The Truth about DDoS Attacks

by Kelly Beardmore, CEO of Carbon$_{60}$ Networks

Distributed Denial of Service (DDoS) attacks are a plague on the Internet. In the last 5 years, the magnitude of DDoS attacks has steadily increased to the point that in a recent NYT interview Matthew Prince from CloudFlare compared DDoS attacks to "nuclear bombs" in their capacity to cause wide-spread damage. While, as the CEO of an Internet security company, Mr. Prince is guilty of some self-interested hyperbole there is no question that the amount of targeted and collateral damage from DDoS attacks continues to escalate. Not just individual websites but major hosting centers and regional Internet services are being compromised by these attacks. At the same time it only takes a small DDoS attack to disrupt the vast majority of websites on the Internet. In this article, we will examine the scope of the DDoS problem and then make suggestions on the best ways of dealing with this threat if you have a business-critical website.

Both the frequency and amplitude of DDoS attacks have risen rapidly in the last five years. Figure A shows the average growth of DDoS attacks as well as the peak size of attacks monitored by Arbor Networks since 2009. Given that 100 Mbps of sustained traffic will grind the vast majority of websites to a halt, today's average-sized DDoS attack of 1 Gbps (i.e. 100 Mbps x 10) is almost always fatal to its intended target.

Both the frequency and amplitude of DDoS attacks have been rising rapidly in the last five years.

Moreover, as an attack grows beyond 100 Mbps the likelihood of collateral damage increases rapidly. Any inbound or outbound attack 100 Mbps or larger is very likely to "sideswipe" other sites sharing hosting infrastructure with the target website. The extent of the collateral damage will depend on the exact nature of the hosting solution and the hosting provider's system architecture. The more shared and the less distributed the architecture is, the greater the risk of collateral damage. Attacks in the 5 to 10 Gbps range are likely to cause significant collateral damage to all sites hosted in the same datacenter. Attacks in the 50-100 Gbps range are likely to cause serious issues for the world's largest hosting facilities as well as regional Internet services.

Figure A: Average Monitored Attack Sizes Month-by-Month (January 2009-Present)[1]

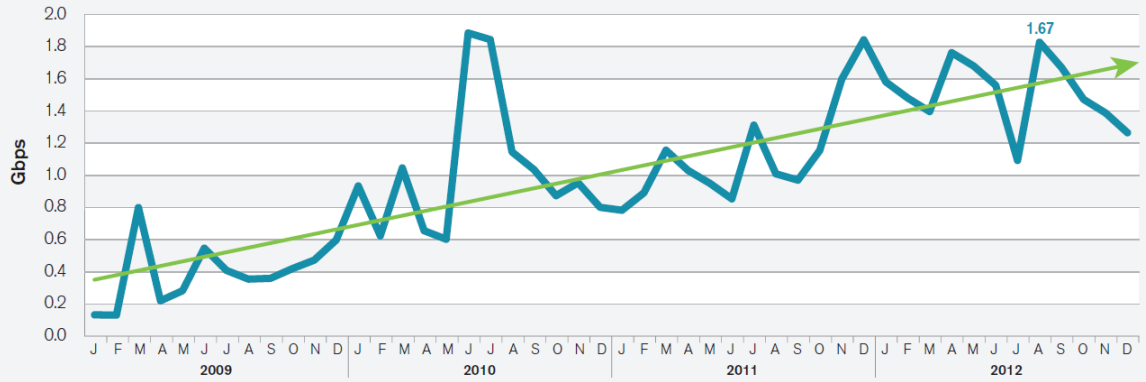

Major Internet hubs routinely handle Tbps of Internet traffic and, therefore, remain responsive through even the largest attacks. However, even these hubs may be threatened in the future as the peak amplitude of DDoS attacks increases through: the ongoing development of more powerful DDoS tools and techniques; the possibility of a state-sponsored attack; and the global explosion of broadband connected wireless devices which offers a powerful new platform for high-powered DDoS attacks.

