
The straight talk on DNSSEC

Peter Focas

October 2009

The DNS (Domain Name System) is a key building block of the Internet. The technology's most important task is translating host names (for example `www.google.com`) to IP (Internet Protocol) addresses (for example `74.125.45.100`).

DNS is used by almost all IP applications such as web browsing, email and file sharing. Applications will forward DNS requests to the DNS server they've been configured to access, typically the one provided by your ISP. The general flow of a DNS request is:

1. A host sends a request to its configured DNS server to enquire what the IP address is for a particular host name. To use web browsing as an example, a request for the IP address of `www.google.com`;
2. The DNS server queries the hierarchical list of Internet DNS Servers (the root DNS server, followed by the 'com' DNS server, followed by the 'google.com' DNS server followed lastly by the 'www.google.com' DNS server) until it finds the mapping for the host name requested and passes back the IP address; and
3. The host completely trusts the response and starts communicating with the received IP address.

What would surprise Internet users, if they knew to think about it, is that there are plenty of ways that the IP address returned by a DNS server can be subverted to point to an attacker's host rather than the genuine one requested. DNS queries and responses are sent 'in-the-clear' so anyone in the path of the traffic has the potential to be able to alter the data going between the DNS server and the requester.

A notable recent attack against DNS, outlined by Dan Kaminsky, involved sending fake packets to inject an attacker's IP address into a target DNS server effectively 'hi-jacking' a server's valid address without even having to intercept DNS queries and responses!

The effect of this type of attack is insidious because users are completely unaware that they are communicating with an unauthorized party. Any traffic including passwords, bank account and credit card details can be captured by the attacker who can either fake some service or relay the traffic on and simply record it on the way through.

In 2000 leading Internet Security company RSA had its DNS entry changed to point to an attacker's website, which looked like the real thing but had been defaced. Whilst there had been no compromise or breach of the RSA website or systems, it certainly looked that way to an outsider. This posed the open question - "Even though RSA's systems were completely secure did RSA suffer reputational damage by this attack on DNS?"

This exposure is the result of inherent weakness in the DNS protocol itself and one of the ways we can fight this unique type of attack is via an enhancement to DNS; *DNS Security Extensions (DNSSEC)*.

The goal of DNSSEC is to digitally sign the exchanges between a requestor and a DNS server so that responses cannot be subverted. This provides total assurance that the end-point IP address being communicated with is genuine.

So, if it's that foolproof, why aren't we all doing it? Simply put, there is a coordinated programme of work required – its not a simple fix. For example:

- Coordination is required by many parties to ensure full security;
- Domain name registries, ISPs and users may be required to upgrade their software and must establish processes to digitally 'sign' their DNS zone information and exchange cryptographic keys;
- All DNS servers in the hierarchy need DNSSEC or it is not completely secure; and
- It's hard to make business cases for security related technologies stack up in the face of competing revenue/profit generating initiatives.

All but the final point can be managed, but developing the business case is where the rubber hits the road. Can you articulate the reputational damage if your Internet enabled services are hijacked? Does your organization take its reputation seriously?

So, what can you do?

A client computer without DNSSEC capability will still work OK with a DNSSEC enabled DNS server but it will not be able to verify a returned address. Organisations cannot normally dictate what client systems will access their publicly available web servers so mandating that clients will have DNSSEC capability will not be practical. Therefore, your planning should be focused towards doing what you can in environments that you control:

- Secure your own internal DNS servers with up-to-date patch levels, access controls and anti-spoofing controls (packets using fake IP addresses) to prevent your DNS server from being attacked and possibly subverted;

-
- Deploy DNSSEC in your DNS zone and request that your DNS registrar has a roadmap to support DNSSEC;
 - Make sure that all internal hosts and servers have DNSSEC capability; and
 - Monitor your domains both inside and outside your network and make sure that you have a security policy in place specifying processes to follow to cope with DNS hijacking if it occurs.

But don't simply put this in the "too hard" basket. DNS attacks occur and losses are suffered virtually every day here in NZ. They go largely unreported, but the reality is you could be next.

If you found this interesting and would like to chat about it, drop me a line at peter.focas@voco.co.nz