

IS-IS: Homenet and Data Centers

Fred Baker



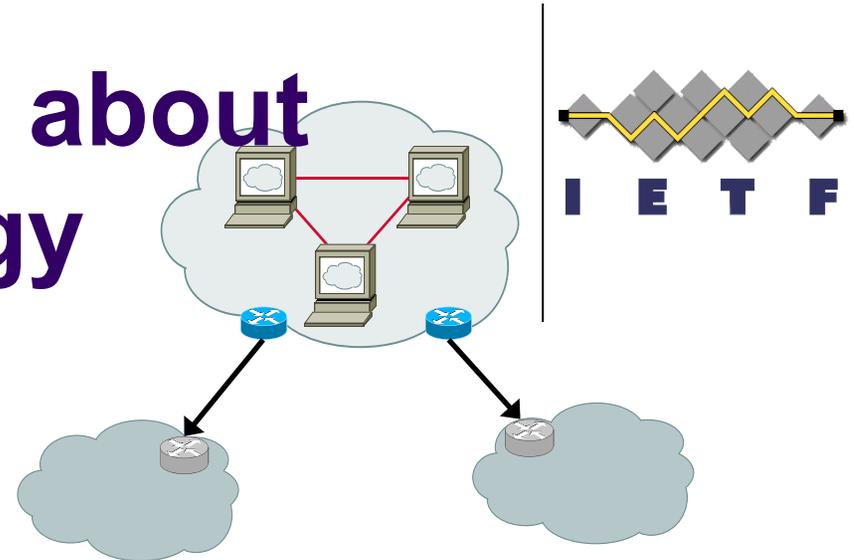


Homenet Requirements

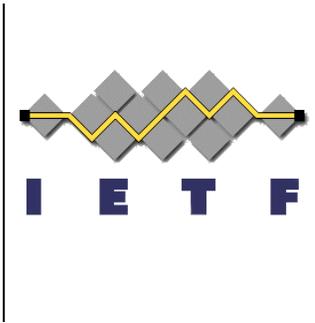
- Homenet is trying to develop supporting technologies for a very simple, but technologically advanced, home
 - Primarily focused on IPv6
 - Zero Configuration if at all possible
 - Interface to Smart Grid technologies including Zigbee/802.15.4
 - **Multi-subnet with routing an option**
 - **Potentially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

I have been asked about ISIS-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 routing cases could be implemented as multi-topology



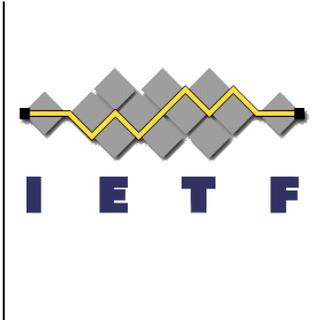
- **Edge routing is routing to a default route that is outside the routing domain**
 - The IS-IS topologies for each PA prefix are *identical*
 - There is no link advertised in IS-IS that might have the indicated metric
- **Edge routing is a *reachability* problem, not a *topology* problem**



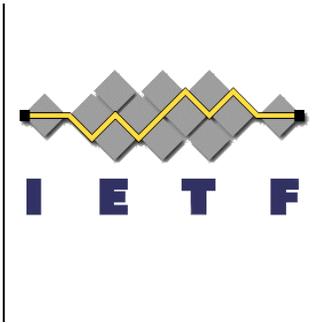
Context

- 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
- I'm looking at
 - Homenet requirements for egress routing
 - Homenet requirements for automated prefix allocation as in draft-ietf-ospf-ospfv3-autoconfig
 - Multi-tenant Data Center requirements for tenant-to-tenant access control

draft-baker-ipv6-isis-automatic-prefix

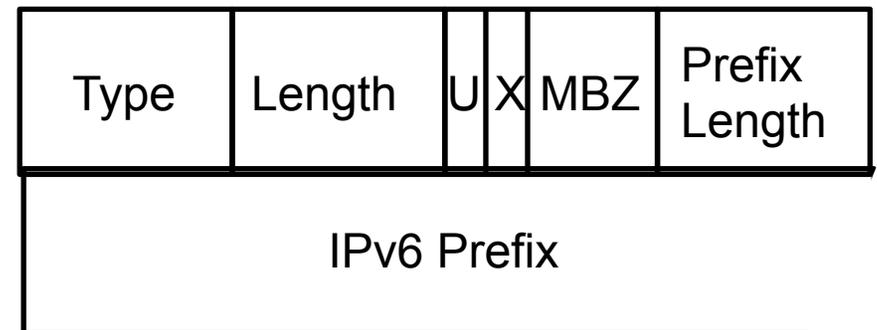


- Concept:
 - A specific system, maybe the CPE router that has received a DHCP-PD prefix allocation, announces the prefix into a network
 - Routers (including pseudonodes) allocate a /64 at random from the prefix
 - If there is a collision, conflicting routers back off a random interval and guess again
 - If the TLV is withdrawn, they forget the derived IPv6 prefix

