
Mitigating DDoS Attacks at Layer 7

**Detect, Localize and
Mitigate using DNS GSLB**

Allan Jude

ScaleEngine Inc.

Introductions

Allan Jude

- 12 Years as FreeBSD Server Admin
 - Architect of the ScaleEngine CDN (HTTP and Video)
 - Host of BSDNow.tv & TechSNAP.tv Podcasts
 - Former Professor @ Mohawk College (2008-2011) teaching Network Engineering and Security Analysis
 - Extensive work with Puppet to manage our 80 odd servers in 28 data centers in 10 countries
 - Lots of work with ZFS to manage large collections of videos as well as extremely large website caches (15+ million objects in ff*ff directories)
-

History

From 2002 - 2012 I ran an IRC shell provider, specializing in hosting smallish IRC networks as well as some protection bots and BNCs (proxies)

During this time my servers were subject to numerous different types and sizes of DDoS attacks and other attempts to disrupt services (take down IRC servers, gain control of rooms on EFNet etc)

Overview

- DDoS Attack Types
 - Defensive Techniques
 - Commercial Solutions
 - Using a GSLB to Localize and Mitigate
 - Implementation Examples
 - Where to go from here
-

What Is Distributed Denial of Service

- "A Denial of Service attack (DoS) is any intended attempt to prevent legitimate users from reaching a specific network resource."
[1]
-

Types of DDoS

- Volumetric / Link Saturation (Bandwidth)
 - Protocol Attacks (SYN flood, fragments)
 - Packet Storm (Excessive PPS)
 - Resource Starvation (CPU, I/O, Memory)
 - Stealth/Creeper (Slowloris, Slow POST)
 - Exploit (Application or OS Specific DoS)
 - DoS L2 - Intentionally trigger defensive mechanisms to block traffic from the (spoofed) source address or subnet
-

Statistics

- Average attack lasts 34.5 hours
- China is #1 origin of DDoS traffic, making up 40-50% of all botnet activity
- 75% of attacks are against Infrastructure (layers 3 & 4, SYN Flood, ICMP/UDP attack) with only 25% against Applications (layer 7)
- 25% of all attacks are under 1 gbps, and 50% of all attacks are under 6 gbps

[Source] Prolexic Quarterly Global DDoS Attack Report 2013Q1 - http://www.prolexic.com/kcresources/attack-report/attack_report_q113_english-version/Prolexic_Quarterly_Global_DDoS_Attack_Report_Q113_041613.pdf

Planning ...



Defensive Techniques

- Simple Failover
 - Null Route (Automated or Manual)
 - Web Application Firewall
 - Anycast
 - Proactive Name servers (a service from your registrar that provides failover to additional name servers, especially useful in the event of a DDoS against your name servers)
 - Commercial Solutions
-

Simple Failover (Hidden Spares)

- When a machine or location is under attack, fail over to another machine/location
 - Works better in Active/Passive mode, because in Active/Active the attacker may attack both locations
 - If you have a sufficient number of hidden spares, you may be able to evade the attacker for a while
-

Null Route (Blackhole)

- You or your provider send /32 routes upstream, routers stop forwarding the attack (and all other) traffic to you, preventing the saturation of your link
 - This allows traffic to adjacent machines or customers that are not under attack to continue normally
 - Some providers implement this automatically to lessen the disruption of a single customer being attacked
-

Web Application Firewall

- Only protect against specific known attacks
 - Actually meant to prevent intrusions and exploits, rather than brute force attacks
 - Are often a bottleneck or failure point because they can handle only very limited packets-per-second and often have underpowered CPUs, can actually amplify the attack and take you down sooner
 - Usually subject to limits on the number of concurrent sessions/connections or other scalability issues
-

Anycast

- Announce BGP routes for a single prefix from multiple locations
 - Traffic is directed based on fewest hops
 - Automatically distributes traffic between locations based on source network
 - Limits damage caused by DDoS attacks to the areas nearest the attackers
 - Requires your own IP space or LoA
 - Often requires 24/7 NOC
 - Much harder to maintain and scale on a limited budget with limited personnel
-


Commercial Solutions

- Filtering Hardware
 - Arbor Networks (PeakFlow) - Profiling and trend analysis, attempts to automatically detect out of character traffic in addition to known attack patterns
 - Checkpoint (DDoS Protector Appliances)
 - Cisco (Guard XT)
 - Fortinet (FortiDDoS)
 - Protected Hosting
 - BlackLotus - Dedicated Servers or (BGP) GRE Tunnel to your own network
 - Staminus - Dedicated Servers
-

Commercial Solutions (Continued)

- Protection for Hire
 - Prolexic (PLXRouter) - Prolexic announces your BGP routes and Inbound traffic is routed through Prolexic 'Scrubbing' facilities then on to your facilities via a GRE tunnel. Asymmetric (outbound traffic comes directly from your facilities)
 - Prolexic (PLXProxy) - HTTP and HTTPS only, traffic is symmetrically routed through scrubbing facilities
 - CloudFlare - Application Proxy. HTTP(S) only
 - DDoSArrest - HTTP(S) only
-

Global Server Load Balancer



Balance is the Key to Life

Using a GSLB


Another option for dealing with inbound DDoS is to mitigate with a GSLB. The solution we use, `gdnsd` (in ports), brings with it a number of useful features.

