IS-IS/OSPFv3 extensions for destination +<something> routing

Fred Baker

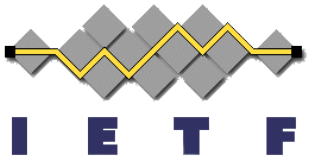

Homenet Requirements

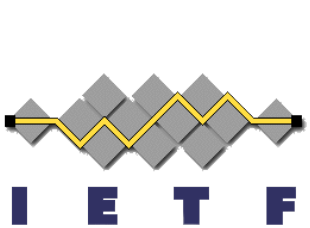

- Homenet is trying to develop supporting technologies for a very simple, but technologically advanced, home
  - Primarily focused on IPv6
  - Zero Configuration if at all possible
  - Interface to Smart Grid technologies including Zigbee/802.15.4
  - **Multi-subnet with routing an option**
  - **Potentially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

Additional issues

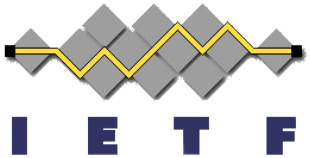

- I have a security problem I want to solve in data centers
  - That's not Homenet's problem, but I would like a corresponding solution, and
  - I'd like to have the debate needed once, not twice
- I think this can be solved in IS-IS or OSPFv3,
  - I have customers likely to want it in either protocol
  - In either protocol, I have *AS-external issues, inter-area issues,* and *intra-area* issues. That implies looking at routing information in all of those areas

Multi-Topology Routing

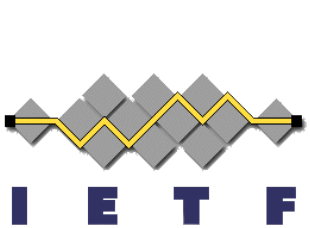

- OSPF and IS-IS Topologies are defined by metrics on links between router interfaces within the routing domain
  - The link does or does not have a metric within the topology
  - Automatically routes around discrepancies between physical and logical topology
- Inter-area and intra-area source/destination routing cases could be implemented as multi-topology
  - draft-xu-homenet-twod-ip-routing is multi-topology

Edge Routing as Multi-Topology Routing

- Edge routing is routing to a default route that is outside the routing domain
  - The IS-IS/OSPF topologies for each PA prefix are identical
  - There is no link advertised in IS-IS or OSPF that might have the indicated metric

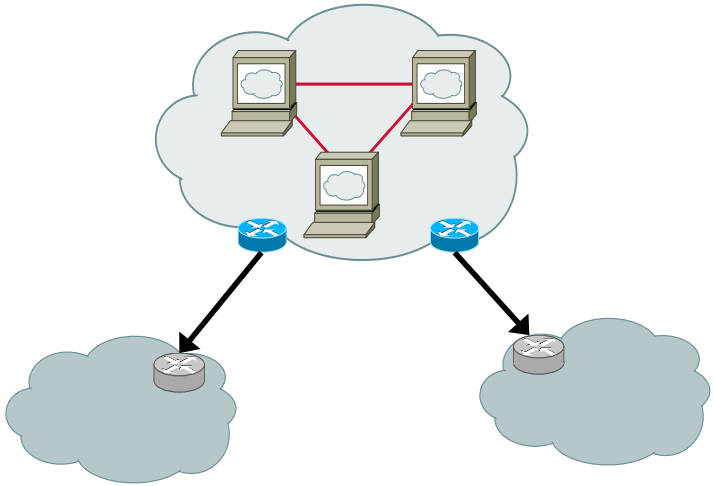

- **Edge routing is a *reachability* problem, not a *topology* problem**

IS-IS Context

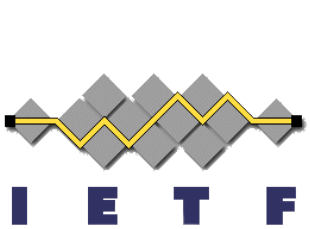

- RFC 5308 defines a Reachability TLV for reachable IPv6 Prefixes
  - It also defines a format for sub-TLVs, which it says may be of value in the future
  - Sub-TLVs add information to a routing decision

draft-baker-ipv6-isis-automatic-prefix

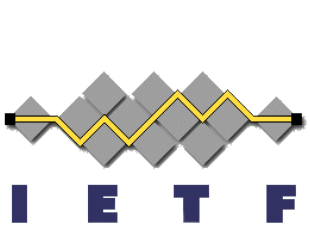

- Problem:
  - Homenet, for zeroconf, needs prefix distribution
  - OK, if I propose IS-IS, I need to solve that
- Concept:
  - A specific system, maybe the CPE router that has received a DHCP-PD prefix allocation, announces the prefix into a network
  - Routers (including pseudonodes) allocate a /64 at random from the prefix
  - If there is a collision, conflicting routers back off a random interval and guess again
  - If the TLV is withdrawn, they forget the derived IPv6 prefix

