RAPID:
Advanced IP/LDP Fast-Reroute

- By Alia K. Atlas
aatlas@avici.com

Outline of Talk

- Motivation
- RAPID
  - Loop-Free Alternates
  - Improving Coverage: U-Turn Alternates
- Network Design Considerations
- Network Examples

Network Disruptions are Daily Events

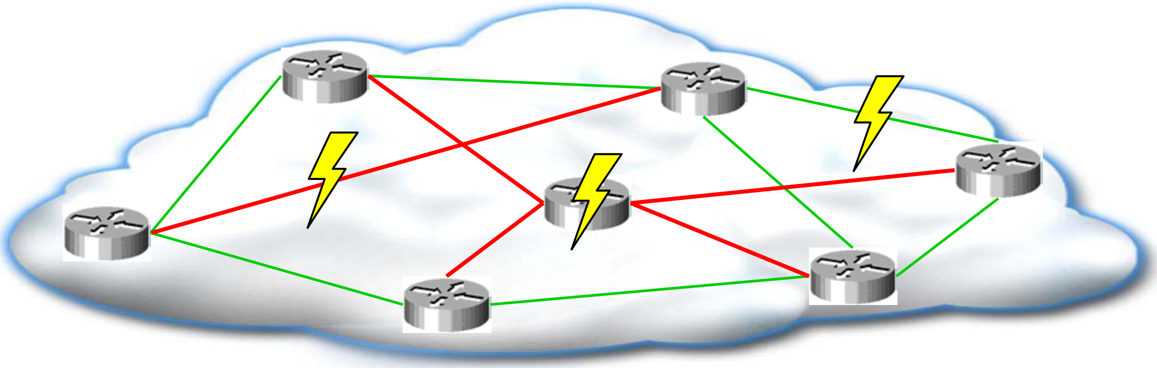

Causes

- Router Failure
- Link Failure
- Disruptive Operations (sw upgrades, config. changes, …)

Service Impact

- Loss of traffic for 10s of seconds
- Disruption of Real-Time Services (voice calls, gaming sessions, video, ATM)

Business Impact

- SLA Penalties
- Customer Service/Maint. Issues
- Customer Churn
- Inability to support High-Margin Real-Time Services

MPLS 2004

AVICI SYSTEMS

Traffic Convergence Goal: < 50 ms

- To support a multi-service network, need to minimize service interruption.
- Network Failures cause service interruption.
  - Node Failure: Avoid disruption with Non-Stop Routing or minimize traffic loss during convergence.
  - Link Failure: Minimize traffic loss during convergence.
- Traffic Convergence
  - IGP Convergence: SPF provides the basis for all other protocols so must be very fast.
  - BGP Convergence: Using forwarding-plane indirection to IGP next-hop allows traffic restoration for BGP learned destination ***before*** BGP re-computation occurs for many failure scenarios.
  - LDP Convergence: Requires IGP SPF results to install new forwarding plane state.

RSVP-TE Fast-Reroute

PROs

- Provides Link, Node, and SRLG Protection
- Provides 100% Coverage except for Ingress & Egress Node Failures
- Understood and Deployed Technology

CONs

- Overlay Network -> Scalability Concerns
- Operator Complexity – Many options & controls
- If No Need for TE, introduces new protocol just for resilience.
- Area & AS border routers are single points of failure (until inter-area/inter-AS TE FRR is practical)
- Failure of Tunnel Ingress and Egress Cannot Be Protected Against

Outline of Talk

- Motivation
- RAPID
  - Loop-Free Alternates
  - Improving Coverage: U-Turn Alternates
- Network Design Considerations
- Network Examples

RAPID Functionality

- Pre-compute an alternate next-hop for each destination.
- On local failure, rapidly switch affected traffic to use alternate next-hop.
- Switch to new primary next-hop(s) when safe.

