IPv6 Multihoming from the IETF perspective

Eric Vyncke, evyncke@cisco.com, @evyncke
Distinguished Engineer, Paris Innovation & Research Lab
December 2017

Agenda

- Problem statement
- Provisioning the host
- Routing to the multi-home exit
- Demo (time allowing)

Problem statement

Hosts and networks are multi-homed

Just a few examples…

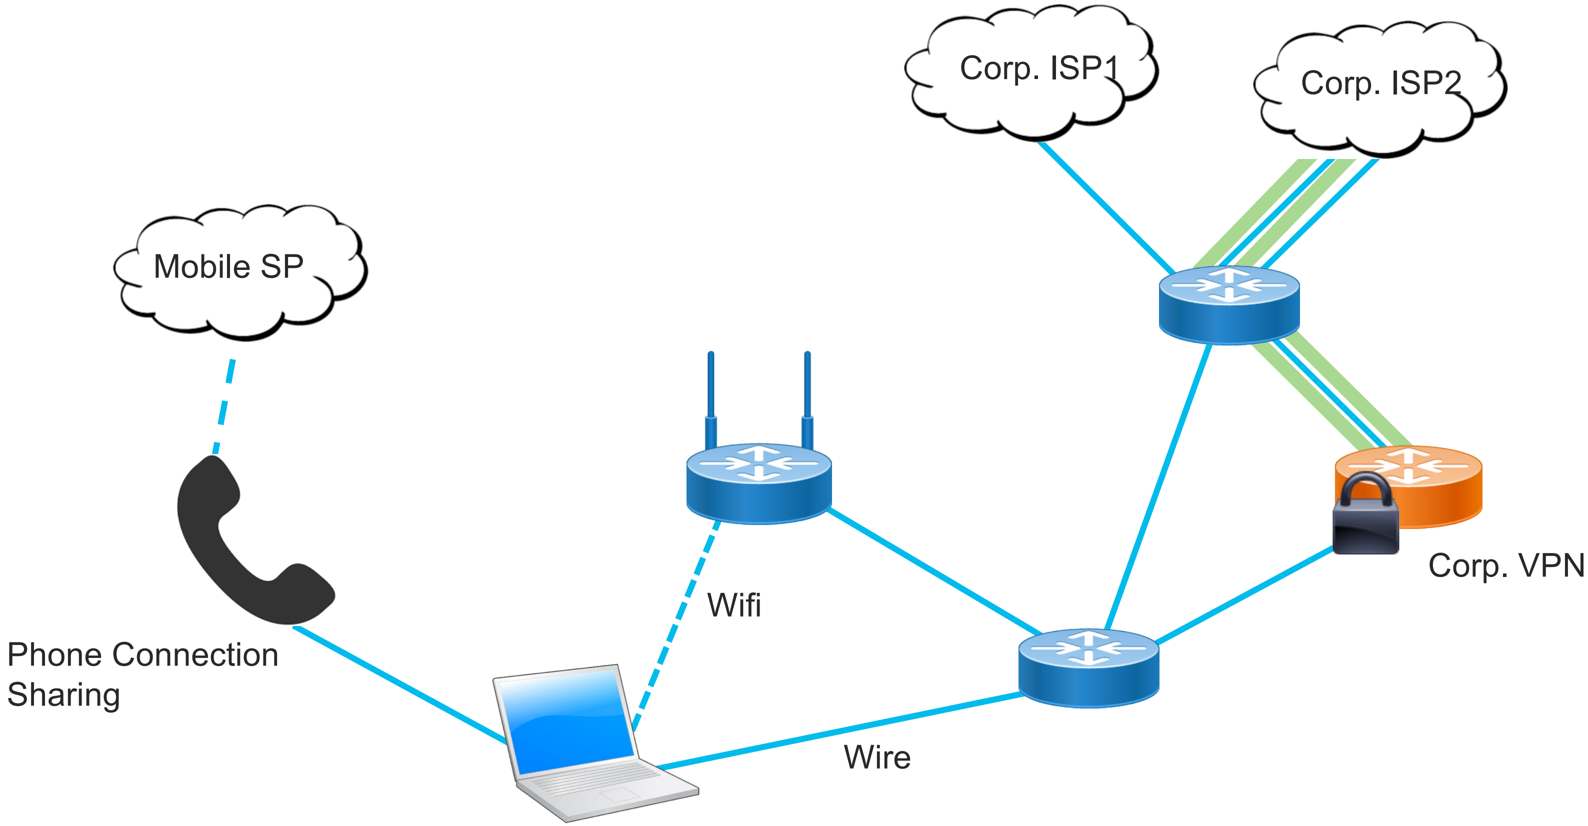

Multi-Homing, the legacy way…

WARNING
This slide has strong and offensive wording: the N word…
BE WARNED

