PBB Network Scalability mp2mp B-VLAN

Maarten Vissers

March 2007

802.1ah PBB PIP, CBP and PBP on I- and B-Components Current design

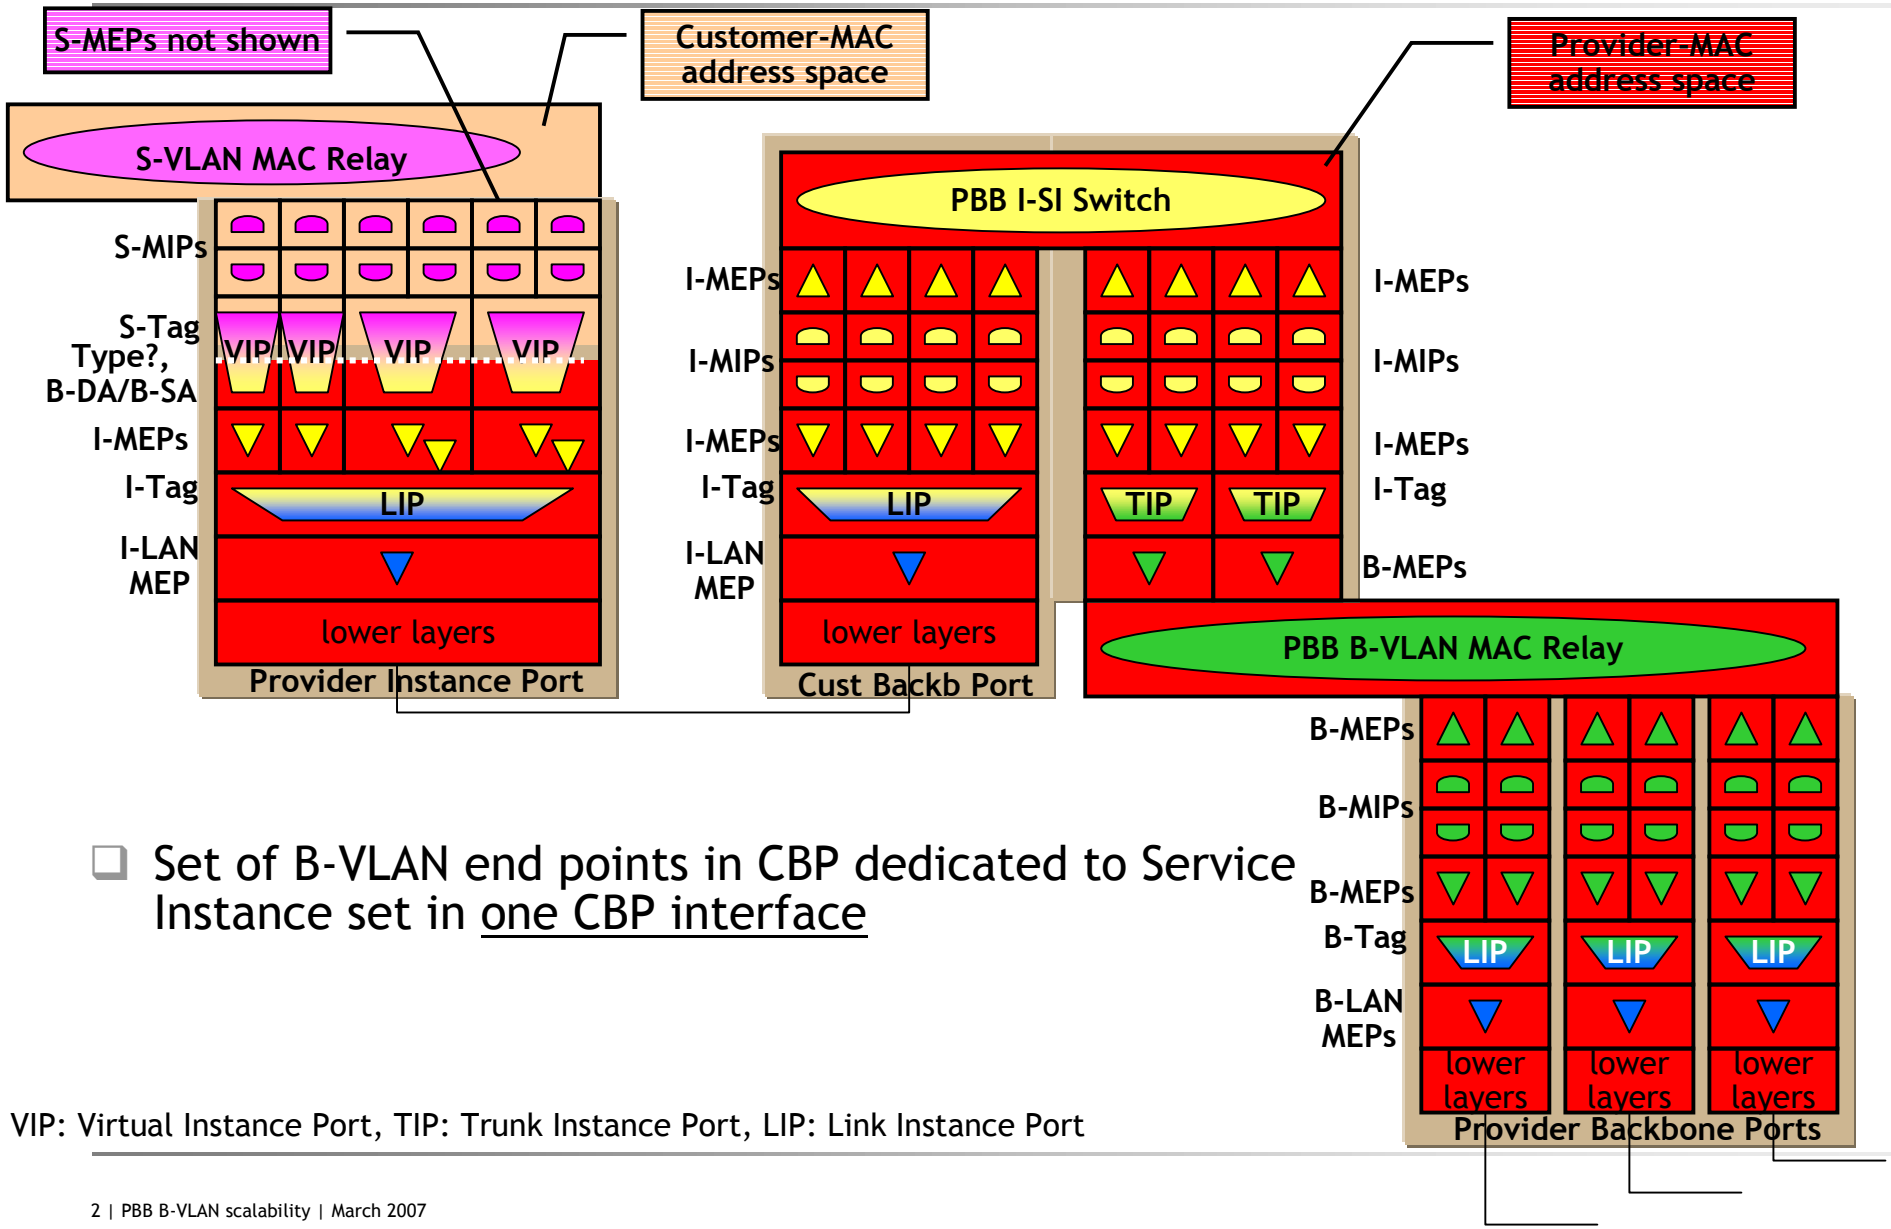

- Set of B-VLAN end points in CBP dedicated to Service Instance set in one CBP interface

