

F5 Friday: So That DNS DDoS Thing Happened



Lori MacVittie, 2013-12-04

#infosec #dns #ddos Smurfs aren't just for ICMP anymore...



Note: This post was originally published in 2014 after a ginormous DDoS attack targetting DNS. Generally speaking, it's still relevant, though the details are not referencing the DYN attack from 10/2016.

DNS, like any public service, is vulnerable. Not in the sense that it *has* vulnerabilities but vulnerable in the sense that it must, by its nature and purpose, be publicly available. It can't hide behind access control lists or other traditional security mechanisms because the whole point of DNS is to provide a way to find your corporate presence in all its digital forms.

It should therefore not come as a surprise that eventually it turned up [in the news](#) as the primary player in a global and quite disruptive DDoS attack.

The gory details, most of which have already been circulated, are nonetheless fascinating given the low technological investment required. You can duplicate the effort with about 30 friends each with a 30Mbps connection (that means I'm out, sorry). As those who've been in the security realm for a while know, that's because these types of attacks require very little on the attack side; the desired effects come due to the unbalanced request-response ratio inherent in many protocols, DNS being one of them.

In the world of security taxonomies these are called "amplification" attacks. They aren't new; "[Smurf attacks](#)" (which exploited ICMP) were first seen in the 1990s and effected their disruption by taking advantage of broadcast addresses. DNS amplification works on the same premise, because queries are small but responses tend to be large. Both ICMP and DNS amplification attacks are effective because they both rely on protocols that do not require a handshake and are entirely uninterested in verifying whether or not the IP address in the request is the one from which the request was received. It's ripe for spoofing with much less work than a connection-oriented protocol such as TCP.

To understand just how unbalanced the request-response ratio was in this attack, consider that the request was: "dig ANY ripe.net @ <OpenDNSResolverIP> +edns0=0 +bufsize=4096". That's 36 bytes. The responses are typically 3K bytes, for an amplification factor of 100. There were 30,000 open DNS resolvers in the attack, each sending 2.5Mbps of traffic each, all directed at the target victim. CloudFlare has [a great blog on the attack](#), I recommend a read. Another good resource on DNS amplification attacks is [this white paper](#). Also fascinating is that this attack differed in that the target was sent a massive number of DNS responses - rather than queries - that it never solicited in the first place.

The problem is DNS is, well, public. Restricting responses could ostensibly unintentionally block legitimate client resolvers causing a kind of self-imposed denial of service. That's not acceptable. Transitioning to TCP to take advantage of handshaking and thus improve the ability to detect and shut down attempted attacks would certainly work, but at the price of performance. While F5's BIG-IP DNS solutions are optimized to avoid that penalty, most DNS infrastructure isn't and that means a general slowdown of a process that's already considered "too slow" by many users, particularly those trying to navigate the Internet via a mobile device.

So it seems we're stuck with UDP and with being attacked. But that doesn't mean we have to sit back and take it. There are ways in which you can protect against the impact of such an attack as well as others lurking in the shadows.

1. DEPRECATE REQUESTS (and CHECKING RESPONSES)

It is important to validate that the queries being sent by the clients are ones that the DNS servers are interested in answering, and are able to. A DNS firewall or other security product can be used to validate and only allow the DNS queries that the DNS server is configured for. When the DNS protocol was designed, there were a lot of features built into the protocol that are no longer valid due to the evolving nature of the Internet. This includes [many DNS query types, flags available and other settings](#). One would be surprised at what types of parameters are available to mark on a DNS request and how they can be manipulated. For example, DNS type 17=RP, which is the Responsible Person for that record. In addition, there are ways to disrupt DNS communications by putting bad data in many of these fields. A DNS firewall is able to inspect these DNS queries and drop the requests that do not conform to DNS standards and do not use parameters that the DNS servers are configured for.

But as this attack proved, it's not just queries you have to watch out for - it's also responses. F5 DNS firewall features include stateful inspection of responses which means any unsolicited DNS responses are immediately dropped. While that won't change the impact on bandwidth, it will keep the server from becoming overwhelmed by processing unnecessary responses.

[F5's DNS Services includes industry-leading DNS Firewall services](#)

2. ENSURE CAPACITY

DNS query capacity is critical to delivering a resilient available DNS infrastructure. Most organizations recognize this and put into place solutions to ensure high availability and scale of DNS. Often these solutions are simply caching DNS load balancing solutions which have their own set of risks, including being vulnerable to attack using random, jabberwocky host names. Caching DNS solutions only cache responses returned from authoritative sources and thus when presented with an unknown host name, it must query the origin server. Given a high enough volume of queries, the origin servers can still be overwhelmed, regardless of the capacity of the caching intermediary.

A high performance in-memory authoritative DNS server such as [F5 DNS Express \(part of F5 BIG-IP DNS\)](#) can shield origin servers from being overwhelmed.

3. PROTECT AGAINST HIJACKING

The vulnerability of DNS to hijacking and poisoning is still very real. In 2008, a researcher, [Evgeniy Polyakov](#), showed that it was possible to cache poison a DNS server that was patched and running current code within 10 hours. This is simply unacceptable in an Internet-driven world that relies, ultimately, on the validity and integrity of DNS. The best solution to this and other vulnerabilities which compromise the integrity of DNS information is [DNSSEC](#). DNSSEC was introduced to specifically correct the open and trusting nature of the protocol's original design. DNS queries and responses are signed using keys that validate that the DNS answer was not tampered with and that it came from a reliable DNS server.

[F5 BIG-IP DNS](#) not only supports DNSSEC, but does so without breaking global server load balancing techniques.

As a general rule, you should verify that you aren't accidentally running an open resolver. Consider the benefits of implementing DNS with an ICSA certified and hardened solution that does not function as an open resolver, period. And yes, F5 is a good choice for that.

Additional Resources:

- [High-Performance DNS Services in BIG-IP Version 11](#) 
- [DNSSEC: The Antidote to DNS Cache Poisoning and Other DNS Attacks](#)
- [F5 DNS Services Infrastructure](#) 
- [The Dynamic DNS Infrastructure | F5 White Paper](#) 



F5 Networks, Inc. | 401 Elliot Avenue West, Seattle, WA 98119 | 888-882-4447 | f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com