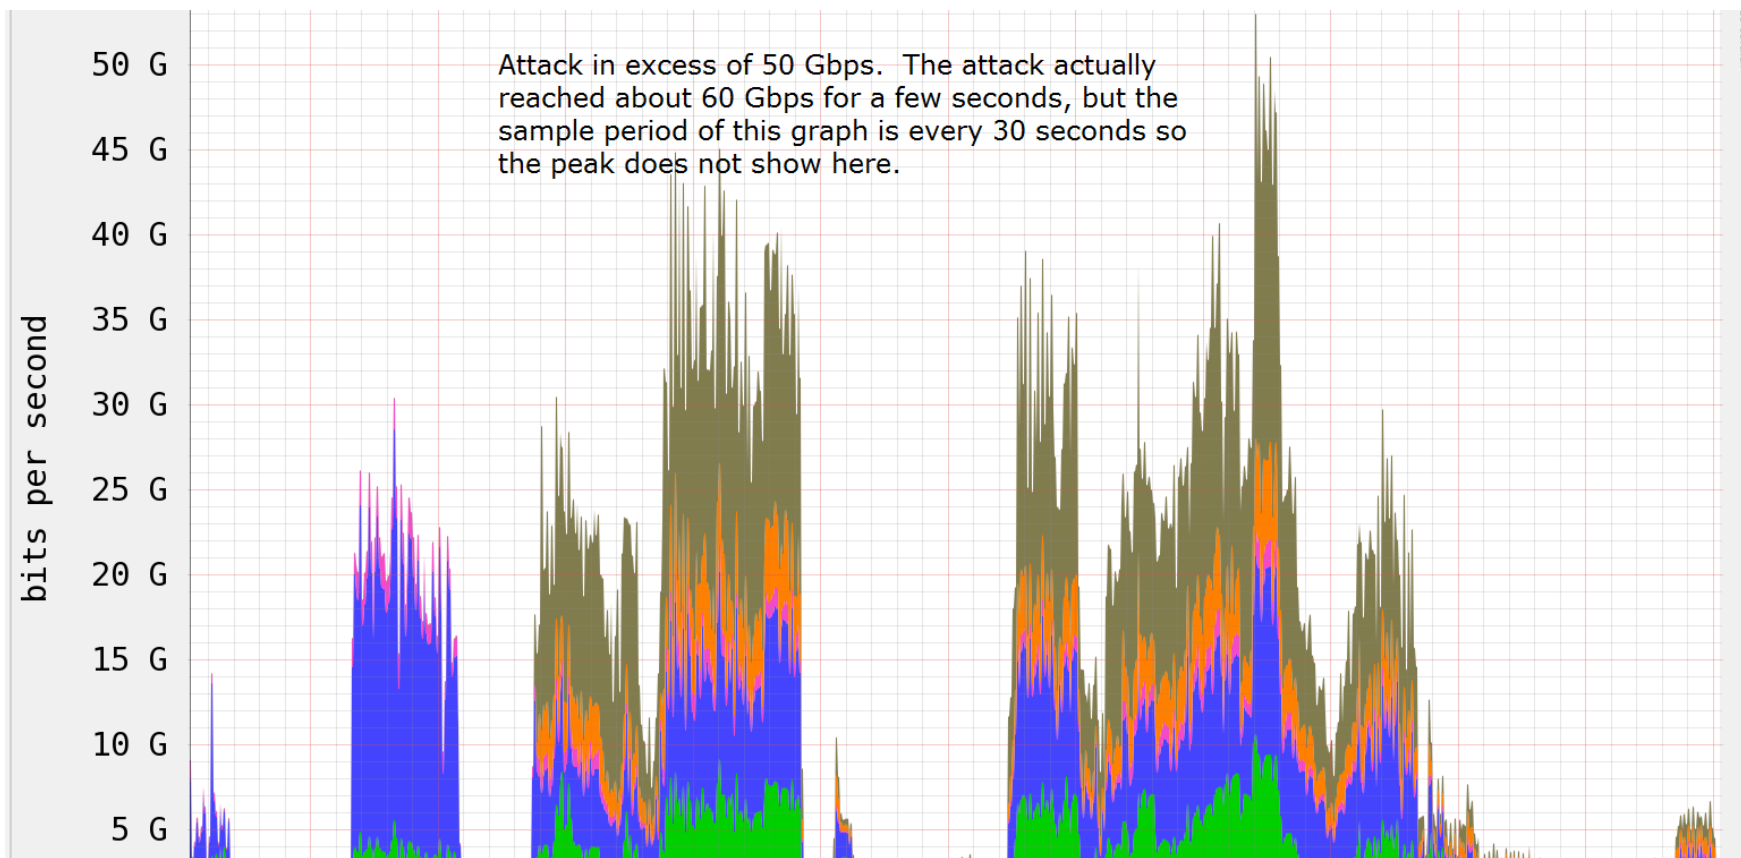
The first of these is its integration with GeoIP. Providing some of the functionality of Anycast, it allows you to localize an attack. If most of the attackers are in Europe, your North American nodes will remain up