VIP: Virtual Instance Port, TIP: Trunk Instance Port, LIP: Link Instance Port

802.1ah PBB CBP and PBP on B-Components and BCBs
Current design

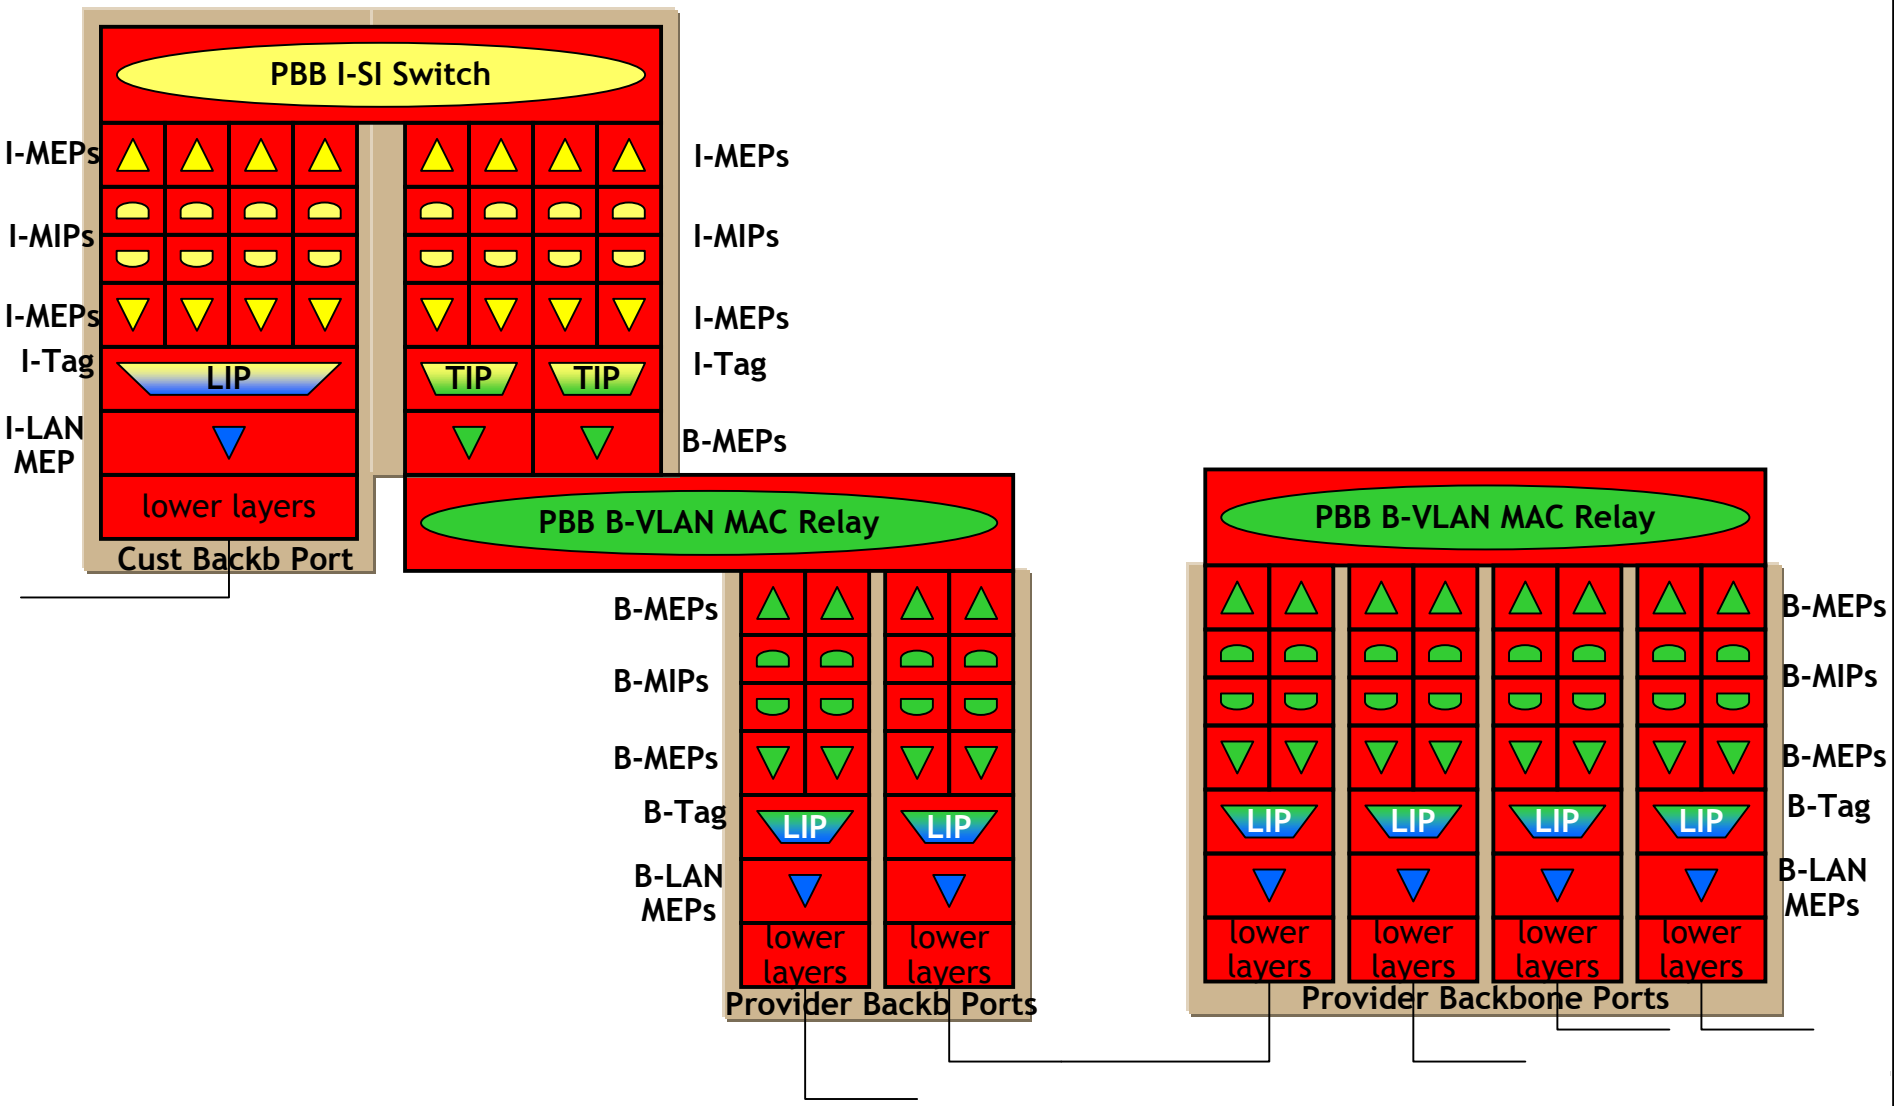

VIP: Virtual Instance Port, TIP: Trunk Instance Port, LIP: Link Instance Port

PBB mp2mp B-VLAN scalability issue 1

Assume that a PBB network domain is bound by 50 B-BEBs. Those 50 B-BEBs are interconnect with each other via intermediate BCBs, while many I-BEBs are connected to the other side of those B-BEBs. Assume that a B-BEB has ten interface cards, each having ten CBP functions. A B-BEB (in this example) may thus have up to hundred CBP functions. The 50 B-BEBs at the edge of the PBB network domain may have up to 5000 CBP functions.
[Note - current interface cards may have up to 48 interface ports.]

Worst case, a B-VLAN will have a service mix which has one end point of its service instances behind each CBP function. Such B-VLAN #X (e.g. X=100) will then have a total of 5000 end points and 5000 B-VLAN MEPs. With 5000 B-VLAN MEPs, each generating CCMs with 100 ms periodicity causes each of those MEPs to receive 4999 x 8 kbit/s = 39.99 Mbit/s of CCM frames. With ten B-VLAN #X MEPs on an interface card, the interface card will receive 10 x 39.99 = 399.9 Mbit/s of CCM frames for B-VLAN #X. Idem for other B-VLANs (#Y, #Z, etc).

Question: How to reduce the number of B-VLAN end points?

PBB mp2mp B-VLAN scalability issue 2

PBB contains a service instance layer and a B-VLAN trunk instance layer. The objective of the introduction and use of a trunk layer in transport networks is to have the nodes within the trunk domains independent of the service instances. Service instance awareness will then be restricted to the trunk domain edges. BCB nodes will then only have to deal with trunk instances, of which there are typically a factor of 50 to 100 less then service instances (e.g. 1 million service instance, 10 to 20 thousand trunk instances with most of those located in the metro domains and just some 10 to 20% in the core domain).

