Administratively Scoped IP Multicast

Van Jacobson
Steve Deering

30th IETF
Toronto, Canada
25 July 1994

Preliminaries

We currently use the TTL field in multicast datagrams for three completely orthogonal purposes:

- To limit the lifetime (in hops) of the datagram.
- To confine traffic to some topological region because of link bandwidth constraints.
- To confine traffic to some administratively defined topological region for privacy or scaling (address reuse) purposes.

This massive overloading of the TTL field isn't working very well.

We propose to perform the 3rd function by defining a region of the multicast address space as 'administratively scoped' and using the address rather than TTL to describe the scope region.

The Plan

A region of the multicast address space (say 239.0.0.0 to 239.255.255.255) will be designated as ‘administratively scoped’ by IANA. These addresses will have roughly the semantics of ‘local use’ Ethernet addresses – they are assigned locally by organizations and *are not guaranteed to be unique* across organizational boundaries.

The Plan (cont.)

A ‘boundary’ attribute would be added to the multicast routing interface config. This attribute would mark the interface as being part of the topological boundary for some range of addresses. E.g.,

```
tunnel foo bar boundary 239.255.0.0/16
```

says that this tunnel is a boundary for addresses in the range 239.255.0.0 to 239.255.255.255. The semantics of a boundary are that no packets in this range are forwarded *either direction* across the boundary.

Typical Use

To put a boundary around some region, you put a ‘boundary’ statement on every link that leaves the region.

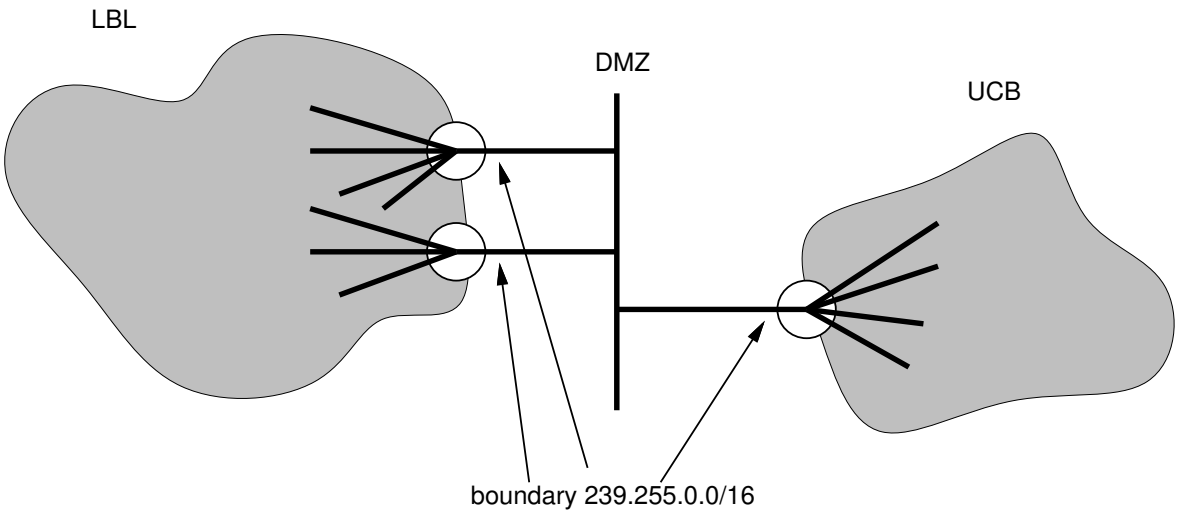

The ‘boundary’ on the LBL–DMZ links makes all 239.255.x.x addresses local to LBL. The ‘boundary’ on the UCB–DMZ link makes all 239.255.x.x addresses local to UCB.