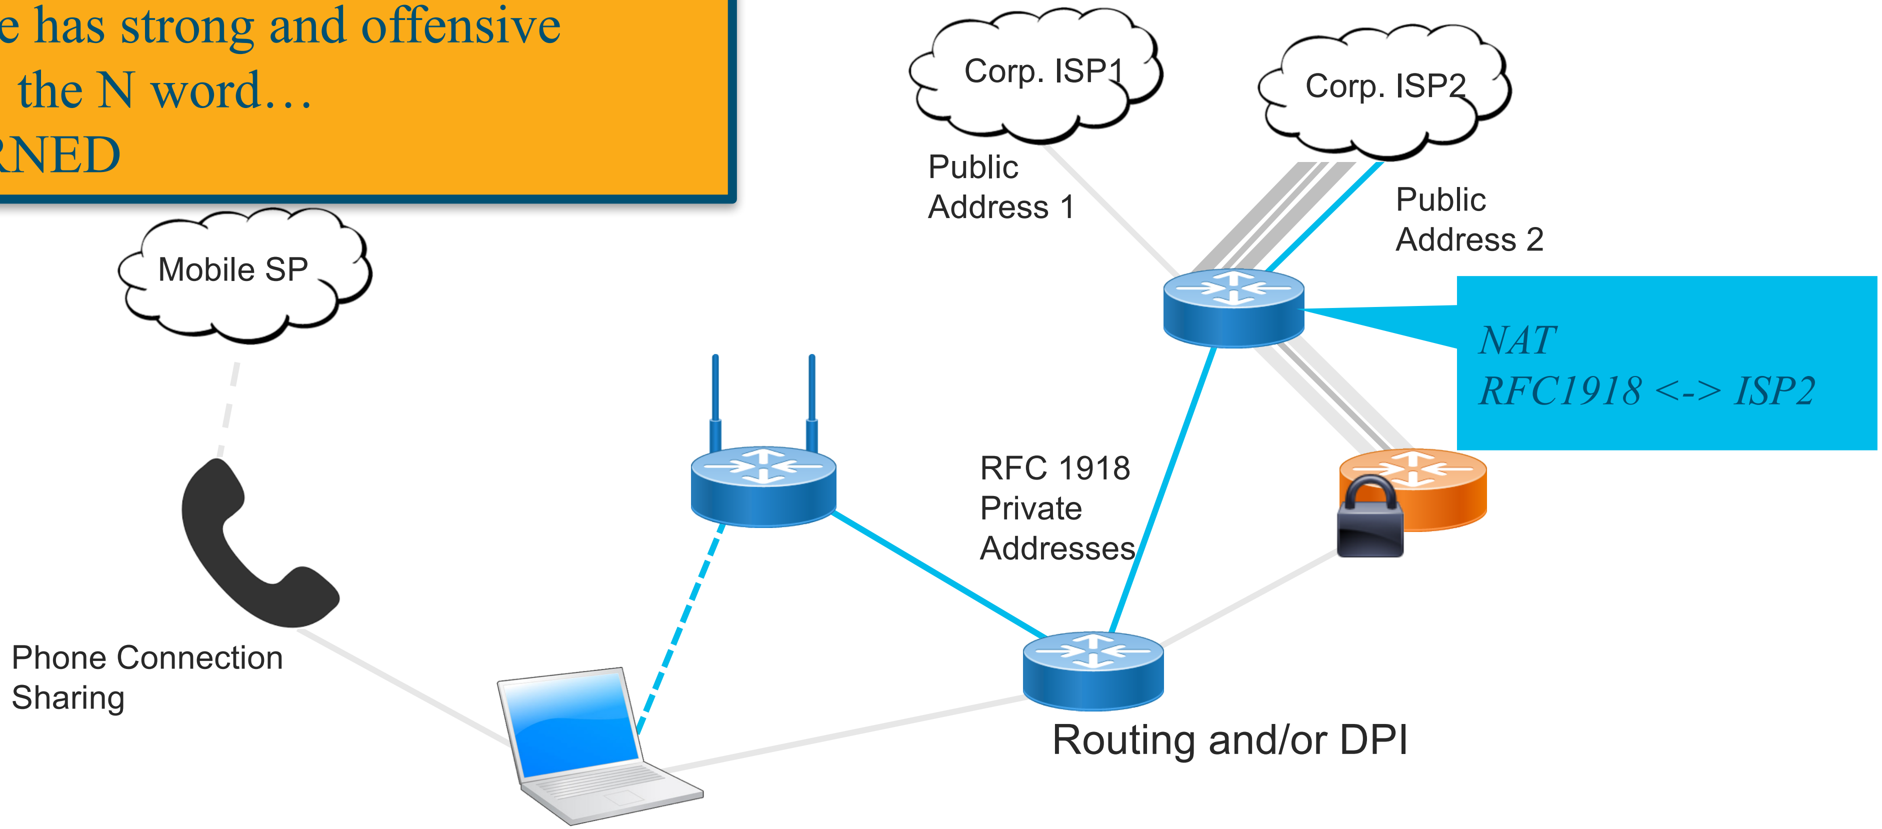

Addressing in Multi-Homed Networks in IPv6

- Assign Provider Assigned (PA) addresses to hosts.
  - Native to IPv6 hosts (RFC4861, ...)
  - HNCP for home networks (RFC7788)
  - draft-ietf-rtgwg-enterprise-pa-multihoming for corporate networks.
- Teach the hosts to pick and use multiple addresses.
  - IPv6 source address selection (RFC6724)
  - Multi-Path TCP (RFC6824), SCTP, QUIC, ...
- Give the host meaningful information about the addresses.