Comparing the above trunk layer objective with the use of the B-VLAN trunk instances inside a PBB network domain, it can be concluded that the B-VLAN trunk layer does not meet the objective. Every B-VLAN needs to take into consideration the service instances carried over it; i.e. it needs to be aware and has to administer unicast MAC addresses and multicast MAC addresses. With at least one multicast MAC address per multipoint service instance, and networks that may have up to 30% multipoint services, it will not be uncommon that 100000 multicast MAC addresses are to be registered within a core domain; i.e. within BCB nodes inside the core domain. When an inner core node fails or a major cable breaks there are then many multicast addresses to re-register within the restored B-VLANs.

Question: How to reduce service instance awareness within the B-VLAN?

Minimize mp2mp trunk instance (B-VLAN) end points

Current:

- One B-VLAN # X end point per CBP interface port
- Maximal “n x m” in B-Component
  - m: # interface ports per CBP interface card
  - n: # CBP interface cards in B-Component
  - k: # B-Components per PBB Network
  - maximal “n x m x k” B-VLAN # X end points per PBB network;
    - e.g. 5000 B-VLAN #X end points and thus 5000 B-MEPs

Alternative I:

- One B-VLAN # X end point per CBP interface card;
- Maximal “n” in B-Component
  - maximal “n x k” B-VLAN # X end points per PBB network;
    - e.g. 500 B-VLAN #X end points and thus 500 B-MEPs

Alternative II:

- One B-VLAN # X end point in B-Component (on PBP interface card)
  - maximal “k” B-VLAN # X end points per PBB network;
    - e.g. 50 B-VLAN #X end points and thus 50 B-MEPs

Current B-BEB equipment architecture

- B-BEB includes
  - one or more PBP interface cards
  - one or more CBP interface cards
  - one or two B-VLAN MAC Relay fabric cards
- PBP interface card includes
  - one or more Provider Backbone Ports
  - (optional) multiplexer to access fabric
- CBP interface card includes
  - one or more Customer Backbone Ports
  - (optional) multiplexer to access fabric

