



Interdomain LDP-BGP VPLS Interworking

**Amit Shukla
Juniper Networks**

Background

- **Virtual private LAN service (VPLS) glues together several individual LANs across a packet-switched network to appear and function as a single LAN**
- **There are two standards for VPLS control plane that are deployed today:**
 - **BGP-based (BGP-VPLS):** uses BGP for auto-discovery and signaling of PWs (RFC 4761)
 - **LDP-based (LDP-VPLS):** uses LDP for signaling of PWs (RFC 4762), with discovery via provisioning
- **The data plane used by both these standards is same**
 - Flooding of unknown unicast, learning/aging MAC addresses, etc.
 - A key notion in the data plane is of split-horizon forwarding
 - MPLS pseudo-wire encapsulation

Business Drivers for LDP-BGP VPLS interworking solution

- To interconnect domains running a different VPLS control plane technologies
- To expand LDP-based VPLS service out of metro domain to WAN in a scalable and efficient manner.
 - By using BGP-VPLS.

Expanding VPLS service using BGP-VPLS


- **Why BGP-VPLS in WAN**
 - Scalable
 - Auto-discovery
 - Inter-provider support

- **Why LDP-VPLS in Metro**
 - Legacy


- **Proposed solution requires no changes in Metro network running LDP-VPLS**

Interconnecting Domains running different VPLS control plane technologies (BGP-VPLS and LDP-VPLS)

- Existing inter-AS VPLS options are not designed for ASes running different VPLS technologies (BGP-VPLS and LDP-VPLS)
 - Option-A can be used but doesn't scale.
- It is apparent that interworking function is required between BGP-VPLS and LDP-VPLS.

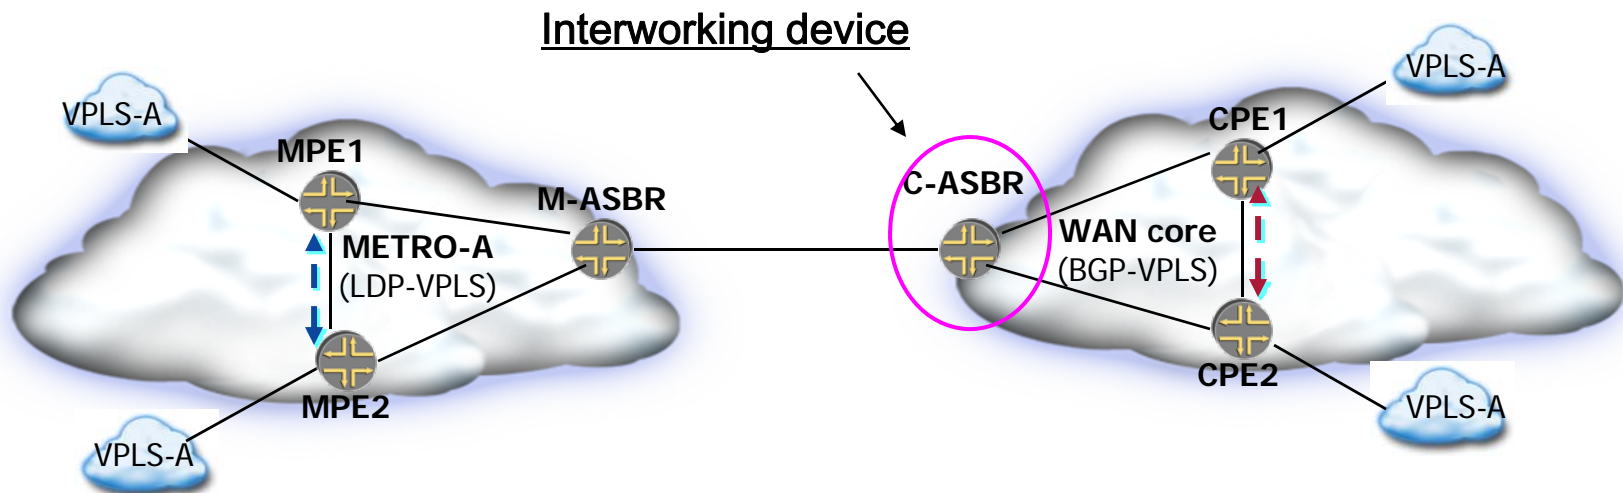


Extending the reach of LDP-VPLS metro-domain to WAN via BGP-VPLS to enable region/nation-wide VPLS service



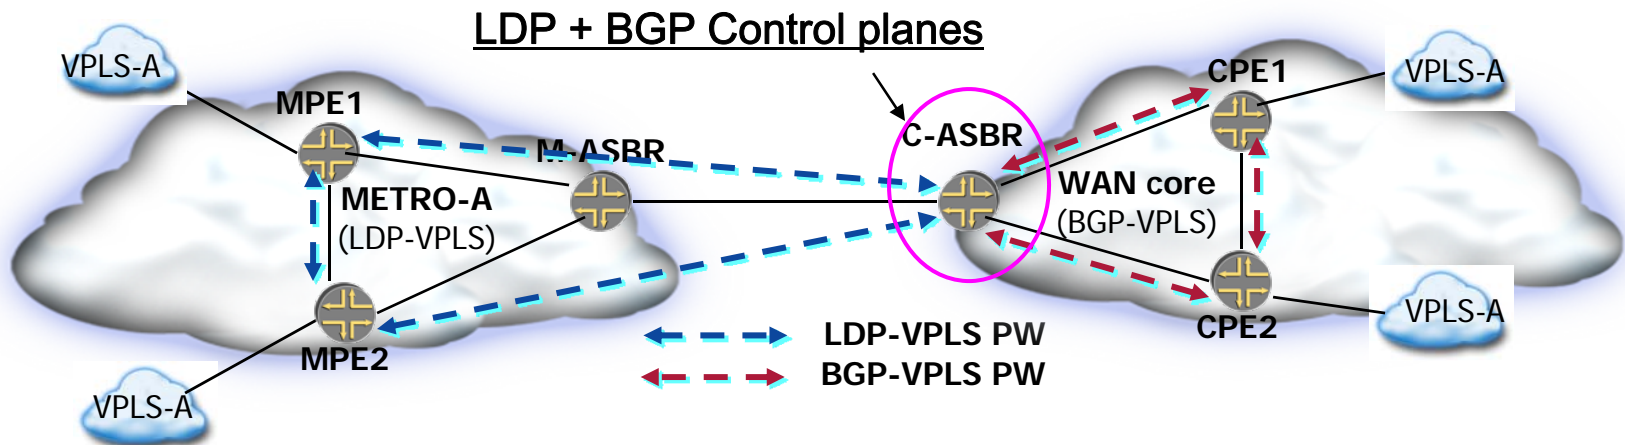
Overview of LDP-BGP VPLS interworking

- Interworking operation localized at WAN border-router (ASBRs)
- Proposed Interworking mechanism require no changes in both
 - LDP/BGP VPLS control plane specifications (and)
 - Current LDP-VPLS Metro network design



