


MIF using reverse DNS (draft-stenberg- mif-mpvd-dns)



Steven Barth
Markus Stenberg

Problem Statement

Multiple Uplinks

1. Different connection types (e.g. DSL, Cable, LTE, ...)
 2. Different latency, speed and reliability attributes
 3. Metered vs. unmetered connections
- Hosts should be able to use the most suitable prefix for a certain application

Special Purpose Prefixes

1. ISP delegates global prefixes which are only usable for special services
 - a. Triple Play, IPTV only, VOIP only addresses...
 - b. ... or access to company VPN!
- Hosts need to distinguish internet-capable and special-purpose prefixes

How to communicate prefix attributes?

Current MIF Proposal: Extend RA & DHCPv6

1. Define PVD Container TLV
2. Enclose RA/DHCPv6 TLVs (Addresses, DNS, ...) in PVD container
3. Add in-band signature for security

Issues

1. Require modifications on **all** intermediate routers (router is responsible!)
2. **Router** needs to have private signature key (for security)
3. **Requires changing** RA + DHCPv6 state machines on **hosts**
4. Increased multicast data
5. **Duplication** for legacy (RA) or shared multi-PVD addresses, information
6. Requires **individual** id and metadata definition for RA, DHCPv6, DHCPv4, ...

Alternatives?

Reuse DNS PTR+TXT records!

1. Existing hierarchy and transparent propagation
2. Prefix owner in control of PVD information
3. No (or merely minimal) modifications to routers

How?

1. ISP adds special well-known reverse DNS records to prefix
 - a. can add arbitrary key=value information pairs
2. (CPE) Routers propagate addresses as-is (optionally indicate PVD presence)
3. Hosts query reverse DNS records before using special addresses
 - a. can happen entirely in userspace (e.g. custom browser resolvers)
 - b. hosts can optionally use DNSSEC to authenticate PVD information
4. Hosts honor PVD information

Host behavior in detail

1. Host has addresses 2001:db8:8765:4321:1234:5678:abcd:beef/64 and 2001:db8:ffff:4321:1234:5678:abcd:beef/64
2. Host queries for PTR records in _pvd.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.1.2.3.4.5.6.7.8.8.b.d.0.1.0.0.2.ip6.arpa and _pvd.0.1.2.3.4.f.f.f.f.8.b.d.0.1.0.0.2.ip6.arpa
3. Receive: PTR internet.acme.example for #1 and iptv.foo.example for #2
4. Host queries for TXT records in internet.acme.example & iptv.foo.example
5. Receives "n=ACME ISP" "tm" "bw=10000" and "n=Foo IPTV" "s" "6=2001:db8::/64" "r=dns-iptv.foo.example"
 - #1 (Internet) connection to "ACME ISP", traffic metered, 10 MBit/s
 - #2 Connection to Foo IPTV, no-internet, access to 2001:db8::/64 only, use dns-iptv.foo.example as Recursive DNS

Note: Multiple PTRs per prefix and multiple TXT per PTR'ed domain possible!

Defining TXT-record PVD metadata

Reachable Services

- (human readable) name, internet accessible: yes/no
- accessible DNS zones, IPv6, IPv4 prefixes

PVD DNS Configuration

- Custom Recursive DNS Server or DNS search domain

Connectivity Characteristics and Limitations

- Maximum bandwidth, latency, reliability, ...?
- Traffic metered: yes/no, Captive Portal: yes/no, IPv4 NAT: yes/no, ...?


Enterprise-Specific Information

- e.g. **x-foobar.com-buzzword=bingo**

Special Purpose Prefixes #1

- Problem: Legacy hosts **MUST NOT** mistakenly use special purpose addresses for internet access.


- RA Extension to Prefix Information Option



- S-Bit works like A-Bit but only if host understands PVD information
 - Legacy hosts **ignore** S-Bit and thus ignore prefix for SLAAC if A=0
 - PVD-aware hosts only use address **after** receiving PVD policy info
 - A=S=1: legacy-capable, **optional** auxiliary PVD information for PVD hosts

Special Purpose Prefixes #2

- DHCPv6 Option Extension



- PVD-aware hosts request OPTION_SPECIAL_PURPOSE in ORO
- Routers **only assign** special-purpose addresses **if ORO exists**
- Special-purpose address / prefixes (IAAddr + IAPrefix) include **nested** option OPTION_SPECIAL_PURPOSE to indicate mandatory PVD

- PVD-aware routers **DO NOT** need to query PVD-information!
 - Only need to propagate special-purpose indication from RA / DHCPv6

Summary

Features

- **No overhead** for legacy hosts & routers
- Little to no changes for RA, DHCPv6/v4
- Scalable (no multicast requirements)
- Single **dual-stack** format definition
- Efficient m:n PVD:Address handling
- Easily **extendable** TXT records
- Reuse existing address / DNS hierarchy

Open Issues

- Do we also need whole special-purpose routers?
- DNSSEC delegation for reverse zones feasible?
- Should PVD-aware hosts strictly prefer routers announcing the PIO when using an address? (draft-6man-multi-homed-host)
- Are all implications of the asynchronicity between getting addresses and PVDs clear?

Thank you for your attention! Questions?

Steven Barth
<cyrus@openwrt.org>