

DNS Deeper Dive: Loosely Coherent... and Loosely Defined...

Geoff Huston <gih@apnic.net>
Wes Hardaker <hardaker@isi.edu>
João Damas <joao@apnic.net>

Note Well

This is a reminder of IETF policies in effect on various topics such as patents or code of conduct. It is only meant to point you in the right direction. Exceptions may apply. The IETF's patent policy and the definition of an IETF "contribution" and "participation" are set forth in BCP 79; please read it carefully.

As a reminder:

- By participating in the IETF, you agree to follow IETF processes and policies.
- If you are aware that any IETF contribution is covered by patents or patent applications that are owned or controlled by you or your sponsor, you must disclose that fact, or not participate in the discussion.
- As a participant in or attendee to any IETF activity you acknowledge that written, audio, video, and photographic records of meetings may be made public.
- Personal information that you provide to IETF will be handled in accordance with the IETF Privacy Statement.
- As a participant or attendee, you agree to work respectfully with other participants; please contact the ombudsteam (<https://www.ietf.org/contact/ombudsteam/>) if you have questions or concerns about this.

Definitive information is in the documents listed below and other IETF BCPs. For advice, please talk to WG chairs or ADs:

- [BCP 9](#) (Internet Standards Process)
- [BCP 25](#) (Working Group processes)
- [BCP 25](#) (Anti-Harassment Procedures)
- [BCP 54](#) (Code of Conduct)
- [BCP 78](#) (Copyright)
- [BCP 79](#) (Patents, Participation)
- <https://www.ietf.org/privacy-policy/> (Privacy Policy)

What's in a Name?

In a network that uses destination-based packet forwarding, a “name” is a human-use alias for an endpoint destination

- An endpoint destination is an IP address used by the network to relay a packet onto its intended destination
- Multiple services on a host are implicitly disambiguated by using Transport Protocol Port addresses – the name itself is the name of the platform, not the name of a particular service
- Names were useful to the client (not necessarily to the host named by the client)

What's in a Name?

Human-use names were intended to be constant across protocol families for a single host

- Host Name conventions became a Lowest Common Denominator issue of matching various protocol-specific and application-specific limitations
- Hence the subtle distinctions between host name and DNS labels
 - Which João will explore in detail later on...

DNS Resource Records

Original name model was based on HOSTS.TXT model

- Simple Name to IP address mapping
- Match the name to 1 or more entries in the file
- Return one of more IP addresses listed in these matched entries
- The “resource” here is the IP address in the hosts file

This is a client-centric name alias model

- Different hosts could use different HOSTS.TXT files without fundamental damage

The DNS shifted this replicated file into a distributed database

- This created a publisher-centric model of names and resources

A structured name space and a distributed database created a requirement for additional “meta” records

- SOA records to control common zone behaviours
- NS records to control navigation through the distributed database to resolve names

Service Names

Then we added service names to host names


It became common practice to use a label prefix to denote the service:

- www.example.com for port 80 HTTP service
- ftp.example.com for port 21 ftp service

This enabled a set of hosts to be used to support the set of services located behind a common host name in a manner that was specific to the service, not the host

It was a small step to then include the service name into the application-level protocol exchange so that a named host could provide services to multiple named services:

- Secondary MX server records for a named mail destination domain

```
foo.example.com MX 5 bar.example.com  "Hi 192.0.2.0, I want to send mail to user@foo.example.com"  
bar.example.com IN A 192.0.2.0
```

Service Names

Service Names become an abstraction **used by application protocols** to allow orchestration of multiple servers to provide a named service to clients

- These points of service delivery may not be visible to users, and need not necessarily be names within a domain of human-friendly use
- Some service names are loaded into the DNS itself through resource records
 - Such as NS and MX RRs
- Others are established through explicit alias records
 - CNAME RRs
- Others are implicit
 - Using the A/AAAA records of the service delivery host against the service name

What's a Name?

- A long-held universally visible stable service point identifier suitable for human use
- A transient label used by applications to resolve temporal rendezvous requirements
- A bridging point to glue together the DNS database
- A fragment of code (query) sent to a code interpreter (server)

Or: All the above and more!

