Point-to-Multipoint Traffic Engineered LSP (MPLS and GMPLS)

D.Papadimitriou (dimitri.papadimitriou@alcatel.be)
R.Aggarwal (rahul@juniper.net)
S.Yasukawa (yasukawa.seisho@lab.ntt.co.jp)

MPLS 2004

Agenda

- What is a P2MP TE LSP ?
- Why RSVP-TE ?
- What is Required ?
- Solution Overview
- Standards Status
- Applicability

MPLS 2004

What is a P2MP TE LSP ?

- Point to Multipoint Label Switched Path (LSP)
  - Efficient traffic replication in the network
  - Application agnostic
- Set up with TE constraints
  - May involve resource reservations throughout the network
  - Determine path of these P2MP TE LSPs
- RSVP-TE Signaling
  - Enhancements to P2P (GMPLS) RSVP-TE

What is P2MP MPLS TE ?

Source : PE1
Destinations: PE2, PE3, PE4, PE5
P1, P2 : Branch nodes

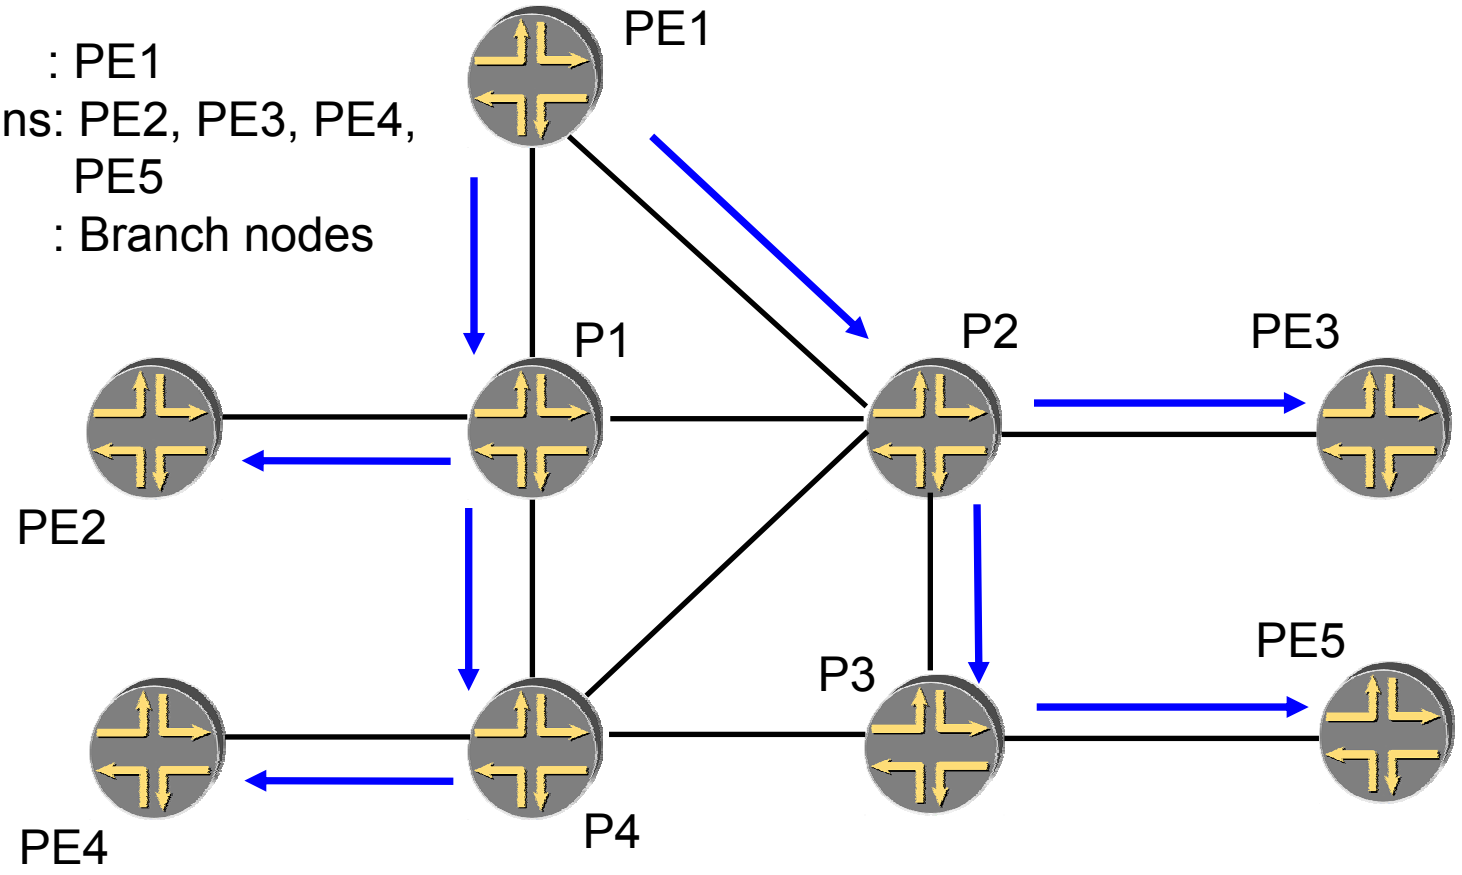

Why RSVP-TE ?

