DCI for SP Cloud Services

Thamir Alhammad, talhamma@cisco.com

@talhamma

DCI for SP Cloud Services

Focus

- The primary focus of these slides is on DCI solutions for SP multi-tenant cloud services (e.g. Virtual Private Cloud), i.e. where SPs offer virtualized data center services to their enterprise customers via the SP VPN network

DCI for SP Cloud Services

Agenda

- Introduction
- VPLS
- EVPN
- OTV
- LISP

SP Multitenant DCI: Interconnection Models

- Standalone DCI network provides interconnection between main SP DCs
  - Owned by SP DC team
  - Addresses SP2SP only
  - Very high bandwidth – packet / optical solution likely the most cost effective
- NGN based DCI solution:
  - Addresses E2SP for workload migration
  - Addresses SP2SP for regional or distributed data centres

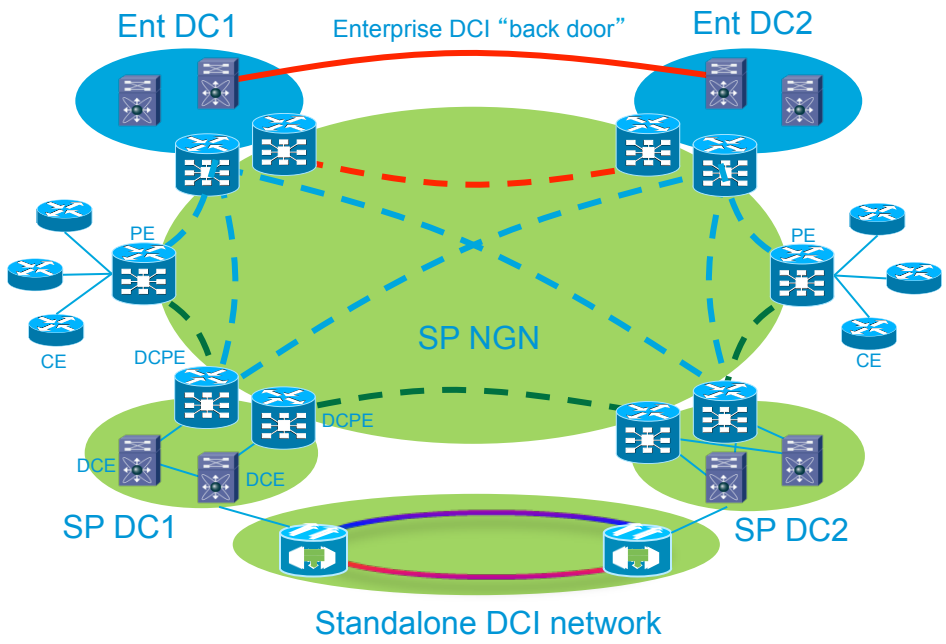

- Interconnection models:
  - Enterprise to Enterprise (E2E)
  - Enterprise to Service Provider (E2SP)
  - Service Provider to Service Provider (SP2SP)

SP Multitenant DCI: Baseline Requirements

- Scales to the level required for SP virtual private cloud
  - 100s of thousands to millions of MAC addresses per data centre
  - Thousands of tenants
  - 10s of data centres
- Optimally forward unicast and multicast
  - Shortest path
  - Loop free
  - Avoiding duplicates
- Uses network resources efficiently
  - All connections active with load balancing
  - Flood minimisation
- Provides control plane isolation between DCs
- Is resilient to all single element failures, i.e. in both NGN and DC
- Fast to converge
- Integrates with SP NGN, whilst honouring any administrative boundaries between DC and NGN
- Supports geo-redundant PEs, i.e Enterprise DCI “back door”
- Works for plain old spanning tree environments,. i.e. legacy Enterprise
- Easy to manage and operate

DCI:
Criteria to compare solutions

- Resiliency model: active:standby or active:active? Per interface, VLAN or flow?
- What transport is supported: IP, MPLS or both?
- NGN PE-based solution or overlay
- Whether requires an L2 or L3 service
- Standards status
- Underlying technology: MAC learning/bridging or MAC routing
- Applicability:
  - Enterprise to Enterprise (E2E)
  - Enterprise to Service Provider (E2SP)
  - Service Provider to Service Provider (SP2SP)
- Scalability: PW, C-MAC, ..
- Convergence
  - Both after failure and VM move
- Broadcast, unknown unicast and multicast handling
- Provisioning / configuration complexity
- DC transparent or not

