of DDoS mitigation uses proxying via DNS. And that might be fine for websites: his company, however, is dealing with VoIP traffic. "So we mitigate at the network level," he explains. "We prefer to pass 100% of network traffic and then we can mitigate an attack on any service. And we do that either by the customer being directly connected to the network, or we can do it over GRE [Generic Routing Encapsulation].[13] In either event, the customer uses its own IP address, but behind our network. An ISP can use us just as it would any other transit provider. But in the event of one of their customers having an issue, they can cease announcing that customer's addresses through other ISPs to effectively force all traffic through us. So they have the ability to turn on mitigation through a subset of customers rather than building it themselves."

As for what Simwood actually does, it's what might be called a 'defence in depth' approach. "We have several layers," says Woodhead, "with dirty traffic one side, clean traffic coming out the other side. Those layers get more fine-grained as we go through. The DDoS element is one of those, although the others all contribute to it."

The front end is what Woodhead calls 'best network practice'. Simwood filters

out traffic coming from bogon sources – spoofed IP addresses or IP addresses that aren't actually in issue or IP addresses that are reserved for internal use. "If every ISP took efforts to not route traffic like this, then the Internet would be a far cleaner place," he says. "So right at the edge of our network we're filtering out all of that rubbish. That makes a big difference, particularly in a DDoS sense, because if you're interested in disturbing service, you don't want packets back. So using a spoofed IP address is common practice."

The next level uses IP reputation service ThreatSTOP.[14] "We're the first people in the world to apply their IP reputation for DDoS mitigation," claims Woodhead. "Given that a DDoS will normally be preceded by a network scan to try to determine points of vulnerability, that will normally be undertaken from a malware-infected machine or a co-lo machine that has a track record of being up to no good. By virtue of blocking traffic from disreputable IP addresses, we potentially prevent a DDoS from ever happening in the first place."

'Traditional' DDoS mitigation is the next layer. "It is network behavioural analysis and it is, generally speaking, passive," he explains. "It is basically monitoring everything that is going through and building up patterns of behaviour. So a typical source and destination pair of IP addresses, or just a destination address, will have a typical level of traffic at certain times of day. There are tens of thousands of metrics it measures, and what it aims to identify is something that is out of character. When it discovers something out of the ordinary it flips into mitigation mode. All that actually means is that it's far more aggressive in ageing connections, so your good traffic will be entirely unaffected, your bad traffic will find it far harder to do any damage."

Underneath that, Simwood has an IPS layer which also provides some DDoS mitigation because it does SYN proxying, rate limiting and so on. And it is using conventional signature-based analysis on everything coming through.

In the cloud

Inevitably, an increasing amount of DDoS mitigation takes place in the cloud. According to Petro, if you need to scrub 10Gbps or 15Gbps of data, "the cloud is the natural place to do that. If a cloud provider can sanitise that traffic and provide only the good queries to your perimeter, then you have the reduced cost of not having CPE, the reduced cost of not requiring ungodly amounts of bandwidth for something that may or may not happen to you, and third you don't have to hire the subject matter experts. The folks who can manage a DDoS attack, who can look at your Juniper router and start to make those changes, or play with the Cisco equipment or whatever, are expensive and high-demand personnel."

It's the cloud approach that allows Prolexic to operate so far upstream. "Our strategy starts in the cloud where we're able to accept that attack traffic regionally before it even gets close to the customer," says Sop. "So we have a global set of Tier 1 connections – it'll be over 400Gbps in 2012. When someone turns this on, it's like snapping in a 400Gbps global capability in front of all of their ISPs, which is distributed to be close to where the attackers are."

But Rapid7's Carey points out that the cloud can be a double-edged sword. "I can spin up 100 machines right now, with someone like Amazon, and do a slow DDoS attack on somebody," he says. "The cloud is the mitigation and it is the force multiplier too."

Risk analysis

Given that DDoS mitigation services are expensive, how do you know if you need them? Carrying out a risk analysis can be very tricky. However, it seems that an increasing number of firms are erring on the side of caution.

"We conducted a survey," says Petro. What was interesting was that of the folks that did not have a DDoS mitigation solution, 71% said they were modelling it into their budget and would be buying one in the next four quarters. Will they follow through on that? Don't know. But, in my opinion, if you'd taken that survey two years prior, it would have been more like 7%."

What it comes down to is the nature of your business. All kinds of organisations can find themselves at the nasty end of a DDoS attacks these days but, clearly, if your business is online gambling, financial services or e-commerce then your expectation of an attack is going to be that much greater. It might even be simpler than that. "I think that the companies that are leveraged more than 10% of their revenue online are now actively aware of DDoS," says Petro, "and it's no longer a case of 'do I buy the insurance?', it's 'how much insurance do I buy?'."

He gives the example of online travel booking firms such as Travelocity. "For them, 100% of their business is online. And it's immediacy. If you're down for a minute, you have customers that are unable to get their tickets, unable to find their confirmation numbers. If you're down for an hour, you've lost customers. If you're down for a day, that might be the end. But if you're GM and the website goes down? People are still in the dealerships buying cars – not much has changed for GM that day."

This is not a security issue, it's a business continuity issue. "The larger organisations already have the infrastructure in place to absorb huge amounts of traffic," says Shulman. "They usually have agreements in place with their ISPs to help them in case of traffic floods. And they have done that not so much with respect to DDoS but just making sure their online business is available at all times."