- What are the choices ?
  - RSVP-TE
  - PIM
- Why is RSVP-TE a better fit ?

RSVP-TE versus PIM for P2MP TE

RSVP-TE

- Has resource reservation mechanisms
- Supports explicit/record routing along paths that may be different from hop-by-hop IP routing
- P2MP LSP is signaled by the root and hence allows flexible P2MP computation algorithms
- Fast reroute and Make before-break capabilities

PIM

- No resource reservation mechanisms
- No equivalent support
- Receiver initiated trees are better in terms of scaling and responsivness but limited in tree computation flexibility
  - Do not support Minimum cost trees
- PIM capabilities are NOT oriented for TE purposes !

Requirements for P2MP TE extensions

- P2MP TE Tunnels should be identified by unique P2MP ID
- P2MP TE LSP Tunnel establishment, teardown, and modification mechanism
  - should support grafting/pruning mechanism
  - non-disruptive (forwarding/control) for other P2MP sub-trees
- P2MP TE LSP Tunnel explicit routing support
  - provide a means of establishing arbitrary P2MP TE LSP Tunnel, e.g. cost minimum tree or delayed bounded tree
  - explicit routing with loose hops and widely scoped abstract nodes
- Record routing support
  - information collected and updated during P2MP TE LSP establishment and modification process

Requirements for P2MP TE extensions

- Failure Reporting and Error Recovery
  - node must report all errors to ingress/branch node to initiate fast recovery around the failure
- Call Admission Control and QoS Control mechanism
  - must support resource sharing and exclusive resource utilization
  - must be applicable to Diffserv-enabled networks and SHOULD satisfy the DS-TE requirements
- P2MP TE LSP Tunnel parameters
  - no variation of attributes along the P2MP LSP TE Tunnels
  - homogenous QoS
- Fragmentation of protocol message(s)
  - when a single protocol message cannot contain all the information

Requirements for P2MP TE extensions

- Re-optimization of P2MP TE LSP Tunnels
  - must support Make-before-break (whole and partial operation)
- Support of Multi-Area/-AS and hierarchical P2MP TE LSP Tunnels
- Routing advertisement of P2MP (node) capability
  - node ability to support branching/act as an egress and a branch
- GMPLS
  - Solution for MPLS P2MP TE when applied to GMPLS P2MP PSC or non-PSC MUST be backward and forward compatible with P2P GMPLS features
- Backward compatibility and interoperability with (G)MPLS TE capable legacy nodes

Solution Overview

- Terminology
- Problem Statement
- Mechanisms
- Examples
- Comments

Solution Terminology

- Ingress LSR: LSR responsible for initiating the signaling messages that set up the P2MP TE LSP (also referred to as source or root)
- Egress LSR: one of potentially many destinations of the P2MP TE LSP (also be referred to as leaf nodes or leaves)
- branch LSR: an LSR that has more than one directly connected downstream LSR
- P2MP TE LSP: A traffic engineered label switched path that has one unique ingress LSR (also referred to as the root) and one or more egress LSRs (also referred to as the leaf)
- P2MP ID: A unique identifier of a P2MP TE LSP, that is constant for the whole LSP regardless of the number of branches and/or leaves.
- P2P sub-LSP: Label switched path from the ingress LSR to an egress LSR