Bundling IP address & DNS resolver

Multihoming and CDNs

- Name lookups for resources stored on CDNs give different answers depending on the network connection
- Host on homenet may look up name using resolver from provider A, then connect to CDN using provider B
- This will generate support requests
- What to do?

Ted Lemon, Homenet WG, IETF-99

Example Topology

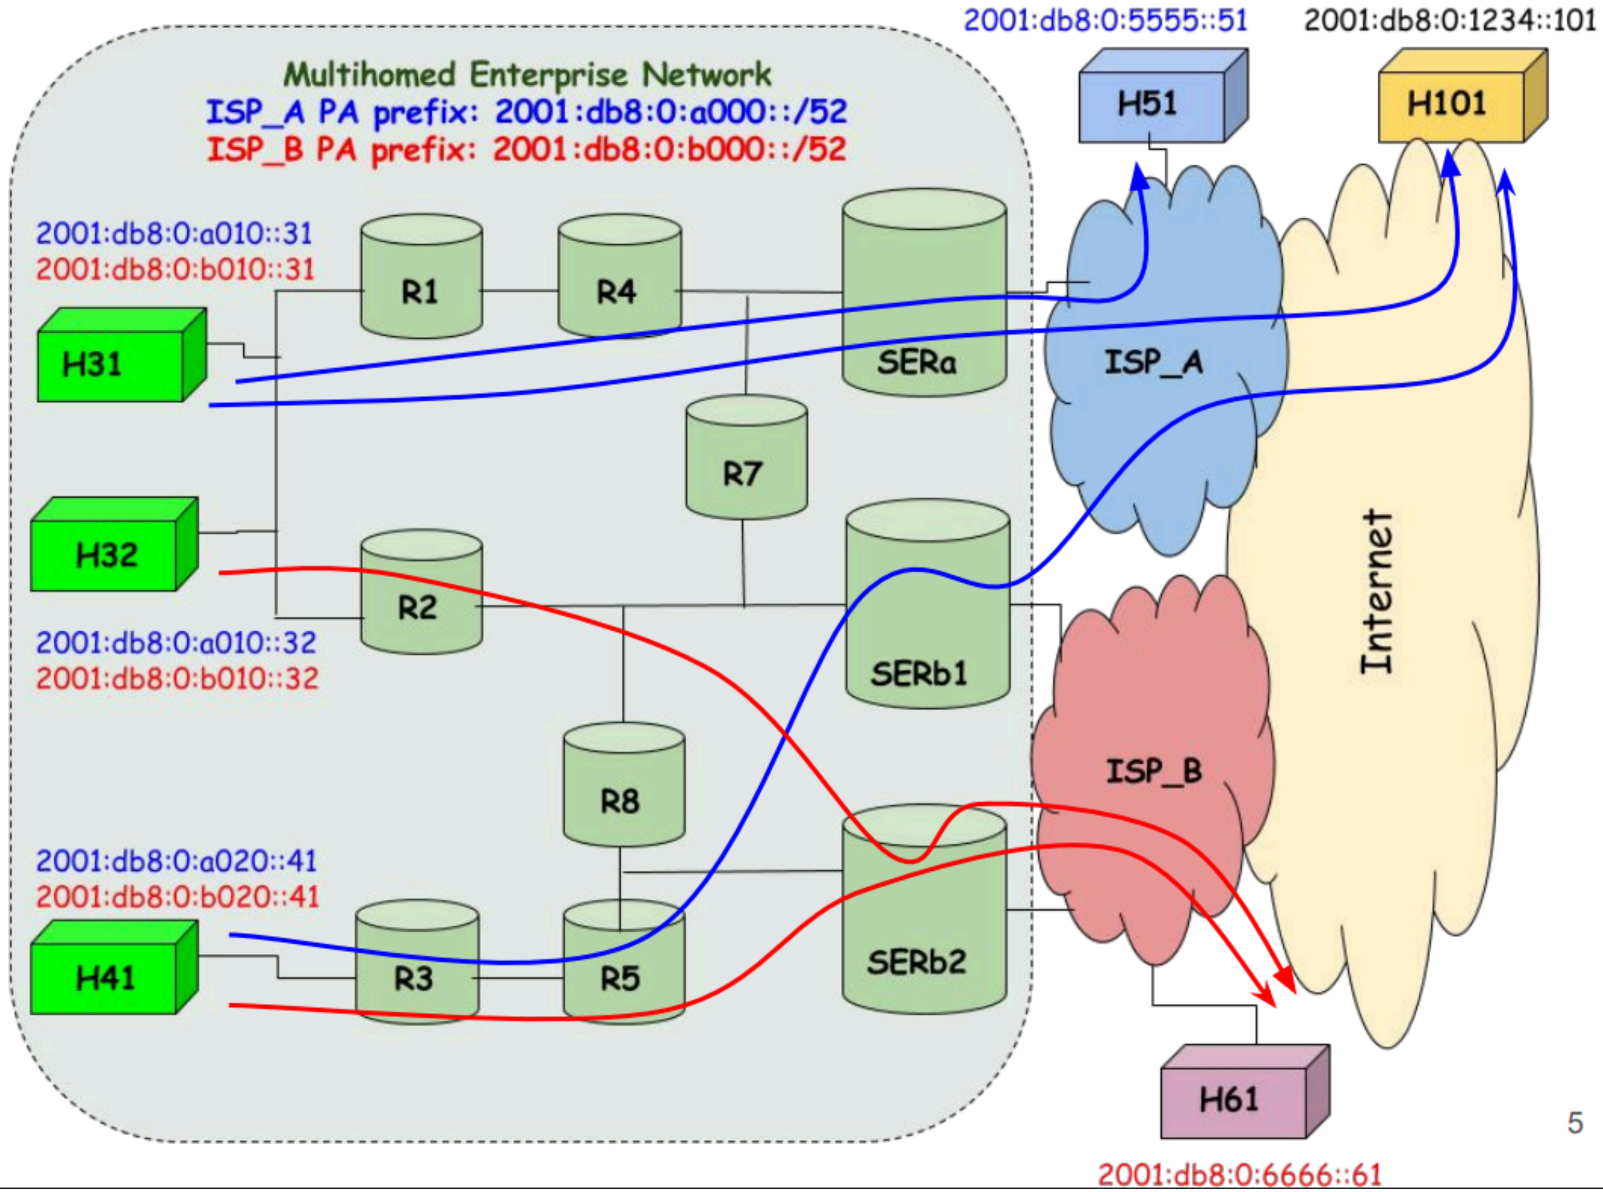

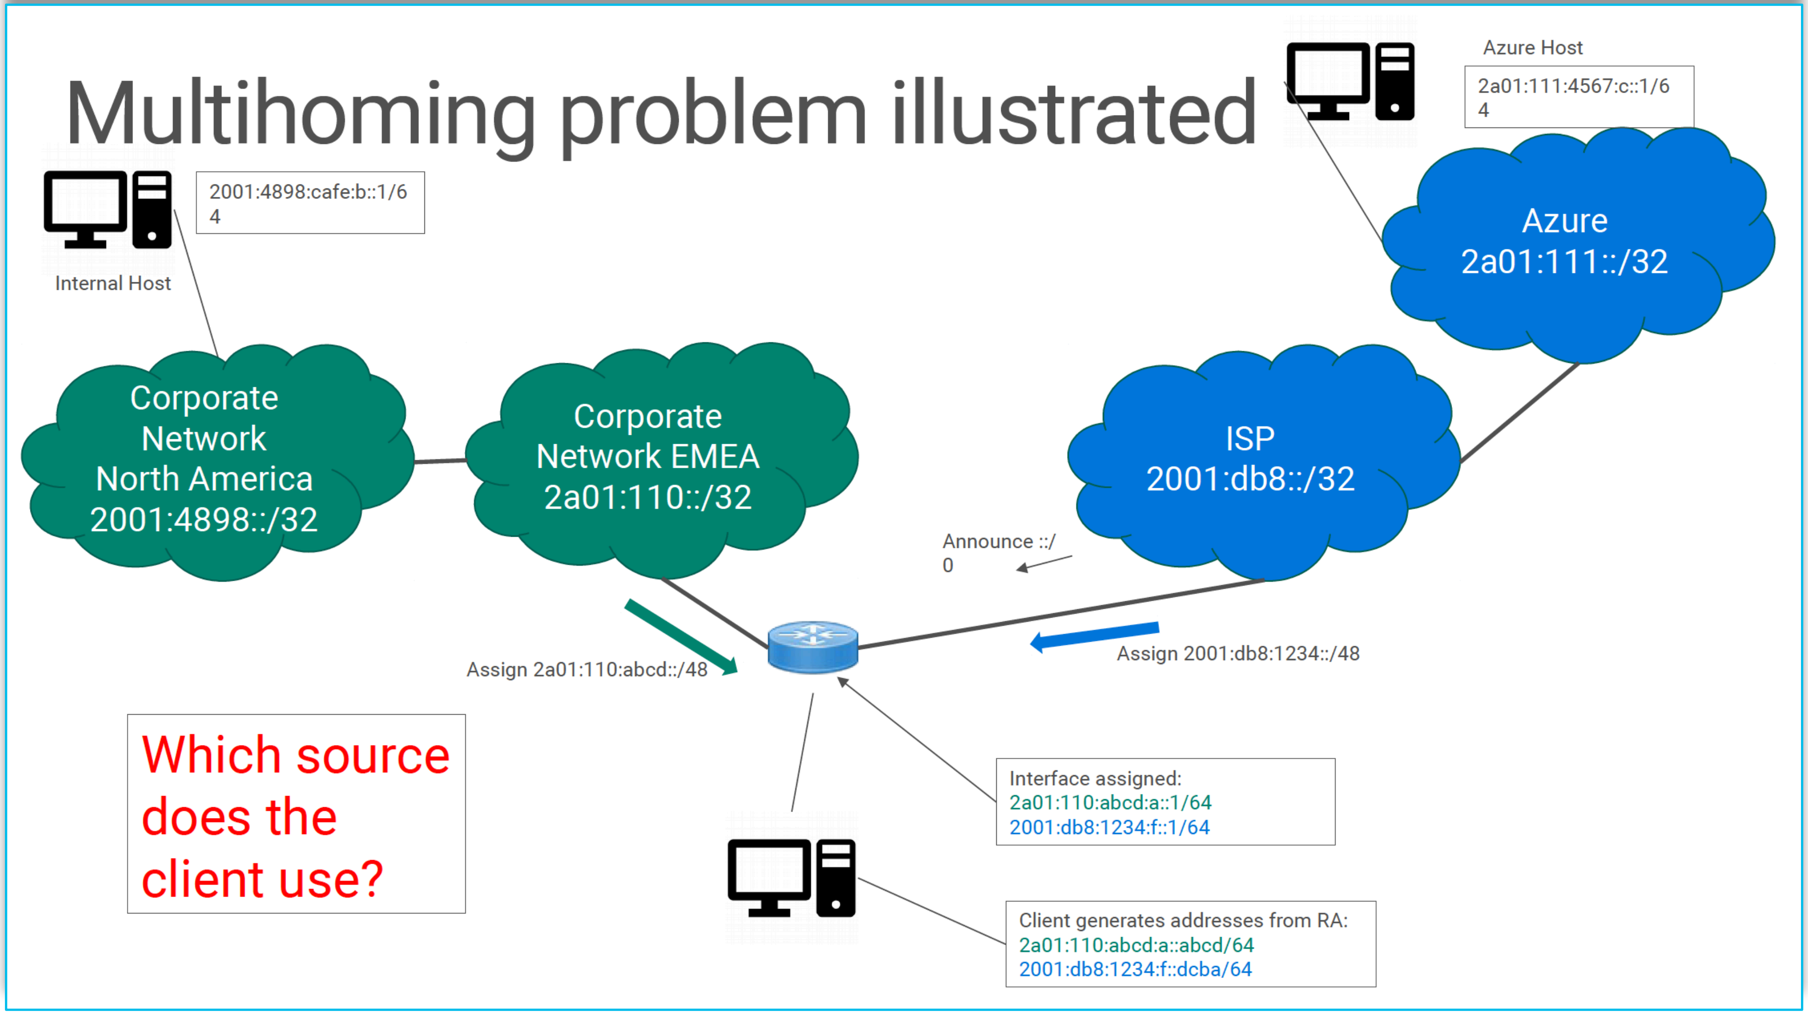

From Marcus Kean, Microsoft IT, at V6OPS IETF-99

Selecting the Service by Source Address

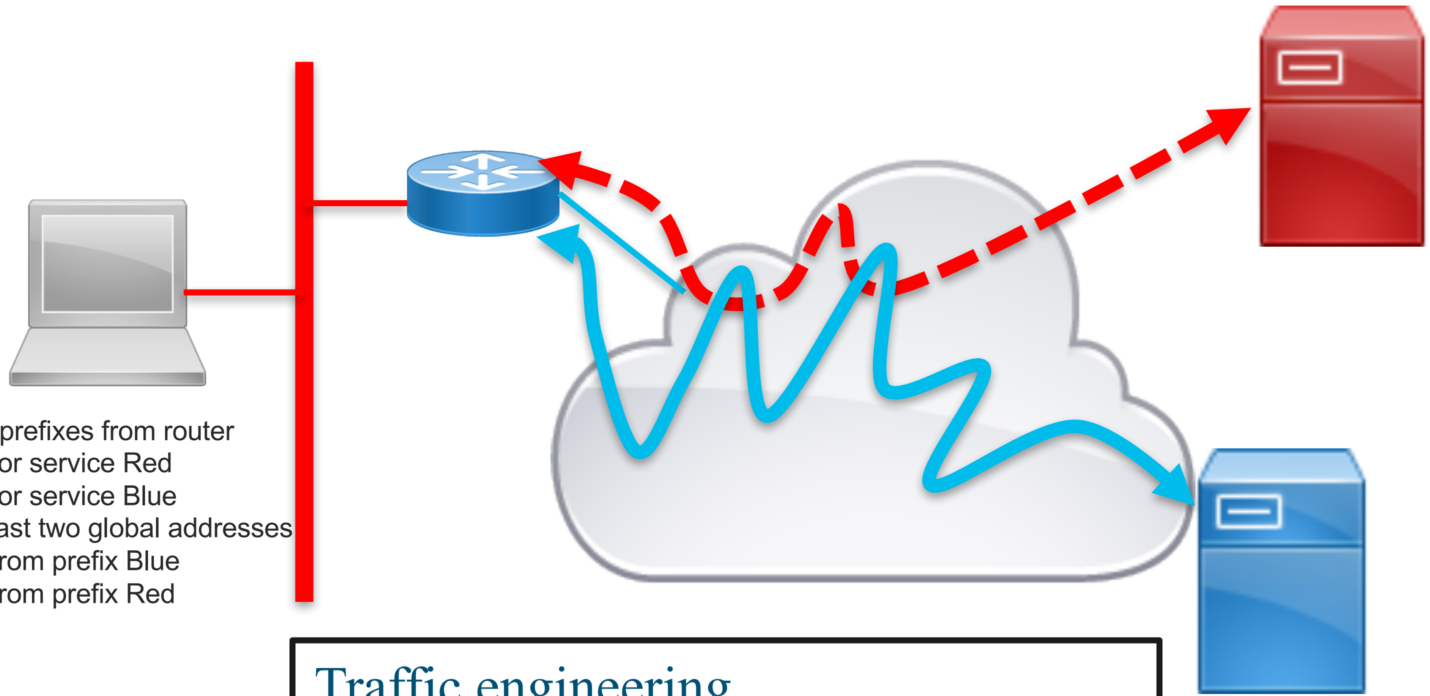

Two prefixes from router

1. For service Red
2. For service Blue

At least two global addresses

1. From prefix Blue
2. From prefix Red

Traffic engineering
Different QoS