Service Identification

- There are now multiple ways in the DNS to associate a service provider with a service
- Each DNS rendezvous mechanism we invent proposes a generic mechanism that can be used by any application
- But each application appears to use a unique specific DNS solution!
 - Place information in the value part of the key/value store
 - Map the query name to a different query name
 - Modify the Query Type to have application specific data

TXT Records

The universal DNS “one size fits all” DNS resource record

- Treats the DNS as a key-value store
- TXT records can be internally structured according to the requirements of the application
- Of course if the same name is used for multiple services the overloading of the TXT record can become a problem

SPF Records

- Sender Policy Framework for Mail
 - Validation of service identity as an authorized agent for the application level service identity
- Implemented as a TXT record with internal structure
 - See previous comment on overloading of TXT record!
- Can use more domain names on the RHS as an include: directive
 - Which allows the formation of extended chains and loops

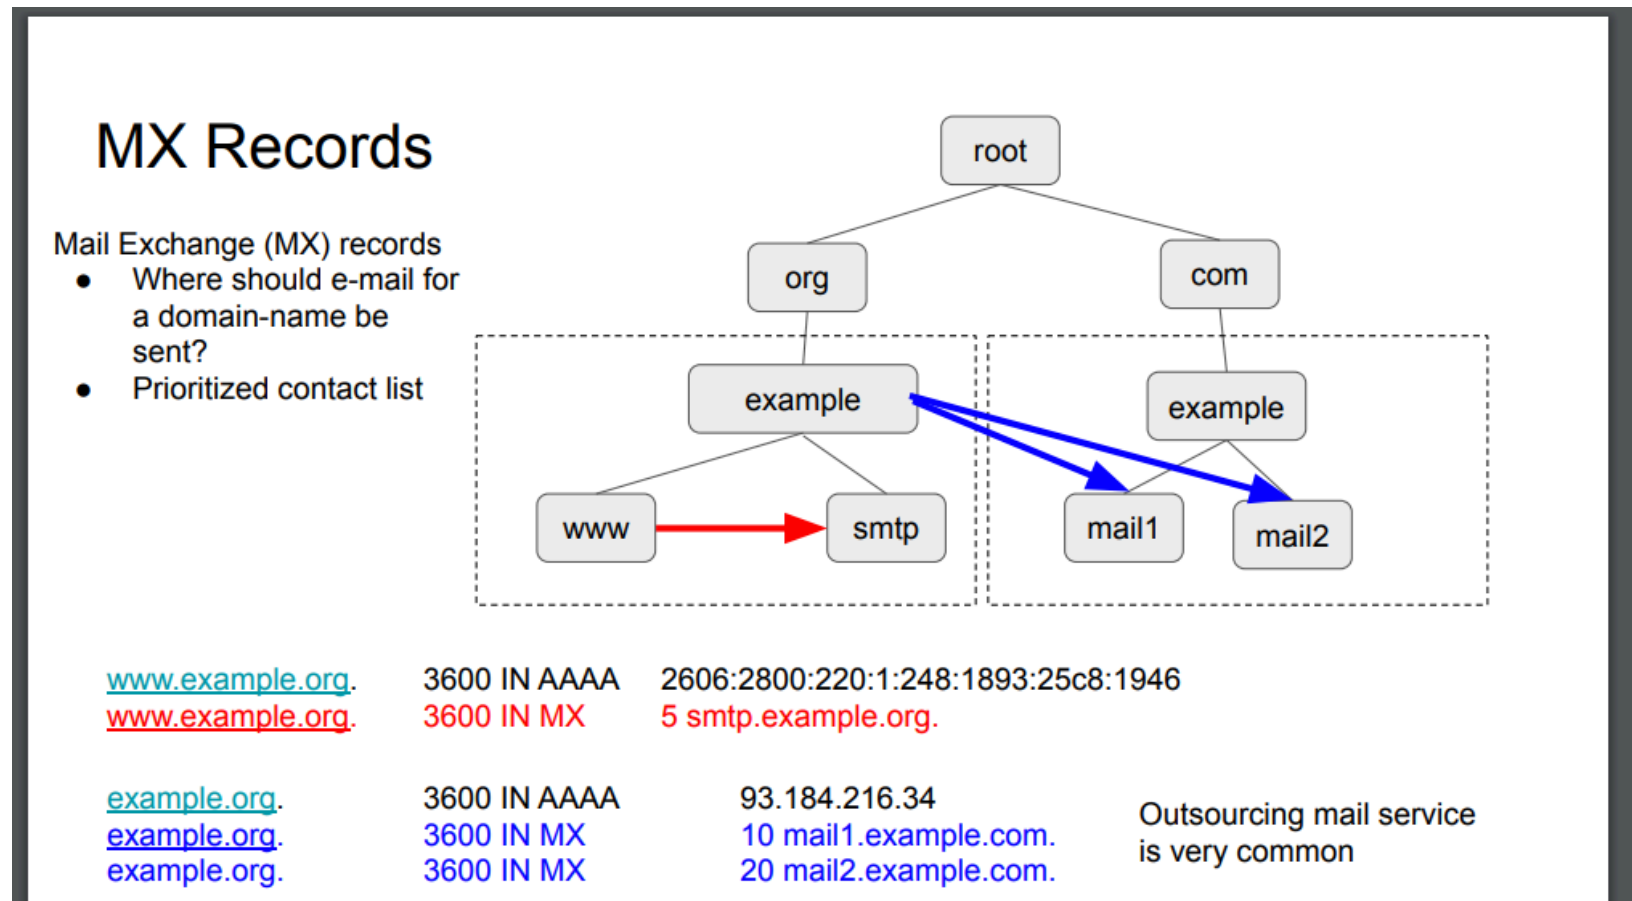
```
example.com TXT "v=spf1 ip4:34.243.61.237 ip6:2a05:d018:e3:8c00:bb71:dea8:8b83:851e include:thirdpartydomain.com -all"
```

