Multicast Support for VPLS

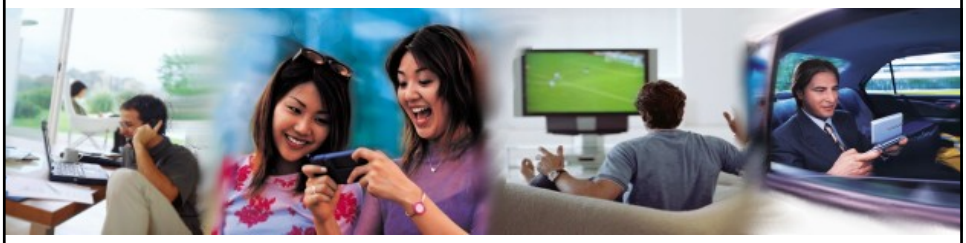

Marc Lasserre

Lucent Technologies
Bell Labs Innovations

Default VPLS Multicast Processing

VPLS Ingress Replication

- Ingress replication for
  - Broadcast
  - Unknown
  - Multicast
- Original design goals
  - Keep the VPLS core stateless
    - No need to run a multicast routing protocol
    - No need to build multicast trees
    - No need to maintain (S,G) state
  - No congruency issues
    - Between unicast & multicast paths

VPLS Multicast

Basic VPLS

- BUM traffic replicated by ingress PE to all PEs serving corresp. VPLS domain
- Without snooping, multicast traffic sent to non listeners
- With AC snooping, multicast traffic sent to listeners only

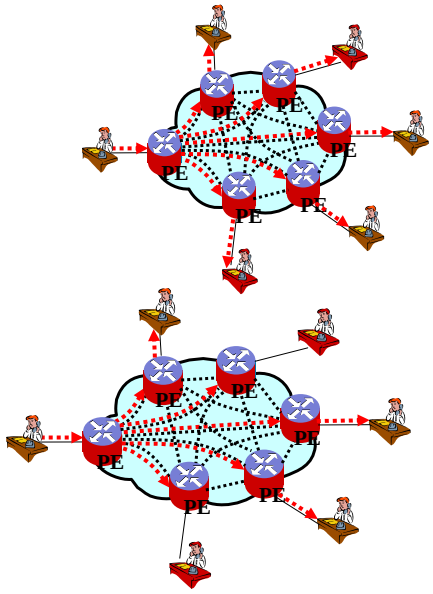

VPLS Multicast

Hierachical VPLS

- With HVPLS, replication distributed between PEs and MTUs
  - MTUs replicate traffic towards their ACs
  - PEs replicate traffic across VPLS PWs and their locally attached MTUs

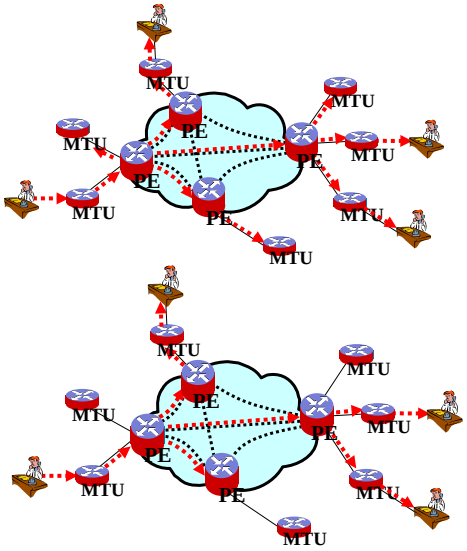

Optimizing VPLS Multicast

Multicast Optimizations

- Traffic Delivery
  - To receivers only
    - Tracking Joins/Prunes
- Bandwidth Usage
  - Minimize number of copies

Traffic Delivery Optimization

- IGMP/PIM snooping
  - Snooping on ACs not an issue
  - Amount of (S,G) state to be maintained is bounded
  - Snooping on core PWs can lead to a large amount of state to be maintained per PE
  - IGMP Snooping
    - Defined in draft-ietf-magma-snoop
  - PIM Snooping
    - Defined in draft-hemige-serbest-l2vpn-vpls-pim-snooping

Bandwidth Optimization

- So far, replication within metro networks has not been an issue
  - Simple topologies (rings or very few P routers between PEs)
  - Average number of sites per VPN typically small (between 5 and 20)
  - Hierarchical VPLS constructs distribute replication across multiple nodes
- With more complex topologies, use of p2mp LSPs leads to better bandwidth utilization

Multicast Optimizations Dependencies

- The amount of multicast traffic dictates:
  - Content location
    - Centralized vs distributed content
    - Core bandwidth usage
  - Snooping location
    - PE-rs
    - MTU-s
    - Access (e.g. DSLAM)
- Multicast optimizations depend upon network topology
  - Number of hops between source and terminating devices