It is not just the growing amplitude of DDoS attacks that is concerning but also their increased frequency. For every high profile DDoS attack reported in the mainstream media – such as the recent massive and sustained attack on Spamhaus – there are thousands of DDoS attacks that go unreported every day. For example, between 2009 and 2011Akamai reported a 2,000% increase in DDoS attacks over its network which handles between 15% and 30% of the Internet's total traffic.[4] A newly released study conducted by the Poneman Institute and Radware claimed that as many as 65% of organizations were the victim of at least three DDoS attacks in the past 12 months.[5] Some network and hosting providers claim to see hundreds of DDoS attacks per month and these are only the ones big enough to get noticed.

This rise in frequency is mainly due to three factors: 1) the ease by which an effective DDoS attack can be launched. You can download a readily available DDoS tool to do it yourself or contract a "hacker-for-hire" to attack the target of your choice for about $5 to $10 per hour; 2) the emergence of DDoS attacks as a form of political protest; 3) and the continued growth of international cyber-crime.[6]

Figure B: Most Common Motivations Behind DDoS Attacks[2]

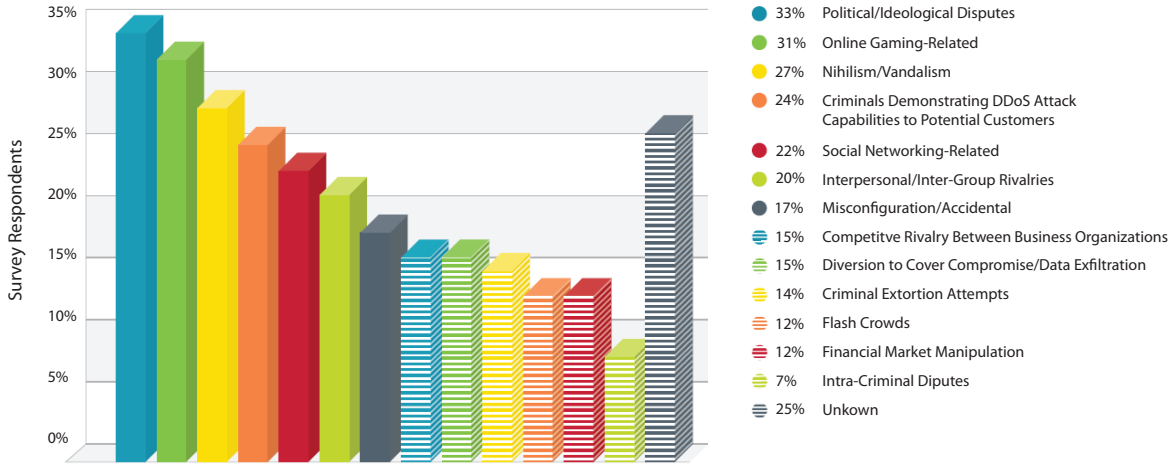

DDoS attack vectors tend to fall into one of three broad categories:

1. **Volumetric Attacks:** These attacks are about causing congestion. They attempt to consume a target's available hosting resources and are typically executed using botnets to generate a high volume of http/s page requests. Attacks on VoIP and authoritative DNS servers are also popular ways to disrupt service. Recently the magnitude of volumetric DDoS attacks has increased significantly by leveraging the recursive function of tens of thousands of misconfigured DNS servers on the Internet to "amplify" attacks.[7] This development represents another escalation in the ability of DDoS attacks to cause wide-spread collateral damage.