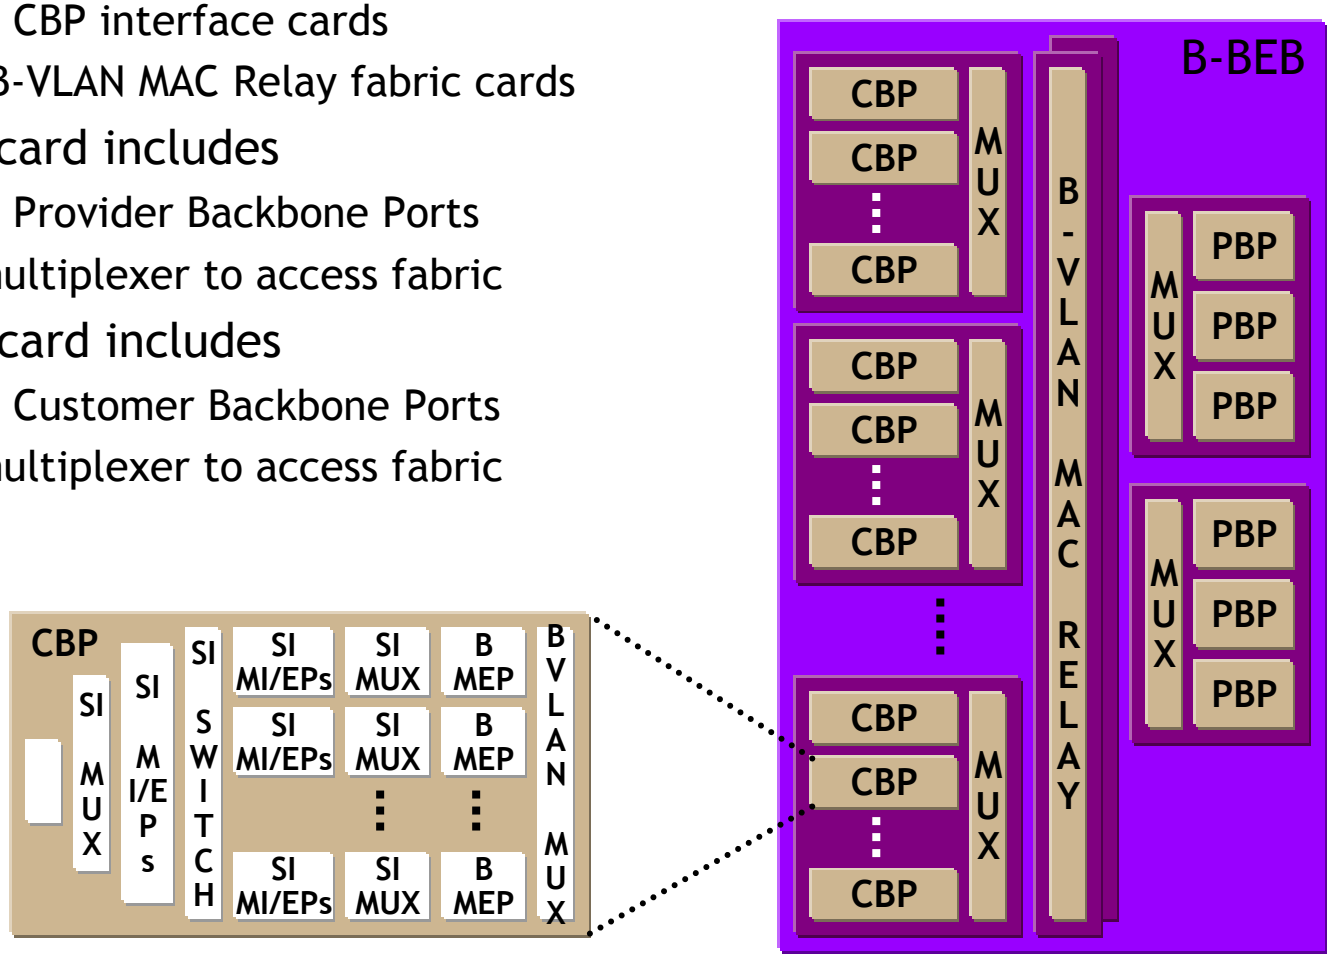

Alternative B-BEB equipment architecture (I)

- Split CBP into two parts
  - CBP-I (I: interface) module
  - CBP-C (C: common) module
- Modified CBP interface card includes
  - one or more CBP-I modules
  - one CBP-C module

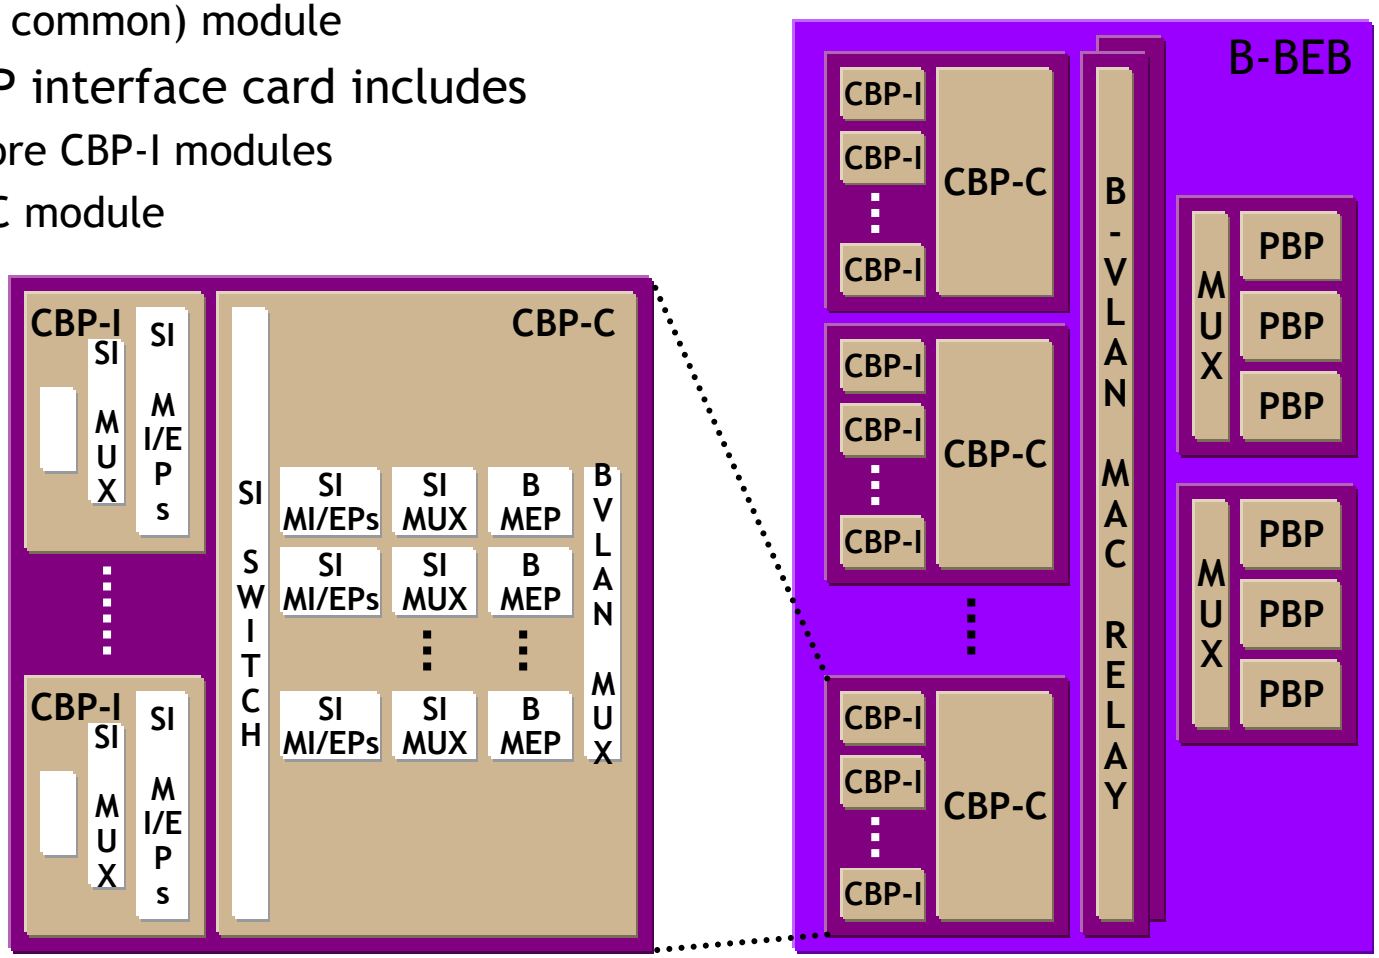

Alternative B-BEB equipment architecture (II)

- Relocate CBP-C module (from CBP interface card to PBP interface card)
- Modified CBP and PBP interface cards
  - CBP: one or more CBP-I modules
  - PBP: one or more CBP-C/PBP modules
- MAC Relay forwards SIs