Multicast Transport LSPs

- RSVP-TE p2mp extensions
  - draft-ietf-mpls-rsvp-te-p2mp
- LDP p2mp extensions
  - draft-minei-mpls-ldp- p2mp
  - draft-minei-wijnands-mpls-ldp-p2mp
    - Includes capabilities to set up p2mp & mp2mp trees
  - draft-boddapati-mpls-pim-ssm-ldp-p2mp
    - Uses a combination of PIM-SSM & LDP
      - PIM-SSM to build mcast trees
      - LDP to distribute labels
- mLDP over mRSVP_TE
  - draft-yasukawa-mpls-ldp-mcast-over-p2mp-lsps

Multicast & QoS

- Multicast applications often have strict QoS requirements
  - E.g. Broadcast Video, Video Conferencing
- RSVP-TE provides
  - Explicit path control
  - Resource reservation
  - Protection
- If mLDP were to be used, it would have to be carried over mRSVP-TE to meet QoS reqs

Multicast Options

- Trade-off between:
  - State maintained in the core
  - Optimization of bandwidth usage
  - Optimality of multicast routes
- L3 Multicast solutions aim to optimize b/w usage
  - draft-rosen-vpn-mcast
- L2 Multicast solutions aim to keep the core stateless
  - draft-hemige-serbest-l2vpn-vpls-pim-snooping
  - draft-ietf-magma-snoop
- Hybrid model
  - draft-ietf-l2vpn-vpls-mcast

VPLS Multicast Drivers

- Broadcast video/radio delivery
  - Carrier based services
  - Broadcast TV, HDTV
- Dedicated multicast streams
  - Business based services
  - Customer video feeds
    - E.g. Bank video advertisements in branch offices
  - Financial information
    - E.g Reuters, TIBCO
  - Video conferencing
    - E.g. NetMeeting

VPLS Multicast Options

Broadcast Trees

Shared Broadcast Tree

- One broadcast tree across VPLS instances
  - Rooted at each VPLS PE or mp2mp tree (shared tree)
- Used to carry all customers' bcast & mcast traffic
- Applicable to both L2 bcast/mcast and L3 mcast
- Minimizes amount of multicast state in the core
- VPLS/VC agnostic
- Requires support of *draft-ietf-mpls-rsvp-te-p2mp*

Suited for residential Broadcast Video/Radio delivery

Dedicated Broadcast Tree

- One broadcast tree per VPLS instance
  - Rooted at each VSI
- Used to carry one customer's bcast & mcast traffic
- Applicable to both L2 bcast/mcast and L3 mcast
- Minimizes amount of multicast state in the core
- VPLS/VC agnostic
- Requires support of *draft-ietf-mpls-rsvp-te-p2mp*

Suited for Business Broadcast Video delivery

Multicast Trees

Dedicated Multicast Trees

- Several trees per VPLS instance
  - Rooted at each source
- Used to carry efficiently customer's specific mcast traffic
- Applicable to IP mcast only
- Requires support of *draft-ietf-mpls-rsvp-te-p2mp*
- Requires a discovery procedure of multicast membership in core
  - To map (S,G) to correct multicast tree

Suited to business customers with multiple multicast streams with high b/w requirements

Aggregate Multicast Trees

- Use of p2mp trees to a defined set of PEs across VPLS instances
- Requires label coordination (upstream allocation)
  - Per mcast group VC label for demultiplexing
- Only applicable to IP traffic
- Requires a discovery procedure of multicast membership in core
  - PIM/IGMP snooping
  - "signaling" protocol to advertize membership (LDP or BGP)

Suited to business customers that need to exchange multicast streams

Discovery Protocol for Multicast Trees

- Several options available
  - Use of IGMP/PIM snooping on core PWs
  - Use of LDP extensions to carry mcast membership information
    - *draft-qiu-serbest-l2vpn-vpls-mcast-ldp*
  - Use of BGP or PIM as defined in
    - *draft-ietf-l2vpn-vpls-mcast*

VPLS Dataplane Changes

- Use of default p2mp tree instead of ingress replication for:
  - All customer broadcast and multicast *data* traffic
    - Customer broadcast & mcast *control* traffic still ingress replicated
- Multicast trees
  - Use of multicast FECs to map customer mcast traffic to appropriate multicast trees
- Aggregate multicast trees
  - Encoding of mcast VC label

VPLS Control Plane Changes

- Broadcast Tree requirements
  - IGMP/PIM snooping on ACs
  - RSVP-TE(/mLDP) multicast extensions
- Add'l requirements for Multicast Trees
  - IGMP/PIM snooping on PWEs for multicast trees

  or
  - BGP/PIM/LDP mcast state signaling over PWEs
  - PIM support in Ps

Conclusion

- Various degrees of complexity to optimize bandwidth usage
  - From simple broadcast trees
  - To more complex multicast trees
- Broadcast Trees require minor extensions to VPLS and suffice for main applications
- Will the extra b/w savings from multicast trees outweigh operational complexity?

Q & A