RAPID Basic Concept

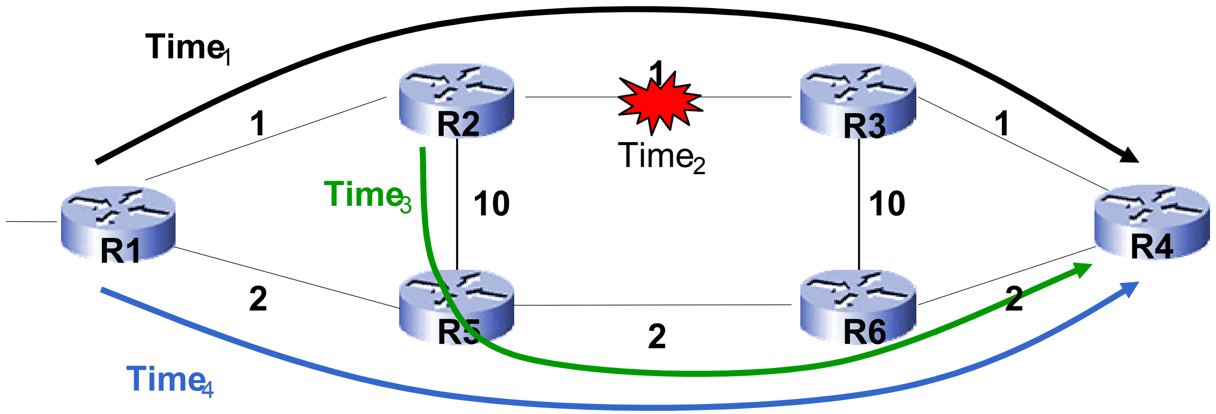

- R2 pre-computes alternate IGP path for R4 traffic in case link fails
- Failure detection triggers R2 to failover to alternate path
  - Failover occurs in **milliseconds** for both IP and LDP
  - R2 also signals failure and runs SPF, but that time does not impact traffic
- Some time later R1 will run a new SPF and send traffic to R5

Loop-Free Alternates

A loop-free alternate is a neighbor of the router R2 whose shortest path to the destination doesn't go through R2.

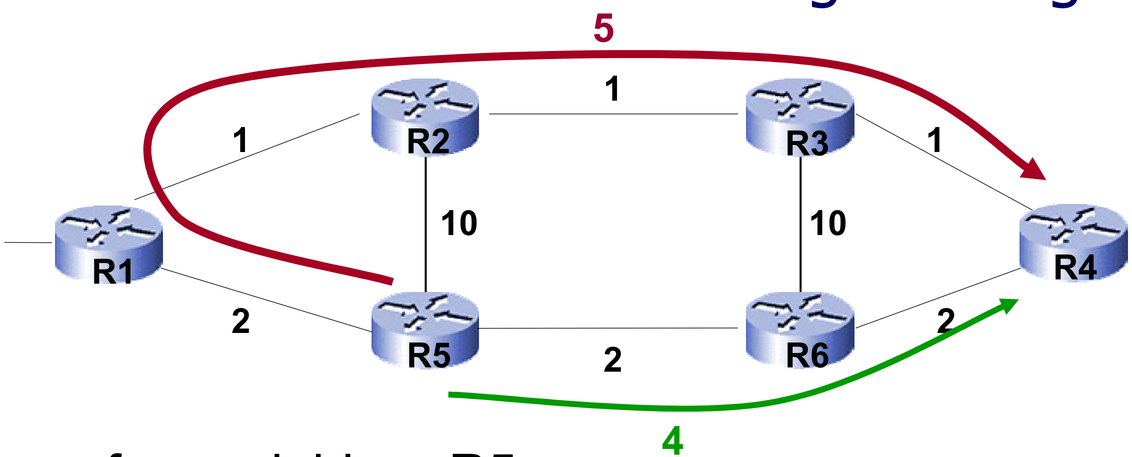

- R2 can find a loop free neighbor: R5
- R5 is loop-free, because the distance from R5 to R4 is less than the distance from R5 to R2 plus the distance from R2 to R4.
- R2 can know all this because it has the full topology

