



OSPFv3: Homenet and Data Centers

Fred Baker

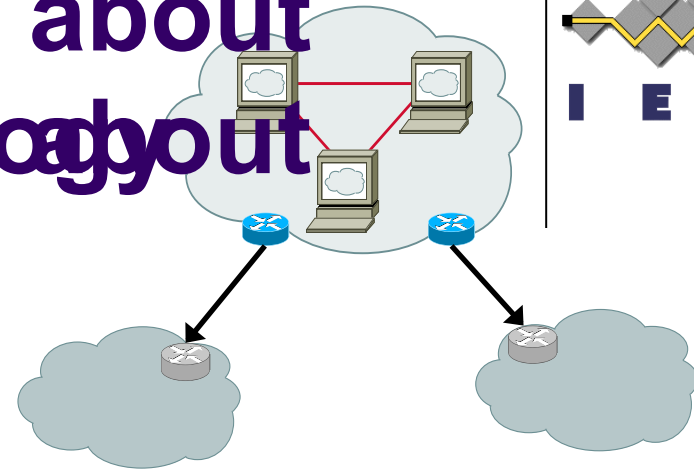






Homenet Requirements

- Homenet is trying to develop supporting technologies for a very simple, but technologically advanced, home
 - Primarily focused on IPv6
 - Zero Configuration if at all possible
 - Interface to Smart Grid technologies including Zigbee/802.15.4
 - **Multi-subnet with routing an option**
 - **Potentially multihomed to multiple ISPs**
 - **Edge Routing to resolve BCP 38 issues**

I have been asked about OSPF Multi-Topology




- Topologies are defined by metrics on links between router interfaces within the routing domain
 - The link does or does not have a metric within the topology
 - Automatically routes around discrepancies between physical and logical topology
 - A number of source/destination indicated metric routing cases could be
 - Edge routing is routing to a default route that is outside the routing domain
 - The OSPFv3 topologies for each PA prefix are *identical*
 - There is no link advertised in OSPFv3 that might have the same metric
 - Edge routing is a **reachability problem, not a topology problem**
 - Edge routing is a **reachability problem, not a topology**



Context


- RFC 5340 defines three prefix LSAs
 - Fixed format, which makes it hard to add information to them
- I'm looking at
 - Homenet requirements for egress routing
 - Multi-tenant Data Center requirements for tenant-to-tenant access control

draft-baker-ipv6-ospf- extensible




- I defined three extensible LSAs, replacements for intra-area-prefix-LSA, inter-area-prefix-LSA, and AS-external-LSA
- I have since been told of Abhay Roy's extensible LSA draft in draft-ietf-ospf-mt-ospfv3 (2007)
- I'll use whatever extensible technology the WG approves