Autoconfiguration TLV

- Fields:
  - Type: IANA
  - Length of TLV
  - U/X as normal
  - No need for sub-TLV flag
  - Prefix Length
  - Prefix, same format as in Reachability TLV

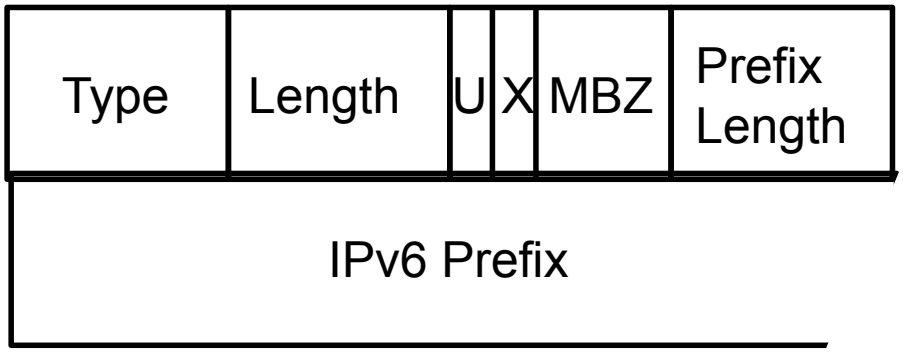

Flow label and Source Prefix sub-TLVs

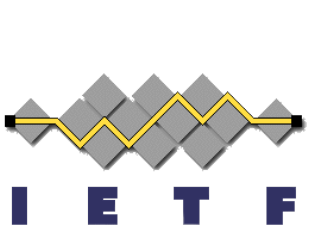

- Drafts I'm describing:
  - draft-baker-ipv6-isis-dst-flowlabel-routing
  - draft-baker-ipv6-isis-dst-src-routing
  - draft-baker-ipv6-ospf-dst-flowlabel-routing
  - draft-baker-ipv6-ospf-dst-src-routing
- Premise:
  - Reachability TLV, with sub-TLV(s), identifies a set of possible messages to send down a route
  - Additional qualifying information while calculating a route, and in the FIB
  - Need comments on route calculation and FIB design

Route Calculation

- Normal OSPF or IS-IS route calculation:
  - *Identifies a sequence of routers and links from calculating router to router advertising reachability*
  - "Router" might be an IS-IS pseudonode or OSPF Network LSA
- TLV, in this case, identifies a destination and a qualification
  - Traffic with a different source address or flow label follows a different route, or no route