Solution Terminology

Ingress LSR: PE1
Egress LSR : PE2, PE3,
PE4, PE5
P1, P2 : Branch nodes

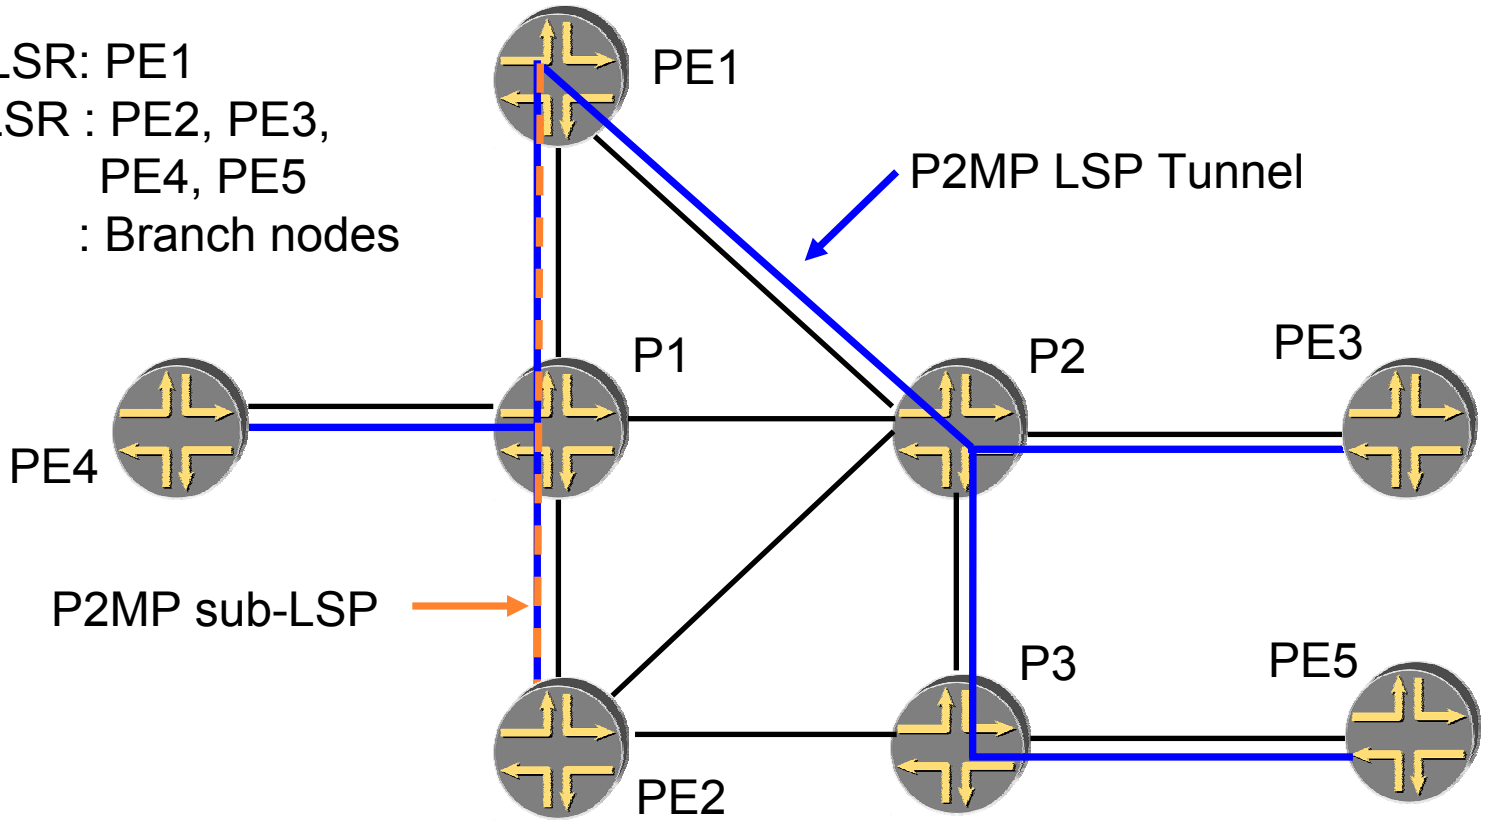

Problem Statement

- The practical problem is to introduce multicast functionality in the MPLS data plane
  - Optimize data plane for high volume multicast
- P2MP TE is performed in the data plane: control plane uses P2P sub-LSPs as building blocks
- Support full and summary refresh mechanisms
- Address message fragmentation (message size > MTU)
- Support aggregated state management and incremental state updates

Problem Statement

- Minimize enhancements to current RSVP-TE
- Operational simplicity
  - P2P RSVP-TE is deployed and understood
  - Leverage the existing control plane model
- Protocol simplicity
  - Minimize complex protocol changes
- Implementation simplicity
  - Minimize changes to existing software: Less Bugs !

Solution Mechanisms

- Building blocks
  - P2MP Tunnel
  - P2MP LSP
  - P2P sub-LSP
- Path Messages
- Resv Messages
- Fast-reroute
- Make-before-break

Solution Mechanism: P2MP Tunnel

- Determines set of destinations terminating the unidirectional traffic flow and for which resource reservation is required
- May comprise multiple P2MP LSP Tunnels (at least one)
- Identified by the P2MP SESSION Object which includes
  - P2MP ID: identifies the destination of the P2MP tunnel
  - Tunnel ID: 16 bit identifier
  - Extended Tunnel ID: local IPv4/IPv6 Address or left unspecified

Solution Mechanism: P2MP LSP Tunnel

- A specific instance of a P2MP Tunnel determined by the source of the traffic flow
- May comprise multiple P2P sub-LSPs
- Identified by the P2MP Tunnel SESSION and P2MP SENDER_TEMPLATE object combination
- P2MP SENDER_TEMPLATE
  - Identifies the sender (ingress)
  - Includes
    - Source IPv4/IPv6 address
    - LSP ID

Solution Mechanism: P2P Sub-LSP

- LSP from the ingress LSR to a particular egress LSR
- A P2MP LSP Tunnel comprises multiple P2P sub-LSPs
- A P2P sub-LSP is represented by
  - P2P sub-LSP object
  - Sub-explicit route object
- P2P sub-LSP Object
  - Identifies a P2P Sub-LSP
  - Egress LSR Destination address
  - P2P sub-LSP identifier (sub-LSP ID)
- Sub-Explicit route
  - Represents the explicit route from ingress LSR to the egress LSR
  - May be compressed

Solution Mechanism: Path message

- One P2MP Tunnel LSP can be signaled using multiple Path message
- Each such Path message can signal multiple P2P sub-LSPs
- Limiting cases
  - A separate Path message for each P2P sub-LSP
  - A single Path message for all P2P sub-LSPs

