draft-ietf-v6ops-tunnel-loops - Update and Status

IETF V6OPS WG - March 31, 2011

Fred L. Templin
Boeing Research & Technology
fred.l.templin@boeing.com

Tunnel Looping Problem Statement

- **Tunnel routers with IPv4-embedded IPv6 addresses use stateless address mapping**
- **When router receives a packet from the IPv6 Internet, tunnels it to the embedded IPv4 address**
- **No way for the router to know whether the holder of the IPv4 address is aware of the tunnel**
- **Holder of the IPv4 address could forward the packet back into the IPv6 Internet**

Example – Two ISATAP Routers with Different Prefixes

v6src = 2001:db8:0:2::5efe:192.0.2.3
v6dst = 2001:db8:0:1::5efe:192.0.2.5

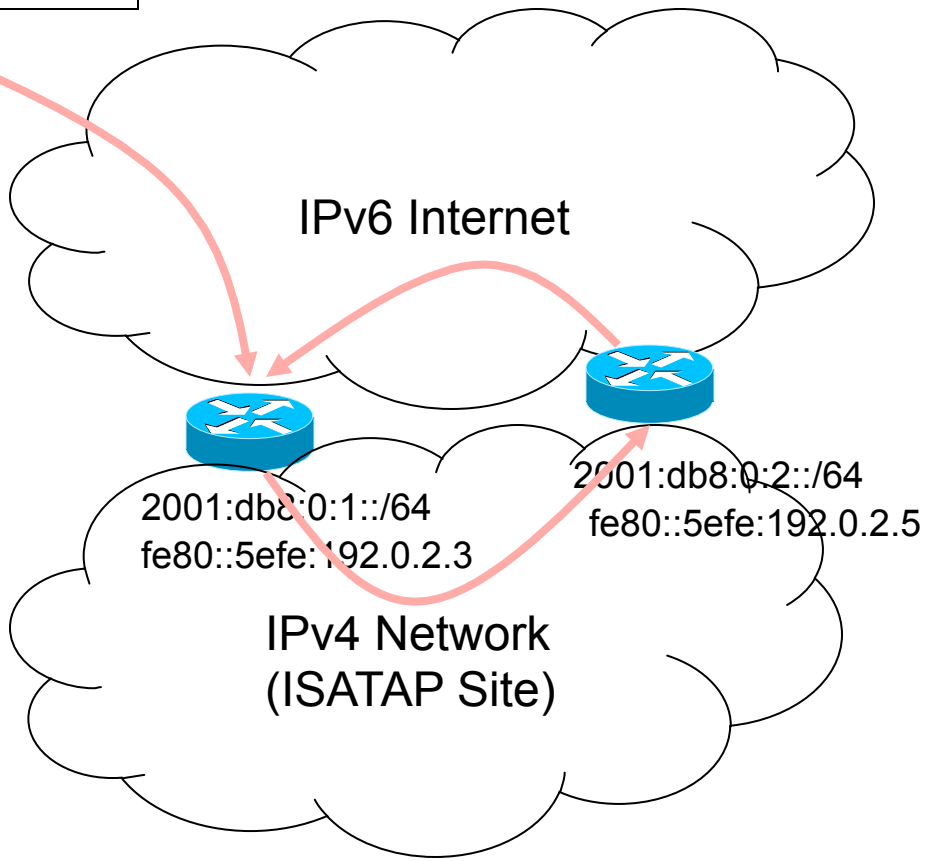

Proposed Mitigations

1. **Verification of endpoint existence**
   a) Neighbor cache check
   b) Known IPv4 address check
2. **Destination and Source address checks**
   a) Check whether the embedded IPv4 address is one of the router's own addresses (if so, drop)
3. **Operational Measures**
   a) IP-Protocol-41 filtering
   b) Operational avoidance of multiple tunnels within the same bounded IPv4 network
   c) Use only a single border router
   d) Use a comprehensive list of all tunnel routers (e.g., ISATAP PRL)
   e) **DON'T USE ON-LINK PREFIXES**

On-Link Prefix Avoidance

- **Can be used for ISATAP, 6rd, 6over4, etc.**
- **IPv6 prefixes assigned to the tunnel interface are rarely used as packet (src, dst) addresses**
- **IPv6 prefixes delegated to edge network links only; tunnel used as transit**
- **IPv6 address can still be assigned on tunnel interface**

Reference Operational Scenario (ISATAP example)

