Network Working Group
Internet-Draft
Intended status: Informational
Expires: May 10, 2013

M. Richardson
SSW

November 6, 2012

Secret Gardens are Better than Walled Gardens
draft-richardson-homenet-secret-gardens-00

Abstract

This document explains a few use cases where operators would like to introduce so-called "walled gardens" into home-networks, including distribution of new DNS anchors. This document proposes an alternative solution involving DNS delegations to access controlled DNS servers. The results are much more scalable, and can be deployed today, using existing operating systems and existing DNS infrastructure.

Table of Contents

1. Introduction: A brief history of split-horizon DNS

DNS settings have been a regular concern since the very early 1990s, when the first firewalls began to partition the Internet. In the earliest works, Cheswick and Bellovin describe the use of split-DNS in an enterprise: all internal machines (including the inside interface of the firewall, and possibly machines on the DMZ/ Service-Network) would use an internal recursive DNS server for names, and if the name was external, the internal recursive DNS server would ask the world.

The above split-horizon configuration survives to today, but it has become very complicated. The first complication was remote access (VPN) to the enterprise. For the computer at home to be able access things, the internal DNS server had to be used. Often short internal names ("smtp", "web", "wiki","printer") would be used, depending upon the fact that all internal machines had the same search path. As the VPN could go on, and off, if it was off, and the end user entered the word "wiki" on the browser, instead of going to the internal resource (which in IPv4 space, had an RFC1918 address), it would either go to an external resource, or cause a search.

Worse, some computer systems originally needed to be rebooted in order to change their DNS settings, as this was really the only way to convince all application to flush name to IP address mappings that they had cached.

Particularly gruesome is the case of the contractor or consulting, who works at enterprise A, and then visits enterprise B. While on the network of Enterprise B (where they may be located for some months), in order to do simple things like reach the printer (using the name "printer"), they need to use the DNS settings for Enterprise B. But, in order to fetch their email, they must have the DNS settings (and VPN) for their home base, Enterprise A. There have been regular reports in the VPN/Remote Access community of situations where a worker needs to have VPNs up with two remote locations, while residing at a third.

The VPN situation is tragic for the technical user, for the non-technical user, it is impossible. For the technical user, typing in longer names, setting up multiple search paths, and running a local recursive name server are possible. For the less technical user, typing in IP addresses rather than names works as long as HTTP Virtual Hosting is not involved (for which the name is important). But the degree to which these things occur has been limited in part due to the inevitable conflict of RFC1918 addresses, which means, that, even if one uses IP addresses, the routing doesn't work.

DNSSEC has put a new twist into things: the best way to run DNSSEC is to have a secure recursive resolver. Now, this recursive resolver needs to be taught about split-horizon names, to talk to the internal name server when it is reachable (and probably to trust it), and to talk directly to the Internet when not reachable. A second problem is that the local recursive name server may not always get flushed. If a query for "smtp.example.com" should result in the internal mail relay when at the office, or connected via VPN, but the same name should resolve in the external address when not connected (possibly, the same machine with multiple interfaces, but not always), then the caching presents a problem. This problem is acute if one tries to run a caching name server on a mobile, multiple-interface devices, such as a stock smartphone that has 3G and wifi interfaces, which may operate concurrently.

Some organizations, where there are multiple data centres, and the network is distinctly non-convex, have solved the caching problem by dispensing with a true split-horizon DNS, and have simply put all external names under "example.com", with internal names under "internal.example.com". Whether or not externally made requests for foo.internal.example.com are resolved by externally facing authoritative name servers is now a security policy question, rather than a routing or caching concern.

There have been further abuses of split-horizon DNS. One of these is by hotels and other places with "captive portals". In order to authenticate the user using a web interface, they must make the portal visible to the user. This involves making the user "captive": no packets may leave the local network until the user has sufficiently authenticated (possibly, involving a financial transaction). Some captive portals have decided that they will intercept all DNS requests, and no matter what the user asks for, they will answer with their own name. This can fail for three reasons: 1) the user does not use the DNS servers provided, instead uses their own (perhaps intending to be reached through a VPN) or or the root name servers directly (perhaps in order to do DNSSEC), 2) the names used in the web interfaces are unqualified ones, and the user has not accepted the search path from the DHCP offer (ref: some BCP on this), 3) even if all of this goes well, the user's browser has now cached an incorrect mapping for the desired web site.