Multiple Path Messages: Example

P2MP Tunnel: ID_1 {PE2,PE3,PE4}

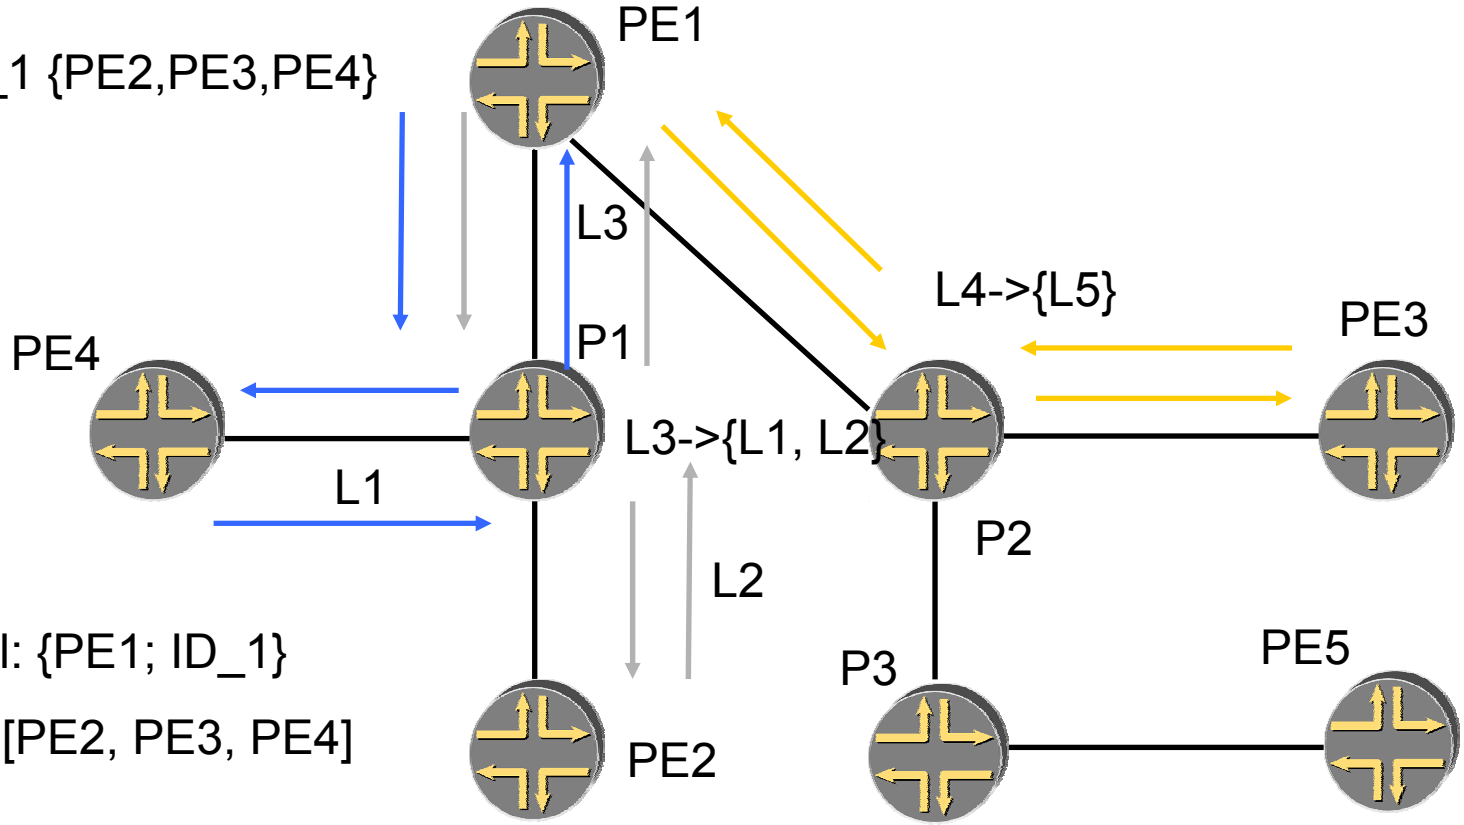

P2MP LSP Tunnel: {PE1; ID_1}

P2P sub-LSP list: [PE2, PE3, PE4]