MPLS 2004

AVICI SYSTEMS

LDP Fast-Reroute Example

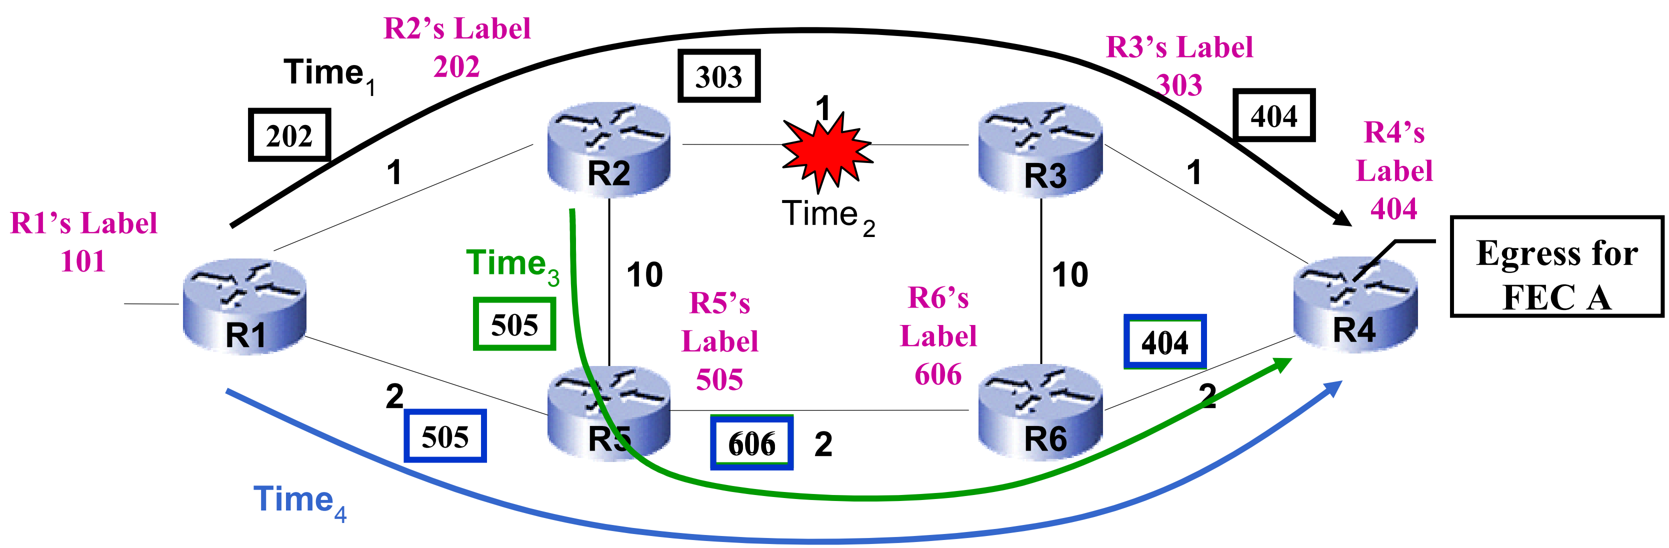

- LDP Must know FEC Label before Failure:
  - Liberal Label Retention Mode
  - Downstream Unsolicited
  - Send Label Mappings to All Neighbors

Network Coverage for Loop-Free Alternates

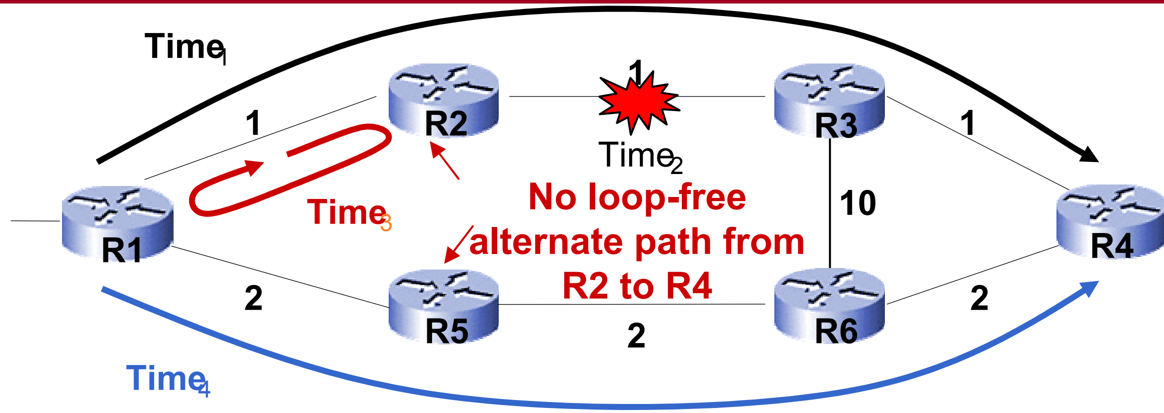

- Many networks don't have alternate links at all points
  - Loop-free IP/LDP Fast-Reroute provides an average 79% failure coverage
  - But 79% of the source/destination pairs does not equal 79% of the traffic – could be a lot less if the 21% unprotected are important source/destination pairs.
- If R2 could use R1 as an alternate, the coverage would increase dramatically

MPLS 2004

AVICI SYSTEMS

Improving Network Coverage

Goals:

- Increase Coverage
- No LDP Extensions or Additional Sessions to Manage
- No Overlay Solution
  - RSVP-TE Fast-Reroute already available
- Minimal Signaling Extensions
  - None at time of failure
- Minimal Operational Complexity
  - No tunnel management overhead

Breaking the loop: U-Turn Alternates

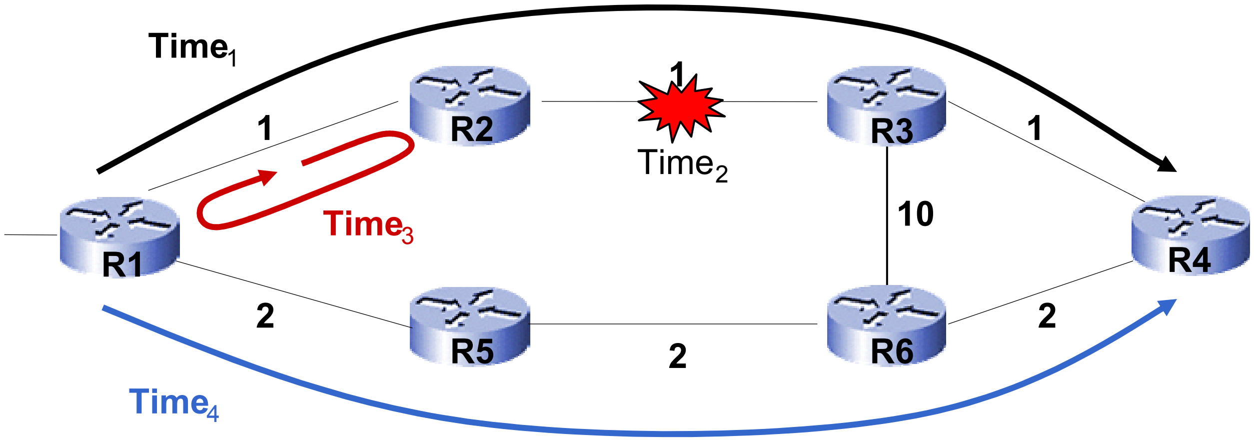

R1 can provide a U-turn alternate to R2 if:

- R1 itself has a loop-free node-protecting alternate path to reach R4
- R1 can break the loop
- R1 is a U-Turn neighbor
  - R2 is R1's primary neighbor to destination R4
  - R1 is capable of breaking the loop

U-Turn Alternates

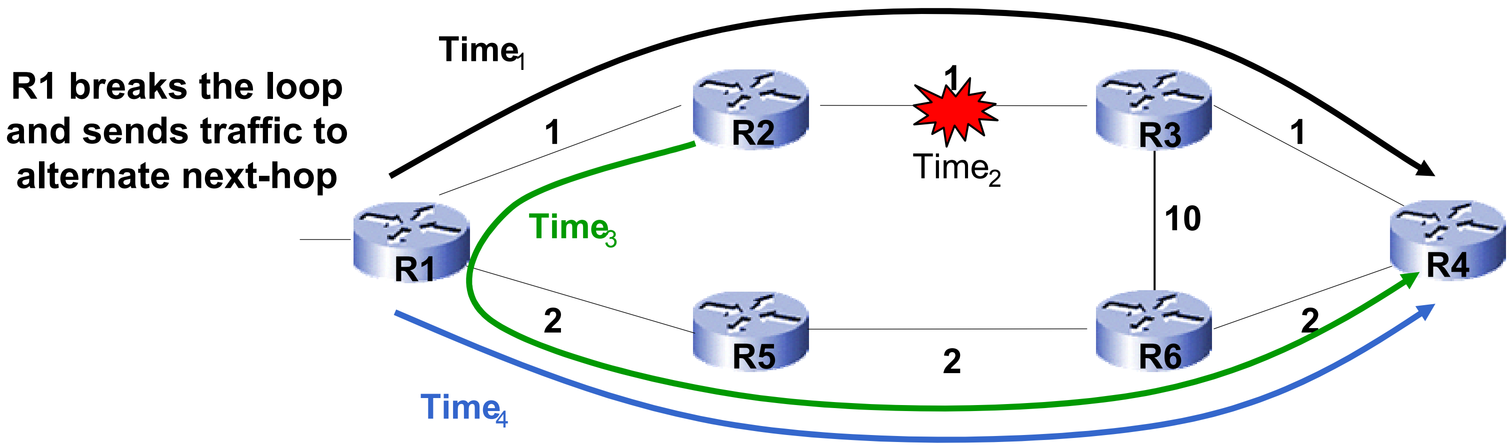

- R1 can break the loop, if its hardware can identify traffic as U-turn traffic
  - Traffic from primary neighbor
  - R1 can require specific well-known marking
- R1 sends U-turn traffic to alternate next-hop
- R1 has to support U-turn alternates.
- Thus new IETF drafts to signal capabilities: OSPF, ISIS

