FOUNDRY NETWORKS

Multi-Provider Ethernet Service Delivery

Ananda Rajagopal

Product Manager, Service Provider Solutions

Foundry Networks

THE POWER OF PERFORMANCE™

Agenda

- Multi-Provider Ethernet Service Delivery
- Ethernet Service Delivery Today
- A Proposed Model
- Comparison with other approaches
- Benefits
- Challenges
- Q&A

Multi-Provider Ethernet Service Delivery

Is There a Need?

- Major drivers for large-scale Ethernet Service deployment across multiple provider networks:
  - Increasing globalization of customers
  - Limited reach of providers: lots of local providers more common than global titans
  - Need to simplify inter-provider hand-offs
  - Business continuity reasons require enterprises to connect critical sites to a second provider
  - Integration of networks after acquisition
    - The Dow 12000 Era!
  - Multi-provider service delivery is/was common in Frame Relay and ATM deployment
    - Customers exploring Ethernet service would have similar expectations

Ethernet Service Delivery Today

- Predominantly point-to-point (E-LINE)
- Multi-point services (E-LAN) delivered as “Transparent LAN Services” (TLS)
  - Limited scalability of 4k VLANs
  - Complexity in management
  - VPLS picking up
- Fast-growing market; yet small compared to alternatives
  - Majority of services continued to be leased line, FR or ATM
- Geographic coverage continues to be the major challenge and obstacle to increasing adoption of Ethernet service. Obstacles include:
  - Reach of fiber
  - Reach of the provider