Multiple versus Single Path Message

P2P sub-LSP PE4 ERO: {P1,PE4}
P2P sub-LSP PE2 ERO: {P1,PE2}

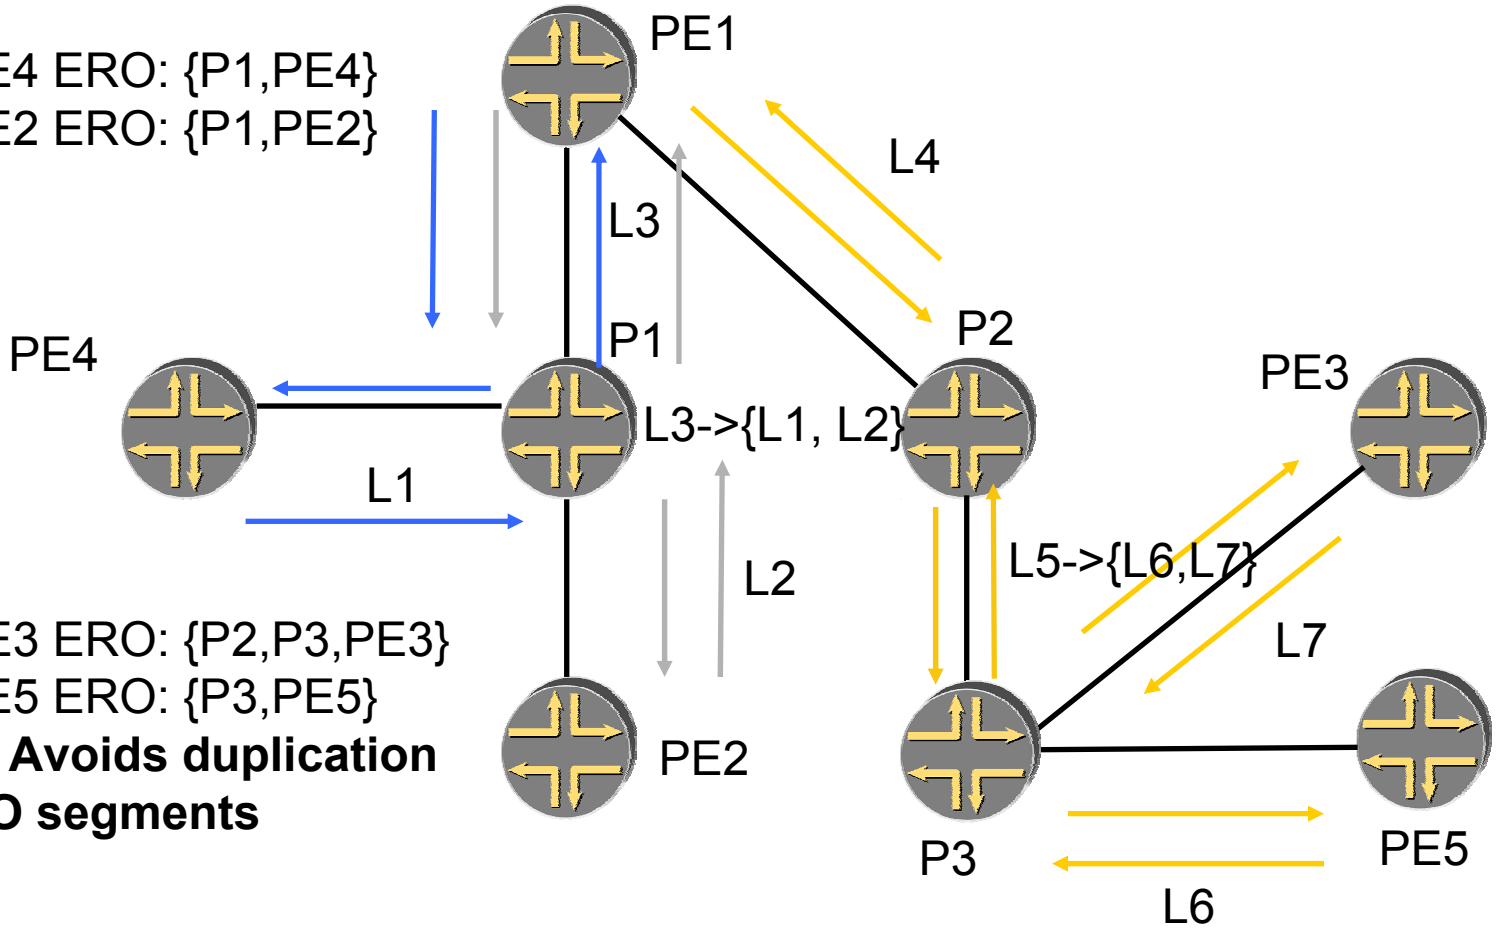

P2P sub-LSP PE3 ERO: {P2,P3,PE3}
P2P sub-LSP PE5 ERO: {P3,PE5}
Compression = Avoids duplication of common ERO segments