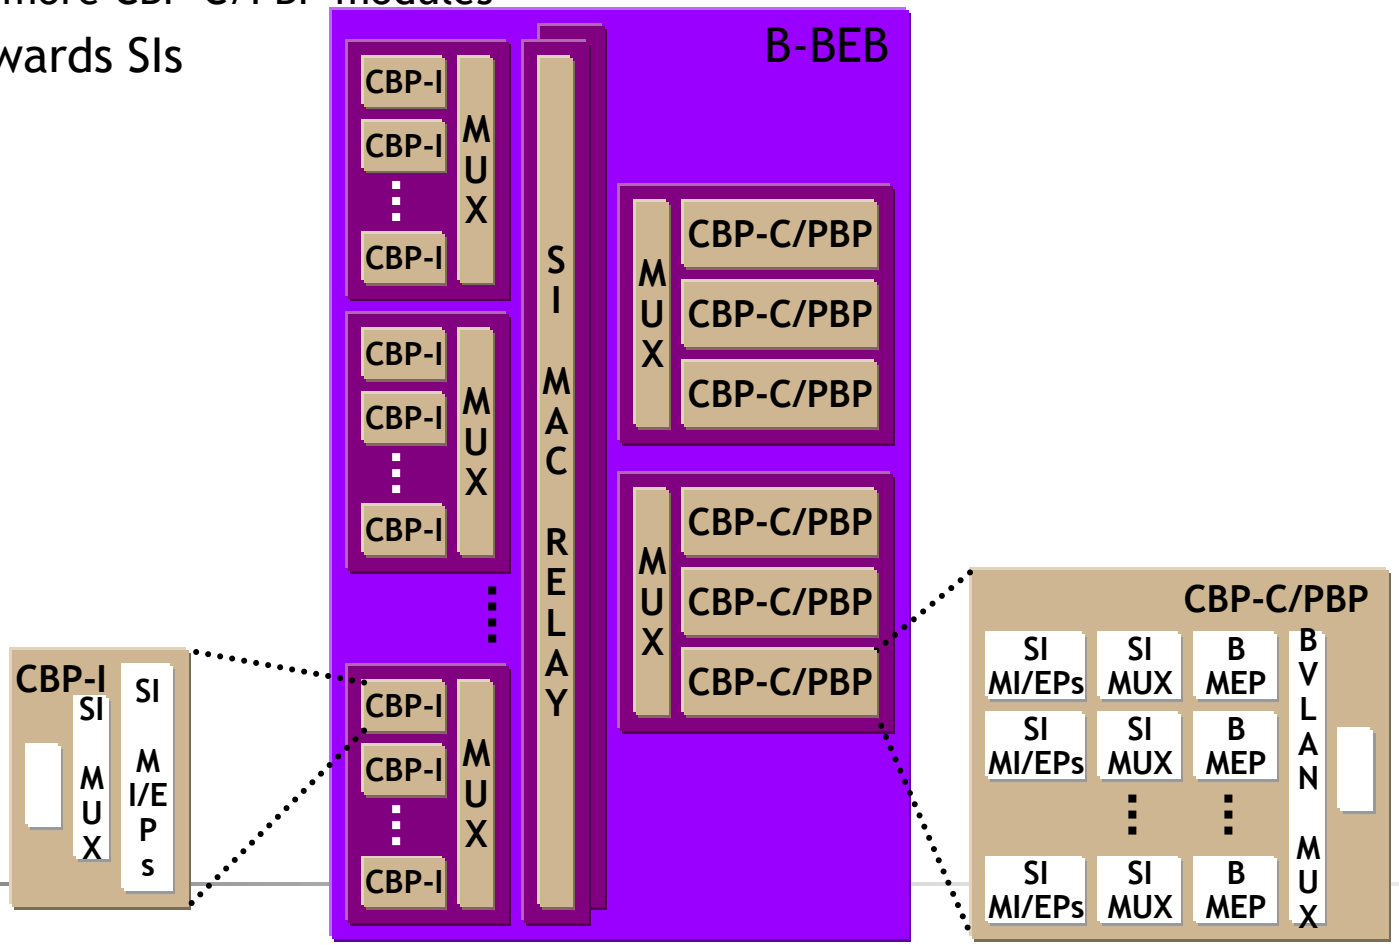

PBB I-SI MAC Relay technology availability

MAC Relay function on basis of I-SID and B-DA

- Every 2.5 – 3 years fabrics support factor 4 more instances
  - 802.1Q-1998: 4k, today (2007) 4k x 64 = 256k should be possible
  - especially as number of B-DAs to learn is limited to one per mp service endpoint
  - 70% of service instances is p2p, not requiring MAC learning
  - as such, restricted number of service instances with restricted number of DAs
- 24-bit I-SID may not be significant in first years of deployment
  - like in early days ATM, restrict use to k-bit (k<24)
- Availability of technology agnostic switch fabrics
  - supports SDH, OTH, MPLS and Ethernet
  - use of 16- to 20-bit instance identifiers, supporting 64k to 1M switched/bridged instances
  - C-VIDs, S-VIDs, B-VIDs, I-SIDs mapped onto equipment specific instance identifier at ingress port and vice versa on egress port

Minimize service instance awareness in B-VLAN trunks

Current:

- individual Provider-MAC (B-MAC) address registration
  - one per PIP
  - one per CBP (if Service Instance MEP active in CBP)
  - registration method: MAC learning
- group Provider-MAC (B-MAC) address registration
  - one per mp2mp or rooted-mp Service Instance
  - registration method: MMRP or other protocol

Alternative:

- zero Provider-MAC address registration
- method: p2p B-VLAN trunks
  - best with Alternative II B-BEB design
    - Metro PBB domain
      - Few hundred trunk instances ➔ no problem to deploy p2p B-VLANs
    - Core PBB domain
      - Few hundred to ten thousand trunk instances ➔ no problem to deploy p2p B-VLANs together with some restricted B-VID Translation

Current service instance awareness in n-port B-VLAN trunks

rationale – consequence – remedy… complexity

General

- multipoint-to-multipoint
  - **rationale**: prevent multiple copies of multicast frame to occur on a link (resource bandwidth optimization)
  - **consequence**: extended flooding scope for every broadcast, multicast and unknown unicast service instance frame (decreases resource bandwidth utilization)
  - **remedy**: install flooding restriction methods:
    - PBB B-VLAN **service instance awareness** (pMAC address administration)
      - additional service instance associated "state" to limit flooding
      - one individual pMAC address per PIP (one or more VIPs)
      - one group pMAC address per n-port service instance
    - PBB service instance pre-conditioning at Trunk Instance Ports (TIP)
      - multipoint service instance: translate all-group and unknown individual/group pMAC addresses into service instance specific group pMAC address
      - point-to-point service instance: translate all-group pMAC address into individual pMAC address of destination VIP
    - Alternative service instance pre-conditioning
      - perform pre-conditioning at Virtual Instance Ports (VIP)
      - additional configuration effort, possibly including additional cMAC⇔pMAC learning
      - superfluous complexity outside n-port trunk solutions

Current service instance awareness in n-port B-VLAN trunks

rationale – consequence – remedy… complexity

Application specific

- rooted-multipoint
  - aggregation service
  - PBB B-VLAN **service instance awareness** (pMAC address administration)
    - additional service instance associated “state” to limit flooding
    - one individual pMAC address per PIP (one or more VIPs)
  - PBB service instance pre-conditioning at TIP or VIP
- point-to-multipoint
  - broadcast TV distribution (root to leaves only)
  - minimized resource bandwidth utilization
  - service instance unaware (broadcast only)

Defeats trunk layer objective, except for point-to-multipoint

Solution (mp2mp trunks) worse than problem (multiple copies on link)