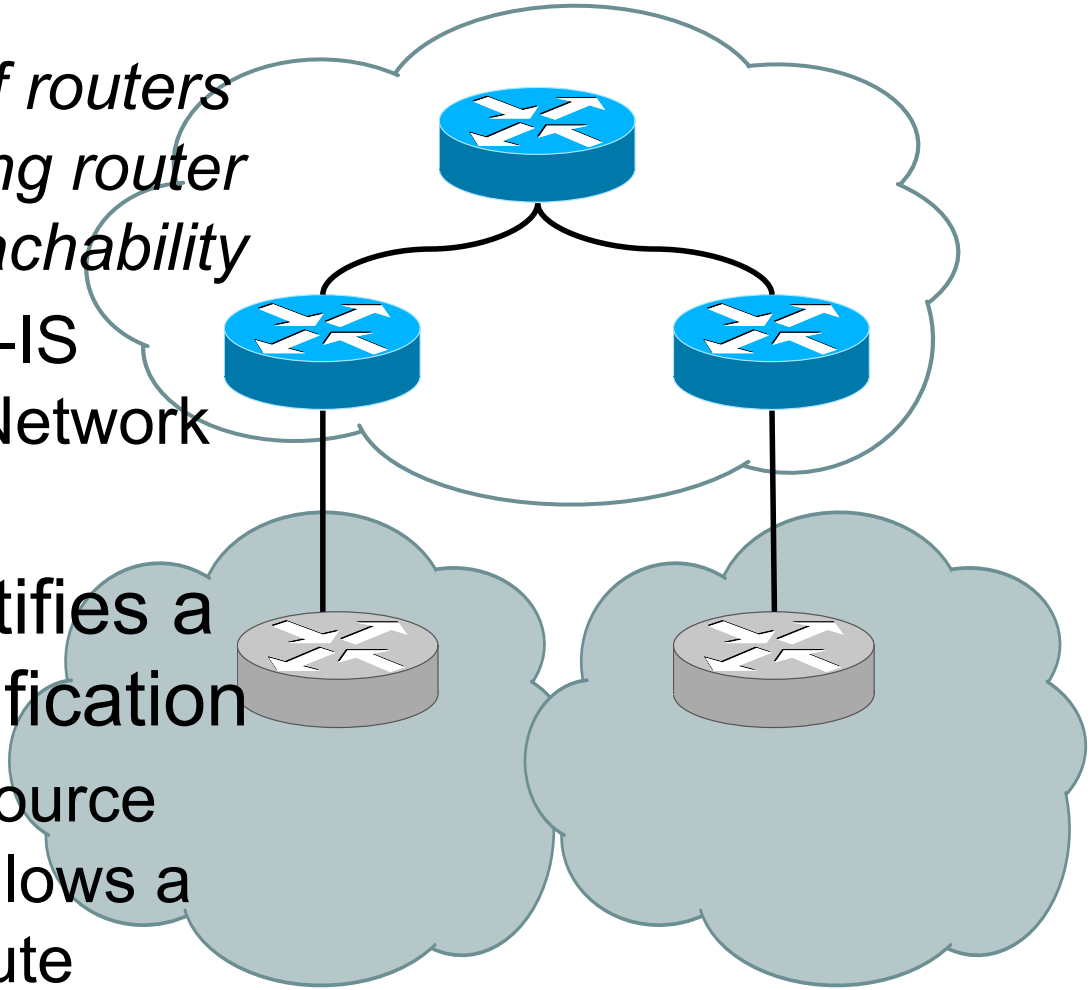

FIB Design

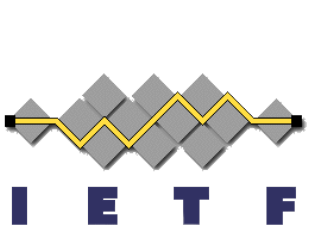

- Not subject to standardization.
- Some suggestions in an appendix
  - Linux (Waikato extensions) has separate FIBs by source prefix.
    - One could insert destination into appropriate FIB, or all FIBs if source not specified
  - PATRICIA tree
    - Allows a discontiguous bit string, differing don't-care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

“So Fred, what’s your problem with OSPF?”

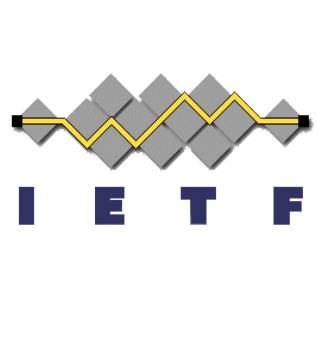

- OSPF (RFCs 2328 and 5340) defines fixed format LSAs for each purpose
  - As opposed to extensible TLVs as RFC 5308 does
  - It also defines separate LSAs for AS-external, intra-area, and inter-area prefixes
    - AS-external-LSA may have additional information beyond the prefix
  - That makes it hard to extend
- To extend OSPF, I need an LSA I can extend

draft-baker-ipv6-ospf-extensible

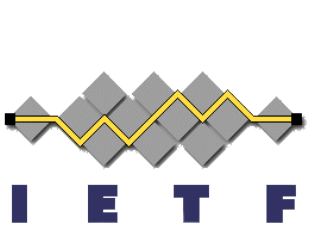

- I defined three extensible LSAs, replacements for intra-area-prefix-LSA, inter-area-prefix-LSA, and AS-external-LSA
- I have since been told of Abhay Roy's extensible LSA draft in draft-ietf-ospf-mt-ospfv3 (2007)
- I'll use whatever extensible technology the OSPF WG approves

Backward compatibility

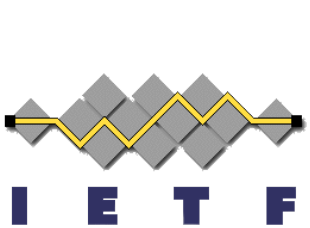

- OSPF WG asked about making this work in networks with RFC 5340 format LSAs as well
- Really not a problem:
  - Definition of source prefix sub-TLV:
    - A zero-length LSA (::/0) can be represented with a sub-TLV whose length is zero or no sub-TLV
  - Definition of flow label sub-TLV:
    - “any” flow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