Single Path Message: Example

P2P sub-LSP PE4 ERO: {P1,PE4}
P2P sub-LSP PE2 ERO: {P1,PE2}

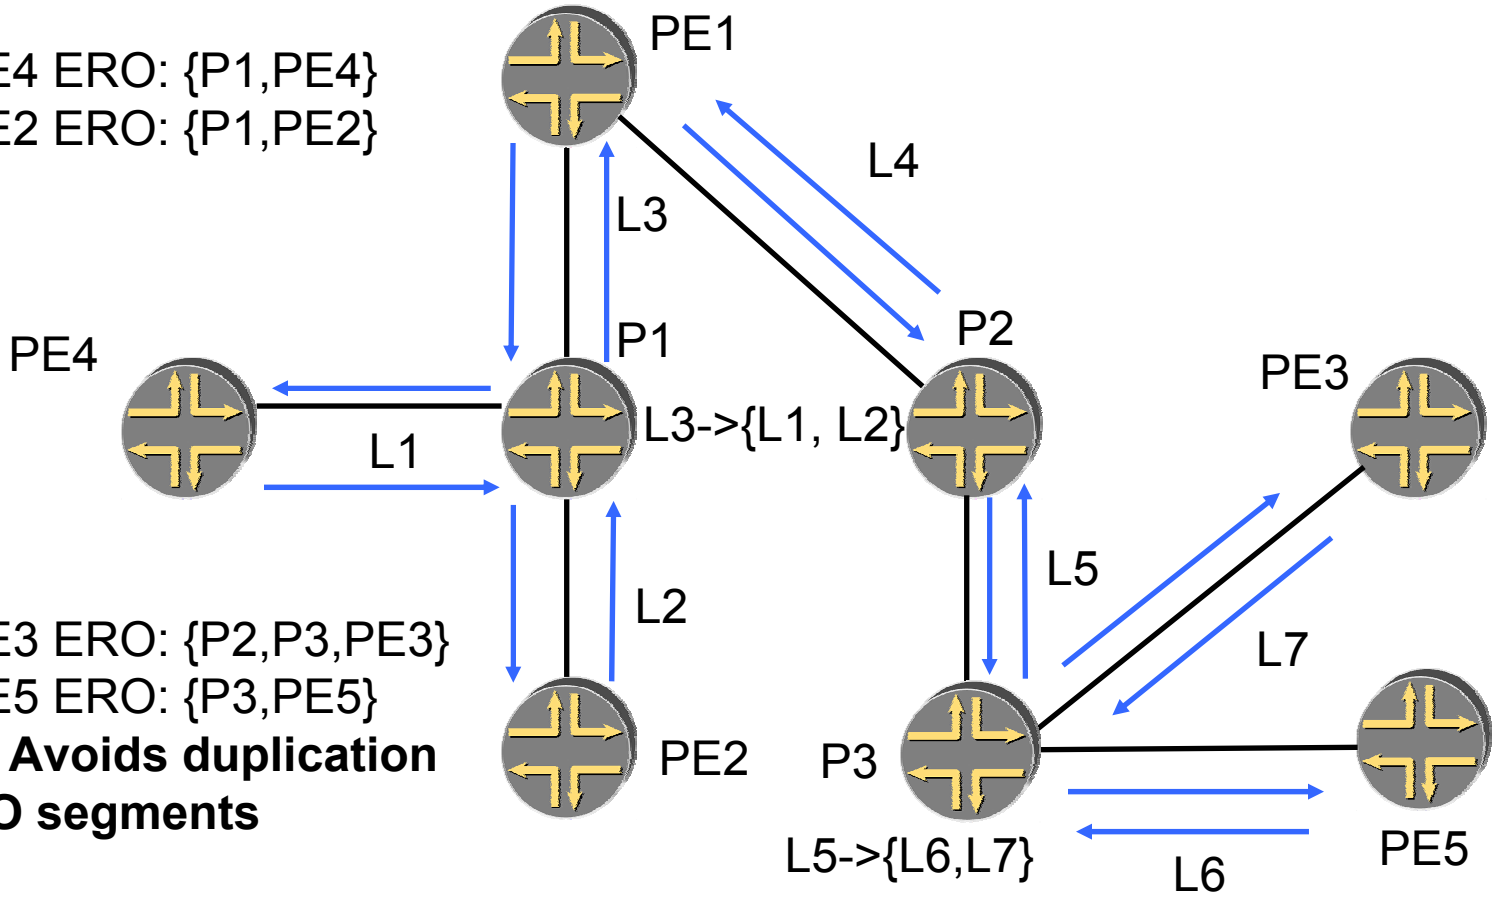

P2P sub-LSP PE3 ERO: {P2,P3,PE3}
P2P sub-LSP PE5 ERO: {P3,PE5}
Compression = Avoids duplication of common ERO segments

Standards Status

- Work done as part of the IETF MPLS WG charter
- Requirement document under last call
  - URL: <http://www.ietf.org/internet-drafts/draft-ietf-mpls-p2mp-requirement-04.txt>
  - revisited version under mailing list discussion
- Solution document (individual status)
  - URL: <http://www.ietf.org/internet-drafts/ draft-raggarwa-mpls-rsvp-te-p2mp-00.txt>
  - virtual team of ~ 30 people working on the document
  - new version to be submitted for the next IETF meeting (Washington DC, Nov'04)