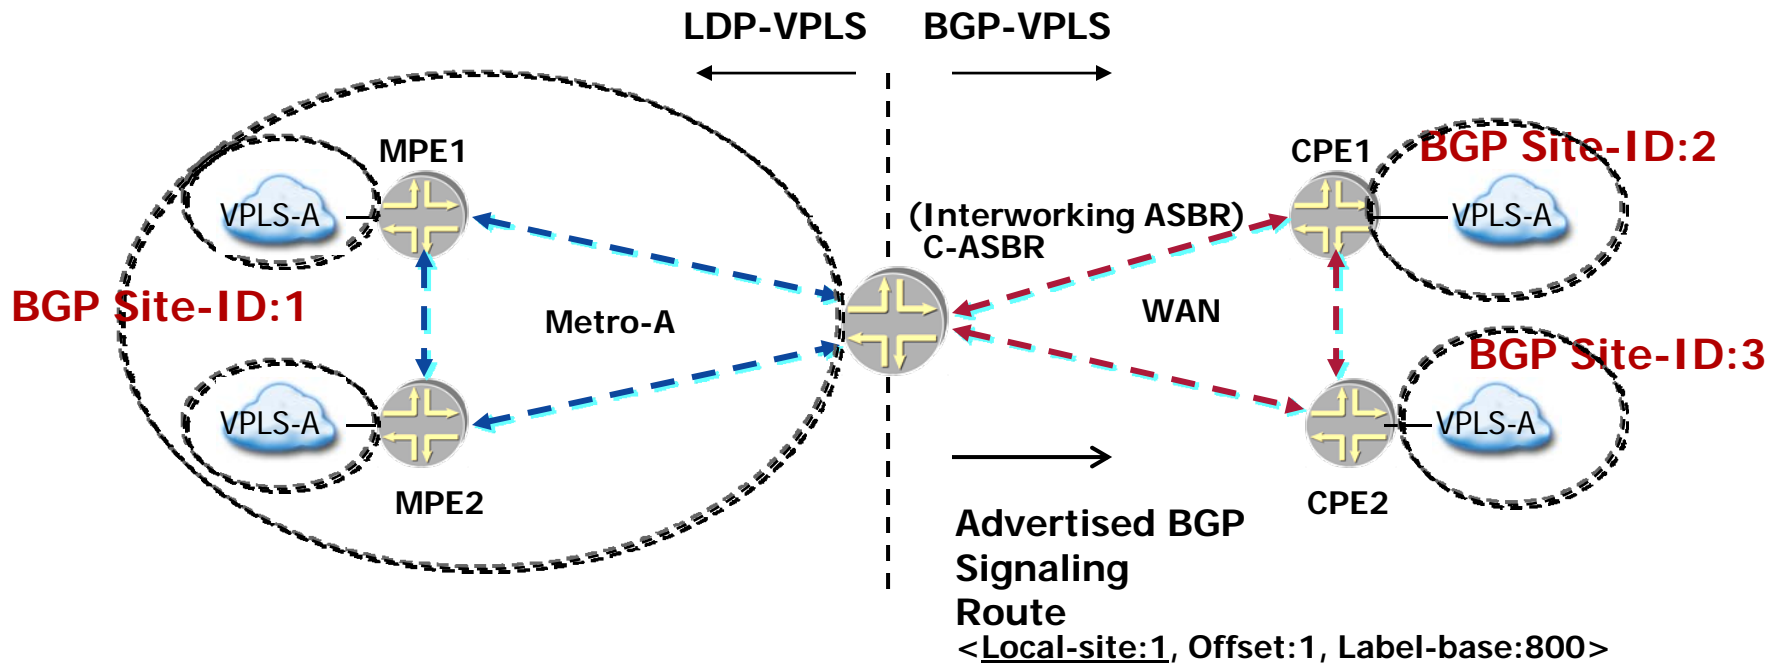
Control Plane Operation

- **Interworking Autonomous System Border Router (interworking ASBR)**
 - Supports both LDP and BGP VPLS control planes
 - Setup full-mesh of LDP/BGP PWs per interdomain VPLS
- **Scalable**
 - Only one additional LDP session provisioned per LDP-VPLS PE
 - Only one additional PW setup on each LDP-VPLS PE per VPLS



How BGP-VPLS view LDP-VPLS

- BGP-VPLS view the LDP-VPLS domain like a VPLS site with a single CE.



VPLS site



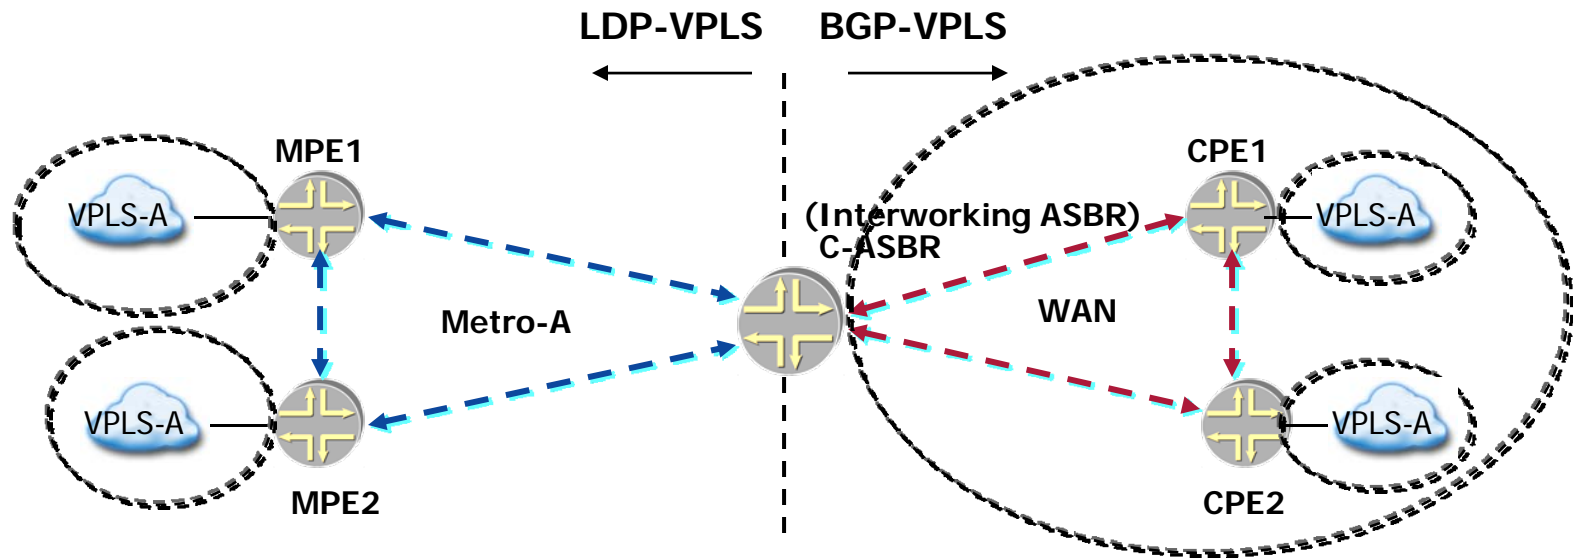
BGP-VPLS pseudowire



LDP-VPLS pseudowire

How LDP-VPLS view BGP-VPLS

- LDP-VPLS view the BGP-VPLS domain like a VPLS site with a single CE.



