More Efficient DNS with resolved from systemd

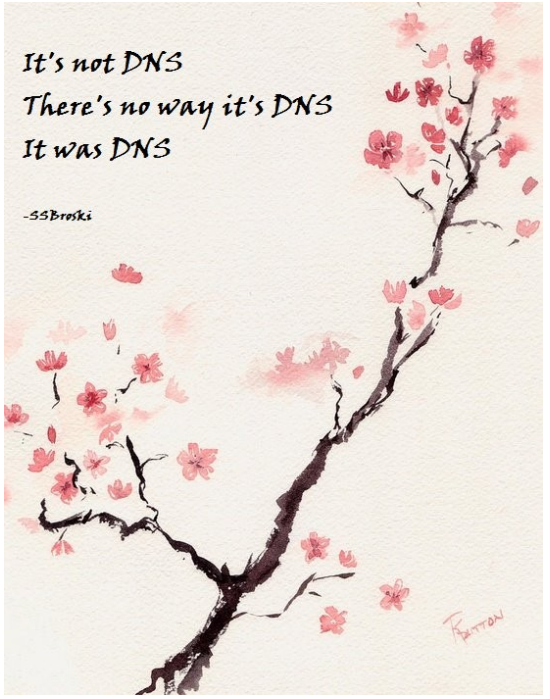

Patrick Ladd
pmladd@gmail.com
people.redhat.com/pladd/

systemd utilities

- systemd utility functions
  - journald
  - logind
  - resolved
  - timesyncd
  - networkd
  - tmpfiles
  - timedated
  - udevd
  - libudev
  - systemd-boot
- General principles
  - Handle common & expected services for modern systems
  - Ease of use for majority use case
  - Ok with:
    - removing functions for simplicity & security
    - removing little used features
    - implementing in more modern fashion
    - extending function in service of simplicity & security

