



Flexible Algorithm (Flex-Algo) for Segment Routing networks

Eduardo Haro – Sr. Sales Engineer

eharo@juniper.net



What is Flex-Algo?

Why Flex-Algo?

Flex-Algo components

Use cases and applications

BGP-CT for inter-domain

What is Flex-algo?

- IGP traditionally computes best effort path based on IGP metric
- Service Providers traditionally use RSVP-TE to steer traffic over a path using different metrics or constraints than shortest IGP path (TE – traffic engineering – path)
- Flexible Algorithm (Flex-Algo) allows IGP themselves to compute TE paths based on various constraints:
 - Metric type: IGP metric, TE metric or latency metric
 - Link color or SRLG constraints
 - Avoid node constraints
- It's possible to create thin network slices based on different requirements (bandwidth, TE or delay).
- Multiple data-planes support such as SR-MPLS, SRv6 (original draft) or IP (more recent draft)
- Being standardized in IETF LSR WG draft-ietf-lsr-flex-algo and draft-ietf-lsr-ip-flex-algo
 - First draft back in 2017 by Shraddha Hegde (Juniper) and Arkadiy Gulko (Thomson Reuters) and evolved to draft-ietf-lsr-flex-algo with Cisco also as co-author.

Why Flex-Algo?

- Easier to implement lightweight TE in a scalable way (no need of static config per LSP)
- Allows configuring TE using segment routing even without a network controller
- Strict TE constraints such as avoid specific nodes or links
- Use ECMP (equal-cost multipath) and TI-LFA (backup path) per Flex-Algo, honoring the constraints
- Only uses IGP (ISIS or OSPF), no need of RSVP or LDP
- Uses single label to create TE paths (helps label stack depth issue)
- Interop with multiple vendors was shown at several events, last being EANTC 2021
- Other alternative solutions offers advantages and disadvantages, depending on the network scenario:
 - Static SR-TE
 - Distributed dynamic SR-TE
 - BGP dynamic colored-SR-TE