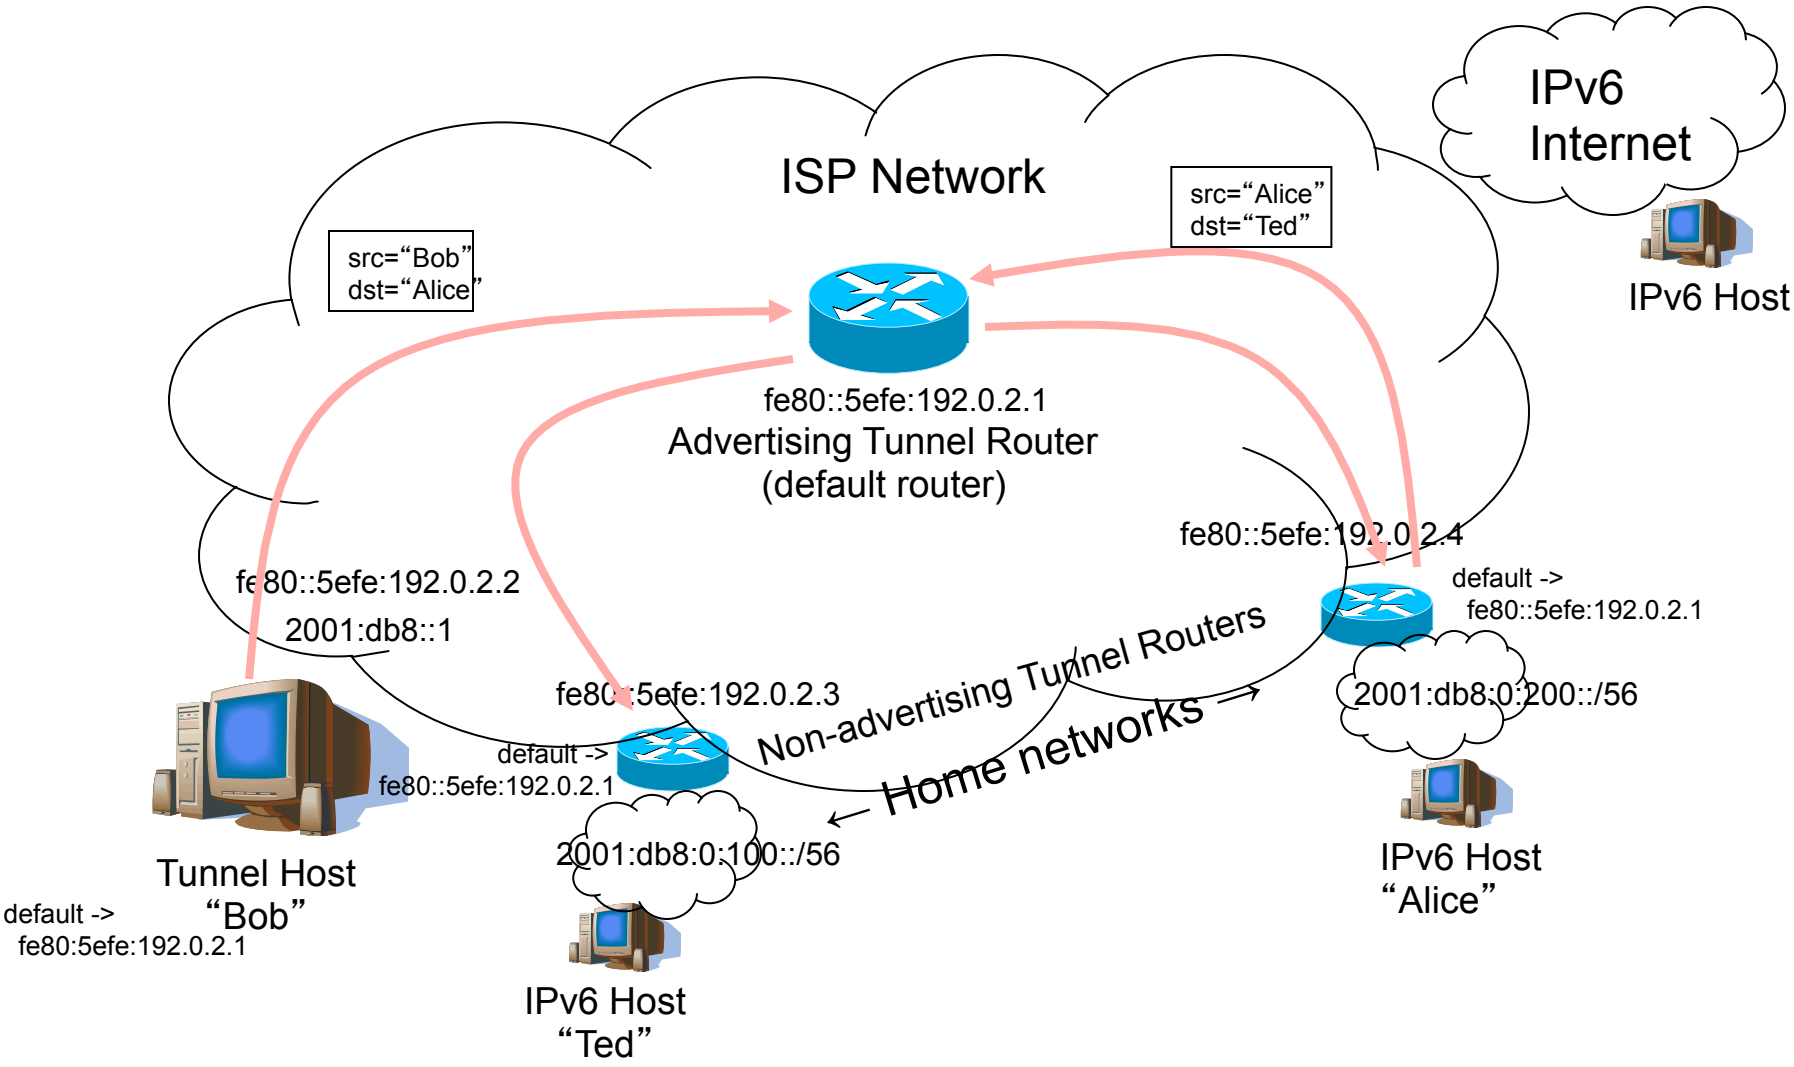

Issues

- **Traffic concentration on ISP infrastructure equipment**
- **Sub-optimal routing between CPE routers**
- **No provisions for ingress filtering, black hole avoidance, etc.**
- **Alternative 1: run IGP between CPE and PE routers**