VPLS site



BGP-VPLS pseudowire



LDP-VPLS pseudowire

Data Plane Operation

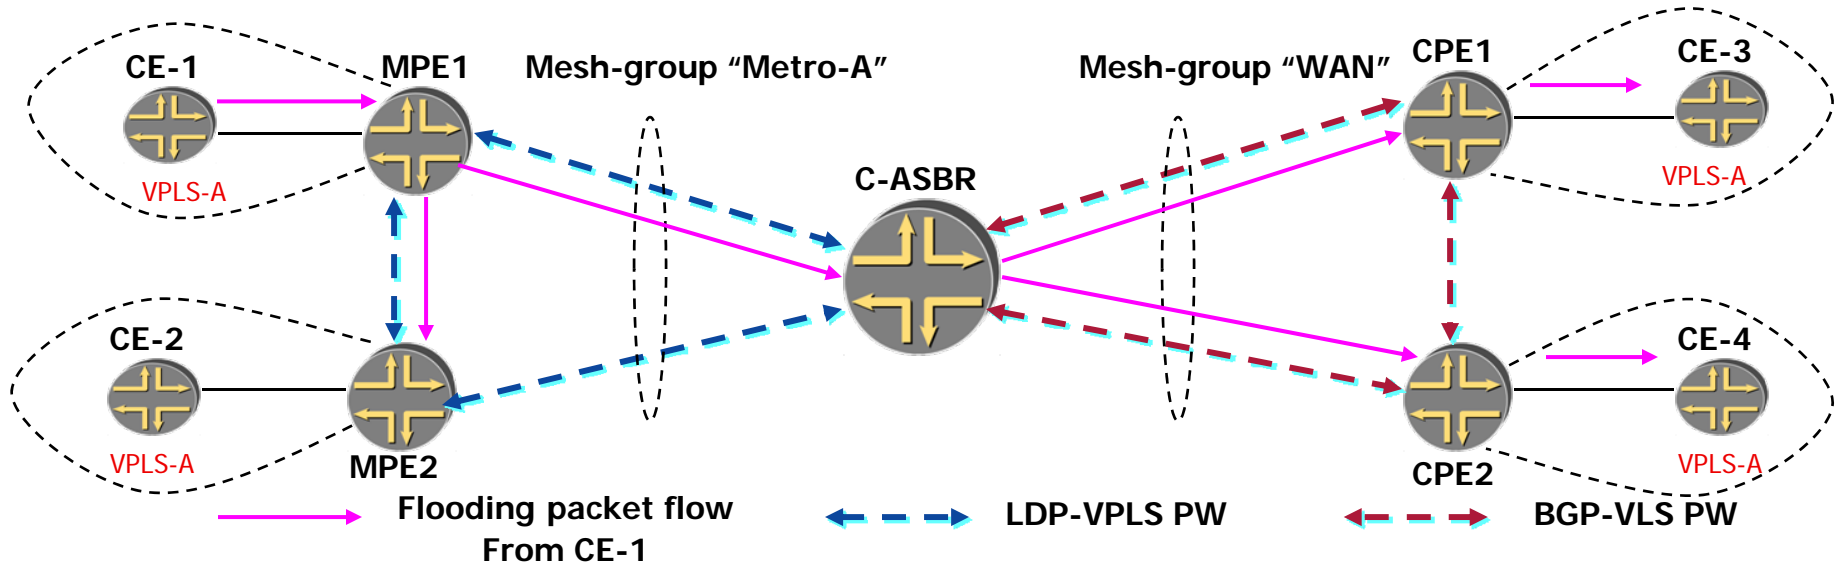
- **Fundamental data operations perform on each incoming frame on a PW:**
 - LDP and BGP VPLS PWs terminates on the interworking ASBR
 - Common MAC-table maintained to stitch the set of fully-meshed LDP and BGP VPLS pseudowires
 - MAC-table populated via learning a source-MAC address from incoming frame
 - Destination MAC-address lookup to determine destination pseudowire

Mesh Group

- **Mesh group concept introduced on an interworking ASBR**
 - Groups all the fully meshed PWs of each domain.

- **Existing split-horizon forwarding rule applicable for each mesh group**
 - Allows forwarding across pseudowires part of a different mesh-group.
 - Prohibits forwarding across pseudowires part of a same mesh-group.