Boundaries can overlap so one net can be part of multiple regions

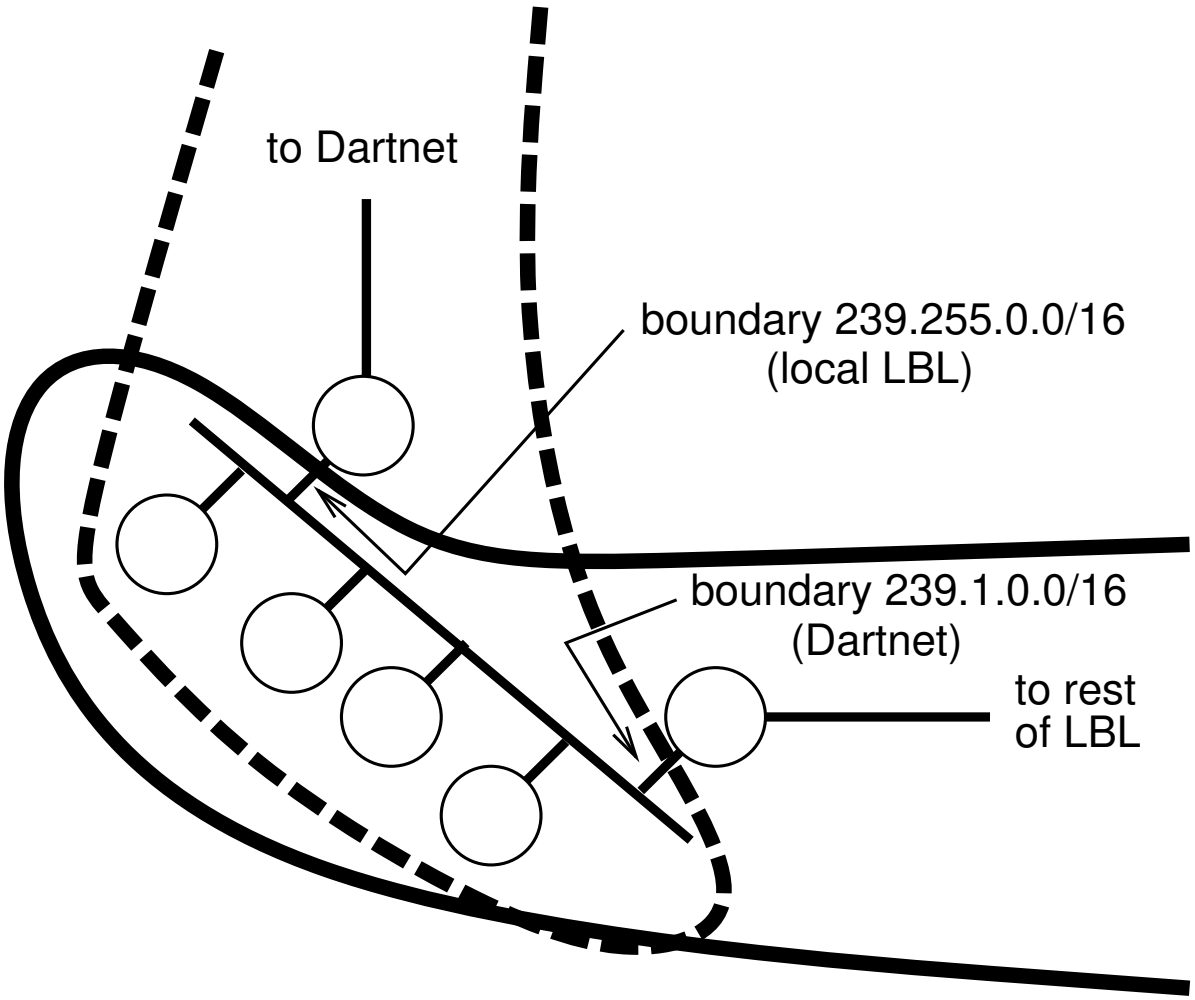

(Solid line shows ‘Local LBL’ scope region. Dashed line shows ‘Dartnet’ scope region. The ethernet segment shown is in both regions.)

Notes

- A boundary is *not* a firewall filter. It has both routing and forwarding semantics (e.g., a router always prunes the boundary).

- The boundary must be impermeable *both directions* across the interface or its primary function, address reuse, won't work.

  (The 'both directions' check required for boundaries on multiaccess links and prevents config errors on point-to-point links.

Topological constraints on boundaries

Bounded region must be *compact*. E.g., between any two nodes in region there must be a path that doesn't cross boundary.

