L2 VPNs.

Pseudowires.
Virtual Private LAN Services.
Metro/Carrier Ethernet.

Petr Grygárek

Layer 2 VPNs

Usages of L2 VPNs

- Server farms/clusters and other L2-dependent applications
  - redundancy and load-balancing implementations dependent on L2 connectivity (single broadcast domain)
- Virtual leased lines
  - Including potential L2 protocol conversion between customer sites
    - e.g. Ethernet - Frame Relay
- Virtual Private LANs (multipoint)
- Overlay networks with customer routing separated from the ISP routing

Comparison of L2 and L3 VPNs (1)

- Information used by ISP to forward packets/frames (L3 or L2 headers)
- Level of customer's control of the routing
  - Customer routing may be integrated or independent on ISP routing

Comparison of L2 and L3 VPNs (2)

- IP-specific (L3) or multiprotocol (L2)
  - GRE may help to carry L2 traffic over L3 tunnels
- Access technology
  - any IP-based line (L3) or specific L2 technology

Note that L3 VPN prevails

- 80% of ISPs' services

Most Common Implementations of L2 VPN Tunnels

- EoMPLS
- L2TPv3
- GRE

L2VPN Services (1)

- Pseudowires
  - P2P, Muxed or unmuxed UNI
    - Muxed UNI allows to terminate multiple (separate) VCs on the same physical interface
    - Muxed UNI possible if L2 framing differentiates between traffic flows
      - 802.1q, FR, ATM
  - Various framing options
    - Ethernet (including 802.1q)
    - Frame Relay
    - HDLC, PPP
    - ATM (AAL5 and Cell Relay)

L2VPN Services (2)

- Virtual Private LAN Service (VPLS)
  - Ethernet Relay
  - Muxed or unmuxed UNI
    - With muxed UNI, user can connect to multiple VPLS instances

L2VPN service classification does not dictate how is the service implemented in the SP core network (EoMPLS, AToM, QinQ, ...)

Any Transport over MPLS (AToM)

- Specifications
  - draft-martini-l2circuit-trans-mpls-07.txt: Transport of Layer 2 Frames over MPLS
  - draft-martini-l2circuit-encap-mpls-03.txt: Encapsulation Methods for Transport of Layer 2 Frames over MPLS
- AToM Technical Overview
  - http://www.informit.com/library/content.aspx?b=Troubleshooting_VPNs&seqNum=61

AToM Usages and Advantages

- Provides traditional L2 connectivity using MPLS core
  - FR/ATM/HDLC/PPP circuits
  - Transparent to users
- All techniques of MPLS TE and MPLS QoS may be applied to reach desirable characteristics of pseudowires
  - Allows the provisioning of QoS-aware virtual leased lines
  - 802.1p, FR DE and ATM CLP may be mapped to MPLS EXP bits

L2 Protocols Supported by AToM

- Ethernet (including 802.1q)
- ATM AAL5 PDUs + OAM cells
- Frame Relay + LMI
- ATM Cell Relay
- PPP
- HDLC
- Protocol Interworking
  - e.g. FR-VLAN
    - See example at http://www.debugall.co.uk/2009/08/03/frame-relay-to-vlan-interworking-atom/

AToM Operation

- Frames encapsulated with 2-level label stack
  - Transport label identifies egress PE
  - VC label identifies outgoing interface on the egress PE
    - Multiple VCs may exist between a pair of PEs
- Directed LDP session between PEs is used to distribute VC labels
  - New LDP TLVs to signal Label-to-VCID mapping and VC type were defined
- 2 unidirectional LSPs

AToM Control Word

- Carried after label(s) instead of the original L2 header
  - Special bits of original L2 headers
    - FECN, BECN and DE for Frame Relay
    - CLP for ATM
  - L2 header is reconstructed on the egress PE
  - May carry sequence number to avoid out-of-order frame delivery
    - Out-of-order frames are discarded
- Mandatory for FR and ATM AAL5, optional for other protocols
  - PEs use new LDP TLV to negotiate whether Control Words will be present

Virtual Private LAN Service (VPLS)

See also http://www.h3c.com/portal/Products__Solutions/Technology/MPLS/VPLS/200701/195598_57_0.htm

Virtual Private LAN

- Ethernet-based any-to-any communication over IP/MPLS core
- Simulates single Ethernet broadcast domain
  - virtual distributed switch that connects together customer's geographically dispersed LANs
  - Behaves as a “real” Ethernet bridge
    - self-learning of MAC addresses, flooding of frames with unknown addresses+broadcasts (+multicasts), MAC address withdrawal after topology change (new LDP TLV)
- Sites are connected by pseudowires (PW)
  - EoMPLS, L2TPv3
  - Much faster convergence in case of failure (LSP rerouting) comparing with STP

VPLS Advantages

- For service providers:
  - May provide a new QoS-aware L2 service on the existing MPLS core
  - Flexible bandwidth allocation
    - Compare with core composed from 100Mb/1 Gb/10Gbps Ethernet links
- For customers:
  - Simple and well-known Ethernet technology
  - The same technology in the carrier network and in customer's LAN

Creation of Virtual Distributed Ethernet Switch

- Full mesh of pseudowires between PE routers
  - PWs signalled using BGP or directed LDP
- Control plane
  - Autodiscovery – finding other routers participating in the same VPN – BGP only
  - Signalling – process of establishing pseudowires – BGP or LDP
  - BGP (RFC 4761)
  - LDP (RFC 4762)
  - Other autodiscovery protocols (DNS, …)