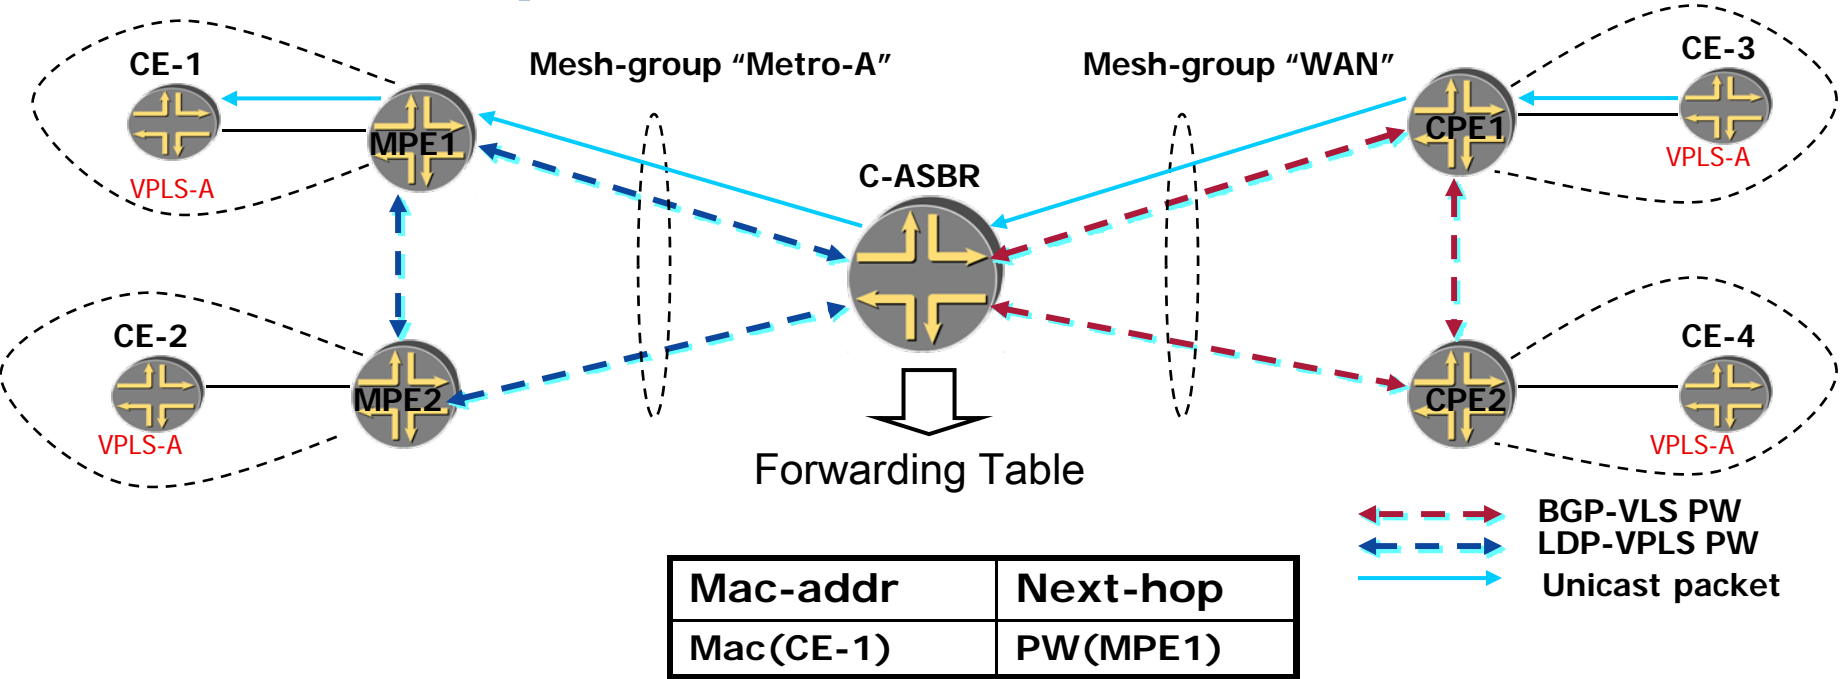
Mesh Groups: Flooding



- Assume CE-1 sends a broadcast ARP request packet. It will be flooded by MPE1, to all PEs in "Metro-A" domain (including C-ASBR) as it's fully-meshed. **C-ASBR receives this packet from MPE1 and forwards it to all the mesh-groups (PWs) except the one in which packet is received, as a result packet is forwarded on all PWs part of mesh-group "WAN".**
- Upon receiving this packet from MPE1, C-ASBR learns the CE-1 Mac-address via MPE1 PW. Destination forwarding table on C-ASBR looks like:

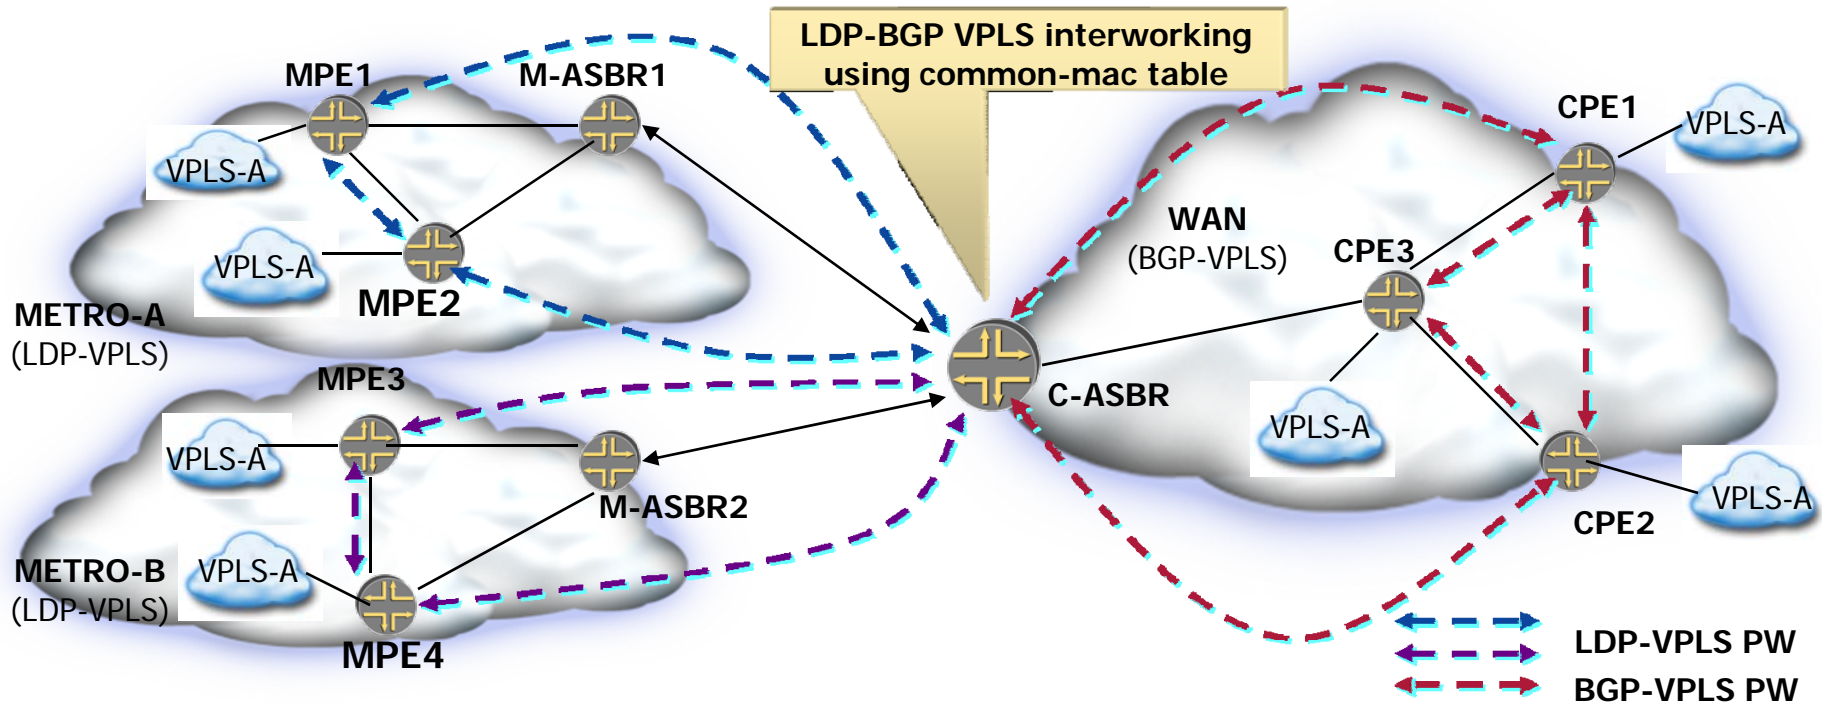
MAC-Addr	Next Hop
MAC(CE-1)	PW(MPE1)

Mesh Groups: “Normal” Known Unicast



- Here, CE-3 is sending a unicast packet to CE-1 and it's assumed that CE-1's MAC address has been learned on all PEs. CPE1 has learned CE-1's MAC via C-ASBR so it forwards the packet to C-ASBR.
- C-ASBR forwards this unicast packet on MPE1 PW since both incoming and outgoing PWs are in different mesh-groups.**