Before Geographic Mitigation



After Geographic Mitigation





“Transit providers are simply not going to transport more than 40 Gbps of dirty traffic across the world; the only solution is to launch multiple geographically distributed scrubbing centers”

[1] <https://blog.staminus.net/mitigation-of-attacks-exceeding-40-gbps>

Sinkhole

Another option with the GSLB is to sinkhole all traffic from a specific region. If the attack is coming from zombies in China, you can point all of those clients to an unroutable address

GSLB also allows 'overrides' to the GeoIP database, allowing you to return a different response for addresses in specific network ranges. This can be used to prevent known bots on some blacklist from reaching your sites

Sinkhole Mitigation - DC Map

```
plugins => { geoip => { maps => { geo_map => {  
  geoip_db => GeoIPCity.dat,  
  datacenters => [US-WEST, EU-WEST, EU-EAST, NULL],  
  map => {  
    EU => {  
      NL => [ EU-WEST ],  
      DE => [ EU-WEST ],  
      PL => [ EU-EAST ],  
      default => [ EU-WEST, EU-EAST ]  
    },  
    NA => { default => [ US-WEST ] }  
  }  
}
```

Identify and Classify Attackers

```
nets => { #Bogons
```

```
  10.0.0.0/16 => [ NULL ],
```

```
  127.0.0.0/8 => [ NULL ],
```

```
#China
```

```
  1.80.0.0/13 => [ NULL ],
```

```
  1.192.0.0/13 => [ NULL ],
```

```
  1.202.0.0/15 => [ NULL ],
```

```
#Autogenerated list of attackers below
```

```
},
```

```
},
```

Adapt DNS results sent to Attackers

```
resources => {  
  prod_www => {  
    map => geo_map  
    service_types => default  
    dcmmap => {  
      US-WEST => 192.0.2.1,  
      EU-WEST => [ 192.0.2.4, 192.0.2.5, 192.0.2.6 ]  
      EU-EAST => {  
        lb01 => 192.0.2.2, lb02 => 192.0.2.3 },  
      NULL => 127.0.0.2,  
    }  
  }  
}
```

Whack-a-mole

If you have various geographically separated nodes, a less graceful approach I have used in the past

- All traffic is sent to node A and everything is fine
 - Attack starts, and overwhelms node A
 - GSLB kicks in and redirects new users to node B
 - Attackers often only do an initial DNS lookups or cache the results, so they keep attacking node A for a while
 - Attackers eventually migrate to node B
 - GSLB shifts load to node C
 - repeat
 - Shift back to node A, which has since recovered
-

Detecting an Attack

In order to respond to an attack, you must first detect when you are under attack

In our case this is especially important. With video streaming, it is not uncommon for our servers to see very large sudden spikes in traffic; this does not necessarily indicate an attack.

Even a large spike of incoming HTTP requests from diverse sources does not presage an attack. We host an advertising network that is used on CBS / CBSSports.com TechRepublic.com, CollegeHumor.com and various others. A large news event can drive a sudden surge of traffic.

Business Rules

What is an expected increase in traffic, and what is an not? Is it an attack or just a surge?

For video streaming, does the traffic level correspond to the viewer count (outbound) or repeater count (inbound)? If not, this may be a sign of an attack

False positives would be a very bad thing, so extra care must be taken here

Flow Analysis

Implementation - Traffic Level

```
service_types => {
  netif => {
    plugin => "extmon", timeout => 5,
    cmd => [
      "/usr/local/libexec/nagios/check_snmp_int.pl",
      "-H", "%%IPADDR%%", "-n", "em0", "-fkBM", "-w",
      "500,500",
      "-c", "640,640", "-d", "300", "-2", "--64bits", "-C",
      "public", "--label"
    ]
  }
}
```

Implementation - Business Rule

```
service_types => {  
  viewers_vs_bw => {  
    plugin => "extmon", timeout => 5,  
    cmd => [ "/usr/local/libexec/nagios/check_ddos",  
            "--viewers",  
            "http://%%IPADDR%%:8086/connectioncounts",  
            "--bandwidth", "-H %%IPADDR%%",  
            "--avgout", "2000"  
          ]  
  }  
}
```

Going Forward

- Implement Flow Analysis
 - Consider additional factors and business rules to more accurately identify attack conditions and attack sources
-

Podcasts

BSDNow.tv is a weekly video podcast featuring News, Interviews and Tutorials about the BSD family of Operating Systems. Hosted by Kris Moore (founder of PC-BSD) and Myself.

TechSNAP.tv is a weekly sysadmin video podcast covering an OS agnostic range of security and production issues of interest to those working , studying or interested in the field.