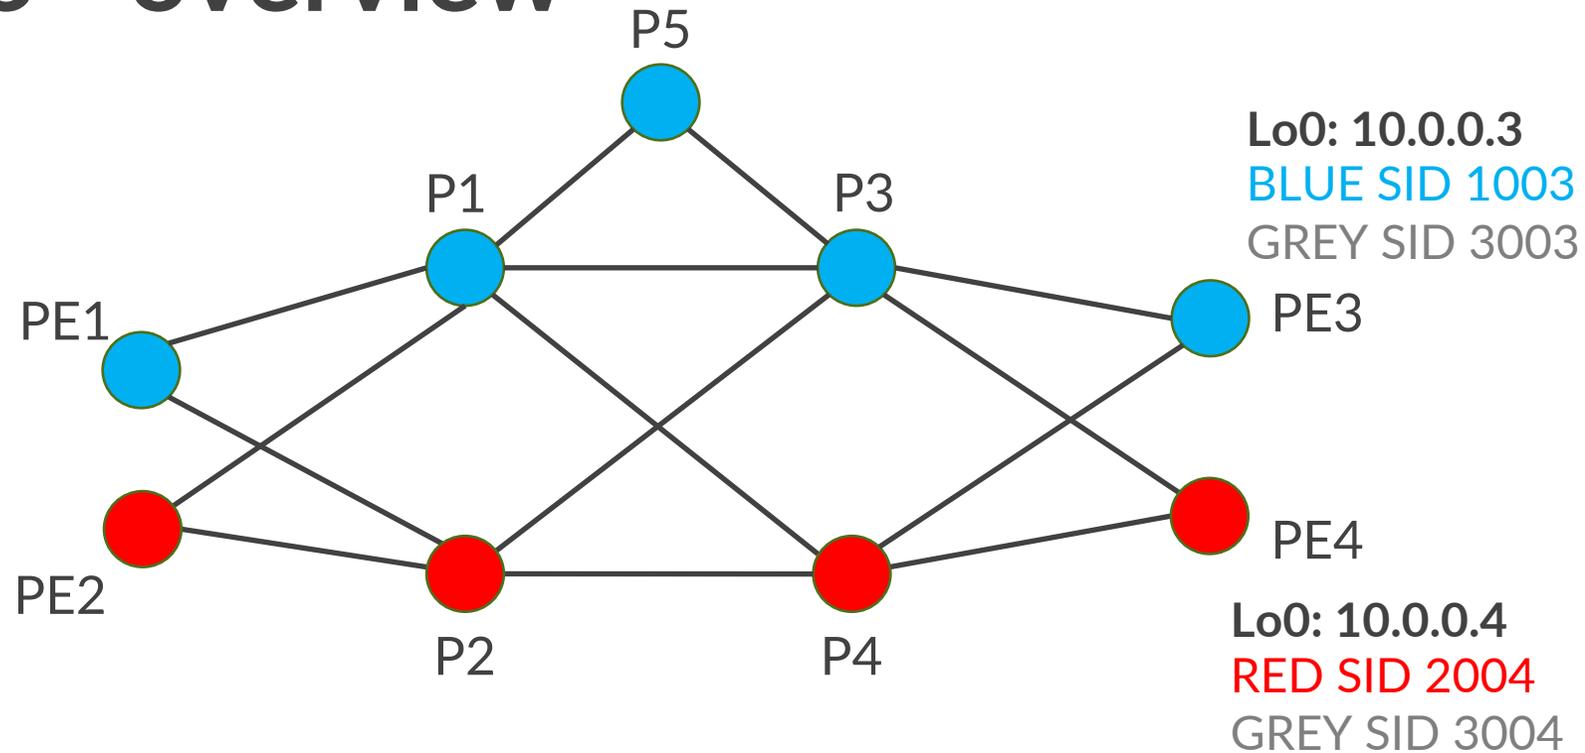
Use cases / applications

- Routing plane separation
 - Multiple routing planes with strict plane separation requirements
- Keep certain traffic isolation between different network domains while merging
- Low latency or high bandwidth routing paths
 - E.g., delay sensitive traffic is routing using a metric proportional to the latency of the link and bulk internet traffic is routing using a metric inversely proportional to the bandwidth of the link
 - In a given network, the two requirements might conflict!
- Path diversity
- Network slicing in transport networks for new emerging use cases (5G, IoT, among other)

Flex-Algo components

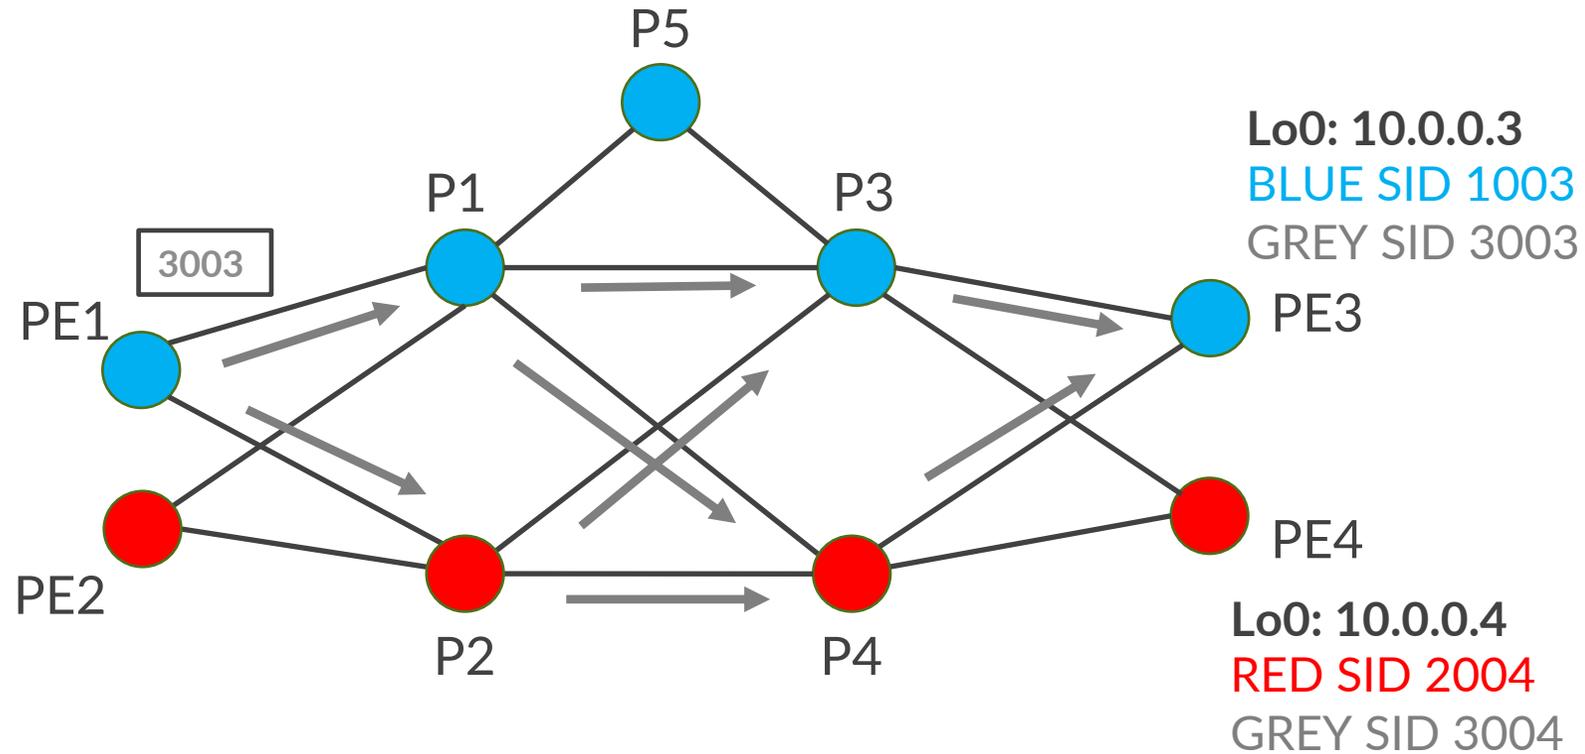
FLEXIBLE ALGORITHM <i>DEFINITION</i>	FLEXIBLE ALGORITHM <i>PARTICIPATION</i>
Used to define and announce topology constraints: link affinity, calculation type, metric type	Used to announce a router's participation in a certain topology
Configured on some devices. Goal is to provide redundancy, but avoid conflicts	Configured on all devices participating in specific algorithm (topology)
Announced in IGP as FAD sub-TLVs	Announced in IGP as SR-Algo sub-TLVs