Example with Dynamic IPv6 IGP between CEs and PEs

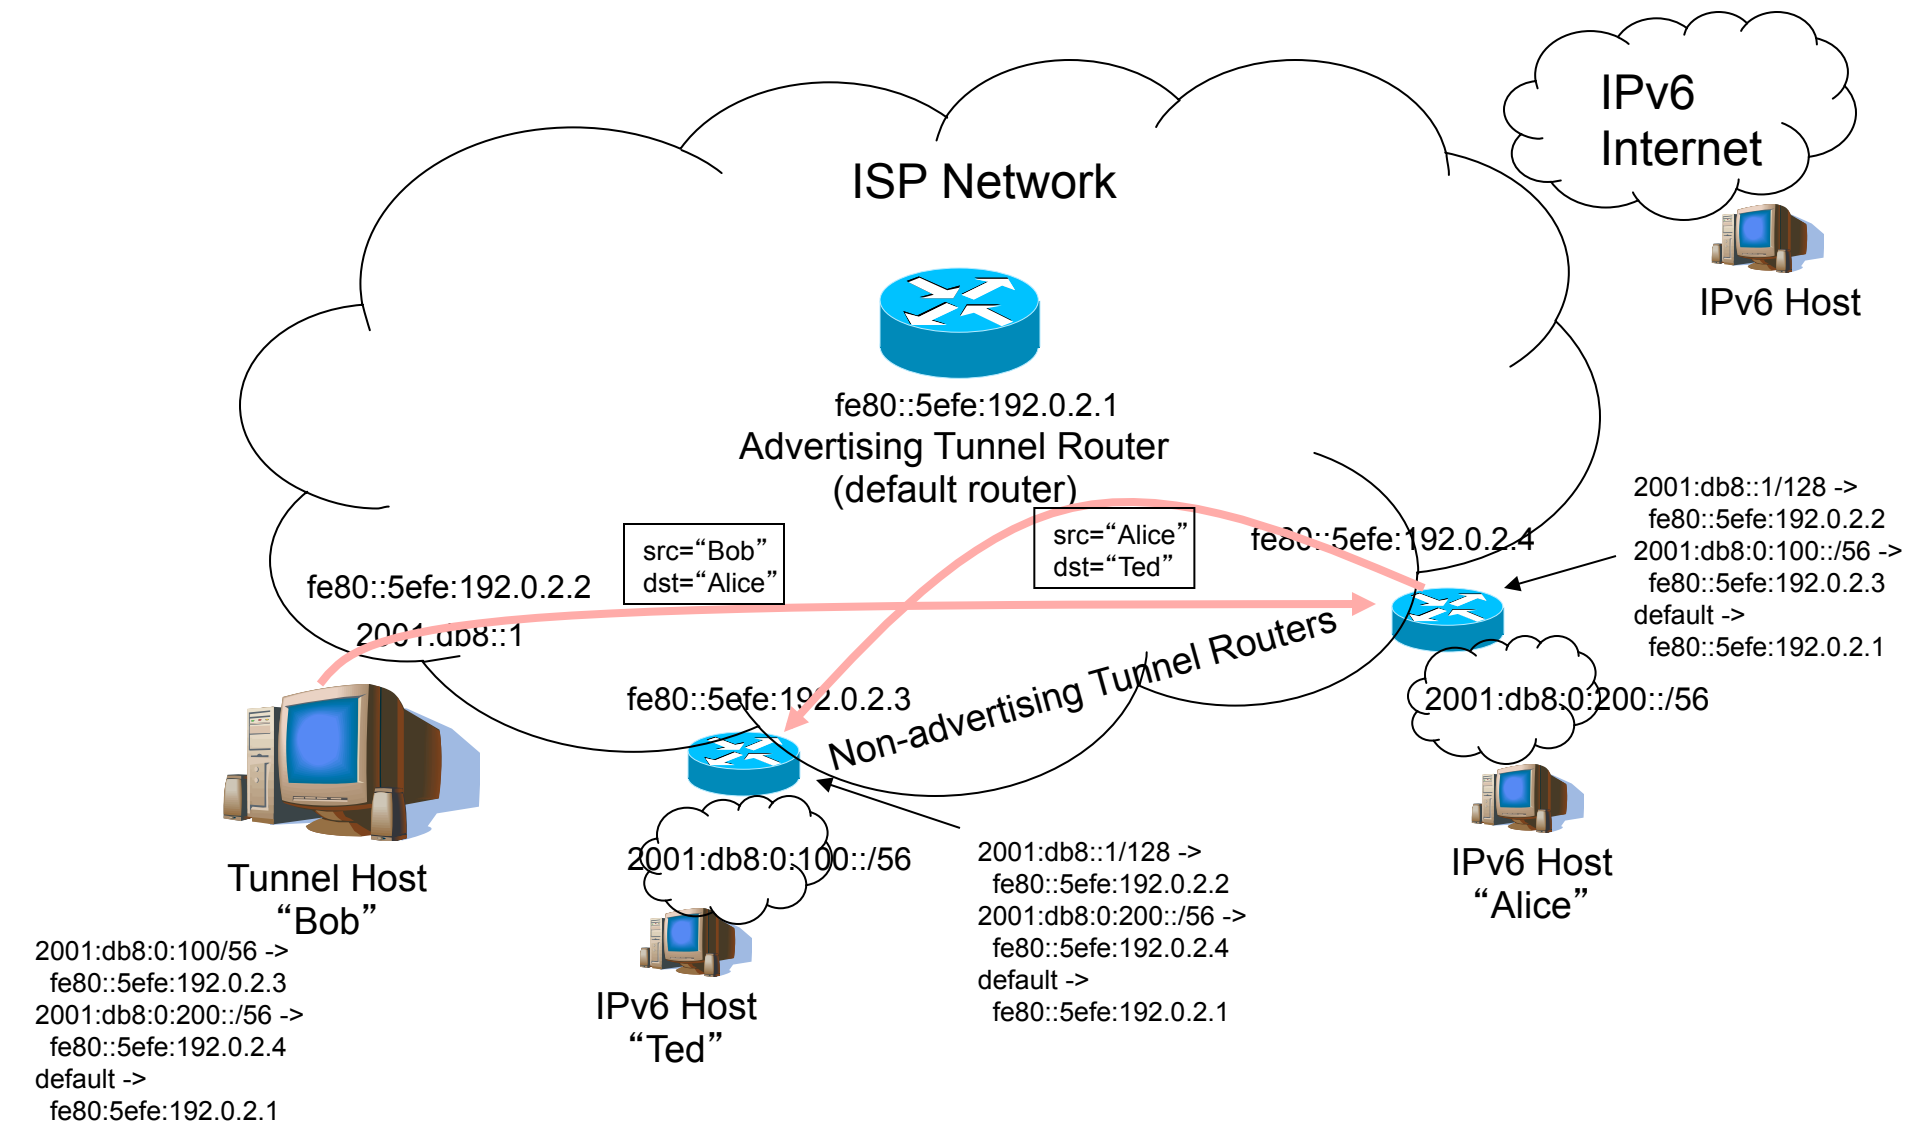

IGP Issues

- **Doesn't scale well**
  - **10K or more routers**
  - **10K or more routes**
  - **excessive control message overhead to keep all routers synchronized**
- **Requires routing protocol configuration**
- **CPEs constantly coming up and going down**
- **CPEs untrustworthy**

Requirements for Zeroconf Dynamic Routing

- R1: Zero configuration on CPE routers
- R2: Security based on chain-of-trust
- R3: Scale to support lots of CPEs
- R4: Off-load performance-critical ISP routers
- R5: CPE-to-CPE route optimization
- R6: Support multiple levels of hierarchy
- R7: Do not circumvent IPv6 filtering
- R9: Do not expose packets to loss due to black holes
- R9: Support IPv6 prefix mobility
- R10: Support the same mechanisms on the LAN side