1.1. Requirements Language

(RFC2119 reference)

2. Use Cases

This section lists a few use cases which homenet believes are within scope. Only the business goals are listed.

2.1. IPv6 TV

A seperate internet connection is provided by the IP TV provider. This can be on a seperate physical wire, or a seperate logical wire. Historically, only the TV set itself is placed on this network, the home owner may be unaware that the wires are other than analog cable TV wires.

The TV set does DNS lookups for the content servers, and is expected to receive answers that are within the IP TV provider's network. The TV set can access the provider's servers as long as it originates connections from the IPv6 prefix that the IP TV provider provided.

The provider does not publish the addresses of the servers publically, and does not want to. Even if the servers can be accessed from the Internet by IP address, they will not serve any TV content that way. It is more likely that the content servers are completely inaccessible from the Internet via firewall rules and/or routing horizons.

2.2. Corporate VPN

The corporate VPN use case involves an employee of an enterprise connecting to the enterprise's network via an IP over IP tunnel. IPsec is a typical technology, and there are various kinds of SSL VPNs, but this use case concerns itself with scenarios where IPv6 pakets will be routed through the tunnel. (Some SSL VPNs provide TCP, usually HTTP-only services. They are not relevant)

As described in the introduction, the Corporate VPN is almost the original walled garden: there is a seperate routing domain (RFC1918 in IPv4, End-User Assigned IPv6, Non-Connected Network or ULA for IPv6), and there is an addition a desire to have an island of DNS.

Some corporate VPNs have a policy that all traffic from the employees' computing device must go through the tunnel, even traffic not addressed to the Enterprise. (That is, the default route on the laptop must point through the tunnel. A host route for the VPN tunnel end point is usually created to permit the encapsulated traffic to travel, or another kind of source-address sensitive policy route is used to create multiple routing tables). These types of VPNs are sometimes called an extruded IP VPN: if the enterprise' network is thought of as a fungible sponge, then the laptop user at

home as an IP address, from the corporate cloud/sponge, extruded, rather like a sea urchin (??) extends a pseudopod.

It is not unusual for the networks of enterprises to grow to be very complex: multiple sites with multiple (sometimes overlapping IPv4 RFC1918) address spaces. With IPv6, enterprises are likely to have fewer blocks of addresses, and none will overlap, but due to acquisitions the number of blocks will still be greater than one.

Frequent teleworkers often have an entire home office connected to the enterprise network. This can include a selection of desktop computers, laptops, tablets, and smartphones (on wifi). The home office network might have several subnets (at least one wired and one wireless). These teleworkers usually have a VPN router as their home gateway, and often these devices are controlled by the Enterprise IT. Just the same, access to the rest of the homenet is desired such that the home worker(s) can share things like printers, talk to home automation from their desk, and to avoid having large-bandwidth cross the corporate network unnecessarily.

More about DNS expectations of enterprises. ActiveDirectory is relevant here.

2.3. Multi-homed to 3G/LTE

There is a problem here: so far it sounds like it's just IP TV over again.

3. Walled Garden DNS

There have been proposals for each ISP/connectivity provider to provide a suffix, such as "mytv.example." or even unregistered names like ".internal" (common with ActiveDirectory) into the homenet.

Along with suffix would be one or more DNS servers, accessible via the semi-private connection, which would resolve names in that zone.

The proposal is for the suffix/server extensions/appendages/ ?-needs-name-? to be made visible up to near the application. At least, if there is a local recursive name server, then it could be the one point which gets updated (this is what some VPN clients do), or it may require applications to be aware, and to use alternative resolver libraries.

4. Using Secret-Gardens to accomplish goals

5. Security Considerations

6. Other Related Protocols

7. IANA Considerations

8. Acknowledgements

9. References

9.1. Informative References

9.2. Normative References

10. Normative references

Author's Address

Michael C. Richardson
Sandelman Software Works
470 Dawson Avenue
Ottawa, ON K1Z 5V7
CA

Email: mcr+ietf@sandelman.ca
URI: http://www.sandelman.ca/