Different routing (€€€ or security)

Provisioning the host

- How can the host discover all network prefixes and services?
- At the network and application layers

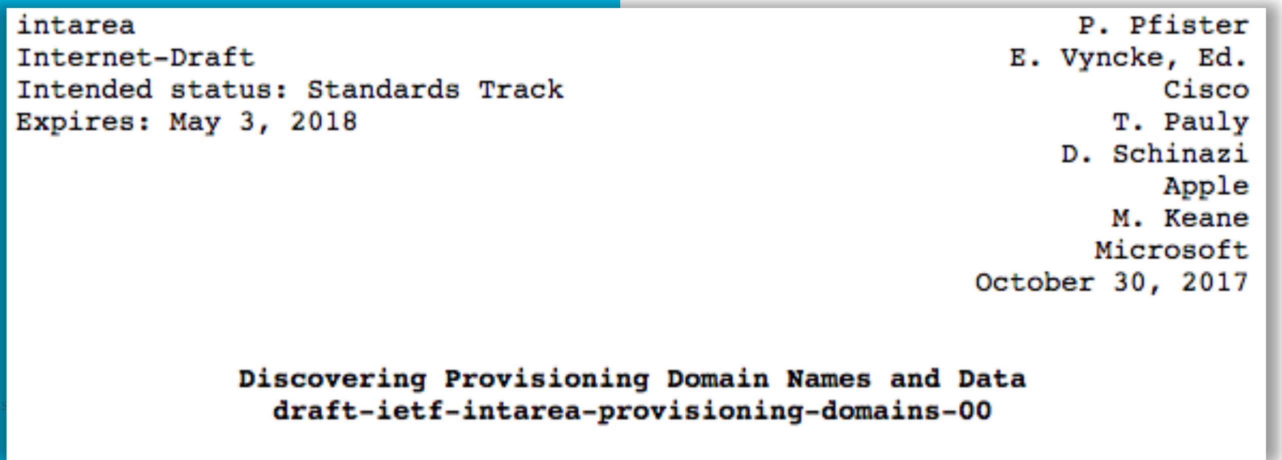

intarea — P. Pfister
Internet-Draft — E. Vyncke, Ed.
Intended status: Standards Track — Cisco
Expires: May 3, 2018 — T. Pauly
D. Schinazi
Apple
M. Keane
Microsoft
October 30, 2017

Discovering Provisioning Domain Names and Data
draft-ietf-intarea-provisioning-domains-00

draft-ietf-intarea-provisioning-domains

1. Identify Provisioning Domains (PvDs)

[RFC7556] *Provisioning Domains (PvDs) are consistent sets of network properties that can be implicit, or advertised explicitly.*

Differentiate provisioning domains by using FQDN identifiers.