DNS resolution flow - original

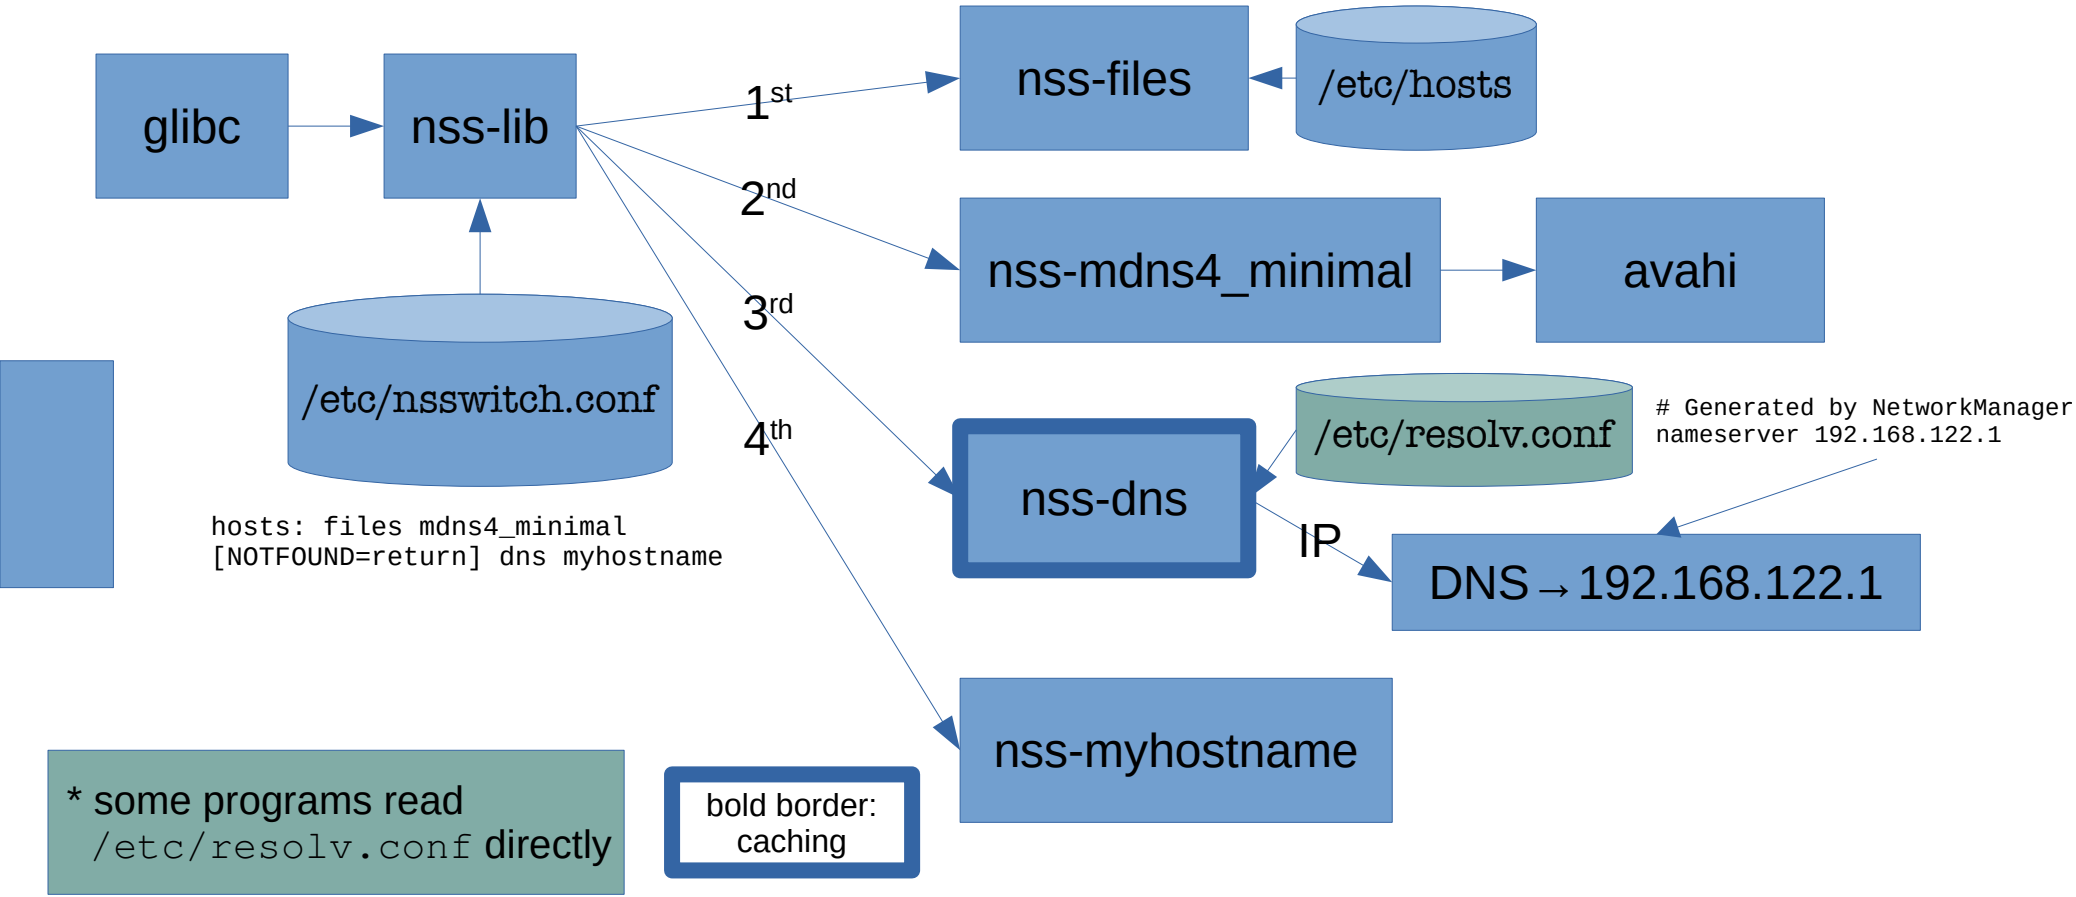

DNS resolution flow – with resolved

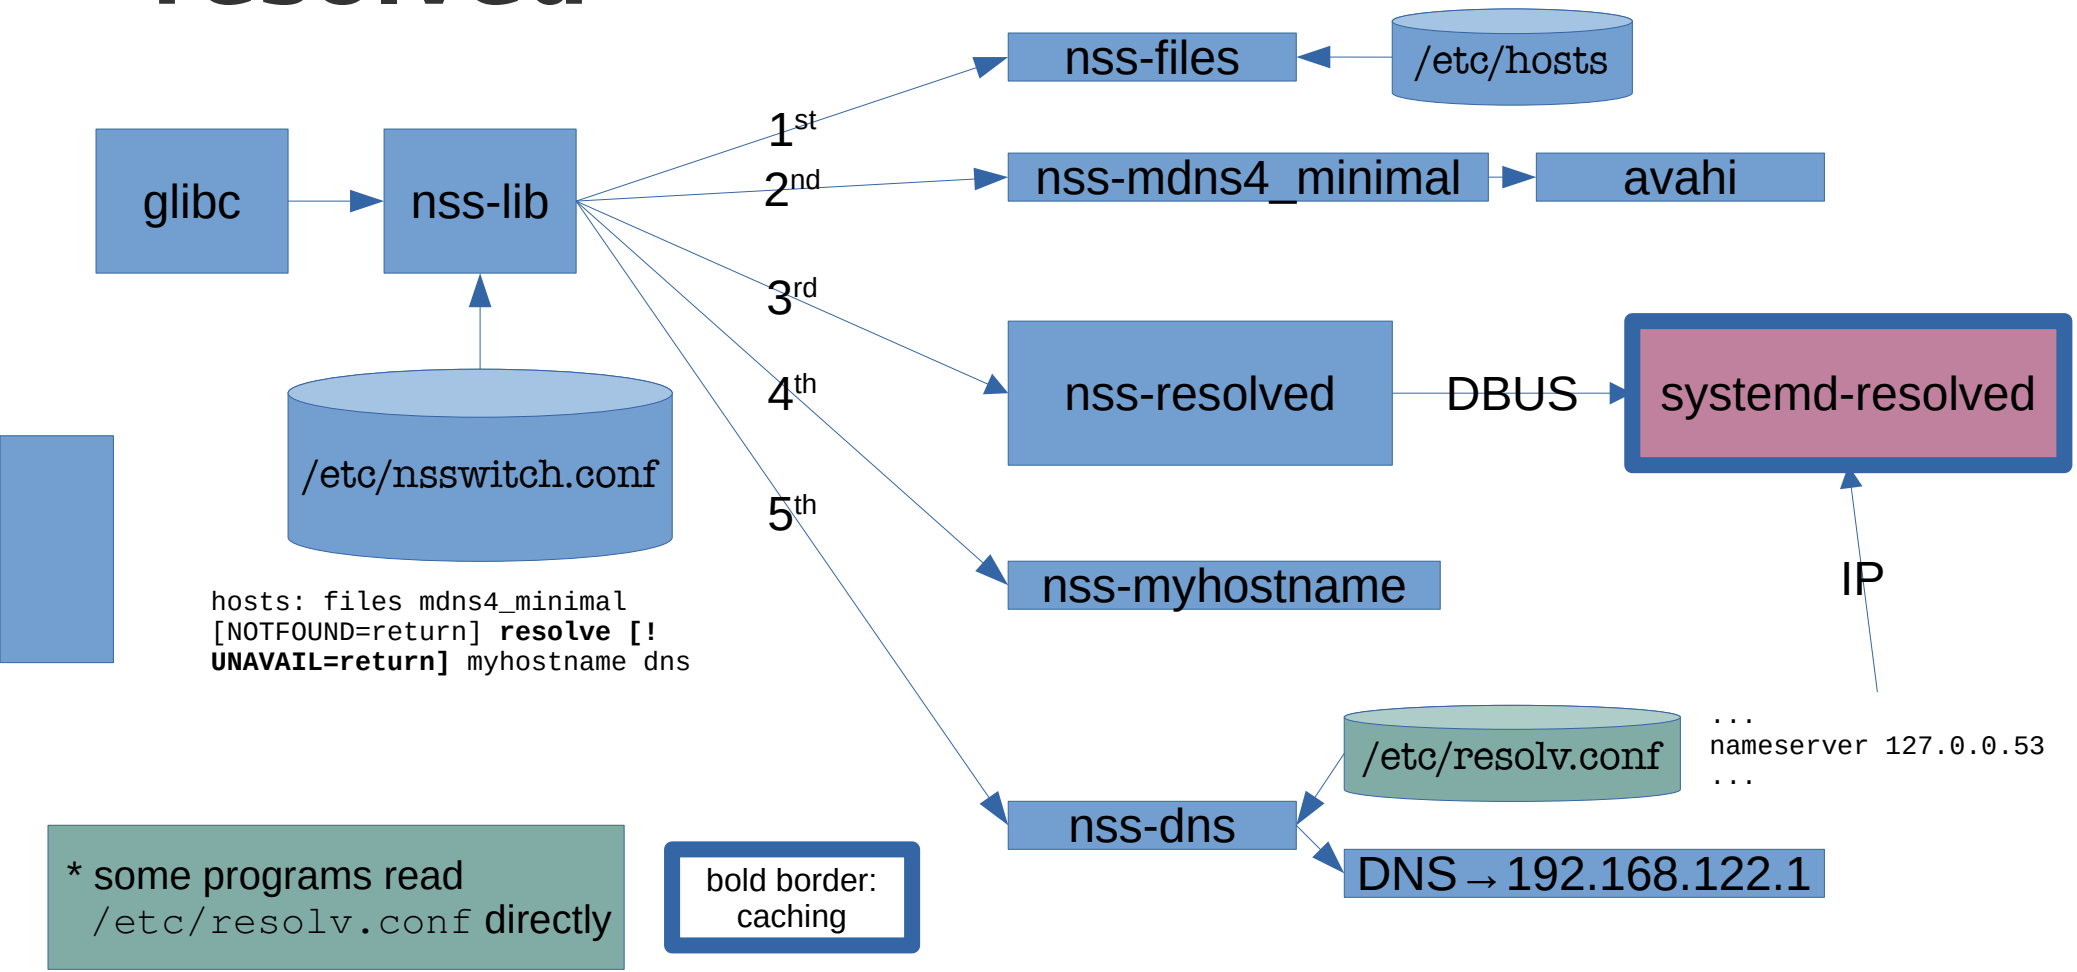

systemd-resolved – replacement for traditional local DNS

- Drop in replacement – majority use case systems should function as they did before
- Implemented differently to improve efficiency
  - **Stub resolver:**
    - Doesn’t resolve queries directly from roots
  - **Cross-process Caching**
    - Can result in far fewer outbound queries and faster response time
  - **Split DNS**: Interface specific resolver options
    - Limit leakage of DNS queries
    - Reduced server load from spurious queries
    - Order of VPN connection unimportant
  - **Modern Features:** DNSSEC, DNS over TLS

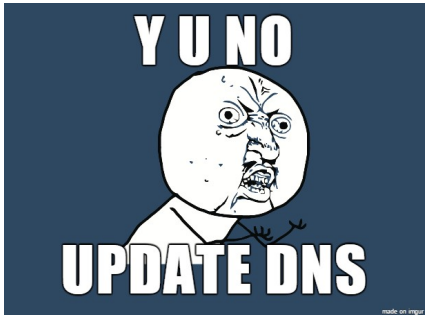

Resolved - Distributions

- Default changed in Fedora 33 (hence today's talk)
- Default in Ubuntu since 16.10 (2016)
  - Doesn't use nss-resolve by default (`/etc/nsswitch.conf` unchanged)
  - Instead of direct connection, IP connection over 127.0.0.53