LDP Example for U-Turn Traffic

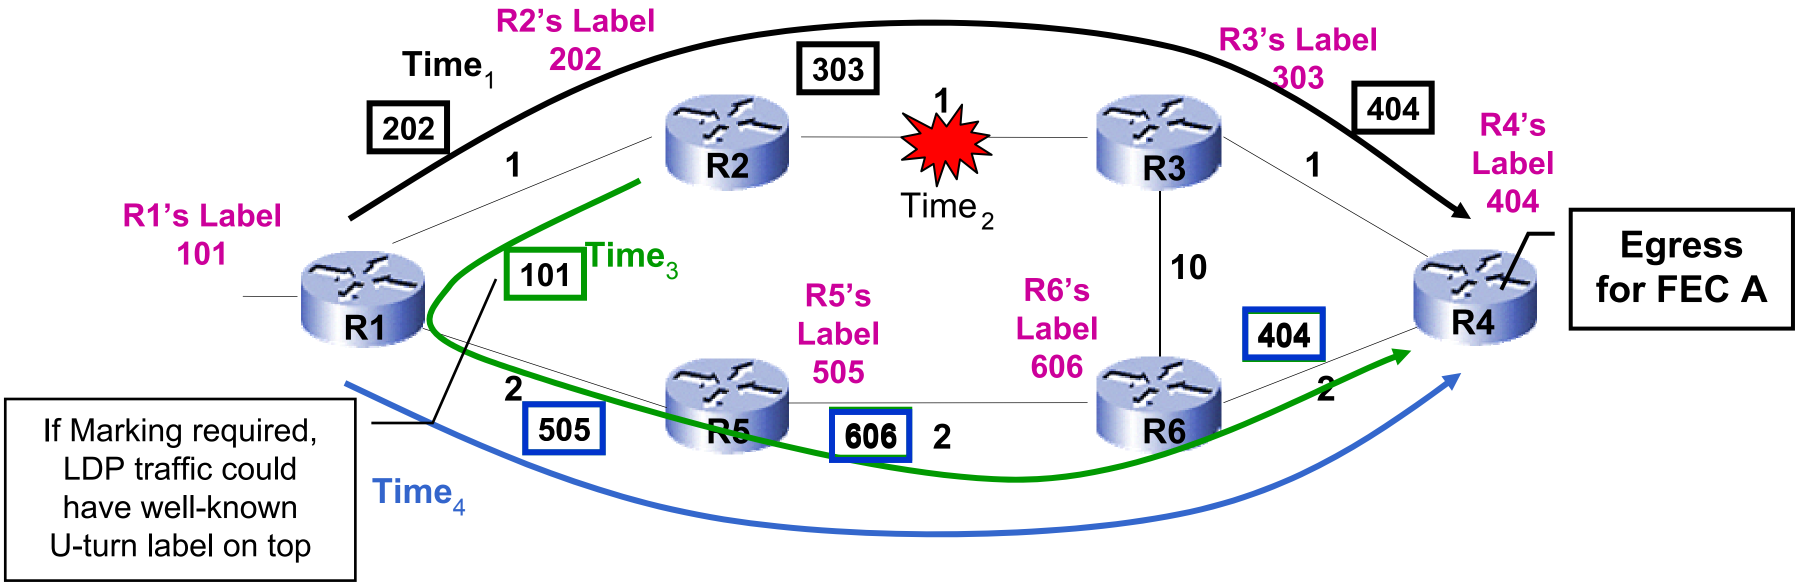

- On local failure, R2 sends traffic for FEC A to R1.
- R1 receives the traffic, identifies it as U-turn traffic, and sends it to R1's alternate, which is R6

Marking U-turn Traffic

Benefits of Explicit Marking of U-turn traffic:

- It is easier for the receiver to identify in hardware. Unmarked traffic can take default forwarding path.
- This covers more general topologies for Nodes on a broadcast interface.
- Traffic which is PHP from a RSVP-TE LSP will not be mistaken for U-turn traffic.

IP/LDP Fast-Reroute Standardization

- Draft-ietf-rgwg-ipfrr-framework-01.txt defines a framework for IP fast-reroute
- Draft-ietf-rtgwg-ipfrr-base-spec-01 defines loop-free alternates and how to use them.
- Draft-atlas-ip-local-protect-uturn-00 defines U-turn alternates and common rules for selecting alternates.
  - Using U-turn alternates increases protection coverage from 79% average, to 98% average
  - IETF drafts define signaling of Router's IP FRR capability, and per link capability for ISIS (draft-martin-isis-local-protect-cap-00) and OSPF (draft-atlas-ospf-local-protect-cap-00)
- Agreement that an additional advanced mechanism is needed beyond loop-free.
  - Other drafts proposing using TE tunnels or tunnels plus directed forwarding to extend coverage.

Outline of Talk

- Motivation
- RAPID
  - Loop-Free Alternates
  - Improving Coverage: U-Turn Alternates
- Network Design Considerations
- Network Examples

Network Design Considerations: Coverage

Consider Each Source to Each Destination

- Importance to Total Coverage Depends on
  - Expected Traffic Flow (Amount Affected)
  - Expected Traffic Service Class
- This can be easy: GUI tools exist

Example:

- If goal is to protect L3VPN traffic between PE routers, then don't worry about non-PE destinations

Potential Traffic Flow for Capacity Planning

- May need to do capacity planning with a Traffic Matrix
  - For each potential failure (link or node), determine added traffic load on alternate paths
  - True for any IP or MPLS FRR mechanism
  - Only important traffic classes need be considered

MPLS 2004

AVICI SYSTEMS

Network Design: Resilience After Failure

- All Fast-Reroute (IP, LDP and RSVP-TE) techniques assume a single failure.
- Consider Expected Mean Time Between Failures
- Plan network for resiliency after sequential failures.
  - Do all routers still have multiple paths after a failure?
  - Are alternates available for all important source to destination traffic?