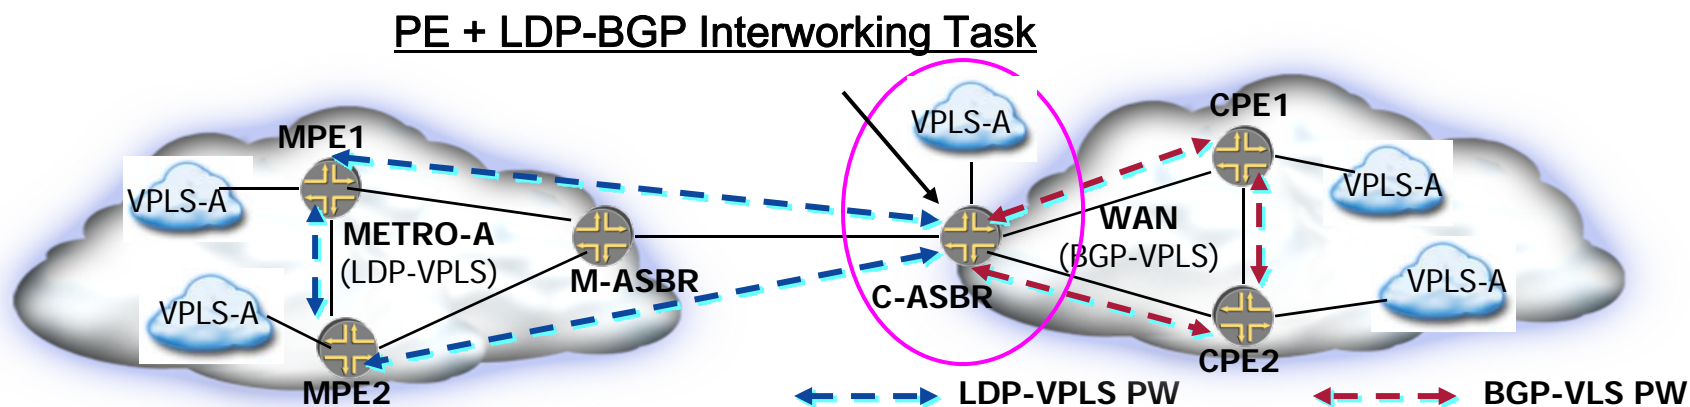
Mesh Groups \Leftrightarrow Full-mesh Domains



- Interconnecting multiple metro domains using a single interworking ASBR.
- Each metro domain is mapped to a dedicated mesh group.

Interworking Device With Attached CEs


- Interworking ASBR can provide PE router functionality at the same time performing the LDP-BGP VPLS interworking task.**



C-ASBR has three mesh groups: one for BGP in the WAN; one for Metro-A; and finally, one for its attached CE(s)


Resiliency

- Interworking ASBR becomes critical for inter-domain VPLS service
- Resiliency becomes critical to complete this solution
 - Solution is redundancy



Background: CE-device resiliency using BGP-VPLS multihoming

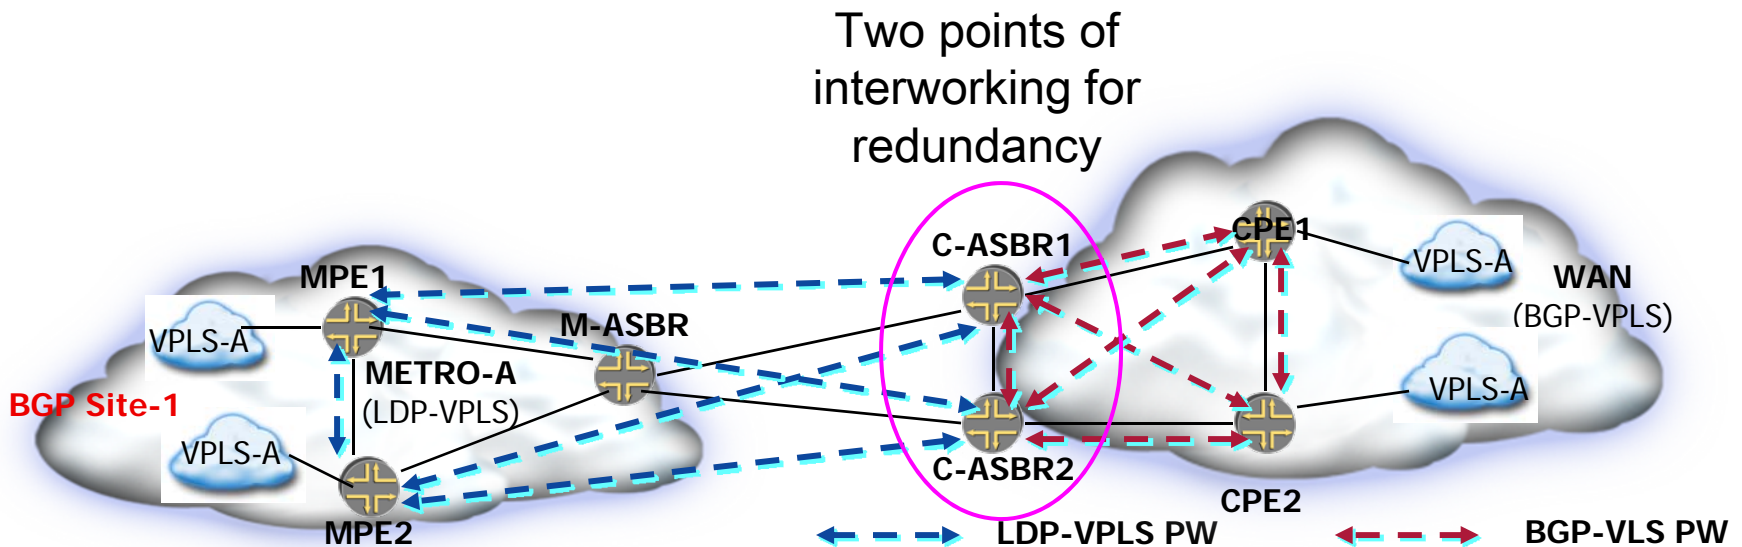
CE device dual-homed to PE devices for redundancy



- BGP-VPLS multihoming procedures enables loop-free redundancy to attached CE devices without relying on Spanning Tree Protocol (STP).