Resolved – Domains & Split DNS

- **Domains**
  - **Search domains**
    - **Added to search path for “single label” queries**
  - **Routing Domains**
    - **Prefixed with ~**
    - **Not added to search list**

```
# resolvectl domain
Global:
Link 2 (enp0s31f6):
Link 3 (wlp4s0):
Link 4 (virbr0): libvirt
Link 5 (virbr0-nic):
Link 6 (anbox0):
Link 7 (virbr1):
Link 8 (virbr1-nic):
Link 12 (wg0):
Link 15 (cni-podman0):
Link 16 (vetha90fc4ad):
Link 18 (tun0): redhat.com ~amazonaws.com
Link 20 (vnet0):
```

Resolved – Search Behavior Difference

- nss-dns
  - Single label names: bare search, then with each search domain
  - Multi-label names: bare search, then with each search domain
- nss-resolved
  - Single label names: same
  - Multi-label names: ONLY bare search, no search domains appended

resolvectl – interact with resolved

- Resolver / reverse resolver for IPv4 & IPv6

```
# resolvectl query yahoo.com
yahoo.com: 2001:4998:44:3507::8001             -- link: wlp4s0
           2001:4998:44:3507::8000             -- link: wlp4s0
           2001:4998:124:1507::f000            -- link: wlp4s0
           2001:4998:24:120d::1:1              -- link: wlp4s0
           2001:4998:24:120d::1:0              -- link: wlp4s0
           2001:4998:124:1507::f001            -- link: wlp4s0
           74.6.143.25                         -- link: wlp4s0
           98.137.11.163                       -- link: wlp4s0
           74.6.231.20                         -- link: wlp4s0
           98.137.11.164                       -- link: wlp4s0
           74.6.143.26                         -- link: wlp4s0
           74.6.231.21                         -- link: wlp4s0

-- Information acquired via protocol DNS in 797us.
-- Data is authenticated: no

# resolvectl query 98.137.11.163
98.137.11.163: media-router-fp74.prod.media.vip.gq1.yahoo.com -- link: wlp4s0

-- Information acquired via protocol DNS in 14.2ms.
-- Data is authenticated: no
```

resolvectl – interact with resolved

- Status

```
# resolvectl
Global
Protocols: LLMNR=resolve -mDNS -DNSOverTLS DNSSEC=no/unsupported
resolv.conf mode: stub

Link 2 (enp0s31f6)
Current Scopes: none
Protocols: -DefaultRoute +LLMNR -mDNS -DNSOverTLS DNSSEC=no/unsupported

Link 3 (wlp4s0)
Current Scopes: DNS LLMNR/IPv4 LLMNR/IPv6
Protocols: +DefaultRoute +LLMNR -mDNS -DNSOverTLS DNSSEC=no/unsupported
Current DNS Server: 192.168.1.1
DNS Servers: 192.168.1.1
DNS Domain: fios-router.home

Link 4 (virbr0)
Current Scopes: DNS LLMNR/IPv4
Protocols: -DefaultRoute +LLMNR -mDNS -DNSOverTLS DNSSEC=no/unsupported
Current DNS Server: 192.168.122.1
DNS Servers: 192.168.122.1
DNS Domain: libvirt

…

Link 18 (tun0)
Current Scopes: DNS LLMNR/IPv4 LLMNR/IPv6
Protocols: -DefaultRoute +LLMNR -mDNS -DNSOverTLS DNSSEC=no/unsupported
Current DNS Server: 10.5.30.160
DNS Servers: 10.5.30.160 10.11.5.19
DNS Domain: redhat.com ~amazonaws.com
```