A Proposed Model

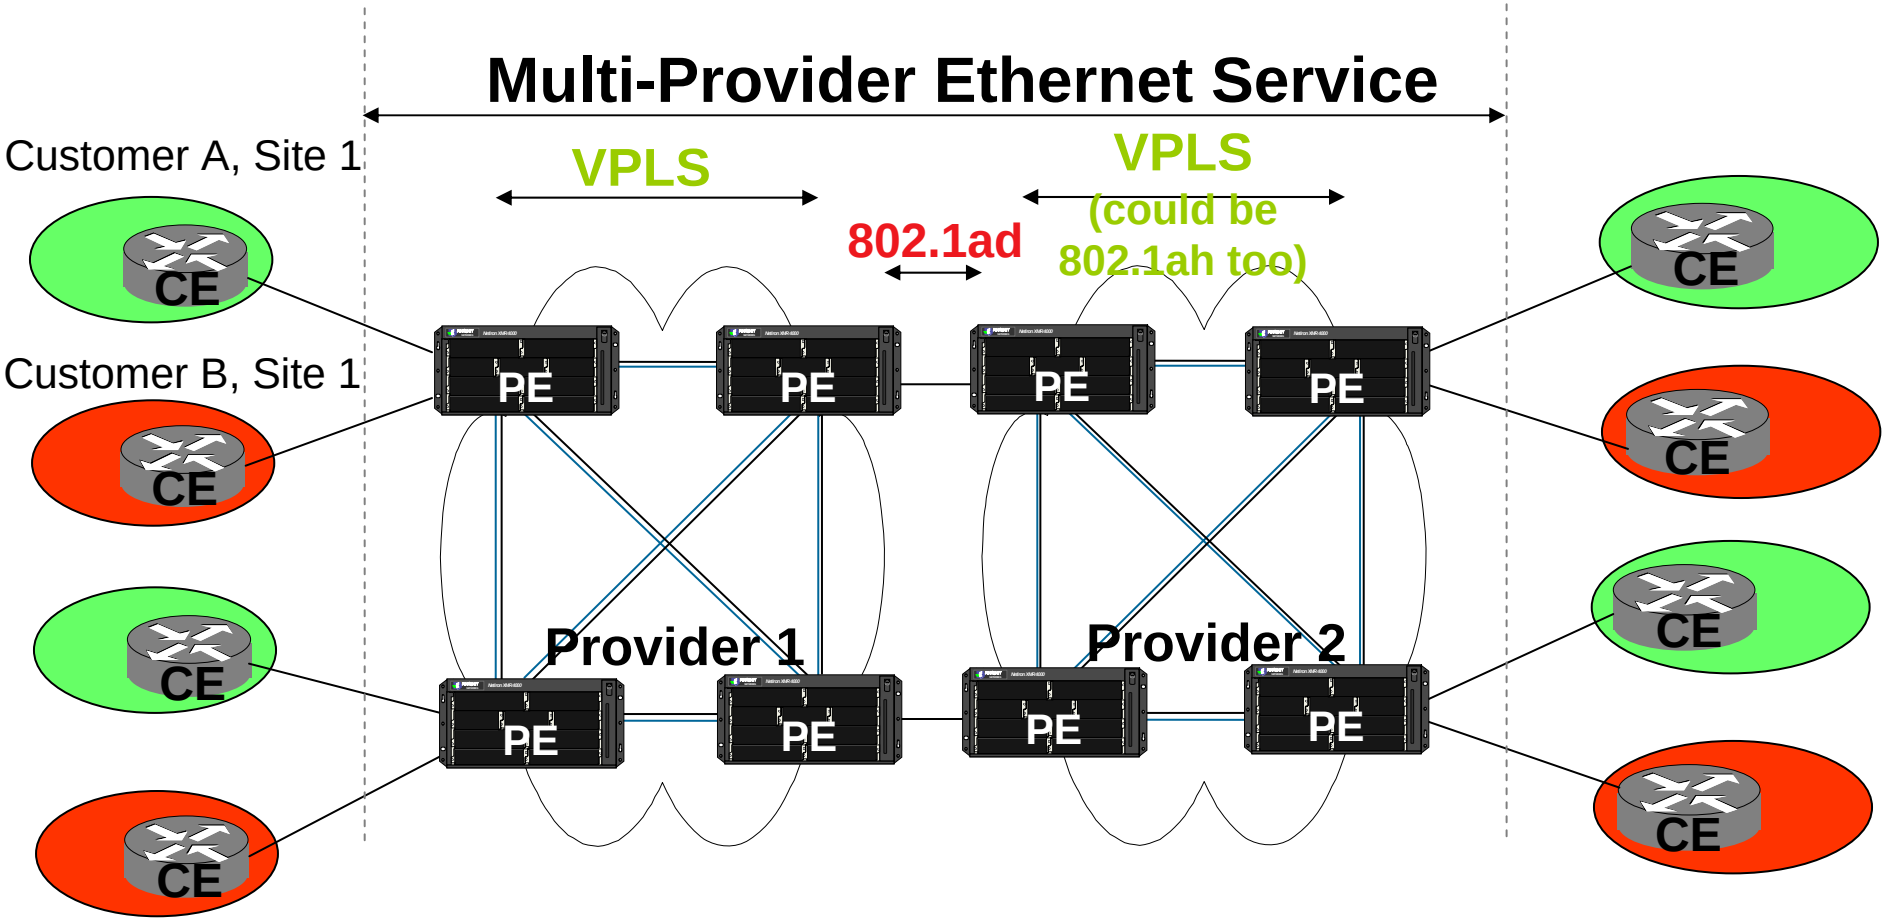

CE: Customer Edge
PE: Provider Edge

The Model in Words …

- Each provider responsible for its network and for "simple" hand-offs to the adjacent provider's network
- Each provider independently chooses the appropriate technology to transport Ethernet service over that domain (e.g. VPLS or VPWS, 802.1ah, L2TPv3 etc)
- Hand-offs to other provider are done using 802.1ad as an E-NNI (External Network-Network Interface)
  - S-VID translation is used to translate the service instance ID if required
  - <S-VID + C-VID> mapping to the partner provider's ELINE or ELAN segment may also be possible