2. **TCP State-Exhaustion Attacks:** These attacks attempt to consume the connection state tables that are present in many infrastructure components, such as load balancers, firewalls and the web application servers. Even high-capacity devices capable of maintaining the state of millions of connections can be taken down by these attacks. Since 2011 there has been a rise in this type of attack on datacenter-level devices in an attempt to maximize collateral damage.[8]

3. **Application-Layer Attacks:** These attacks target a weakness in a particular web application. They are the most sophisticated, stealthy-type of DDoS attacks because they can be very effective without generating abnormal amounts of traffic. This "low and slow" approach makes the attack very difficult to detect using traditional volumetric detection mechanisms. Recently,

Kevin Kennedy, Senior Director of Product Management at Juniper Networks, noted in a blog post: "Forget armies of bots, a single PC was enough to generate a small, well targeted attack that took down one of the e-tailers in Europe within 2 minutes. And precisely because it was so small, it was lost in the noise of legitimate user traffic, taking a full day to identify and remediate and putting $10M of sales at risk."[9]

In a recent survey by Arbor Networks, almost double the number of respondents reported multi-vector DDoS attacks (27% to 46%) in 2012 over 2011. This is a dangerous trend as multi-vector attacks put additional strain on security resources and requires an expertly managed "defence-in-depth" security strategy and response plan to mitigate effectively.

While collateral damage is rising with the increased amplitude and frequency of DDoS attacks, the scope of DDoS targets also remains broad. Although the favoured targets remain e-commerce and gaming sites, all types of sites are attacked and often for no discernible reason. See Figure C.

At the same time, a strong majority of organizations recognize that any service disruption for any reason would have a significant impact on their business. See Figure D. Taken together, these factors – the increased frequency and amplitude of DDoS attacks, the wide scope of targets, and the increased sensitivity of organizations to DDoS attacks – mean increased business risk for most organizations and a thriving DDoS mitigation industry.

We will now explore how a "defence-in-depth" approach is necessary to protect your business-critical website from both infrastructure and application-level DDoS attacks.

Figure C: Targeted Customer Types[3]