Applicability – Layer 2 Multicast over P2MP TE LSP Tunnel

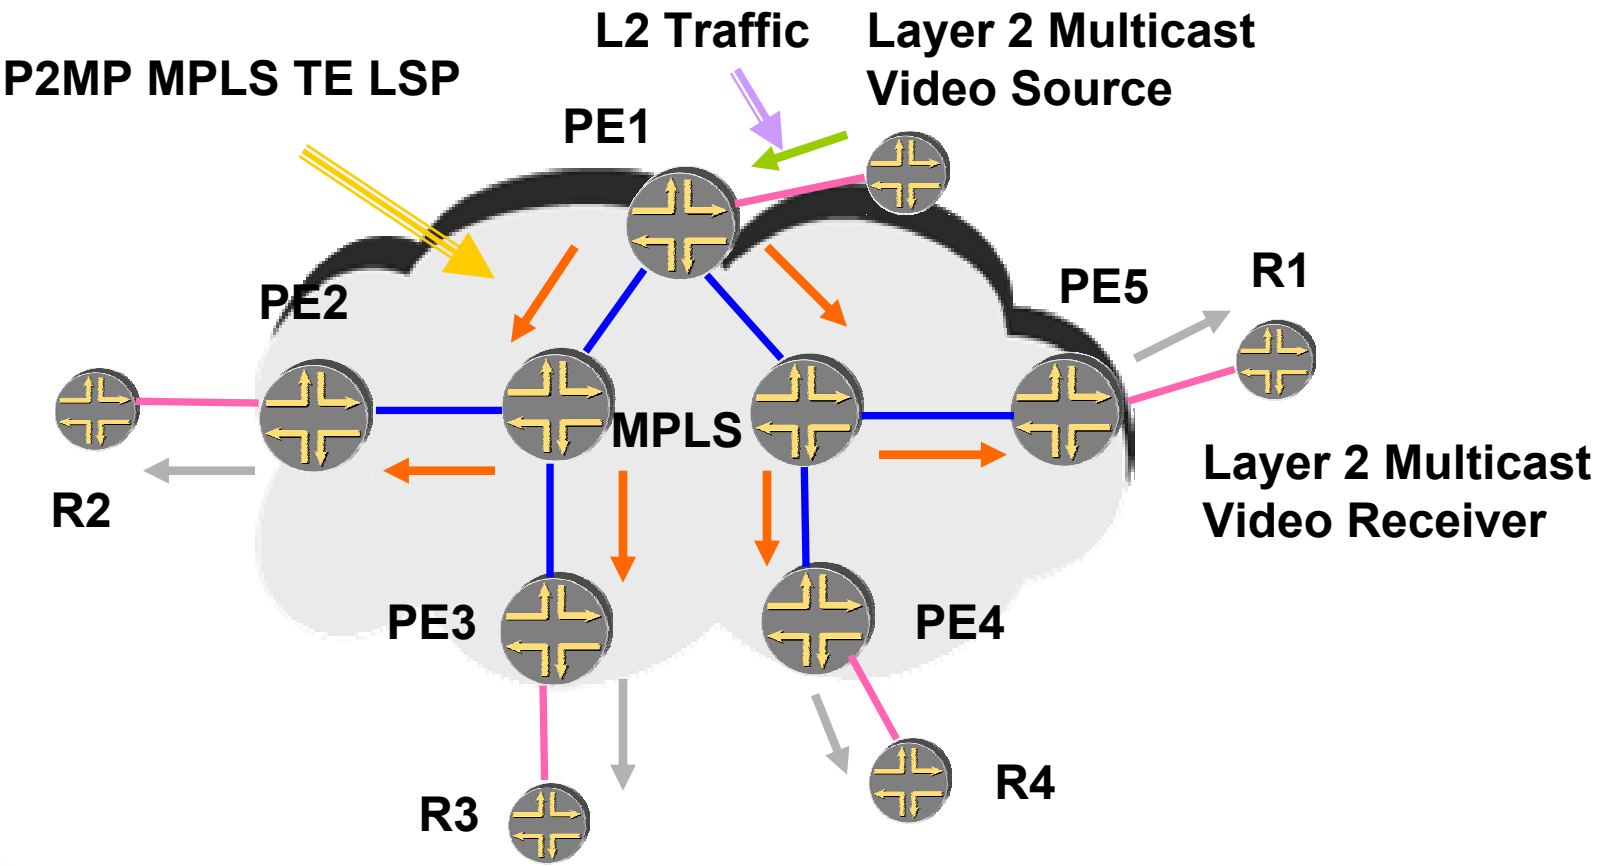

Applicability – Layer 2 Multicast over P2MP TE LSP Tunnel

- Goal is to retain all the functionality available to layer 2 services as they migrate to IP/MPLS
  - P2MP functionality is offered by ATM networks
  - P2MP TE is a missing piece in the layer 2 service migration to IP/MPLS
- A Layer 2 interface can be cross-connected to a P2MP LSP
- TE requirement
  - QoS guarantees: strict SLAs for broadband video traffic
  - Protection: Fast reroute