Zero service instance awareness in 2-port B-VLAN trunks

Completely independent of PBB service instances

Trunk instance identifier administration only

- no Backbone-MAC addresses to register (and learn/age) in BCB nodes

Complies with trunk layer objective

Minimized VIP/TIP configuration and processing

- unrestricted use of all-group Backbone-MAC address ⇒
  - no need for Service Instance specific group Backbone-MAC address (per VIP)
    - no need for 802.1ah OUI
  - no need for Default Backbone-MAC address entry in service instance table in CBP
    - no need to replace all-group Backbone-MAC address by Default Backbone-MAC address
  - no need to change CFM group MAC address usage

Best scalable PBB B-VLAN architecture

Alternative II B-BEB architecture with 2-port B-VLAN trunks

Backup

Multiple PIPs in I-BEB, multiple CBPs in B-BEB

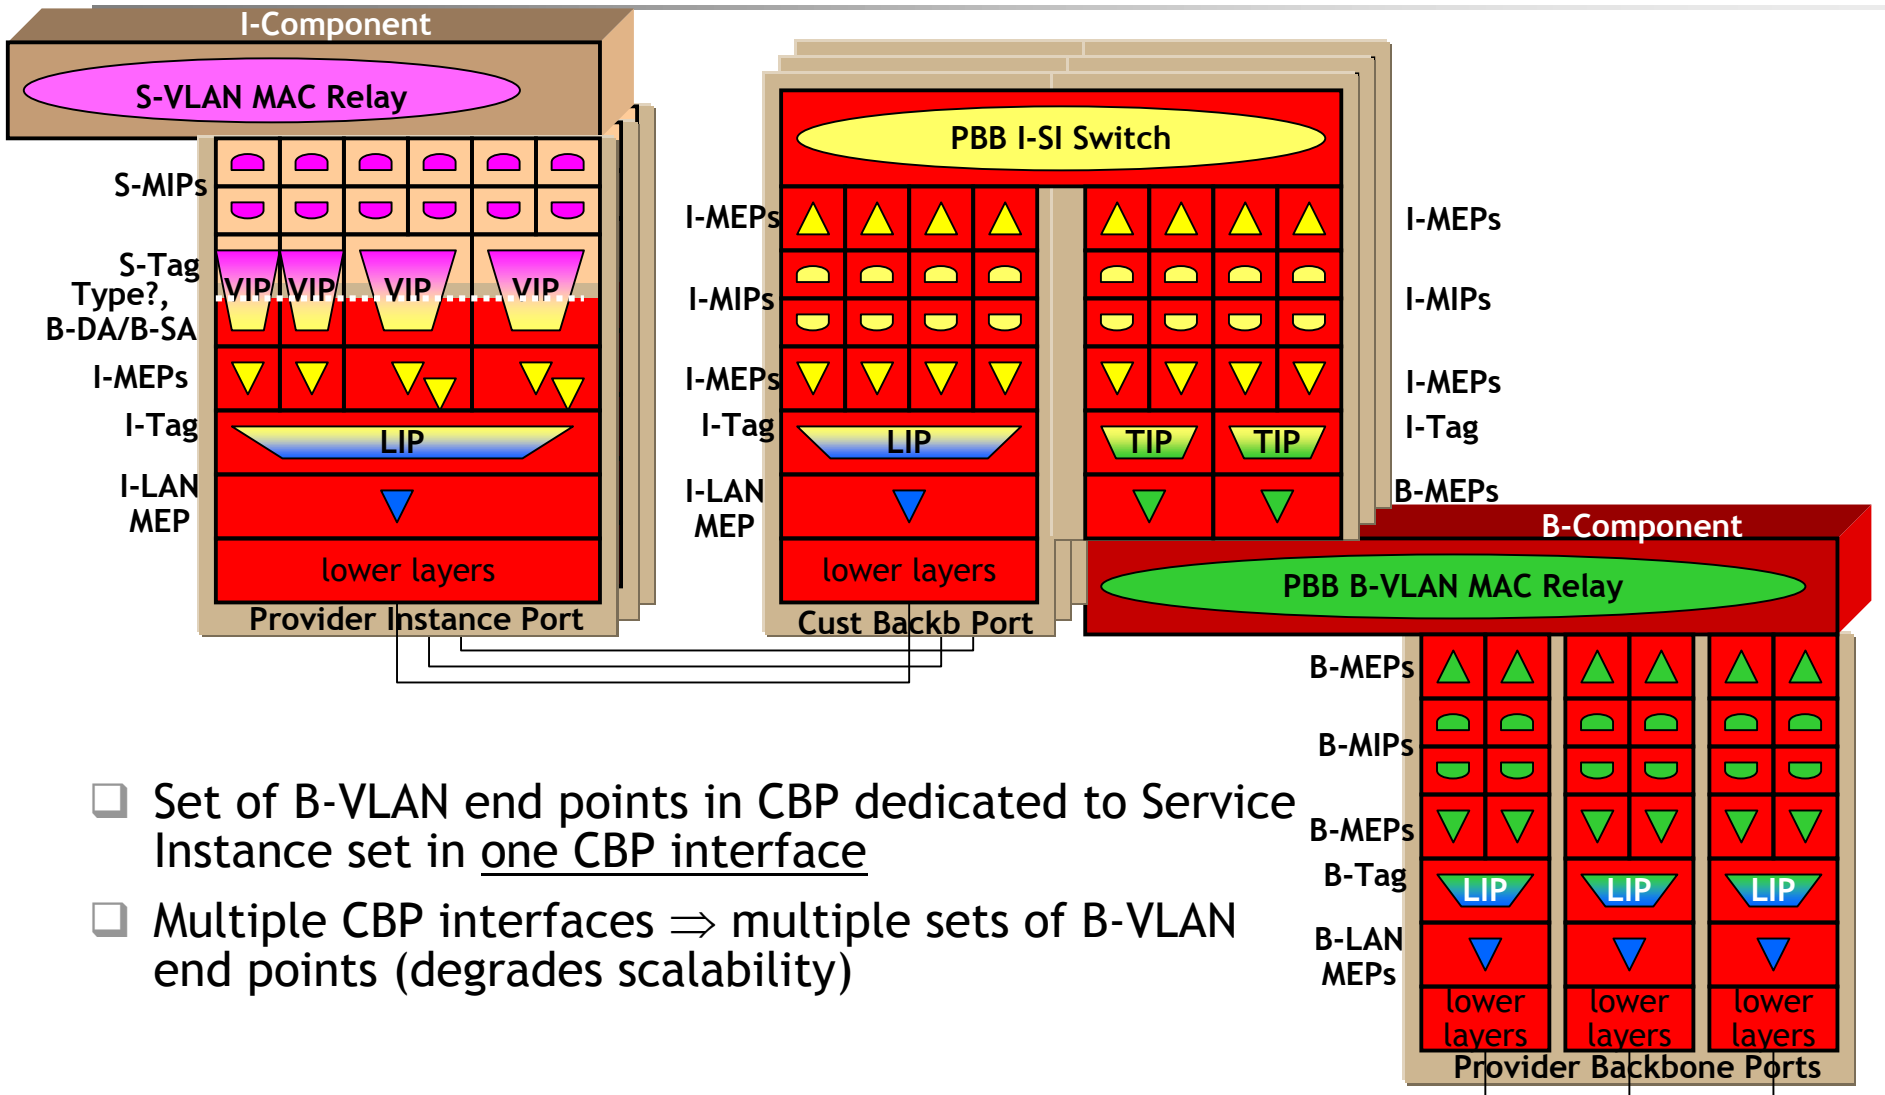

- Set of B-VLAN end points in CBP dedicated to Service Instance set in <u>one CBP interface</u>
- Multiple CBP interfaces ⇒ multiple sets of B-VLAN end points (degrades scalability)