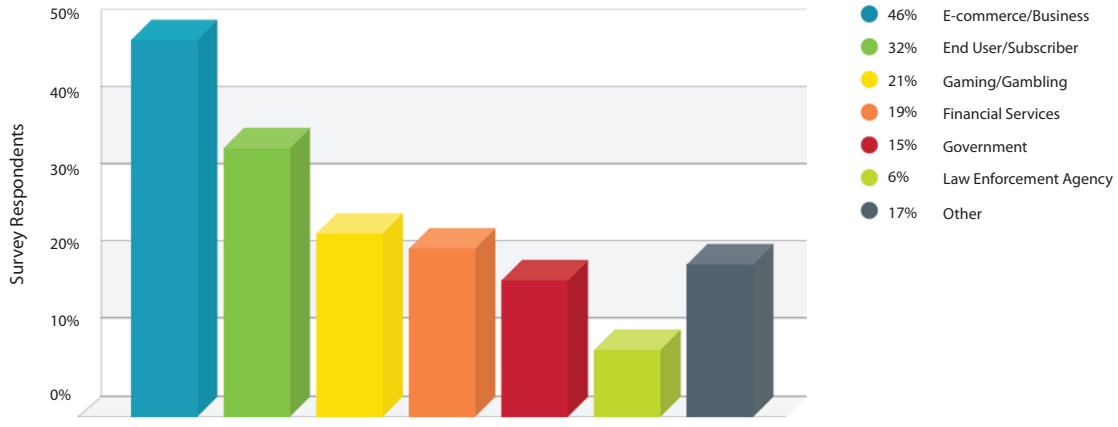

Figure D: Business Impact from Interruption of Services

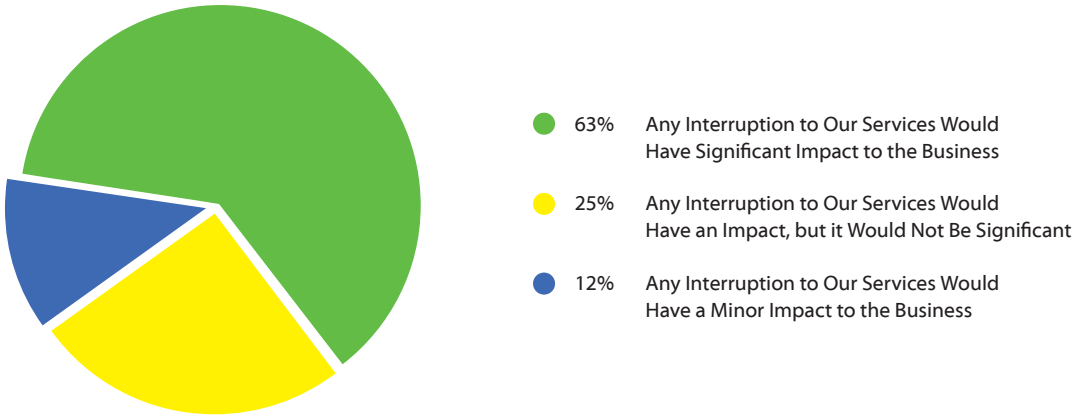

To avoid the direct or collateral damage from even a modest DDoS attack you need a highly scalable and robust multi-origin hosting infrastructure. Or, in other words, you need global edge routing and caching services (e.g. Akamai) to deflect infrastructure level (Layer 3 & 4) attacks and to dissipate application-level (Layer 7) attacks by directly serving requests for cached content and distributing the remaining valid http/s requests between multiple geo-distributed hosting origins.

A global edge network shields your hosting origin(s) a number of different ways. First, an edge network significantly reduces the "attack surface" of your site by eliminating any direct path of attack on your hosting origin(s). This is done by masking your origin IP(s) and firewalling out any traffic not originating from the edge network. Without the ability to reach your origin directly, low level network attacks (e.g. DNS, SYN, ACK and ICMP-type DDoS attacks) will fail. They will be deflected by the edge network since it will only forward valid http/s requests. This is very significant because according to Prolexic's Q1 2013 report approximately 75% of all DDoS attacks are infrastructure-level attacks that would be fully deflected by the edge network. See Figure E for a breakdown of attack types and relative frequency.

To avoid the direct or collateral damage from even a modest DDoS attack you need a highly scalable and robust multi-origin hosting infrastructure.

Second, an edge network can reduce the risk from DDoS attacks through its routing logic. For example, edge routing rules can be configured to deflect http/s requests from

Figure E: Q1 2013 Reported DDoS Attack Types.[10]