VPLS

Virtual Private LAN Services (VPLS)

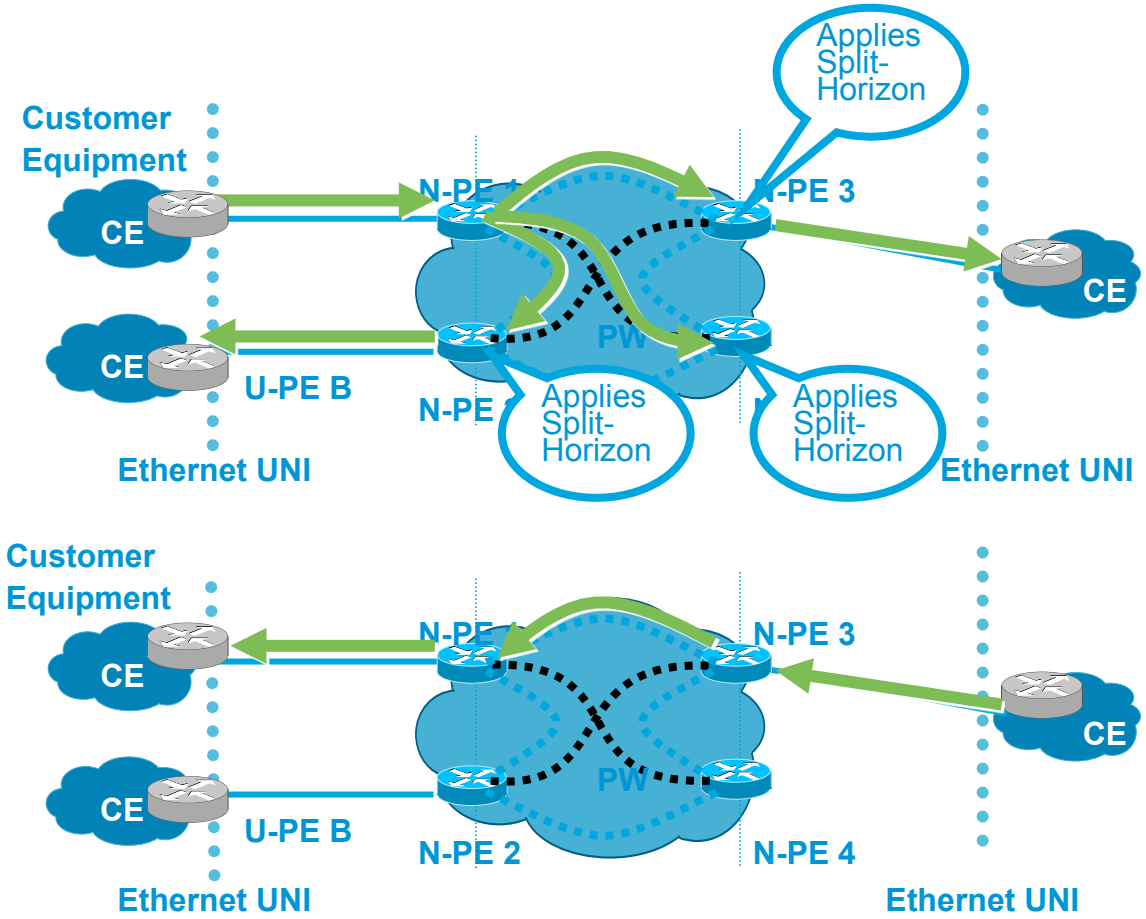

VPLS defines an architecture to provide connectivity between geographically dispersed customer sites as if they were connected using a LAN.