Pseudowire Implementation

- stack of two MPLS headers
  - Outer (transport) label identifies target PE
  - Inner label identifies pseudowire
    - PEs associate it with particular VPLS instance (Virtual Switching Instance)
      - A local switching table related to particular virtual distributed switch
      - Similar concept as VRF
    - Multiple VSIs may exist on the same router
      - customer separation

VPLS Forwarding Loop Avoidance

- A frame received from one PE is never forwarded to another PE
  - only to attachment circuits (to CEs)
  - analogy of Split Horizon rule
  - requires full mesh of PWs
- Spanning Tree may be applied as an alternative
  - not recommended

Problems of VPLS Scaling

- Full mesh of PWs between PEs is needed
  - The same is true for control plane
    - route reflector may help for signalling via IBGP
    - a static configuration of LDP directed sessions is always unscalable
- Signalling and packet replication overhead
- A solution is to establish a hierarchy, i.e. divide a VPLS VPN into 2 tiers
  - Multiple customers are aggregated in 2-nd level and connected to the same PE router

Hierarchical VPLS

H-VPLS

- 2-tier architecture
  - analogical to a star topology of spoke switches connected to a core switch
- High-performance core tier
  - Limited number of PEs
  - Full mesh of virtual circuits
  - Packet replication occurs only in the core
- MPLS or (cheaper) QinQ Ethernet-based access tier in POPs
  - U-PE faces to the customer
  - N-PE faces to the core

H-VPLS Advantages

- Limited size of the PW full-mesh in the core
- Cheaper QinQ-based Metro Ethernet technology in POPs' access networks
- Expansion of POP network does not require configuration change of core PEs

802.1q and MPLS Tags in H-VPLS

- Customer tag
  - Optional, for customers that needs to transport 802.1q-tagged traffic
- Service-provider tag
  - Appended by ingess QinQ access-layer Ethernet switch
  - Converted to (inner) MPLS tag on ingres core PE router
    - Identifies VFI on the target PE router
- Transport tag
  - Identifies egress core PE router

Metro Ethernet (Carrier Ethernet)

Metro Ethernet Forum

- Industry alliance
  - manufacturers of ME provider devices
- Defines
  - L2 services delivered over native Ethernet-based metro networks or other transport technologies (like MPLS/IP)
  - Technologies of carrier-class Ethernet-based transport networks
    - Architectures, Ethernet OAM extensions
- Develops technical specifications for Carrier Ethernet implementations and interoperability (MEF standards)

Ethernet Operation, Administration, and Management

- Necessary for provider-class Ethernet-based links
  - WAN links, Metro Ethernet
- Virtual Circuit Connectivity Verification, Label Switched Path ping, E-LMI etc.
- See http://www.cisco.com/en/US/prod/collateral/routers/ps368/prod_white_paper0900aecd804a0266.html for more details

Metro Ethernet Network Terminology

- User to Network Interface (UNI)
  - Demarcation point between CE device and MEN
  - Uses standard 802.3 PHY and MAC
- Ethernet Virtual Connection (EVC)
  - Connects 2 or more subscriber UNIs
  - P2P or multipoint
- Bundling
  - 2 or more customer VLANs mapped into a single EVC

Metro Ethernet Services Classification

- P2P or multipoint service
- Multiplexed / non-multiplexed UNI

2 x 2 service options give 4 services types in total

Metro Ethernet Service Types (1)

- E-Line - P2P
  - Ethernet Private Line
    - Dedicated UNIs (single EVC per UNI)
  - Ethernet Virtual Private Line
    - Multiplexed UNIs allow customer to connect to multiple EVCs by a single physical line
    - Replacement of FR and ATM
- E-LAN - multipoint L2 VPN
  - Ethernet Private LAN Service
    - Dedicated UNI
  - Ethernet Virtual Private LAN Service
    - Multiplexed UNI

Metro Ethernet Service Types (2)

- E-Tree – P2MP services (broadcasting)
  - Ethernet Private Tree Service
  - Ethernet Virtual Private Tree Service
  - Restrict communication between leaves

ME Service Framework (Service Attributes)

- Characteristics of the service are defined by attributes
  - Does not prescribe the way how the ME core implements the desired behavior
  - Serves as contract specification between customer and service provider
- UNI Attributes
- EVC Attributes
- L2 Control Processing Attributes

UNI Attributes (1)

- UNI ID (arbitrary string)
- Speed (10/100/1000,...)
- Duplex mode
- Service multiplexing
  - multipexed/dedicated UNI
- Ingress Bandwidth Profile
  - Per-UNI, per-EVC, per-CoS
  - CIR, EIR, Bc, Be

UNI Attributes (2)

- CE-VLAN-ID to EVC mapping
  - Customer's 802.1q tags may be either preserved, rewritten or removed
  - All VLANs may be bundled into one EVC

EVC Attributes

- EVC ID (arbitrary string)
- EVC Type (E-Line/E-LAN)
- CE VLAN Preservation (Yes/No)
- CE CoS Preservation (Yes/No)
- Unicast/Multicast/Broadcast frames delivered
- EVC Performance – QoS parameters
  - availability, delay, jitter, frame loss

L2 Control Processing Attributes

- Define how L2 control protocols are tunneled over MEN or interact with control protocols in the MEN core
  - STP, 802.3x, LACP, 802.1x, GARP, proprietary protocols (PAgP, VTP, CDP, ...)
- Processing Options:
  - Pass
  - Discard
  - Peer

Special Capabilities of Metro Ethernet Devices

- Advanced manipulation with 802.1q headers
  - push/pop/match+rewrite
  - works with single tags or with sequences of tags
- ME switches allow to divert a group of VLANs from a trunk to a specific port (Flexible QinQ)
- Ethernet OAM