resolved – Statistics

- Statistics

```
# resolvectl statistics
DNSSEC supported by current servers: no

Transactions
Current Transactions: 0
Total Transactions: 438283

Cache
Current Cache Size: 7
Cache Hits: 534
Cache Misses: 15835

DNSSEC Verdicts
Secure: 41
Insecure: 231
Bogus: 0
Indeterminate: 0
```

resolvectl – interact with resolved

- Inspect / reconfig

```
resolvect dns [LINK [SERVER...]], domain [LINK [DOMAIN...]], default-route [LINK
[BOOL...]], llmnr [LINK [MODE]], mdns [LINK [MODE]], dnssec [LINK [MODE]], dnsovertls
[LINK [MODE]], nta [LINK [DOMAIN…]]
```

```
# resolvectl status wlp4s0
Link 3 (wlp4s0)
Current Scopes: DNS LLMNR/IPv4 LLMNR/IPv6
Protocols: +DefaultRoute +LLMNR -mDNS -DNSOverTLS DNSSEC=no/unsupported
Current DNS Server: 192.168.1.1
DNS Servers: 192.168.1.1
DNS Domain: fios-router.home

# resolvectl dns wlp4s0 8.8.8.8 8.8.4.4 1.1.1.1
# resolvectl dnssec wlp4s0 yes
# resolvectl dnsovertls wlp4s0 yes

# resolvectl status wlp4s0
Link 3 (wlp4s0)
Current Scopes: DNS LLMNR/IPv4 LLMNR/IPv6
Protocols: +DefaultRoute +LLMNR -mDNS +DNSOverTLS DNSSEC=yes/unsupported
Current DNS Server: 8.8.8.8
DNS Servers: 8.8.8.8 8.8.4.4 1.1.1.1

# resolvectl revert wlp4s0
# resolvectl dns wlp4s0 192.168.1.1
# resolvectl domain wlp4s0 fios-router.home
```

resolved – Logging

- Logs via journal

```
# resolvectl log-level
info

# resolvectl log-level debug

# journalctl -fu systemd-resolved.service
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Got message type=method_call sender=:1.248147 destination=org.freedesktop.resolve1
path=/org/freedesktop/resolve1 interface=org.freedesktop.resolve1.Manager member=ResolveHostname cookie=2 reply_cookie=0 signature=isit error-name=n/a error-
message=n/a
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: idn2_lookup_u8: yahoo.com → yahoo.com
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Looking up RR for yahoo.com IN A.
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Looking up RR for yahoo.com IN AAAA.
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Sent message type=method_call sender=n/a destination=org.freedesktop.DBus path=/org/freedesktop/DBus
interface=org.freedesktop.DBus member=AddMatch cookie=1018576 reply_cookie=0 signature=s error-name=n/a error-message=n/a
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Sent message type=method_call sender=n/a destination=org.freedesktop.DBus path=/org/freedesktop/DBus
interface=org.freedesktop.DBus member=GetNameOwner cookie=1018577 reply_cookie=0 signature=s error-name=n/a error-message=n/a
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Got message type=method_return sender=org.freedesktop.DBus destination=:1.0 path=n/a interface=n/a member=n/a
cookie=4294967295 reply_cookie=1018577 signature=s error-name=n/a error-message=n/a
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Removing cache entry for www.google.com IN A (expired 0s ago)
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Cache miss for yahoo.com IN AAAA
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Transaction 3083 for <yahoo.com IN AAAA> scope dns on wlp4s0/*.
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Using feature level UDP+EDNS0 for transaction 3083.
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Using DNS server 8.8.8.8 for transaction 3083.
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Sending query packet with id 3083 of size 38.
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Cache miss for yahoo.com IN A
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Transaction 8761 for <yahoo.com IN A> scope dns on wlp4s0/*.
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Using feature level UDP+EDNS0 for transaction 8761.
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Using DNS server 8.8.8.8 for transaction 8761.
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Sending query packet with id 8761 of size 38.
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Got message type=method_return sender=org.freedesktop.DBus destination=:1.0 path=n/a interface=n/a member=n/a
cookie=4294967295 reply_cookie=1018576 signature= error-name=n/a error-message=n/a
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Match type='signal',sender='org.freedesktop.DBus',path='/org/freedesktop/
DBus',interface='org.freedesktop.DBus',member='NameOwnerChanged',arg0=':1.248147' successfully installed.
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Processing incoming packet on transaction 8761 (rcode=SUCCESS).
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Added positive unauthenticated cache entry for yahoo.com IN A 97s on wlp4s0/INET/8.8.8.8
…
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Added positive unauthenticated cache entry for yahoo.com IN A 97s on wlp4s0/INET/8.8.8.8
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Transaction 8761 for <yahoo.com IN A> on scope dns on wlp4s0/* now complete with <success> from network
(unsigned).
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Processing incoming packet on transaction 3083 (rcode=SUCCESS).
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Added positive unauthenticated cache entry for yahoo.com IN AAAA 82s on wlp4s0/INET/8.8.8.8
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Added positive unauthenticated cache entry for yahoo.com IN AAAA 82s on wlp4s0/INET/8.8.8.
…
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Added positive unauthenticated cache entry for yahoo.com IN AAAA 82s on wlp4s0/INET/8.8.8.8
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Transaction 3083 for <yahoo.com IN AAAA> on scope dns on wlp4s0/* now complete with <success> from network
(unsigned).
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Sent message type=method_return sender=n/a destination=:1.248147 path=n/a interface=n/a member=n/a
cookie=1018578 reply_cookie=2 signature=a(iiay)st error-name=n/a error-message=n/a
Mar 03 16:56:01 pladd-laptop systemd-resolved[1159]: Freeing transaction 3083.
```