Flex-Algo - overview



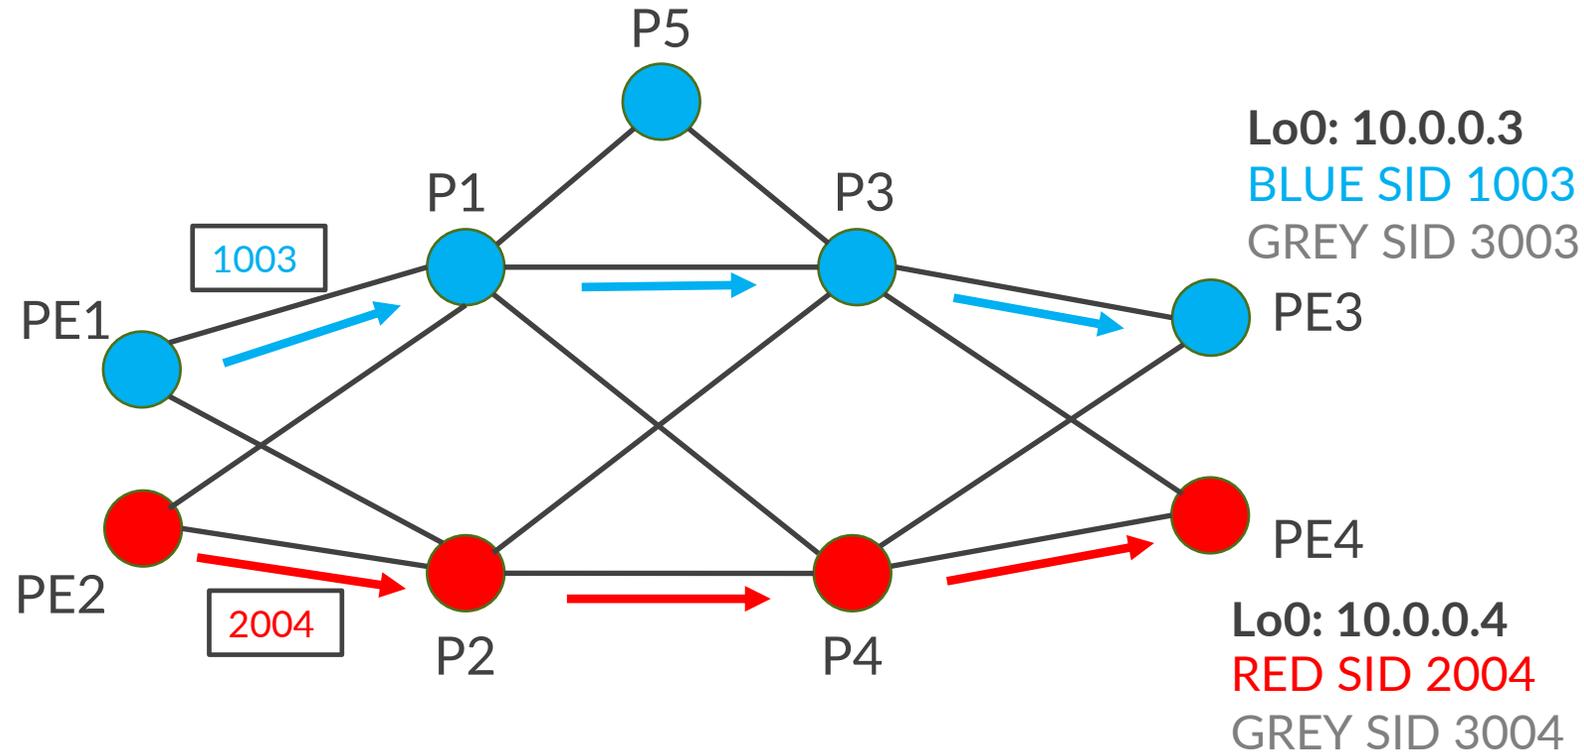
- A node can be a member of multiple algos announcing different node SID for each algo
- Separate SPF per algo.
- In the diagram, all nodes are members of the Grey algo (Flex-Algo 0).
- Nodes in one plane are members of Red algo, nodes in the other plane are members of the Blue algo.
- An algo can be given a color, to allow auto-mapping of traffic to the correct algo.