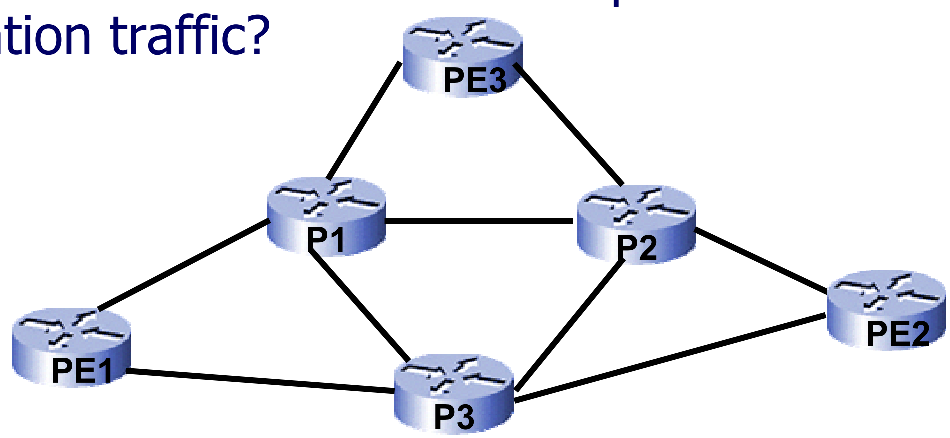

When a link from a PE to a P fails, no alternate possible if link from same PE fails.

Network Example: IP/LDP Backbone

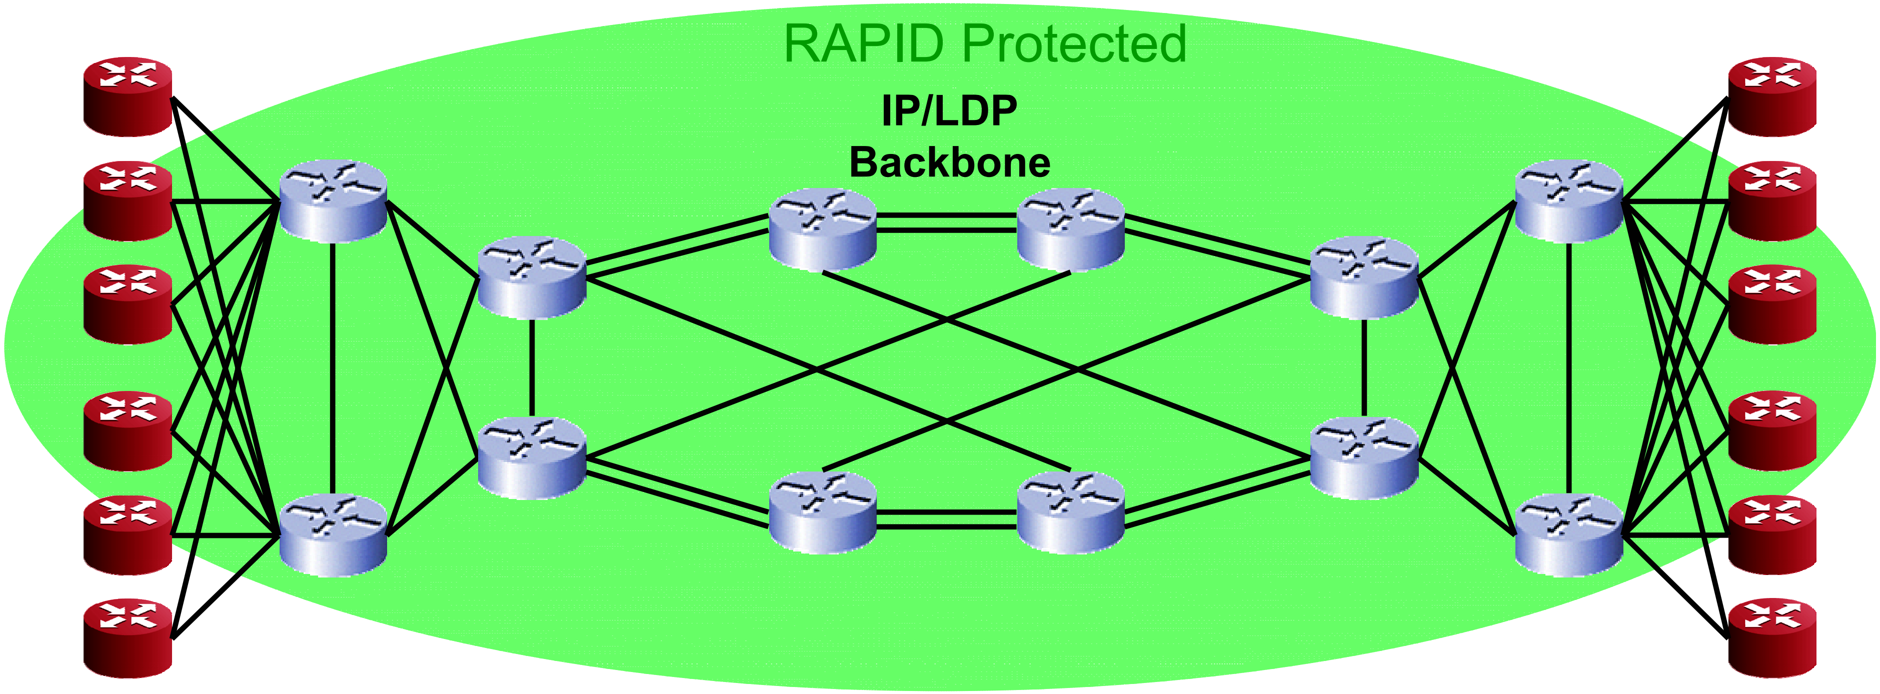

- A highly redundant IP/LDP Backbone – no MPLS-TE
- RAPID provides protection
- Coverage is very good, if the link redundancy is sufficient

