Overlay Transport Virtualization (OTV)

Dino Farinacci
Hasmit Grover
Victor Moreno
Dhananjaya Rao

Introduction

- OTV is a L2/L3 Virtualization Solution for Enterprise environments
- Transparent L2 extension for enterprise sites
- L2 and/or L3 connectivity for site devices
- Multi-site multi-point connectivity
- Core transport infrastructure agnostic
- Extremely simple provisioning and management

Overview

- MAC Routing
  - Uses control plane advertisements instead of data plane learning
  - Remote site MACs learnt via routing protocol
  - No unknown unicast flooding through core
- Inter-site data encapsulated in IP
  - Routed across core to destination site
  - No pre-built tunnels

Overview

- OTV forms an overlay network across core
- Dynamically discovers member Edge Devices
- EDs exchange L2 routing information
  - Unicast MACs of hosts and routers in site
  - Active Multicast Sources in site
  - Interested Multicast Groups
- OTV functionality only in edge devices
  - Transparent to core and site devices
- STP terminated at each site

Overlay Network

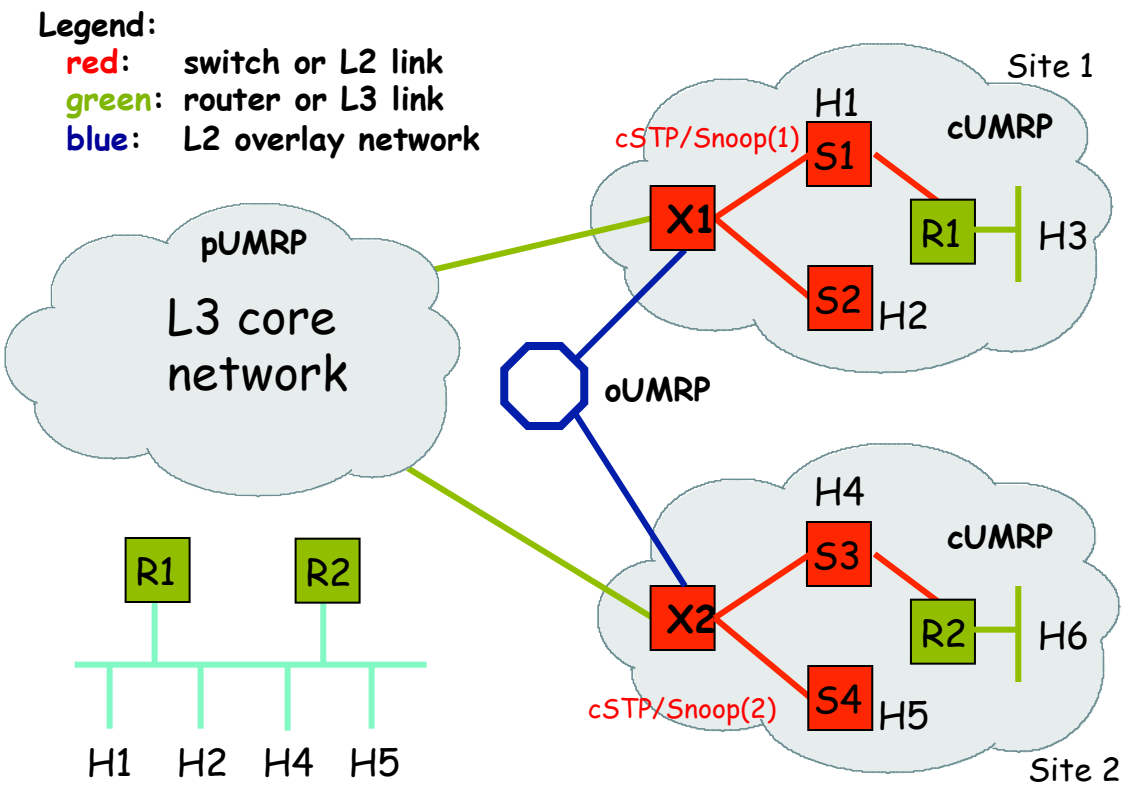

Note: Subnets span across all sites.
VLANs span across all sites.
Each site has its own Spanning Tree.
No L2 flooding or learning on overlay.