2. Extend PvD with additional information

For the applications: name, characteristics, captive portal, etc…

Step 1: Identify PvDs

With the PvD ID Router Advertisement Option

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     Type      |    Length     |H|L|         Reserved          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|          Sequence             |                             ...
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+                             ...
...                       PvD ID FQDN                         ...
...             +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
...             |                  Padding                      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

           PvD ID Router Advertisements Option format
```

- At most **one occurrence in each RA**.
- **PvD ID is an FQDN** associated with options in the RA**.**
- **H bit** to indicate **Additional Information is available with HTTPS**.
- **L bit** to indicate the **PvD has DHCPv4 on the link**.
- Seq. number used for **push-based refresh**.

Step 1b: Identifying PvD (Cont.)

- Information in a RA without PvD ID is linked to an implicit PvD (identified by interface & link-local address of router)
- Option in RA can change of PvD when they are received in a RA with a different PvD ID
- DHCPv6 information MUST be associated to a PvD ID received on the same interface from the same link-local address

Step 2: Get the PvD Additional Application Data

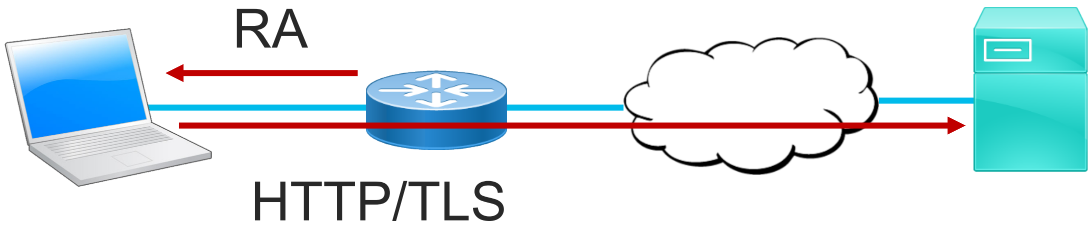

When the H bit is set:

GET https://<pvd-id>/.well-known/pvd

Using network configuration (source address, default route, DNS, etc…) **associated with the received PvD**.

Step 2: Get the PvD Additional Data


```
{
  "name": "Foo Wireless",
  "expires": "2017-07-23T06:00:00Z",
  "prefixes" : ["2001:db8:1::/48", "2001:db8:4::/48"],
  "localizedName": "Foo-Hôtel à Paris Wifi",
  "dnsZones": ["example.com","sub.example.com"];
  "characteristics": {
        "maxThroughput": { "down":200000, "up": 50000 },
        "minLatency": { "down": 0.1, "up": 1 }
  }
}
```


Some other examples (see also https://smart.mpvd.io/.well-known/pvd) :

```
noInternet : true,
metered : true,
captivePortalURL : "https://captive.org/foo.html”
```

Implementation status

Linux - https://github.com/IPv6-mPvD

- **pvdd**: user-space daemon managing PvD IDs and additional data
- **Linux Kernel** patch for RA processing
- **iproute** tool patch to display PvD IDs
- **Wireshark** dissector
- **RADVD** and **ODHCPD** sending PvD ID

Implemented in one commercial vendor router

neat

A New, Evolutive API and Transport-Layer Architecture for the Internet: https://www.neat-project.org/

European H-2020 project

10 partners (Cisco, Mozilla, EMC, Celerway…)

Integration to NEAT code: https://github.com/NEAT-project/neat/pull/80

Asking the user to choose with relevant criteria and simple UI

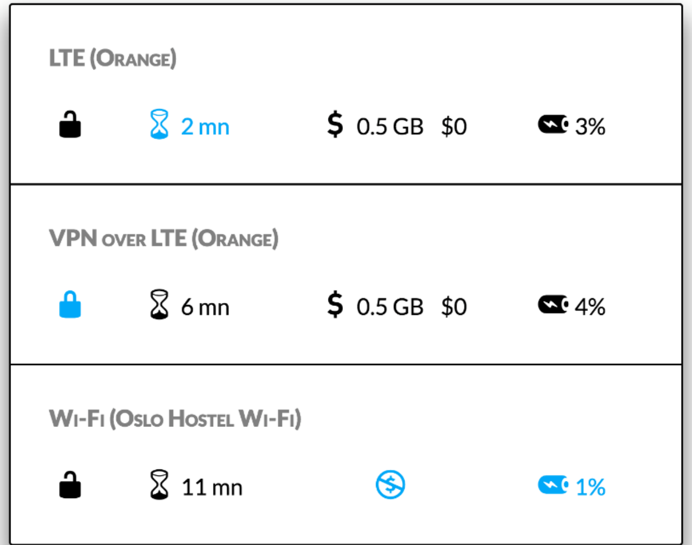

Source Address Dependent Routing (SADR)

- Forwarding based on the SOURCE rather than the destination as usual?
- Based on source scoped Forwarding Information Base (FIB)

```
rtgwg                                                        D. Lamparter
Internet-Draft                                                    NetDEF
Intended status: Standards Track                              A. Smirnov
Expires: May 3, 2018                                 Cisco Systems, Inc.
                                                        October 30, 2017