VIP: Virtual Instance Port, TIP: Trunk Instance Port, LIP: Link Instance Port

VIP, TIP, LIP, I-MEP, B-MEP, LAN-MEP, S-MIP, I-MIP, B-MIP, S relay, I switch, B relay, MAC Service Boundary

Multi-interface CBP

Discussed in Monterey on Friday

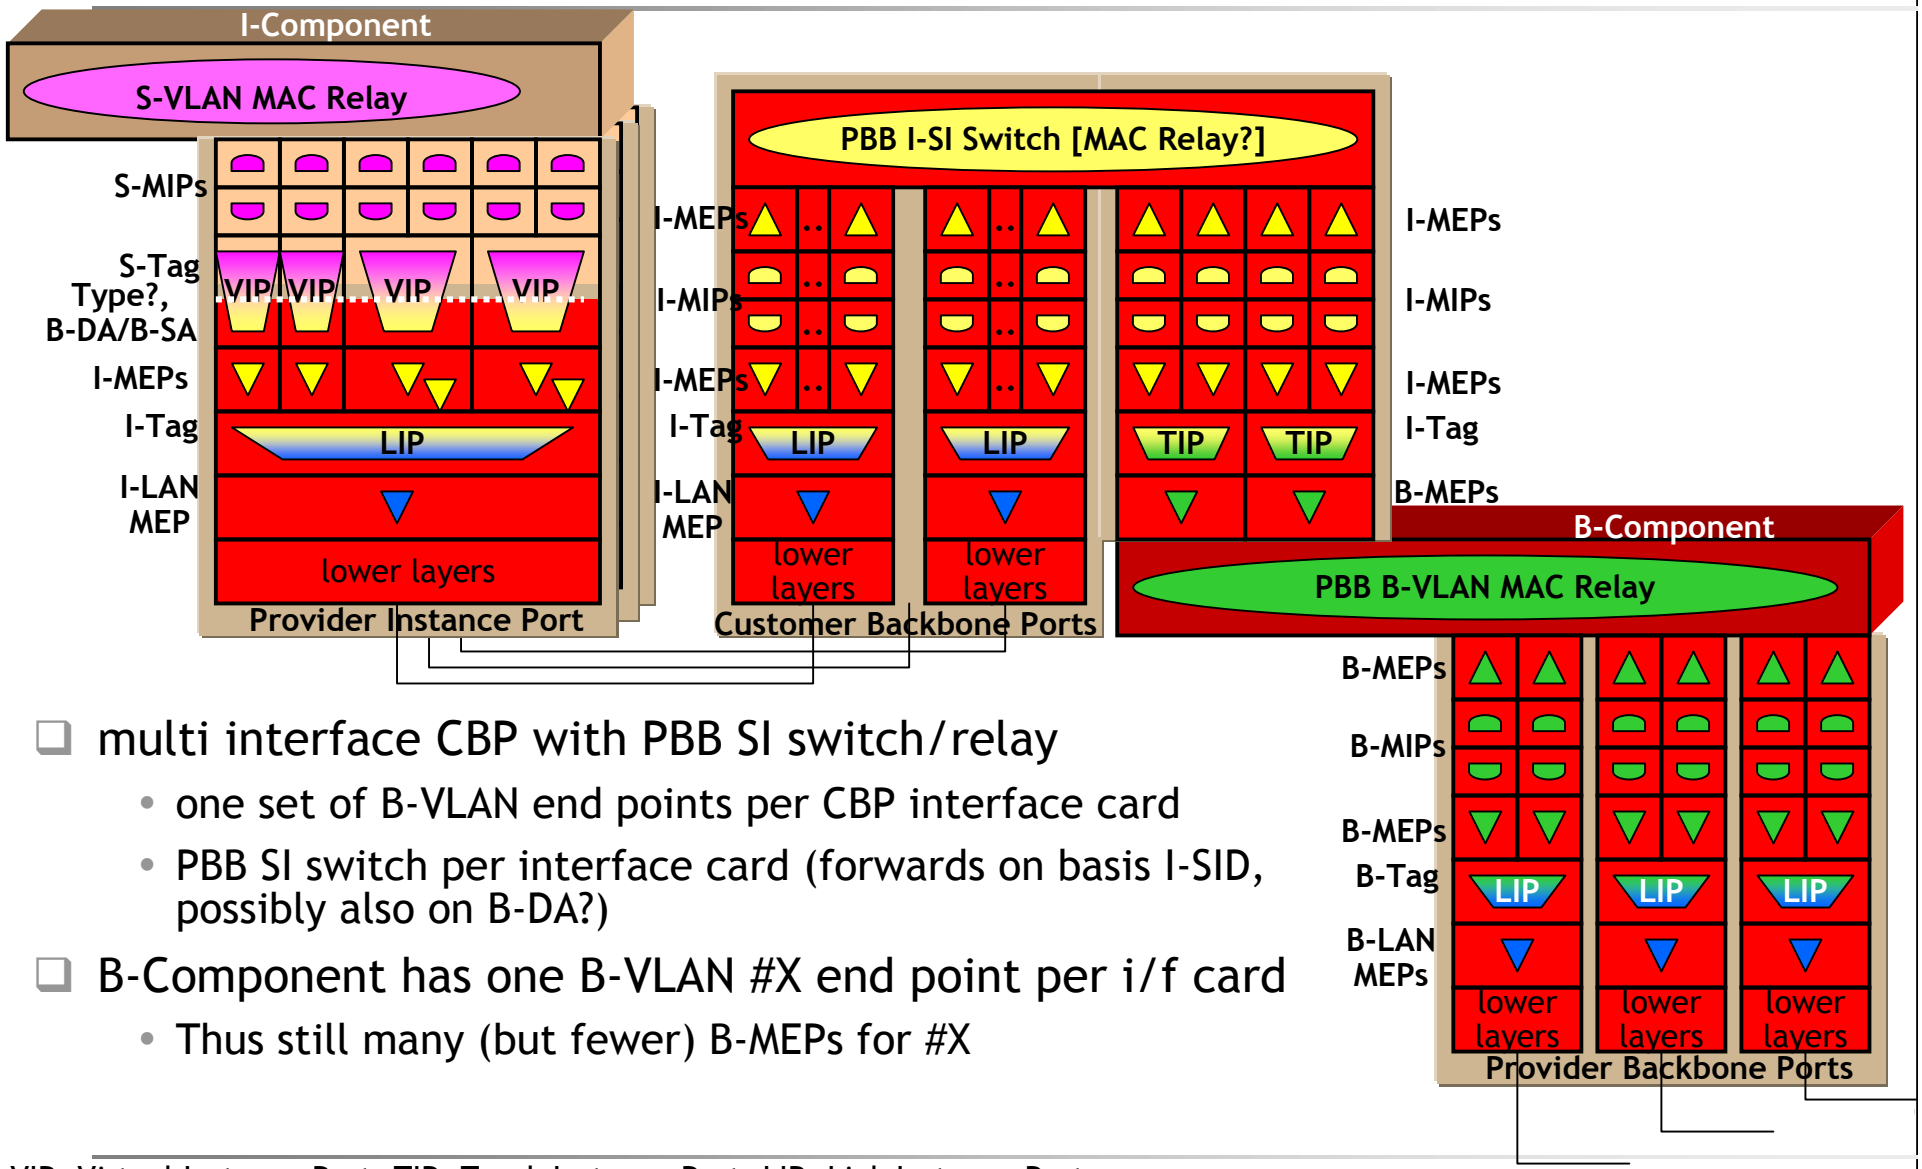

- multi interface CBP with PBB SI switch/relay
  - one set of B-VLAN end points per CBP interface card
  - PBB SI switch per interface card (forwards on basis I-SID, possibly also on B-DA?)
- B-Component has one B-VLAN #X end point per i/f card