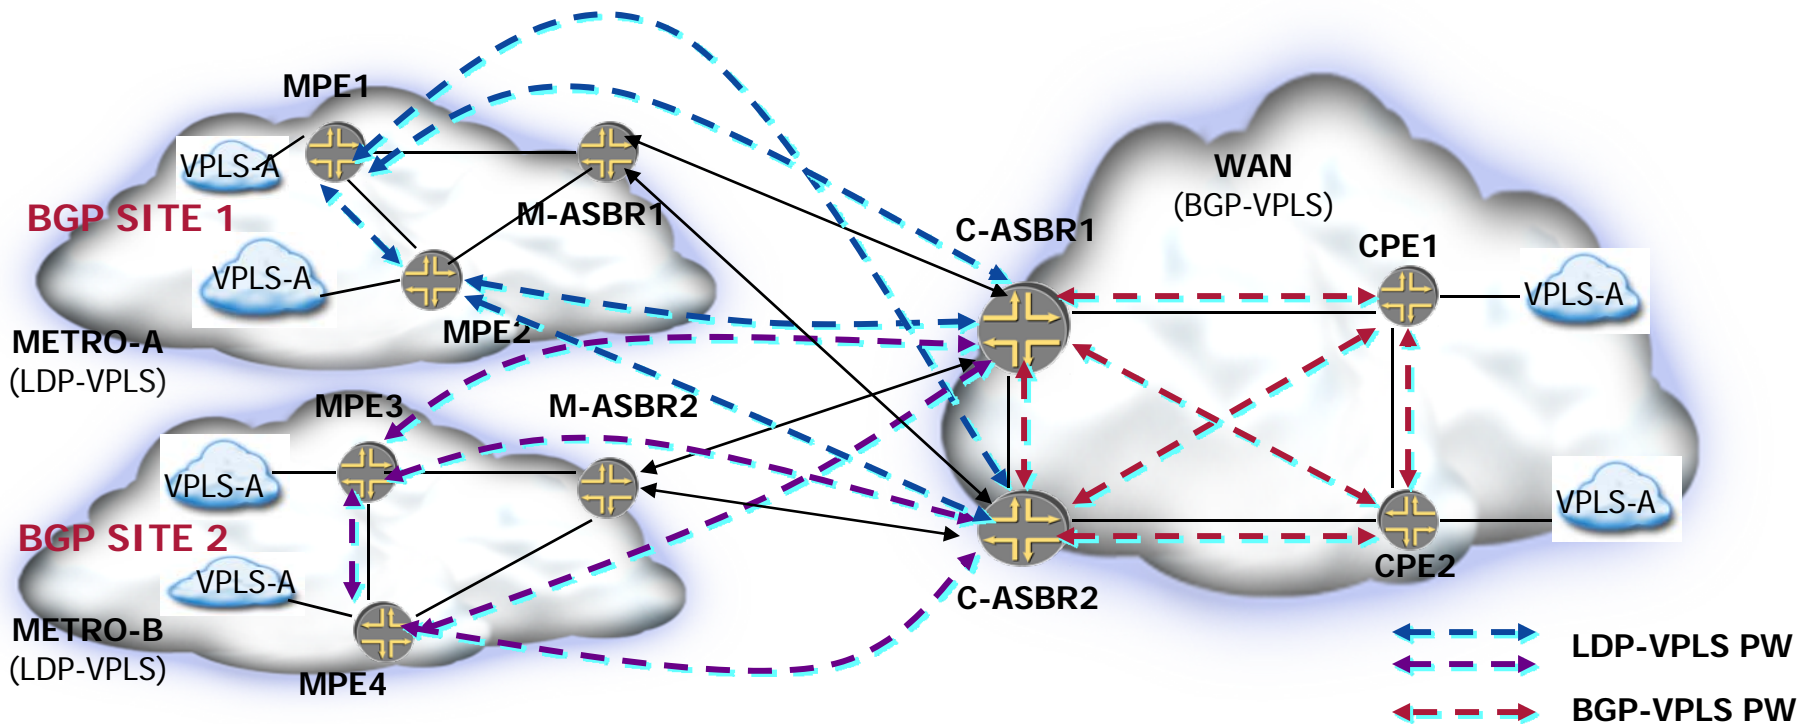
Redundancy using BGP-VPLS multihoming

- Existing BGP-VPLS multihoming procedure used with no changes to enable redundancy for LDP-VPLS metro.