Potential loop!



MX Records

- Application specific redirection in the name space



CNAME and DNAME Records

- Alias name form to create a generic service redirection through query name mapping as part of the name resolution process

CNAMEs and DNAMEs

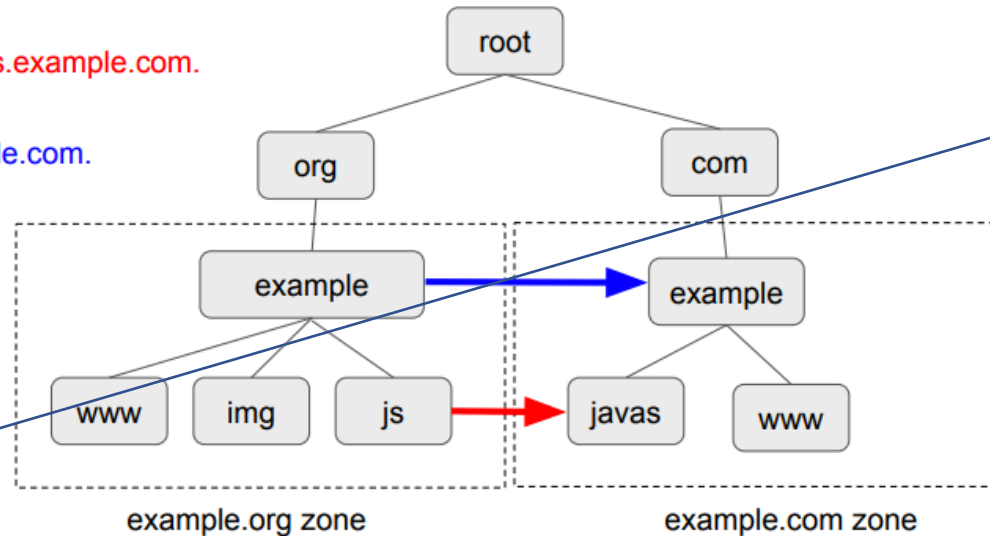
`js.example.org. 3600 IN CNAME javas.example.com.`

CNAMEs cannot occur at the apex

`example.org. 3600 IN DNAME example.com.`

CNAMEs are aliases for other tree elements (can be in the same zone or in another)

DNAMEs are aliases for zones themselves



Zones? NO - Subtrees!