Flex- Algo - overview



“Vanilla” traffic from PE1 is mapped to Grey algo, can go anywhere

Flex- Algo - overview

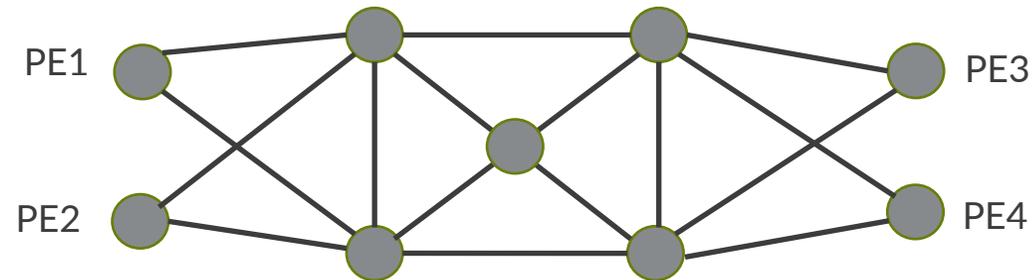


- We have some traffic from PE1 to PE3 that needs to be diversely routed from other traffic from PE2 to PE4
- Use Blue and Red algos respectively

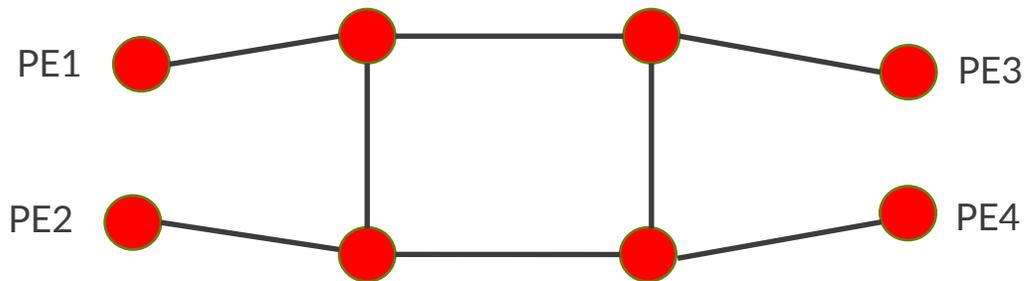
Routing plane separation – multiple topologies

- Nodes and links can be member of multiple Flex-Algos thus enabling multiple topologies