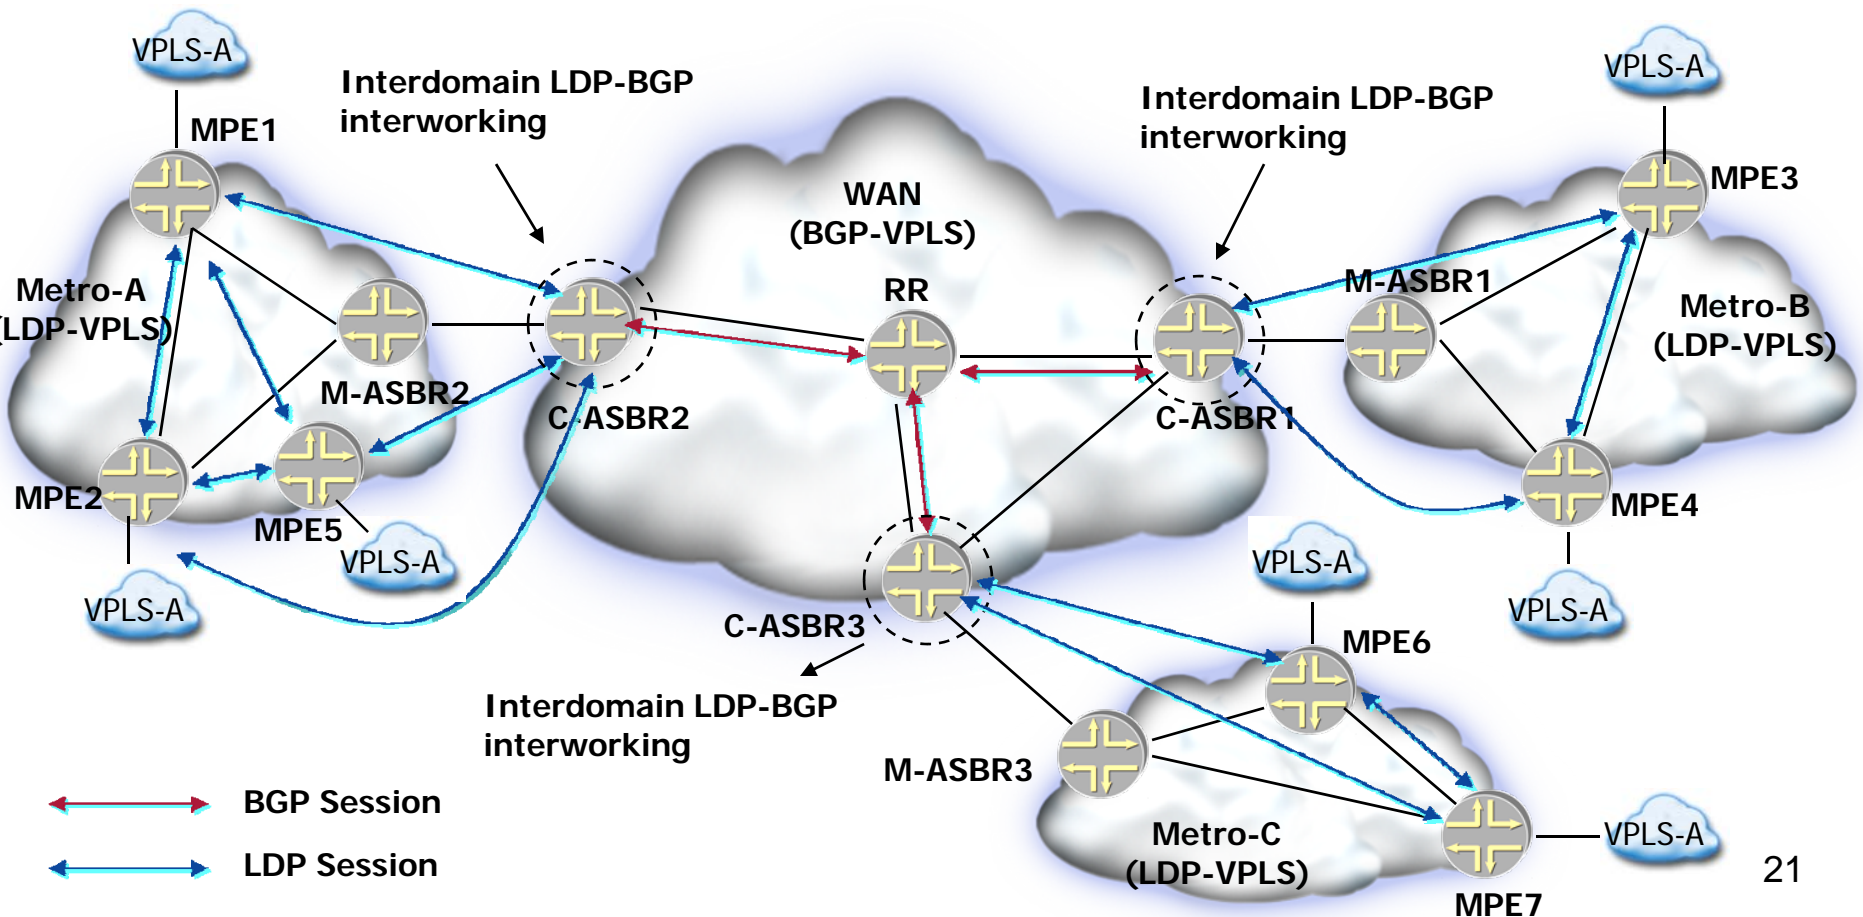
Resiliency for multiple LDP-VPLS Metros

- **Why use BGP-VPLS multihoming**
 - Scalable: Supports multiple metro domains
 - Flexible: Precise control over designated forwarder
 - Load balancing: Distribution of VPLS customers



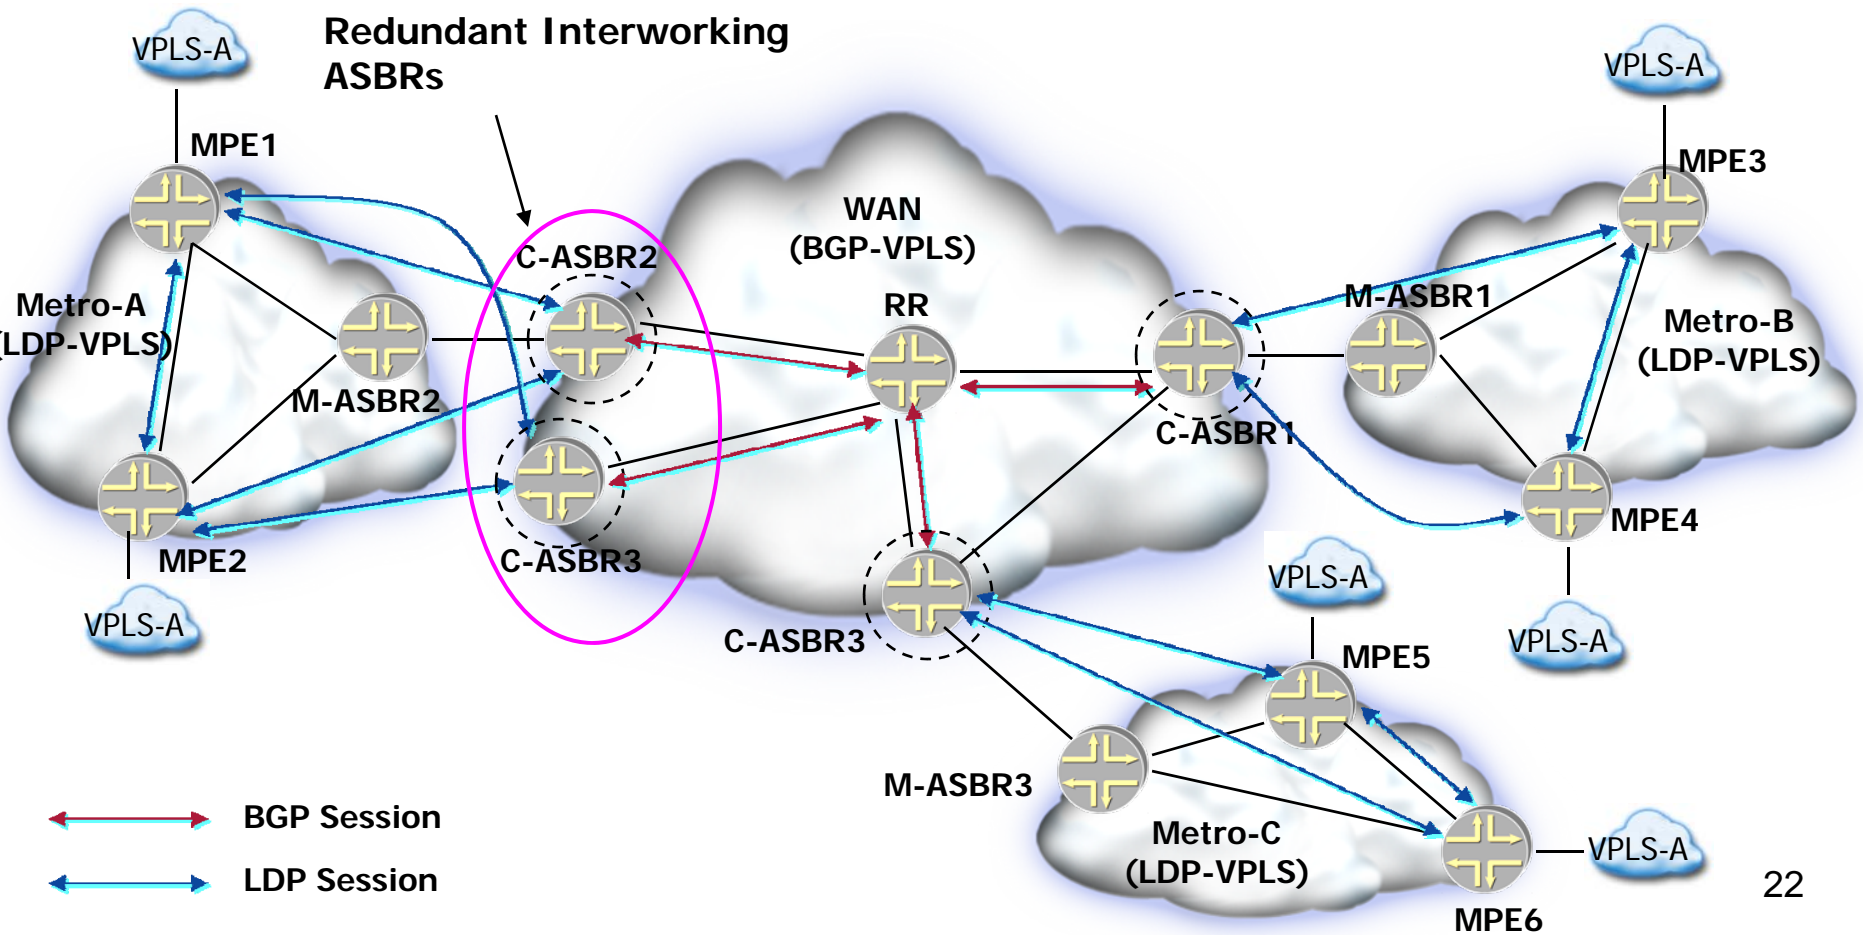
Extending the reach of LDP-VPLS metro-domain to WAN via BGP-VPLS

- Scalable way to enable region/nation-wide VPLS service using BGP-VPLS in WAN.