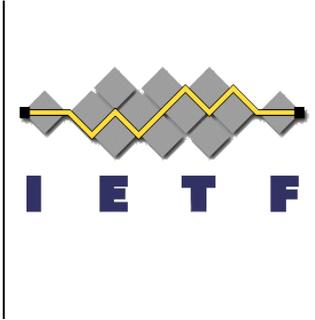


Autoconfiguration TLV

- 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 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

Backward compatibility

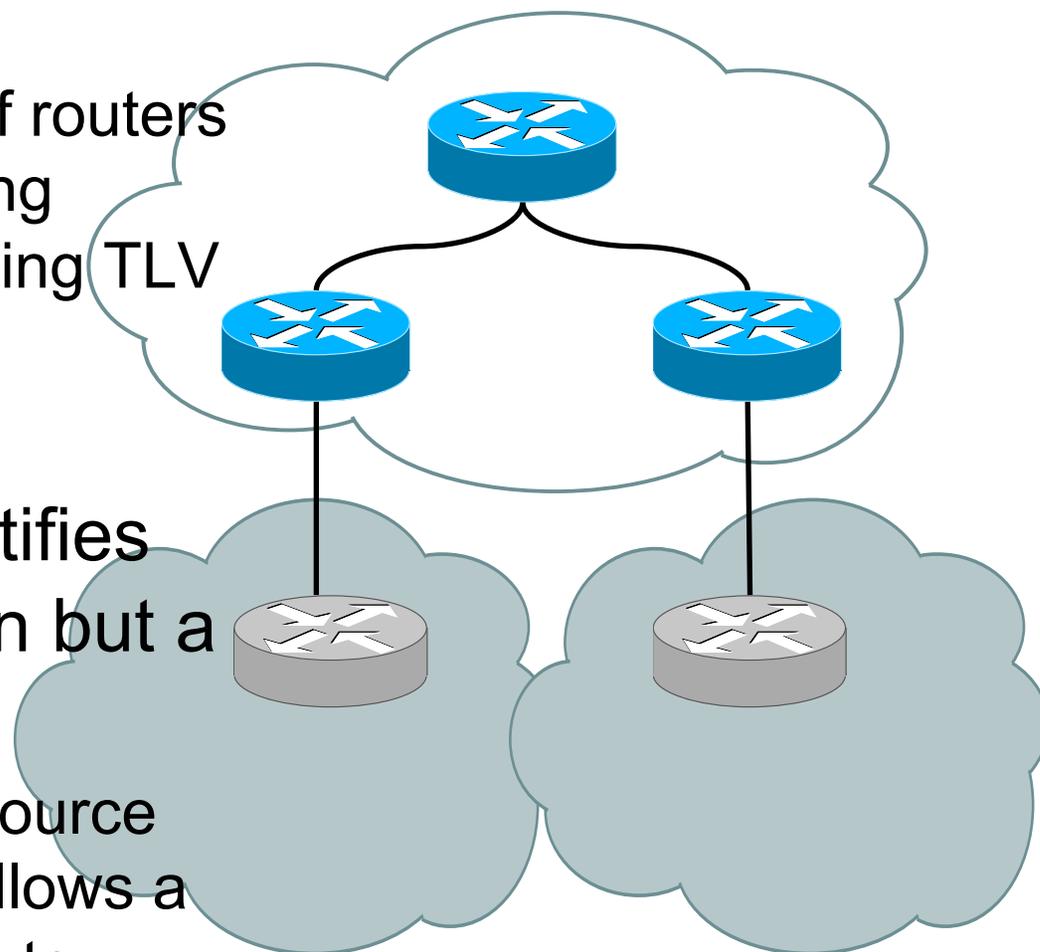


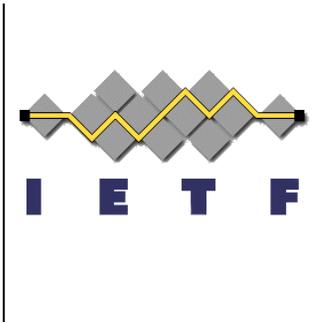
- OSPF WG asked about making this work in networks with RFC 5340 format LSAs as well; IS-IS probably has same question
- Really not a problem:
 - Definition of source prefix sub-TLV:
 - A zero-length LSA (::/0) can be represented with a sub-TLV whose length is zero or no sub-TLV
 - Definition of flow label sub-TLV:
 - “any” flow label is specified by leaving the sub-TLV out
- RFC 5308 TLV by definition leaves those sub-TLVs out. Semantically equivalent.

Route Calculation



- Normal IS-IS route calculation:
 - Identifies a sequence of routers and links from calculating router to router advertising TLV
 - “Router” might be a pseudonode
- TLV, 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





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”