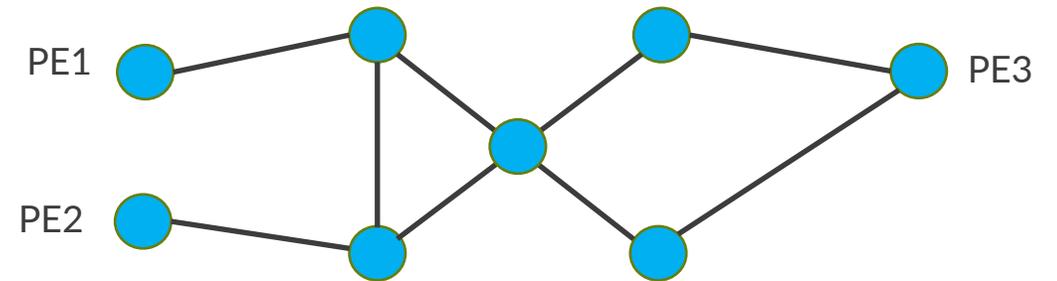
Physical topology (Flex-Algo 0)



Flex-Algo 128

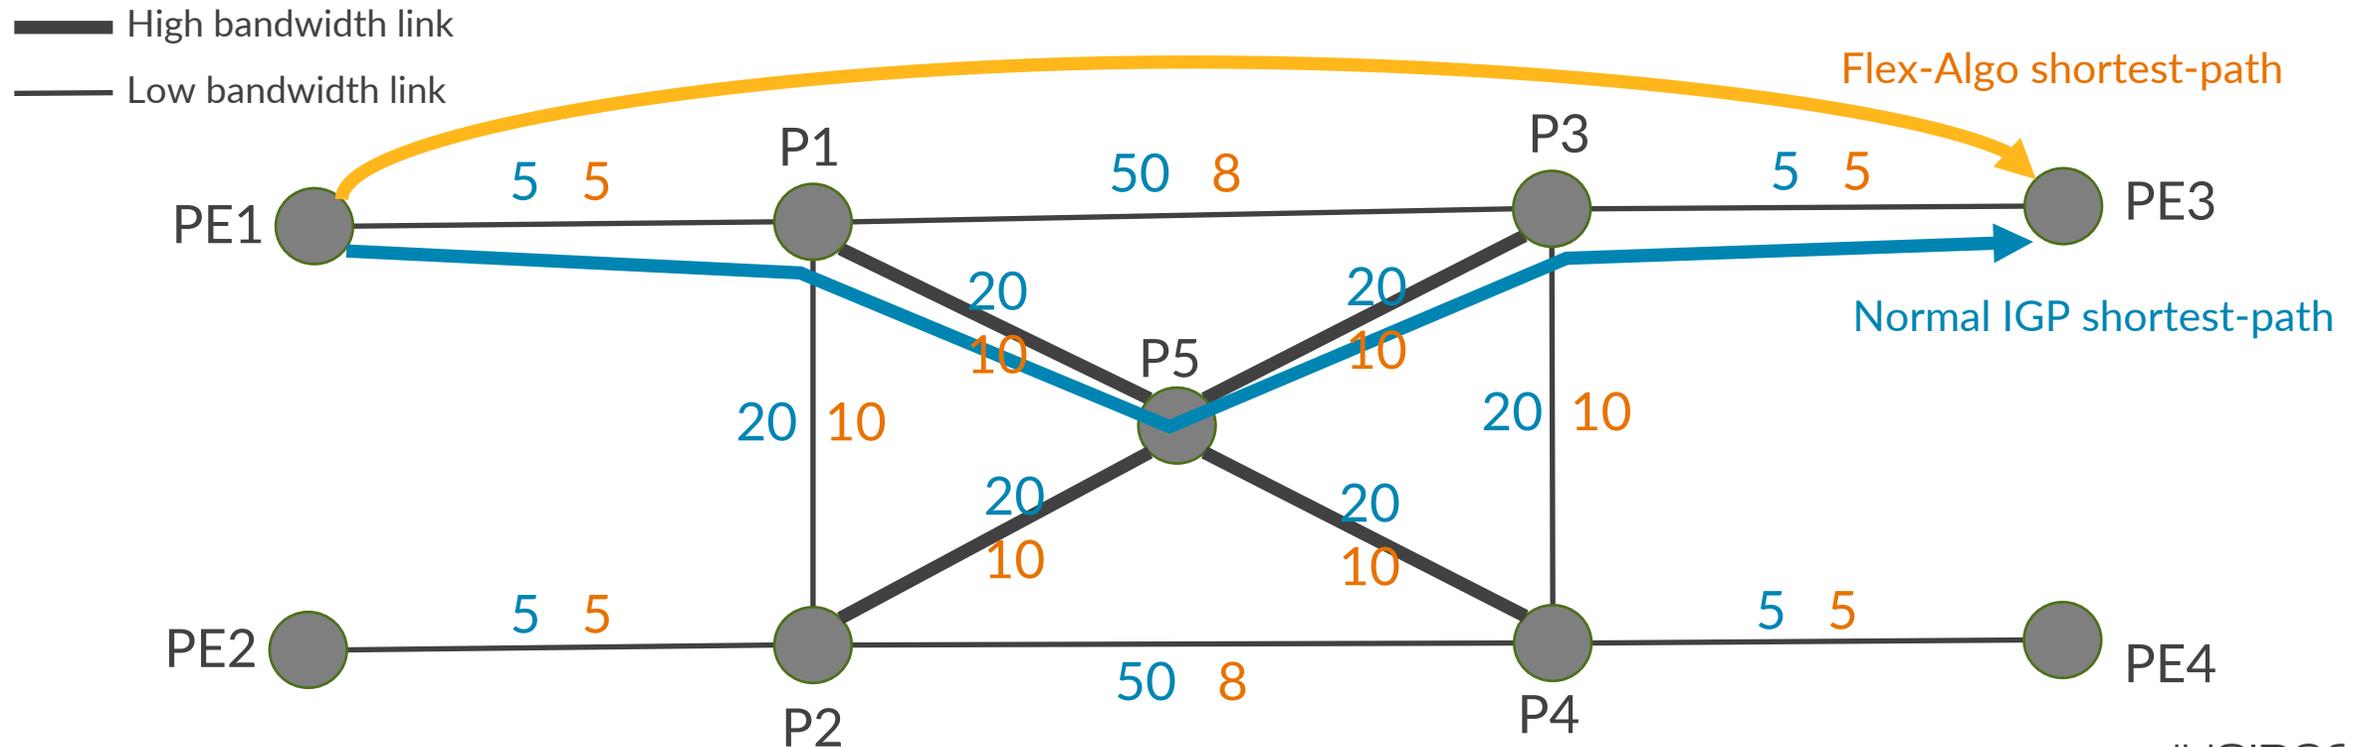


Flex-Algo 129



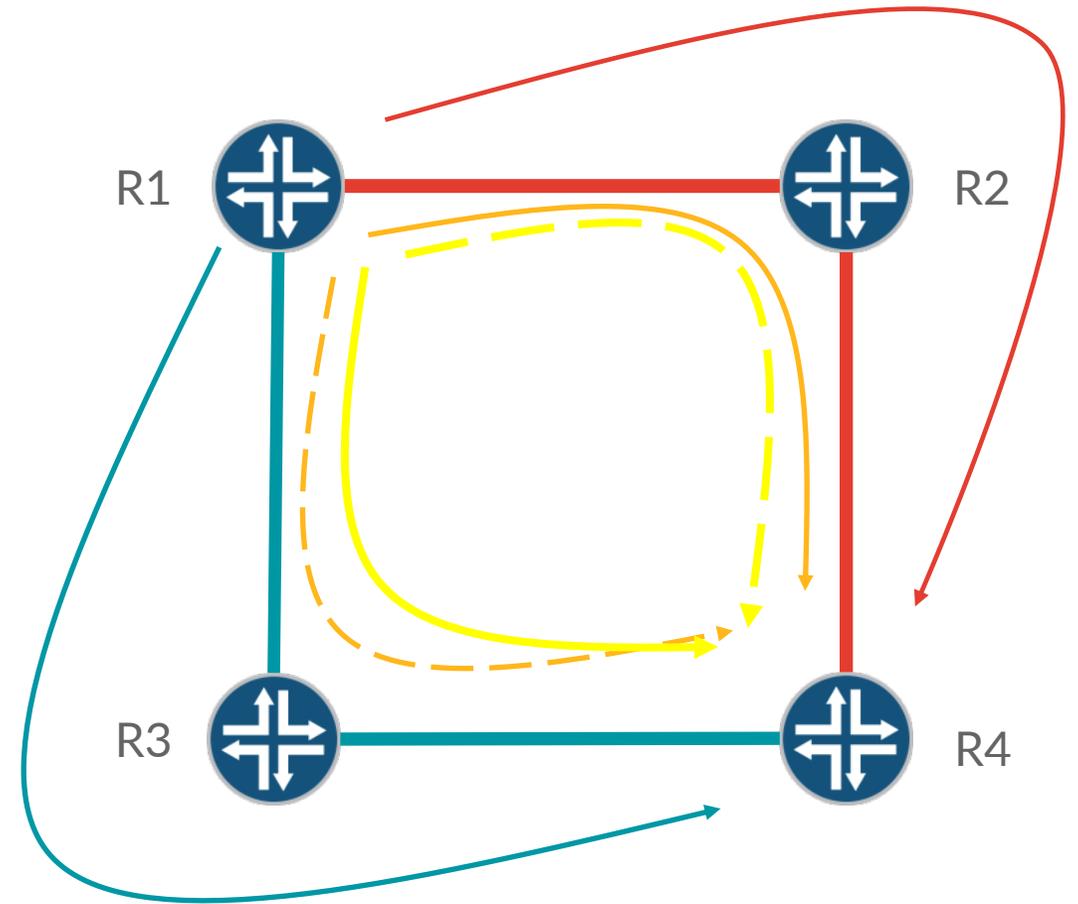
Low latency or high bandwidth path

- Different types of services need to use