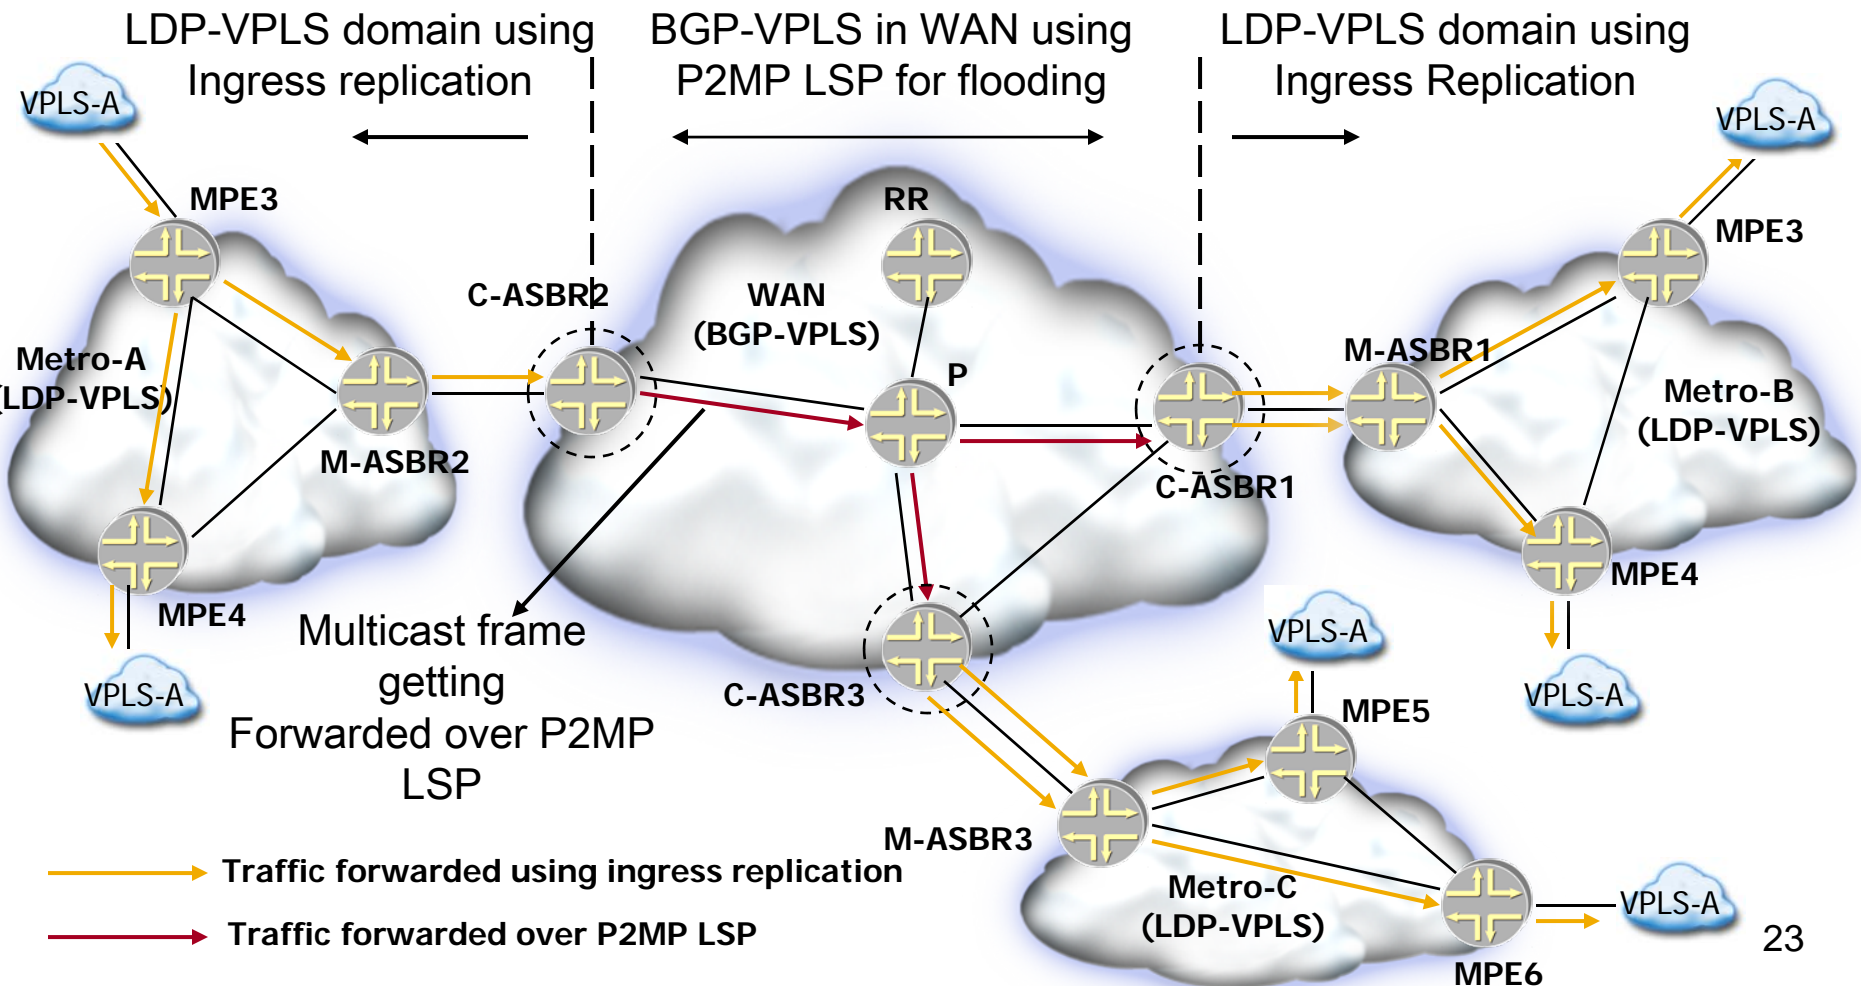
Resiliency using BGP-VPLS multi-homing

- Enabling resiliency for LDP-VPLS domain "Metro-B"



Efficient flooding and broadcasting using P2MP LSP in WAN running BGP-VPLS control plane

- Path of multicast traffic originating from site A3



Summary

- **It works with existing LDP-VPLS standard**
- **Scalable (and minor changes to metro design)**
 - Only one (or couple) additional LDP sessions per PE in metro
 - Only one (or couple) additional PWs setup per PE per VPLS
- **Easier to manage**
 - Adding new sites or PE in a metro area does not require any provisioning on other remote Metro area
 - BGP auto-discovery simplifies provisioning task in WAN

Summary (Cont...)

■ Resilient

- Native BGP-VPLS multi-homing support provides loop-free redundancy without the need of Spanning-Tree Protocol

■ Efficient

- Interworking ASBR provides dual functionality of PE router and LDP-BGP interworking
- Multiple metro-domains can be interconnected via single interworking ASBR
- BGP-VPLS with point-to-multipoint LSP instead of ingress replication for multicast/broadcast unknown



Thank You !