Compact

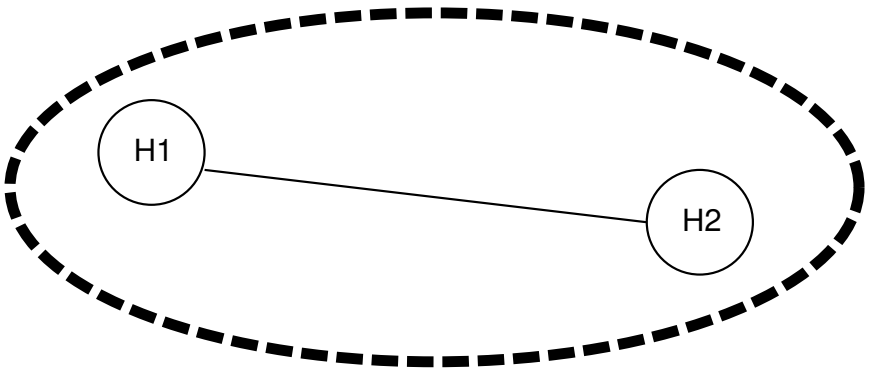

Not

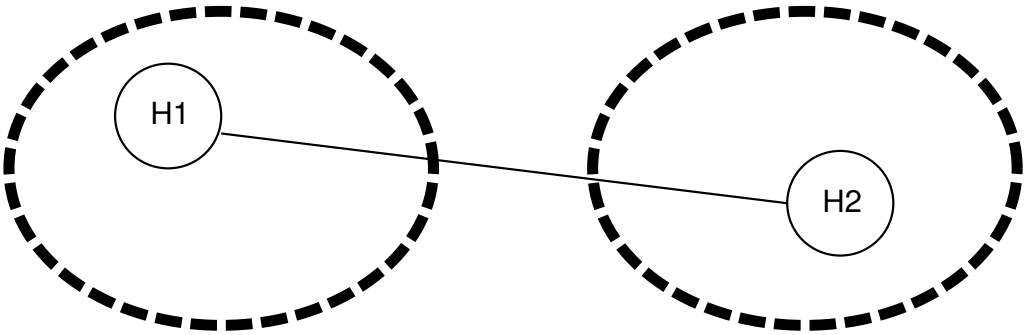

Topological constraints on boundaries (cont.)

To prevent black holes, bounded region should also be *convex*. E.g., no path between any two points in region should cross boundary.

Convex

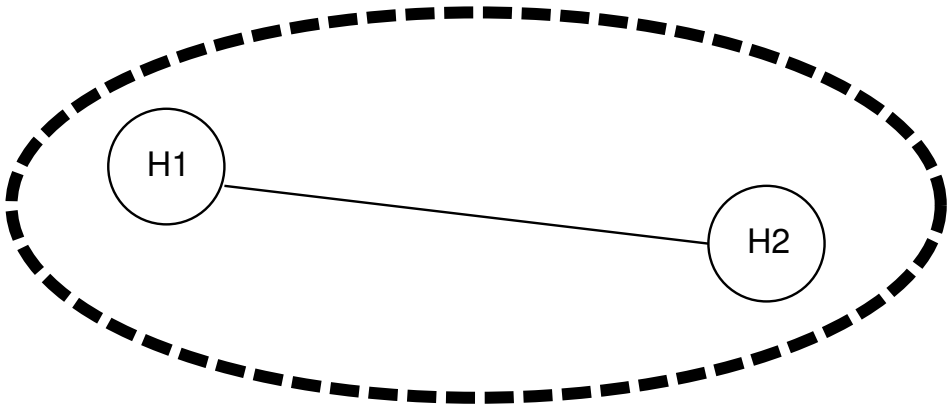

Not

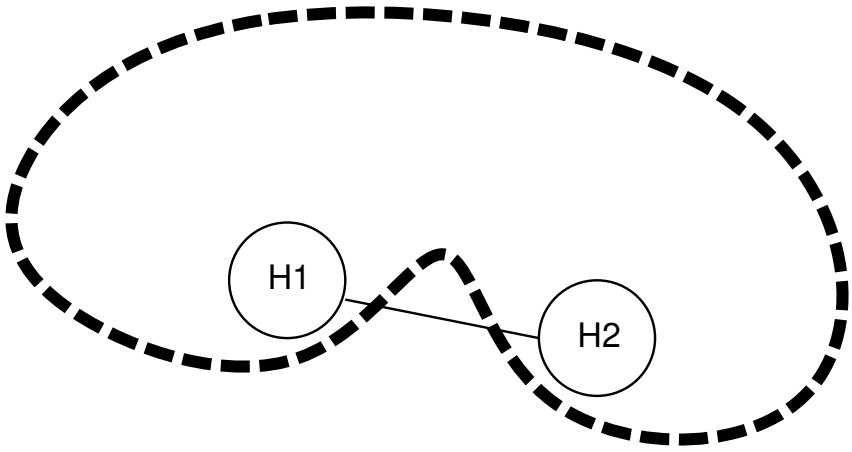

(Note that these same constraints apply to using TTL thresholds to create boundaries.)