GUEST POST: WE NEED TO TALK ABOUT DNS

[*NB: This is a guest post by long-time community member WilliamOckham. Give him a shout in comments. /~Rayne*]

For most people the Domain Name System (DNS) is one of the most boring topics imaginable. However the Department of Justice's Special counsel John Durham – through a frothy mixture of technical incompetence and apparent malice in his published court filings – generated unusual interest in DNS from a lot of folks who've never thought about it before.

To understand DNS better, here's an explanation simple enough even for lawyers who would like to keep their bosses from embarrassing them in federal court.

DNS is used to match and link domain names to Internet Protocol (IP) addresses. When one device needs to connect to another device via the internet, it needs to know the other's IP address. Humans generally prefer to use names. Remembering a person's or business's name is much easier than recalling a string of numbers ranging from 12 to 32 digits (32-bits for older IPV4 addresses and 128-bits for newer IPV6 addresses).

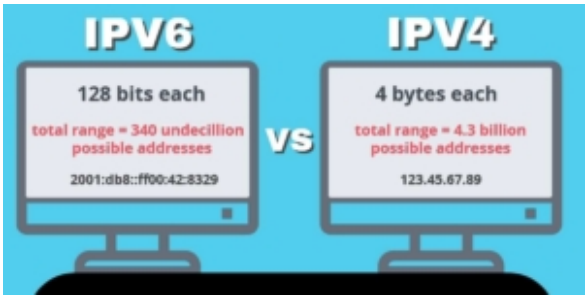

Image: Comparitech.com c. 2019

I'll use "example.com" to illustrate a domain name. As you might guess, example.com is a special-use domain which isn't resolved normally; it can be used to demonstrate how domain names work without inadvertently

generating unnecessary DNS lookups.

It’s a lot easier to input www.example.com instead of 2606:2800:220:1:248:1893:25c8:1946 and certainly a lot easier to remember. However your device can’t possibly store the IP address of every damn server in the entire world just to make data entry easier.

Instead, every device on the internet stores the address of one of the thousands of DNS servers. Devices are usually configured to use a DNS server maintained by the internet service provider which provides connectivity for that device.

When your device needs to connect to www.example.com, it sends a DNS lookup request to its primary DNS server. That server doesn’t store the address of every server on the internet either. If you or someone else using that DNS server has asked for that address recently, the DNS server might know the address and send it back to you.

However if it doesn’t have an IP address for example.com, it will issue requests to other DNS servers, looking for one that does know the address. In a worst case scenario, the request ends up going to one of the root DNS servers. They can reach a DNS server for any domain name on the internet.

During the time period subject to Durham’s investigation, **virtually all DNS lookups happened in the open, unencrypted**. They were recorded by DNS servers. Each time a website address was typed into a browser’s address bar, a DNS server logged the IP address of the device requesting the IP address for some other server. DNS lookup data isn’t proprietary or secret.

Gathering, collating, and analyzing DNS lookup requests, however, is expensive and valuable. It’s a massive amount of data. Billions of DNS requests are issued every day. There are a few companies specializing in managing incredibly large amounts of DNS data. During the time period covered by Durham’s filings, Michael

Sussman’s technology executive client (Tech Executive-1) at a U.S.-based Internet company (Internet Company1) worked for such a firm.

Having access to DNS data had nothing to do with hacking servers, spying, surveillance or anything else nefarious. It was part of Tech Executive-1’s job.

Tech Executive-1’s responsibilities included monitoring anomalies in Internet Company1’s DNS database. As one of Durham’s filings indicated, Tech Executive-1’s firm found “that between approximately 2014 and 2017, there were a total of more than 3 million lookups of Russian Phone-Provider-1 IP addresses that originated with U.S.-based IP addresses.”

Contra Durham, 3 million DNS requests for a related IP addresses over a four-year period means these requests are very rare.

For comparison purposes, my best estimate is that my family (7 users, 14 devices) generated roughly 2.9 million DNS requests just from checking our email during the same time frame. That’s not even counting DNS requests for normal web browsing.

If you’re going to make a federal case out of this, at least make some attempt to understand the topic.