MPLS 2004

AVICI SYSTEMS

Network Example: RSVP-TE Backbone

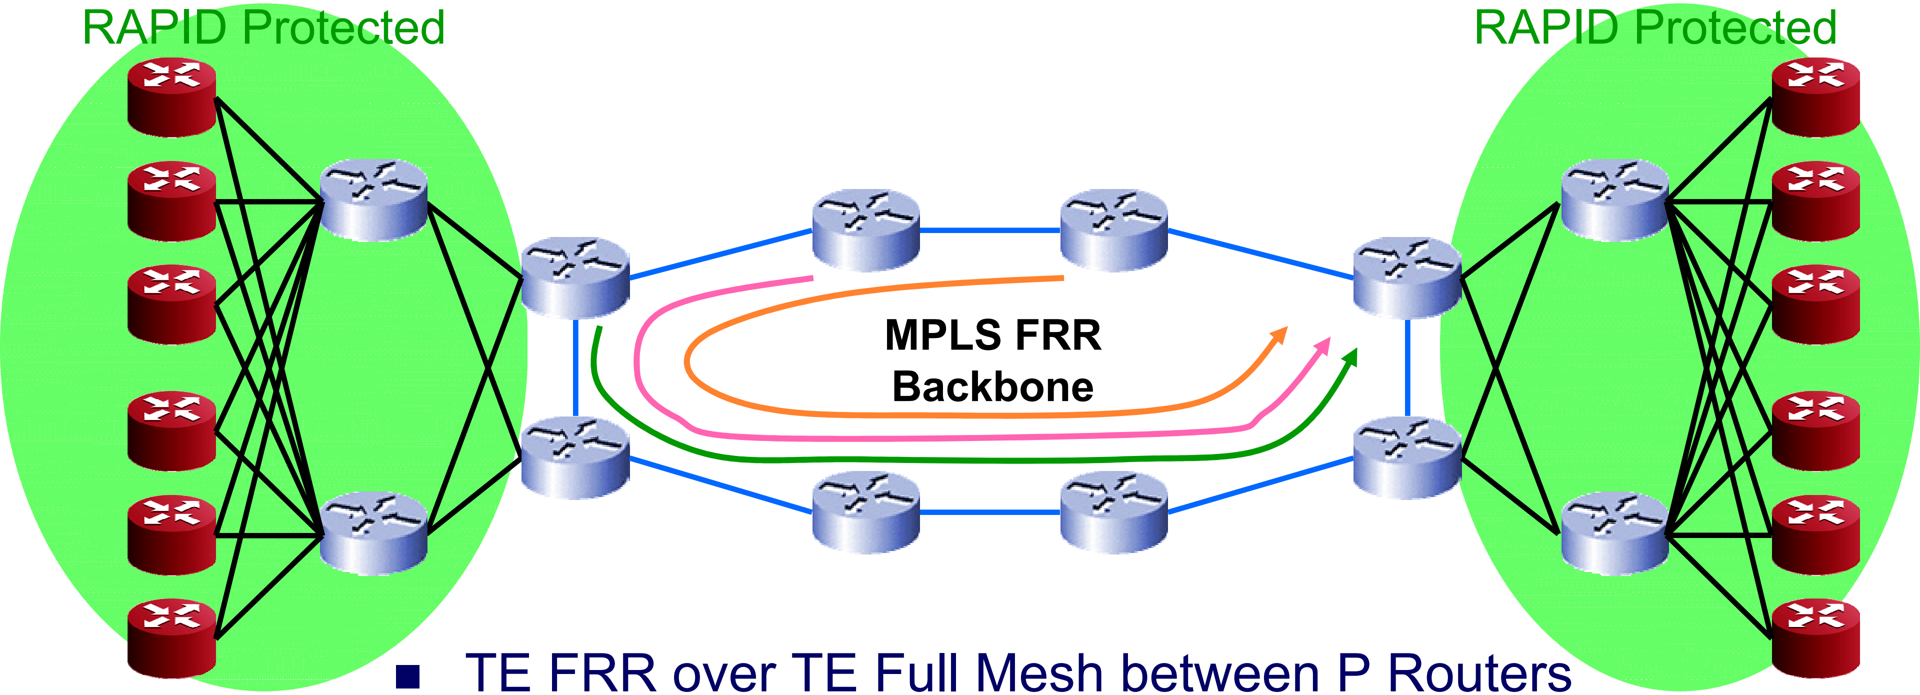

- TE FRR over TE Full Mesh between P Routers
- PE to P connections protected with RAPID
- Non-Stop Routing to Protect Against Egress Node Failures