Anything below a DNAME is not reachable!
There is NOTHING below a CNAME!

IMPORTANT: CNAMEs MUST exist alone at a name (minus DNSSEC entries)
IMPORTANT: CNAMEs point to ALL records at the other name (A, AAAA, NS, MX, etc)

CNAME and DNAME loops

The use of a DNS name in the resource record can form a DNS loop if the name needs to be resolved itself using the same Qtype query

- There are two issues with such loops: cache stuffing of intermediate resolution outcomes and resolver time when following a loop

```
dog.com CNAME cat.com  
cat.com CNAME dog.com
```

```
dog.com DNAME cat.com  
cat.com DNAME dog.com
```

```
dog.com CNAME cat.com  
cat.com DNAME dog.com
```

NAPTR

- What happens if we take a TXT value and interpret it as a regular expression to apply to the query name, and add additional fields to specify order and preference among multiple such NAPTR records for a single label?
 - Then we get a NAPTR record!

```
$ORIGIN 2.1.2.1.5.5.5.0.7.7.1.e164.arpa.  
IN NAPTR 100 10 "u" "sip+E2U" "!^.*$!sip:information@tele2.se!" .  
IN NAPTR 102 10 "u" "mailto+E2U" "!^.*$!mailto:information@tele2.se!" .
```

Black Holes and Loops

- Resolver Forwarder loops
 - A set of *forwarder* clauses in resolver configurations may cause queries to loop across the resolvers in the loop
- DNS Data loops and holes
 - NS
 - Normally a parent zone NS record will contain glue records for the IP address of the listed name server
 - If not then the resolver has to pause the primary resolution task and resolve the name of the name server
 - Black holes: NS<x>.example.com IN NS NS<x+1>.example.com
 - Loops: ns.dog.com IN NS ns.cat.com
ns.cat.com IN NS ns.dog.com

This will consume resolver and server resources if not handled sensibly
 - CNAME, DNAME and NAPTR loops
 - Any DNS structure that includes a name that needs to be applied recursively into the DNS has the potential to create loops

SRV

- Generic form of MX record
- Identifies the service by name and the transport protocol / port address
- SRV maps the service name to one or more service hosts through the use of service prefix labels to the DNS service name
- The RHS host name is the canonical name

```
_sip._tcp.example.com. 3600 IN SRV 10 60 5060 sipserver.example.com.  
_sip._tcp.example.com. 3600 IN SRV 10 20 5060 someone.elses.sip.server.org.
```

SVCB

- When SRV is not enough
 - ALTSVC, ANAME, ESNIKEYS, etc all assembled into a single DNS entry
- An “Alias Form” where operational control of a resource is designated elsewhere
- A “Service Form” to define the parameters of the operational service

```
_8443._foo.api.example.com. IN SVCB 0 svc4.example.net.  
svc4.example.net.          IN SVCB 3 ( svc4.example.net. alpn="bar" port="8004" esnikeys="..." )
```

foo://api.example.com:8443” is aliased to use ALPN protocol “bar” service endpoints offered at “svc4.example.net” on port 8004.

HTTPSSVC

- SVCB for the HTTPS scheme
- Equivalent of the HTTP Alt-Svc response header – shifted into the DNS

Alt-Svc: h3="svc.example.net:8003"; ma=3600; foo=123, h2="svc.example.net:8002"; ma=3600; foo=123



www.example.com. 3600 IN HTTPSSVC 2 svc.example.net. (alpn=h3 port=8003 foo=123)
HTTPSSVC 3 svc.example.net. (alpn=h2 port=8002 foo=123)

DNS and Information Leakage

- Historically, the DNS is not the most secure of protocols
 - And while effort is underway to improve this situation in some ways, the DNS is still extremely chatty – promiscuously so!
- To what extent do these service records exacerbate these issues by “informing” the DNS of the end user’s immediate intentions of subsequent service-level activity?
- Do we need to undertake more analysis of the issue of the trade offs between adding additional rendezvous attributes to the DNS and the extent to which the DNS leaks even more information as to a user’s immediate intentions to an eavesdropper?

And now...

- Over to Wes...