Role of “Boundary” Nodes

- Packets from neighboring provider network:
  - Map the 802.1ad packets to the appropriate tunnel in the provider domain
  - E.g.:
    - Map 802.1ad to a VPLS instance
    - Encapsulate 802.1ad frame within an 802.1ah frame
- Packets going to neighboring provider network:
  - Terminate the local tunnel
  - Encapsulate in 802.1ad and send to neighboring provider

802.1ad Frame Format

DMAC	SMAC	0x88a8	S-TAG	0x8100	C-TAG	Payload

Tagged frame format

C-TAG format

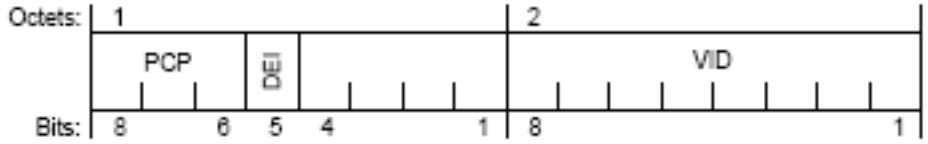

S-TAG format

Traffic Engineering and QoS

- Each provider network determines the best path *within* that network
- Optimized within each domain
- QoS is honored within each provider network
- At inter-provider boundaries, EXP<->802.1ad PCP mappings are done to preserve end-to-end QoS

Redundancy

- Resiliency mechanisms used by each provider independent of what other providers may be doing
  - E.g. MPLS Fast Re-Route (FRR) if VPLS/EoMPLS is in use within a provider's network
- Dual homing may be used for the provider-provider interconnect for inter-provider resiliency

Provisioning

- A provider just needs to:
  - Know the S-VID of the partner provider
  - Submit a request to the partner provider:
    - Partner provider's operations staff maps that S-VID/C-VID combination to a L2-VPN instance internal to their network.
- Wholesaling is an option
  - Do not care about the C-VID; S-VID acts as a "tunnel" that is transported across Provider B's network
  - In this case, "S-VLAN Bridge" functionality of 802.1ad (i.e. no C-VLAN component) is all that is required
  - Maybe simpler to manage based on traffic patterns
- Existing provisioning systems used by the provider can continue to be used

Troubleshooting

- Troubleshoot a reported incident within a provider domain first
  - Can I reach the endpoint of this service instance within my provider domain?
  - Within the domain: OAM utilities such as LSP ping/802.1ag etc can continue to be used to troubleshoot the L2-VPN instance
- Troubleshooting SLA violations would require more cooperation among the providers. Examples:
  - Round-trip delay exceeded
  - Connectivity loss for end-to-end service

Comparison With Other Approaches

- H-VPLS (with spoke VLL)
  - Expects that different provider networks through which the EVC is delivered all support VPLS
- Multi-segment pseudo-wire (MS-PW)
  - Still at the conceptual stage in IETF
  - Little vendor support today
  - Suitable for very large-scale deployments
  - Facilitates interconnection of different PW types
  - Method proposed in this presentation subset of MS-PW
    - Inter-Provider Switching Using Attachment Circuits

Benefits of This Approach

- Hierarchical solution to offering multi-provider services
  - Hierarchy aligned with organizational boundaries
- Uses technologies that are already standardized / well-known
  - 802.1ad ratified earlier this year
  - Avoids vendor lock-in and gives providers a range of options to choose from
- Technology used in Provider A's network independent of what is used in Provider B's network
  - E.g. Provider B may use 802.1ah instead of VPLS
  - What is important is the hand-off

Challenges (1)

- Inter-provider relationships
  - Who “owns” the customer?
  - Customer owned by the provider with whom first contact was made?
  - Ownership of geographical regions?
  - Revenue-sharing arrangements based on above criteria
    - Other thoughts?

Challenges (2)

- Customer (end-user) ownership
  - Provisioning
    - Should provisioning systems of providers be integrated to simplify end-to-end provisioning?
  - Billing and accounting
  - SLA guarantees
    - Certainly not trivial!
  - Troubleshooting an end-customer outage
    - Who takes responsibility for resolution?
    - Provider who sends the bill to the customer!

Thank You!