Network Example: Transit MPLS Backbone

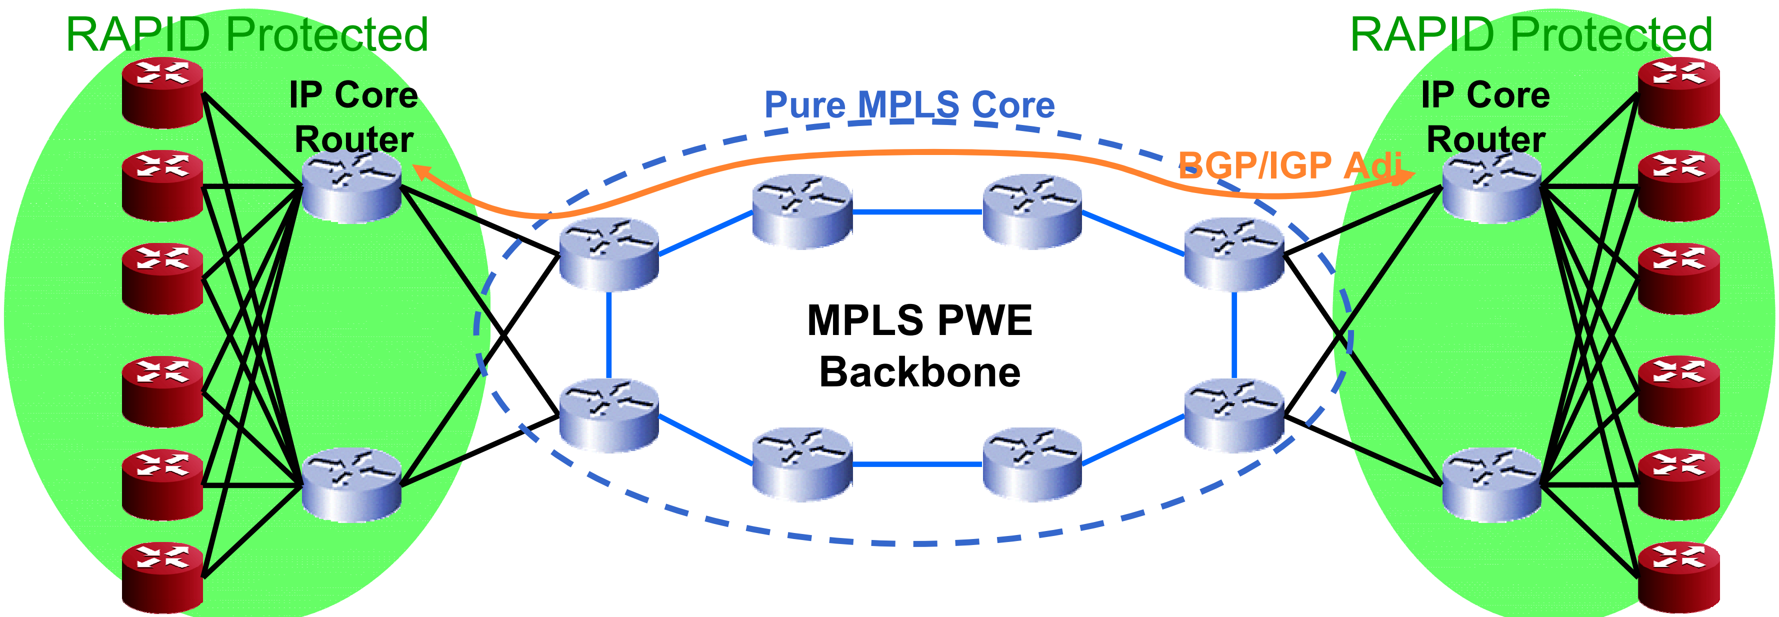

- Separate Transit MPLS PWE/L2-VPN Core design
- IP Core routers do not "see" this MPLS core – they think they have direct connections to the other IP Core routers
- MPLS Backbone can be protected by FRR or Pre-Signaled standby tunnels
- RAPID protects IP Core Routers (not Transit Backbone Routers) and Edge

MPLS 2004

AVICI SYSTEMS

RAPID Summary

- Provide < 50ms traffic convergence in the event of a node or link failure for IP and LDP traffic
- Not all routers have to support RAPID to get the benefits of the basic mechanism (loop-free alternates)
- U-Turn Alternates expand the potential failure coverage on networks
- Simple to configure, manage and interoperate
- Can be incrementally deployed – the benefit of U-Turn Alternates will be seen as more routers are deployed with this feature in the network
- Can be used in conjunction with other mechanisms (e.g., RSVP-TE FRR, Non-Stop Routing, etc.)

Goal is to simplify operations and provide local protection