resolved – Cache Dump

```
# killall -USR1 systemd-resolved

# journalctl -fu systemd-resolved.service
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]: [Scope protocol=llmnr interface=virbr0 family=AF_INET]
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]: [Scope protocol=dns interface=virbr0]
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]: [Scope protocol=llmnr interface=vnet0 family=AF_INET6]
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]: [Scope protocol=dns interface=wlp4s0]
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]: CACHE:
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]:         api.wetter.com IN AAAA -- NODATA
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]:         slack.com IN A 18.214.242.166
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]:         slack.com IN A 54.211.89.16
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]:         slack.com IN A 54.87.197.95
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]:         mail.google.com IN CNAME googlemail.l.google.com
…
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]: [Server 1.1.1.1 type=link interface=wlp4s0]
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]:         Verified feature level: n/a
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]:         Possible feature level: TLS+EDNS0+D0
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]:         DNSSEC Mode: yes
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]:         Can do DNSSEC: yes
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]:         Maximum UDP packet size received: 512
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]:         Failed UDP attempts: 0
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]:         Failed TCP attempts: 0
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]:         Seen truncated packet: no
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]:         Seen OPT RR getting lost: no
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]:         Seen RRSIG RR missing: no
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]: [Server 8.8.8.8 type=link interface=wlp4s0]
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]:         Verified feature level: UDP+EDNS0
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]:         Possible feature level: UDP+EDNS0
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]:         DNSSEC Mode: yes
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]:         Can do DNSSEC: no
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]:         Maximum UDP packet size received: 512
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]:         Failed UDP attempts: 0
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]:         Failed TCP attempts: 0
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]:         Seen truncated packet: no
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]:         Seen OPT RR getting lost: no
Mar 03 17:46:45 pladd-laptop systemd-resolved[1159]:         Seen RRSIG RR missing: no
```

resolved – More resources

- https://fedoramagazine.org/systemd-resolved-introduction-to-split-dns/
- https://blogs.gnome.org/mcatanzaro/2020/12/17/understanding-systemd-resolved-split-dns-and-vpn-configuration/