  - Thus still many (but fewer) B-MEPs for #X

VIP: Virtual Instance Port, TIP: Trunk Instance Port, LIP: Link Instance Port

Single Switch/Relay B-Component

Minimized number of B-VLAN end points

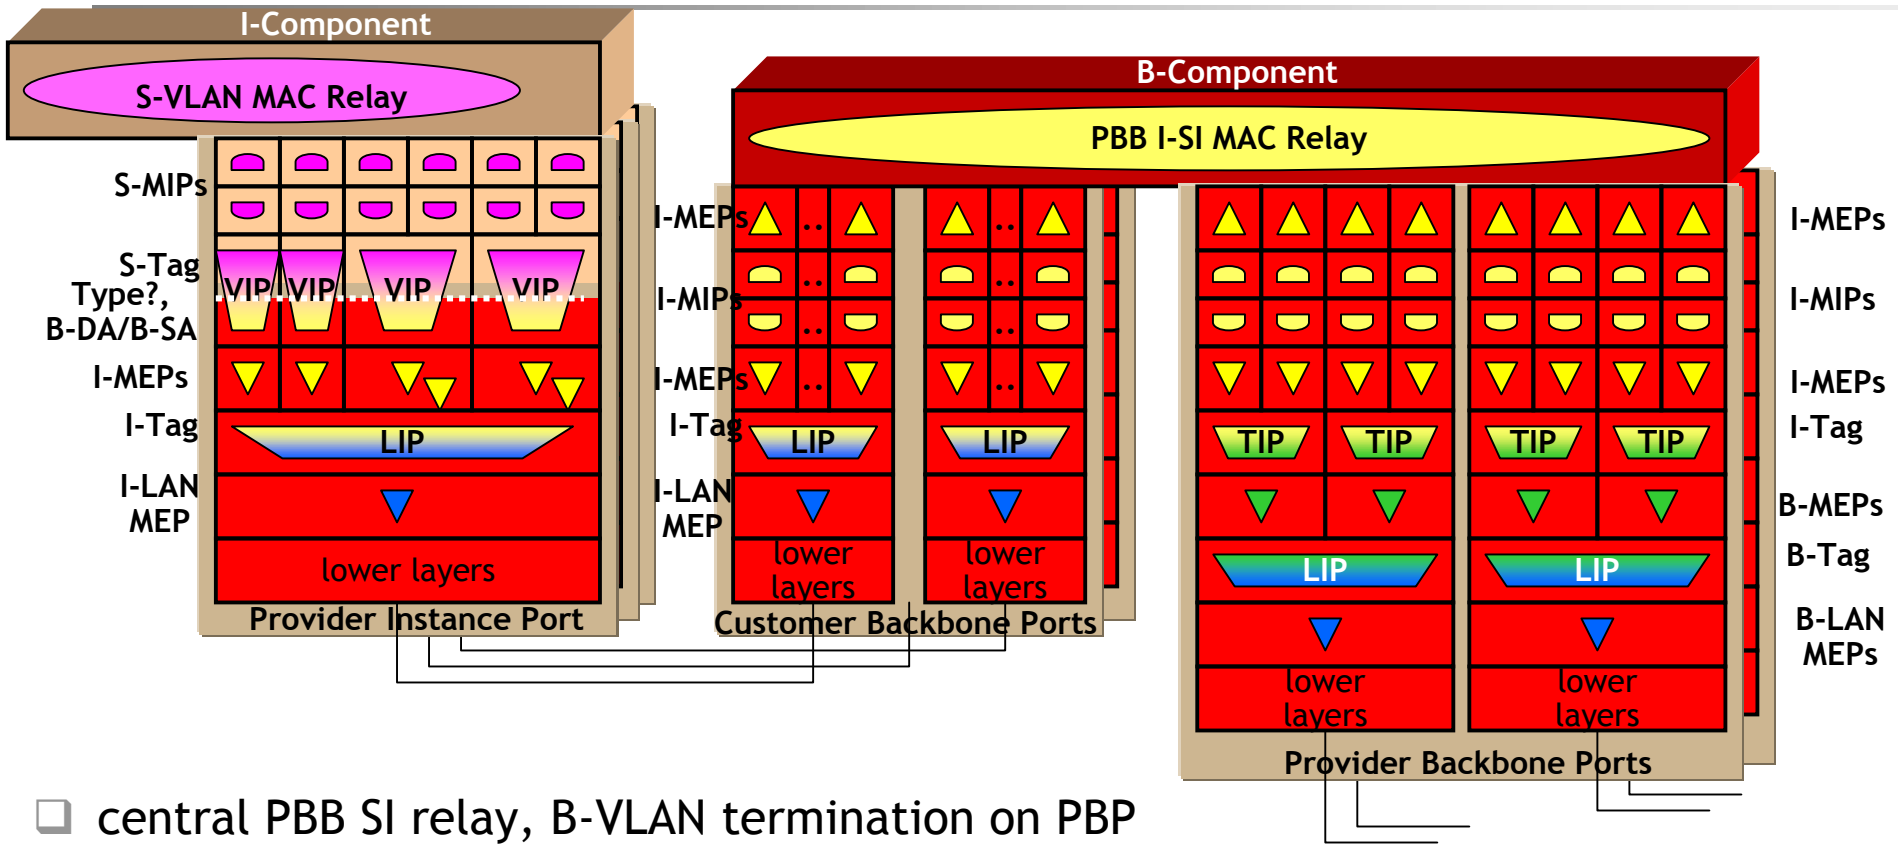

- central PBB SI relay, B-VLAN termination on PBP
  - set of B-VLAN end points per node
  - PBB SI relay per node (forwards on basis I-SID and B-DA)
- B-Component has one B-VLAN #X end point per node
  - Minimized number of B-MEPs

VIP: Virtual Instance Port, TIP: Trunk Instance Port, LIP: Link Instance Port