It's likely, however, that many firms continue to view DDoS as a security concern, even though an effective attack rarely results in data breaches or intrusions. And the fact that this is really about resilience rather than security can take some by surprise. When Rapid7 launches a DDoS simulation, how often does it show that the client company would be brought down? "It's 100%," says Carey. And how often are customers surprised by this? "They see DDoS in the news, so they understand that it can happen, but they don't understand, sometimes, how their organisation could be brought down within minutes. It's always a surprise when you spend a lot of money on bandwidth, servers, load balancers and all that stuff and then I just launch an attack and your website is crippled in a matter of minutes. It's always a shock."

The fact that so many more types of organisation are being attacked should act as a wake-up call. According to Woodhead, firms need to stop thinking of DDoS as a 'black swan' event – in other words an event that, in spite of its severity and impact, needn't be considered because it's impossible to predict. "We're very firmly of the view of thinking of DDoS as a black swan event is misplaced," he says. "The black swan aspect may be the scale of the DDoS, and the degree of the denial created by it, but the kind of attack people associate with DDoS happens every day on most networks."

About the author

Steve Mansfield-Devine is a journalist who has covered the IT industry for more than 30 years. He specialises in infosecurity and is the editor of Network Security *and its sister publication* Computer Fraud & Security. *He is also a Certified Ethical Hacker.*

References

1. 'Prolexic Attack Report Q3 2011'. Prolexic, Nov 2011. Accessed Nov 2011. <http://ww.prolexic.com/l/9892/2011-11-15/68YX>.
2. Leyden, John. 'DDoS bot infests food processing firms'. The Register, 1 Feb 2011. Accessed Nov 2011. <http://www.theregister.co.uk/2011/02/01/ddos_bot_targets_industrial_control_firms/>.
3. 'The Hacker's Choice releases SSL DoS tool'. The Hacker News, 24 Oct 2011. Accessed Nov 2011. <http://thehackernews.com/2011/10/hackers-choice-releases-ssl-ddos-tool.html>.
4. CVE-2011-3192. Common Vulnerabilities and Exposures. Accessed Nov 2011. <http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2011-3192>.
5. 'Apache httpd Remote Denial of Service (memory exhaustion)'. Full Disclosure list. Accessed Nov 2011. <http://lists.grok.org.uk/pipermail/full-disclosure/attachments/20110820/848b4dca/attachment.obj>.
6. 'Internet Crime Complaint Center's (IC3) Scam Alerts'. IC3, 14 Jul 2011. Accessed Nov 2011. <http://www.ic3.gov/media/2011/110714.aspx>.
7. Leyden, John. 'Bargain-basement botnet kit – yours for just €5'. The Register, 22 Sep 2011. Accessed Nov 2011. <http://www.theregister.co.uk/2011/09/22/aldi_bot/>.
8. 'Anatomy of a Botnet'. Arbor Networks. Accessed Nov 2011. <http://www.arbornetworks.com/white-papers-global-network-security-topics.html>.
9. Anderson, Nate. 'Massive DDoS attacks target Estonia; Russia accused'. Ars Technica, May 2007. Accessed Nov 2011. <http://arstechnica.com/security/news/2007/05/massive-ddos-attacks-target-estonia-russia-accused.ars>.
10. 'Chinese DDoS bots lack sophistication, stealth'. Threat Post, Kaspersky, 5 Oct 2011. Accessed Nov 2011. <http://threatpost.com/en_us/blogs/chinese-ddos-bots-lack-sophistication-stealth-100511>.
11. '2010 Report on Distributed Denial of Service (DDoS) Attacks'. Berkman Center for Internet & Society, 20 Dec 2010. Accessed Nov 2011. <http://cyber.law.harvard.edu/publications/2010/DDoS_Independent_Media_Human_Rights>.
12. 'Poor firewall implementations pave way for DDoS attacks'. InfoSecurity, 1 Feb 2011. Accessed Nov 2011. <http://www.infosecurity-magazine.com/view/15596/poor-firewall-implementations-pave-wave-for-ddos-attacks>.
13. 'Generic Routing Encapsulation'. Wikipedia. Accessed Nov 2011. <http://en.wikipedia.org/wiki/Generic_Routing_Encapsulation>.
14. ThreatStop. <http://www.threatstop.com/>.

Could 'wait and see' be the best IPv6 strategy?

Jérémy D'Hoinne,

Jérémy D'Hoinne, NETASQ

The year 2011 was supposed to be the year for IPv6. The depletion of version 4 addresses from the top level provider (IANA), and the announcement of IPv6 World Day on 8 June, set the tone. The message was that we need IPv6 and we need it now. While tech-Nostradamuses have never been in short supply in the IT sector, this particular message gained traction in the media. Put simply, it was a nice story: the exponential growth of the Internet has continued to the point where we now need more IP addresses than there are people on earth.

There are a few facts we can't deny here: in an ever-growing e-world, IPv4 addresses will soon run out. In fact, certain Internet providers in Asia, already facing this challenge, have successfully moved to IPv6, most notably for their IP-TV offerings. This has given rise to a whole section of the Internet that cannot be accessed from IPv4 addresses.

However, if we must advise IT professionals in Europe and the US, we can't rely solely on this macroscopic analysis. If you'll permit an analogy with global warming, this is a problem for everyone,