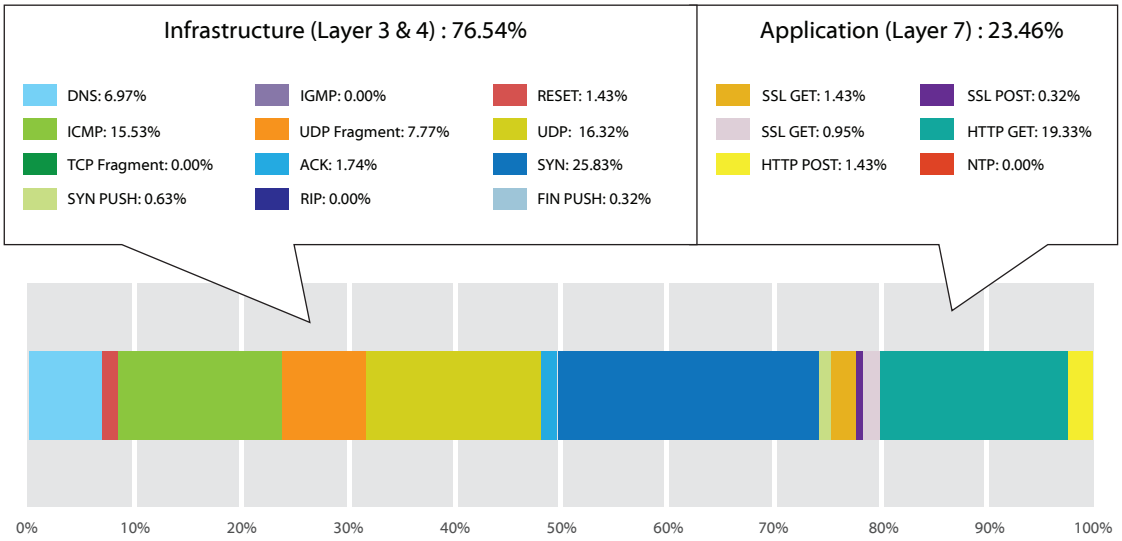

high-risk regions like China and South East Asia by either dropping them altogether or redirecting them to a secondary site safely away from your primary site. For organizations with a truly global audience, edge routing rules can also direct http/s requests to geo-specific versions of your site hosted at different origins. In this way, there is a better chance any DDoS attack will be defused as http/s requests are dispersed across multiple origins or localized to a single region depending on the source(s) of the attack. Figure F shows where DDoS attacks originated from in Q4 2012. It should be noted that this distribution can change month by month as new botnets – the engines for DDoS attacks – are brought online in different regions at different times. For example, a recent exploit aimed at hosting providers in the US and Canada saw a surge in DDoS attacks from North America last quarter while India and Russia – historically high threat regions – saw a significant drop in activity.

Third, the robustness of edge network services makes them an ideal anti-DDoS weapon against application-layer attacks. For example, Akamai's global edge network is composed of more than 120,000 servers across 1,100 networks globally and is capable of handling almost a third of all Internet traffic. By effectively leveraging this scalability, you can shield your site from even the largest http/s-type DDoS attacks or surges in legitimate web traffic.

Leveraging the immense scalability of edge network services against application layer DDoS attacks requires intelligently caching

Figure F: Q4 2012 Reported Source of DDoS Attacks[11]

Just over 18% of observed attack traffic originated in North and South America, just under 25% originated in Europe, and 56% originated in the Asia Pacific/Oceania region. The remaining 1% of attack traffic originated in Africa.

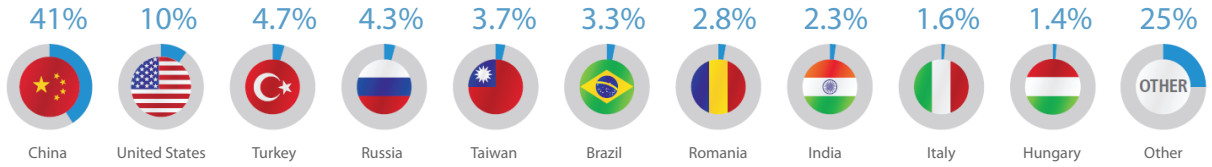

The blue areas represent each country's percentage of the overall total amount of attack traffic observed by Akamai

as much of your web content as possible on the edge network. For the vast majority of sites, a large portion of their web "objects" (html files, javascript files, media files, etc.) are cacheable on the edge network. In fact, many dynamically created web objects are also cacheable. With proper tuning Carbon$_{60}$ often sees edge offload rates (i.e. the percentage of edge served requests to total web requests) above 80%. If the benefits of this edge caching were distributed evenly across a website then, all things being equal, an 80% offload rate would translate into a five-fold increase in the number of concurrent http/s requests a website could manage without any increase in response time. A five-fold improvement in scalability is good in fending off application level DDoS attacks.[12] However, real world performance is often far better.

If there is a flood of valid http/s requests aimed at a URL where all its constituent objects are cacheable – which is often the case – then the attack, regardless of size, will crash against the edge network. Of course, the opposite is also true. A DDoS attack targeted against a very transactional page will let the weight of the attack fall more directly on the more vulnerable hosting origin. As a result, you need more than edge network services to fend off the most targeted, most skilled attacks.[13]

The strength of edge network services against DDoS attacks is the result of its highly distributed nature. To further strengthen your DDoS defences against application-layer DDoS attacks and limit the collateral damage from DDoS attacks against others, you must also distribute your origins. The best way to do this is through a multi-origin cloud hosting solution that works in combination with your edge routing and caching services. This is for a couple of reasons. First, scalability is intrinsic to cloud computing. A cloud hosting provider can scale your compute capacity by adding more computing power to existing virtual machines or by adding more virtual machines to existing web,

application, or database clusters, to compensate quickly for surges in traffic during an html/s type DDoS attack. Better yet, you only pay for this reserve computing power while you need it.