Data Forwarding

- Unicast data sent to "next-hop" EDs
  - Packets load-balanced across core ECMPs
- Multicast uses *Delivery Groups* across core
  - Source ED encapsulates site data in a (DS,DG)
  - Core optimally replicates to interested EDs
- Broadcast data sent as IP multicast
  - All Edge Devices join this core multicast tree

Multi-homing

- OTV provides loop-free multi-homing
- Authoritative Edge Device (AED) per site
  - Edge Devices in the site elect AED
- Only AEDs forward traffic on overlay
  - Avoids loops and duplicates
- Site traffic load-balanced among EDs
  - Per-VLAN AED

MAC Mobility

- MAC moves supported by control plane
  - MAC advertised with default metric
  - When MAC moves, ED in new site advertises MAC with lower metric to indicate MAC move
  - When original advertiser sees this, withdraws its own advertisement
  - New site ED then readvertises with default metric

OTV UDP Encapsulation

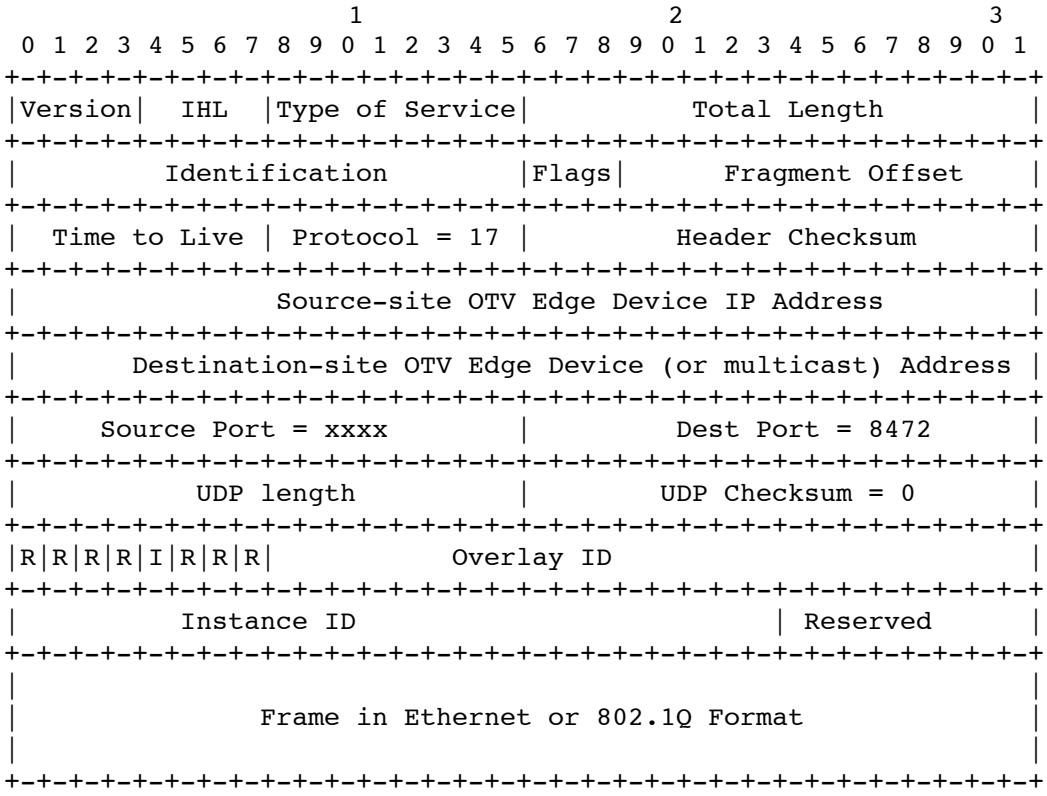

Overlay Routing Protocol

- Routing protocol for OTV control plane
  - Discovers overlay members
  - Forms adjacencies on overlay
  - Exchanges unicast and multicast routes
- IS-IS used as oUMRP
  - Overlay forms a logical LAN over the core
  - Edge Devices run IS-IS at L2 on overlay
  - Leverages Layer-2 IS-IS extensions

I-Ds

- Overlay Transport Virtualization

  http://www.ietf.org/id/draft-hasmit-otv-03

- IS-IS Extensions to support OTV

  http://tools.ietf.org/html/draft-drao-isis-otv-00

Comments ?

- Authors would like to solicit feedback and suggestions