different type of metrics, e.g., delay sensitive traffic may use latency as metric, while bulk internet traffic may use link bandwidth
- Normal IGP metric in *blue* font (internet traffic); delay metric (or TE metric) in *orange* font



Path diversity example

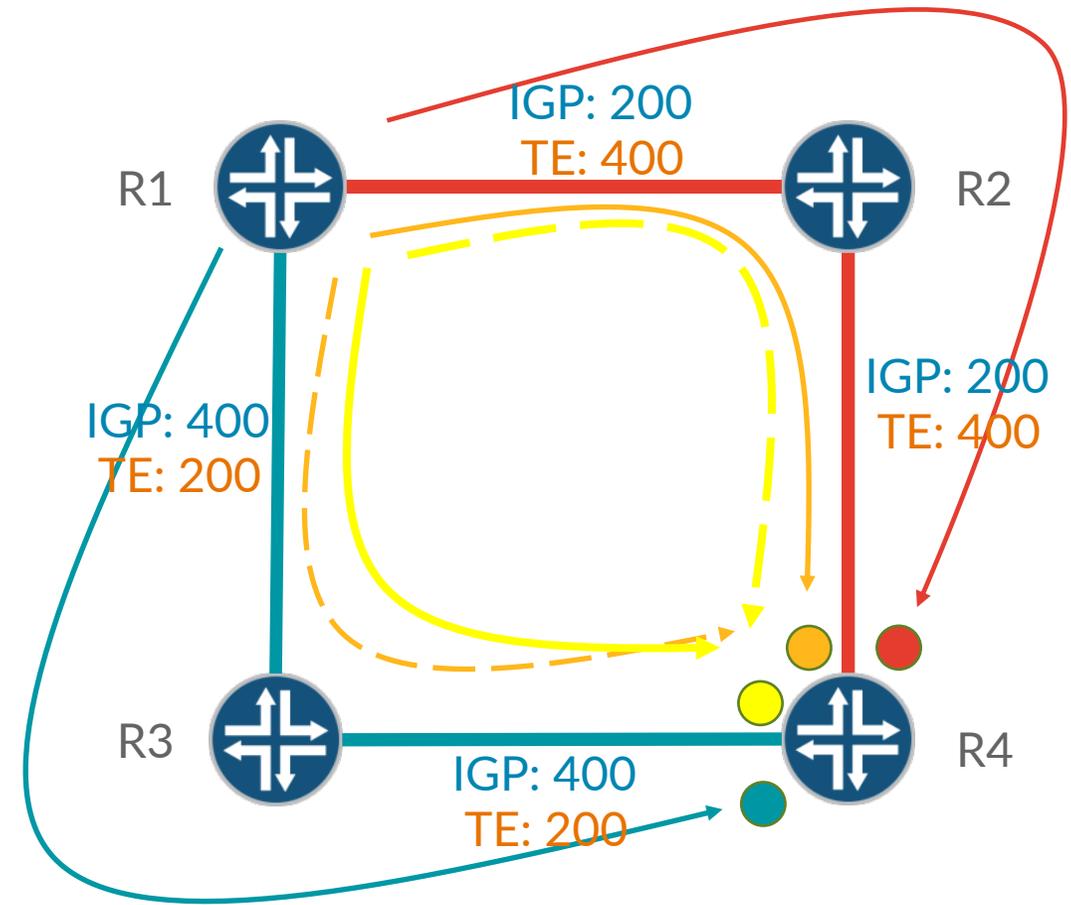
- Red flows traverse red links
 - And no others
- Blue flows traverse blue links
 - And no others
- Orange flows prefer red links
 - But can fail over to blue links
- Yellow flows prefer blue links
 - But can fail over to red links