- Data plane mac learning
- VPLS uses split-horizon and full-mesh of PWs for loop avoidance in core
  - SP does not run STP in the core
- RFC 4761 (BGP-Based VPLS); RFC 4762 (LDP-Based VPLS)
- Hierarchical VPLS option for LDP based VPLS
  - 802.1ad based
  - MPLS PW based

VPLS Challenges

- Multihoming
  - Loop avoidance
  - Load balancing
  - Resilience / Failover
  - Geo-redundant PEs
- Scaling

Scaling VPLS

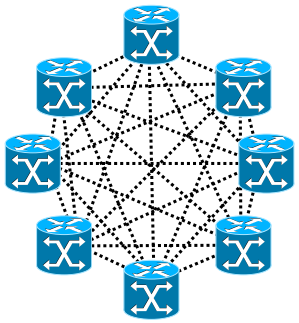

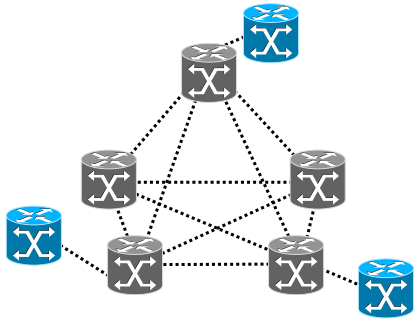

- Challenges
  - MAC-address scalability at the PE

    MAC address scalability is an issue on the DC WAN edge for solutions that use L2 extension

    Example: 250k CMAC addresses in a single SP data centre

    This is multiplied across interconnected SP data centres and Enterprise data centres

    MAC scale problem applies to EVPN also
  - PW scalability

    Full mesh per VSI per customer

    Scales O(n^2)
- Approach: PBB / 802.1ah / TRILL with VPLS
  - C-MAC hiding at the WAN edge node
  - Reduces # of PWs
  - http://tools.ietf.org/html/draft-ietf-l2vpn-pbb-vpls-interop-00

PBB-VPLS

- Leverages Provider Backbone Bridging (PBB) with VPLS
  - PBB = IEEE 802.1ah a.k.a mac-in-mac
  - C-MACs are tunnelled in CMACs
  - B-MACs can represent a PE, a PE linecard, or potentially a CE
- Benefits:
  - Significantly improves PW scaling
  - Can significantly improves MAC-address scalability at the PE

    Depending on the PBB model used

EVPN

EVPN

- Ethernet VPN
  - BGP MAC routing – BGP distributes client-MAC-address reachability and client-multicast group information between the edge devices
  - The coming together of Cisco R-VPLS and Juniper MAC-VPN drafts
- IETF Standards track:
  - http://tools.ietf.org/html/draft-sajassi-raggarwa-l2vpn-evpn-req-00
- Reuses several building blocks from existing BGP-MPLS technologies
- Requires extensions to existing BGP-MPLS technologies…

EVPN Concepts

1. EVPN PEs use data plane learning to identify the MAC addresses of locally connected CEs and hosts
2. EVPN PEs learn MAC addresses of remote CEs and hosts via the control plane using BGP
3. Route Targets (RT) are used to define the membership of a MAC VPN, i.e. by selectively importing/exporting RTs
4. MPLS is used for forwarding between PEs with a label stack:

   Inner label: EVPN label advertised by destination PE

   Outer label: Label for LSP to destination PE

OTV

What about OTV?

- Overlay Transport Virtualisation:

  ISIS distributes client-MAC-address reachability and client-multicast group information between the edge devices

  Mac-in-IP encapsulation

  IETF draft – http://tools.ietf.org/html/draft-hasmit-otv-01

- Not currently optimized for SP multi-tenant deployments – primary applicability today is for enterprise overlay deployments

LISP

What about LISP?

- Locator/ID Separation Protocol (LISP)
  - IP mobility solution rather than VLAN extension
  - IETF experimental: http://tools.ietf.org/html/draft-farinacci-lisp-12
- LISP currently addresses 2 use cases in the context of DCI
  - VM mobility
  - L3 tenant virtualisation, i.e. L3 VPN
    - Not envisaged to replace MPLS VPN but could potentially be used for intra-DC L3 virtualisation, i.e. In addition to VM mobility
- LISP does not currently address server clustering or other VLAN extension requirements
  - Although potentially could be supported in future
- Potential solution for DCI requirements if L2 extension not needed

Thank You