Flow label and Source Address sub-TLVs

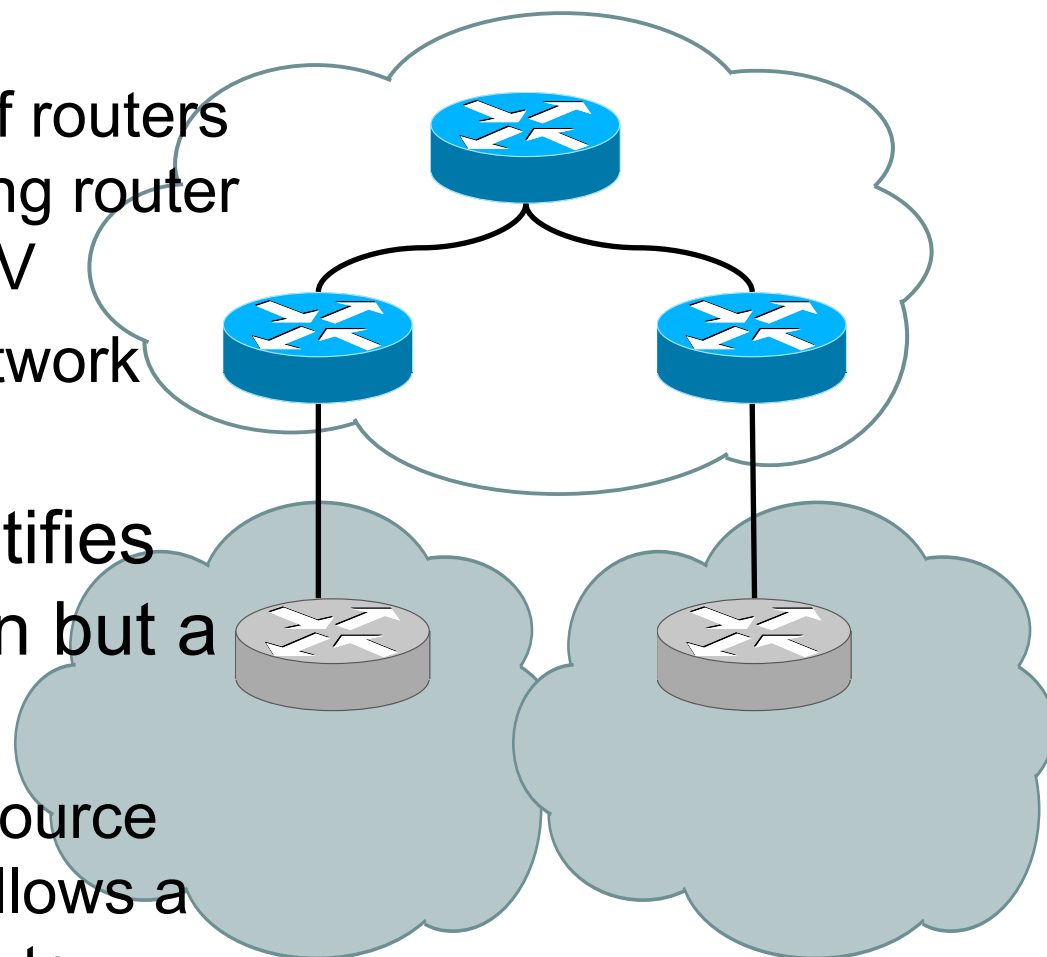



- Premise:
 - Reachability TLV, with sub-TLV(s), identifies a set of possible messages to send down a route
 - Need comments on route calculation and FIB design

Route Calculation




- Normal OSPFv3 route calculation:
 - Identifies a sequence of routers and links from calculating router to router advertising TLV
 - “Router” might be a Network LSA
- LSA, in this case, identifies not only the destination but a qualification
 - Traffic with a different source address or flow label follows a different route, or no route





Backward compatibility

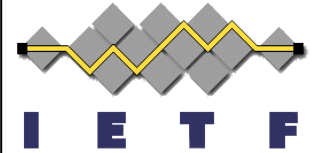
- Acee asked about making this work in networks with RFC 5340 format LSAs as well
- Really not a problem:
 - Definition of source prefix sub-TLV:
 - A zero-length LSA (::/0) can be represented with a sub-TLV whose length is zero or no sub-TLV
 - Definition of flow label sub-TLV:
 - “any” flow label is specified by leaving the sub-TLV out
- RFC 5340 LSA by definition leaves those sub-TLVs out. Semantically equivalent.



FIB Design

- Not subject to standardization.
- Some suggestions in an appendix
 - Linux (Waikato extensions) has separate FIBs by source prefix.
 - One could insert destination into appropriate FIB, or all FIBs if source not specified
 - PATRICIA tree
 - Allows a discontinuous bit string, differing don't-care sets
 - Recursive descent following most useful bits
 - Final answer compared to entire specification

Possible use cases



Source Prefix

- Egress Routing
 - Most TLVs in network destination-only
 - Default routes to upstream specify PA source prefix
- One could imagine more general uses, such as dynamic “ACL”

Flow Label (RBAC model)

- Long discussion about use of the Flow Label in the IETF, with many suggestions
- One could also use it as a tenant id in a multi-tenant data center
 - IPsec or TLS still required for proper end-to-end security
 - Tagged route limits attack possibilities to neighbors that know the “password”