Applicability - IP Multicast Over P2MP MPLS TE LSP Tunnel

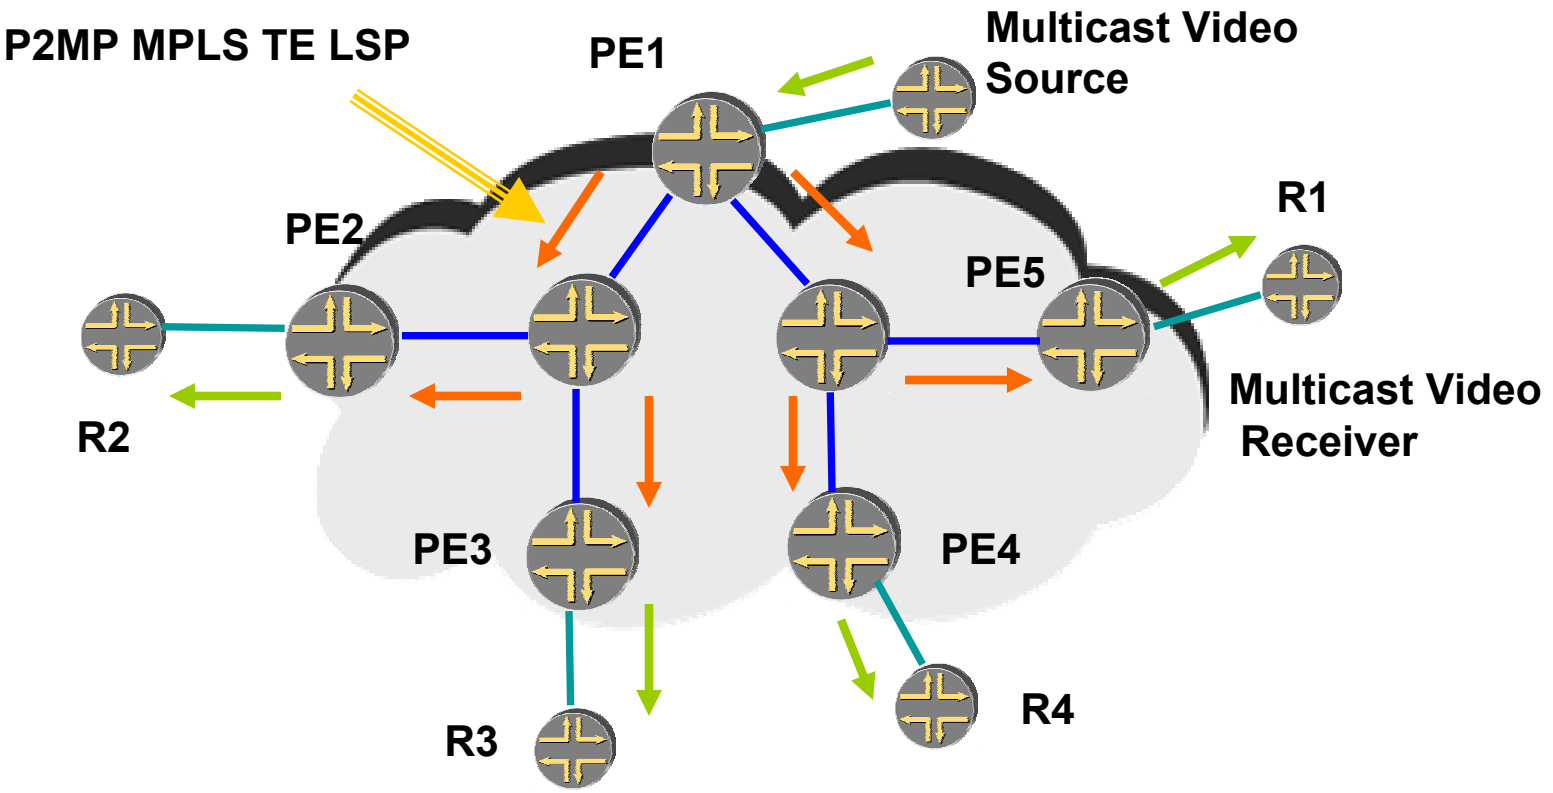

Applicability - IP Multicast Over P2MP MPLS TE LSP Tunnel

- TE for broadband video multicast traffic
  - QoS for content distribution
  - Protection: Fast Reroute
- Multicast (PIM-SM) free core
  - Keeping multicast routes out of the core
- Eliminates the need to use BGP in the core to distribute unicast routes used by multicast RPF
  - Particularly useful if the core is BGP free for unicast routing (e.g. by running RSVP-TE)

Thanks for your attention

… Questions

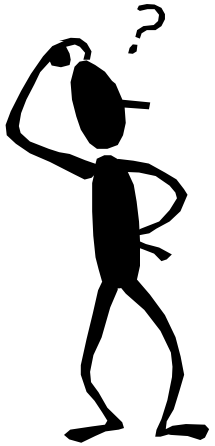

MPLS 2004

References

- D.Awduche et al., *RSVP-TE: Extensions to RSVP for LSP Tunnels*, RFC 3209, December 2001.
- L.Berger (Editor) et al., *Generalized Multi-Protocol Label Switching (GMPLS) Signaling – Resource Reservation Protocol - Traffic Engineering (RSVP-TE) Extensions*, RFC 3473, January 2003.
- S.Yasukawa (Editor) et al., *Requirements for Point to Multipoint Traffic Engineered MPLS LSPs*, Internet Draft, Work in progress, draft-ietf-mpls-p2mp-requirement-04.txt, September 2004.
- R.Aggarwal, D.Papadimitriou, S.Yasukawa (Editors) et al., *Extensions to RSVP-TE for Point to Multipoint TE LSPs*, Internet Draft, Work in progress, draft-raggarwa-mpls-rsvp-te-p2mp-00.txt, July 2004.