Second, only hosting your site from a single origin makes your site vulnerable to any service impacting events at that origin. This includes DDoS attacks against other sites sharing the same cloud compute, storage, or network infrastructure. While it is always advisable to host in a good neighborhood, i.e. away from high risk sites, it is better to geographically distribute the hosting of your site. This can be done using an edge network's routing logic to balance traffic between multiple origins and automatically re-directing traffic away from an unavailable origin. This works particularly well when inbound requests are routed according to their geographic source to multiple hosting origins within a corresponding region.

Another approach is to use the edge network itself as a standby origin by leveraging the edge network's own storage service to host a static, "DDoS-proof" version of your website. While not capable of serving dynamic content or processing orders, this solution at least maintains your web presence and permits emergency updates to keep customers informed.

In general, the cost of multi-origin hosting solutions has reduced dramatically in recent years and the supporting technologies for distributed computing have improved to support a broader-range of applications. While it is beyond the scope of this article, it is worth investigating these technologies when assessing your next web application framework.

In short, leveraging edge network services along with a multi-origin cloud hosting solution delivers a very robust solution to the escalating magnitude and frequency of single and multi-vector DDoS attacks. At the same time, this solution provides significant benefits related to site reliability, performance, scalability, and security that are not provided by the many dedicated DDoS filtering services on the market today. Of course, selecting a robust, DDoS-hardened hosting solution is only one aspect of an overall site security strategy. Any high value, business-critical web property requires a comprehensive security strategy that encompasses all aspects of web development and delivery.

1 - Figure A source: Arbor Network's Worldwide Infrastructure Security Report, 2012 Volume VIII

2 - Figure B source: Arbor Network's Worldwide Infrastructure Security Report, 2012 Volume VIII

3 - Figure C source: Arbor Network's Worldwide Infrastructure Security Report, 2012 Volume VIII

4 - See http://www.circleid.com/posts/20120131_ddos_attacks_increased_by_2000_percent_in_past_3_years/

5 - See http://www.securitybistro.com/blog/?p=3683

6 - It should be noted that most DDoS attacks are not reported to law enforcement because of a lack of time and resources, low confidence in the efficacy of law enforcement, and corporate policy.

7 - See http://news.techworld.com/security/3407339/open-dns-resolvers-used-to-amplify-ddos-attacks-hide-original-source/

8 - See http://pages.arbornetworks.com/rs/arbor/images/WISR2012_EN.pdf

9 - See http://forums.juniper.net/t5/Security-Mobility-Now/It-s-Not-Size-But-Sophistication-That-Matters/ba-p/185087

10 - Figure E Source: Prolexic Q1 2013 Quarterly Report

11 - Figure F Source: Akamai Q4 2012 State of the Internet Report

12 - Conversely, the same site could theoretically be able to handle the same number of requests with a five-fold decrease in resource usage at the origin – which, of course, is how companies cost-justify their edge services.

13 - Thankfully, only a small percentage of DDoS attacks are performed by highly skilled and motivated attackers that use DDoS tools as anything more than blunt instruments. For example, Prolexic believes the largest attacks are the work of a relatively small core of veteran mercenaries and the majority of attacks are the work of "script kiddies".

Carbon$_{60}$ is a safe haven for private and public sector organizations that need a trusted partner to manage their business-critical IT workloads. By embracing the hosting solution stack end-to-end, we deliver extraordinary value and assume full responsibility for the quality of service delivered from the datacenter to the end user's doorstep.

Carbon$_{60}$ Networks
17705 Leslie Street - Suite 9
Newmarket, ON,
Canada, L3Y 3E3

sales@carbon60.com
www.carbon60.com
1.888.227.2666