                       Destination/Source Routing
                  draft-ietf-rtgwg-dst-src-routing-06
```

SADR in a nutshell

- All FIB entries are associated with a source prefix
  - ::/0 for entries without a source prefix
- draft-ietf-rtgwg-dst-src-routing
- Algorithm
  1. `PotentialRoutes :=Longest match(es) on destination prefix`
  2. `SourceRoute := longest match on the packet source in the PotentialRoutes`
  3. `If not found, then back to 1) with a shorter match`

- Other implementations are possible

Trivial SADR Example

- SADR FIB

Source	Destination	Next - Hop
::/0	::/0	R3
2001:db8::/32	::/0	R3
2001:db8:2::/64	::/0	R4

- Packet SRC = 2001:db8:1::1 to DST = 2001:db8:cafe::babe
- Packet SRC = 2001:db8:2::1 to DST = 2001:db8:cafe::babe

Incremental Deployment

- SADR **only** on edge routers
- Best effort forwarding:
  - R3 can have a SADR route to R4 for ISP2 source prefix
- SADR on R1 / R6 would only improve
- If R3 and R4 are not adjacent, then SRv6 (or a tunnel) is required

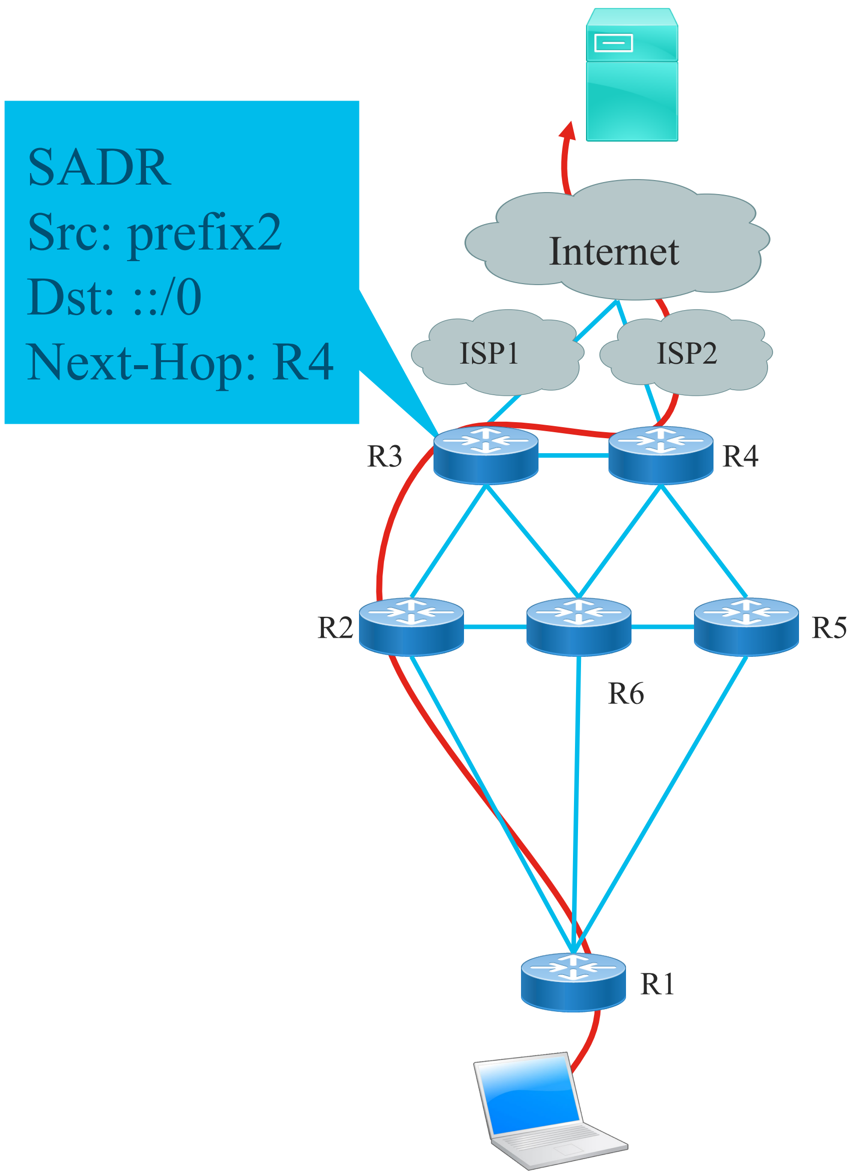

Summary of SADR for multi-homing

- SADR allows network to send packets to the “right” egress point
- SADR can be deployed incrementally
  - MUST be enabled on the edge
  - Tunnels may be required until complete deployments
- Routing protocols can be extended to SADR

Demo Time

- Only if time is available

- Source Address Dependent Routing

- Provisioning and Captive portal

Conclusion

- Multi-homing in IPv6 is vastly different than in IPv4
- Several addresses per interface
- Several interfaces per host in 2017
- Host must select the right bundle of DNS, address, next hop
- Network must route according to the host-selected address
- Implementations exist
- Huge momentum at IETF