Path diversity example

FAD	Metric Type	Calculation Type	Constraints
Red	IGP	SPF	Exclude blue
Blue	TE	SPF	Exclude red
Orange	IGP	SPF	Include all
Yellow	TE	SPF	Include all

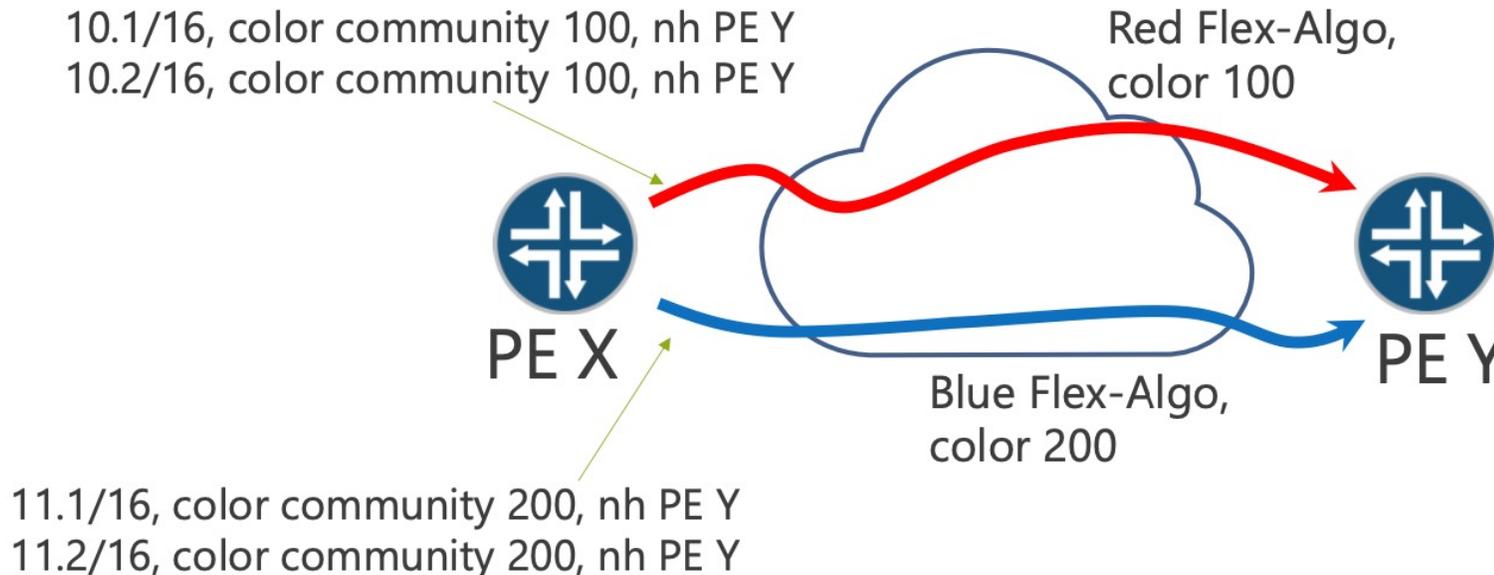
Link	IGP Metric	TE Metric	Admin Group
R1-R2	200	400	Red
R1-R3	400	200	Blue
R2-R4	200	400	Red
R3-R4	400	200	Blue



Loopback Red: Flex-Algo 128
 Loopback Orange: Flex-algo 129
 Loopback Blue: Flex-algo 130
 Loopback Yellow: Flex-algo 131

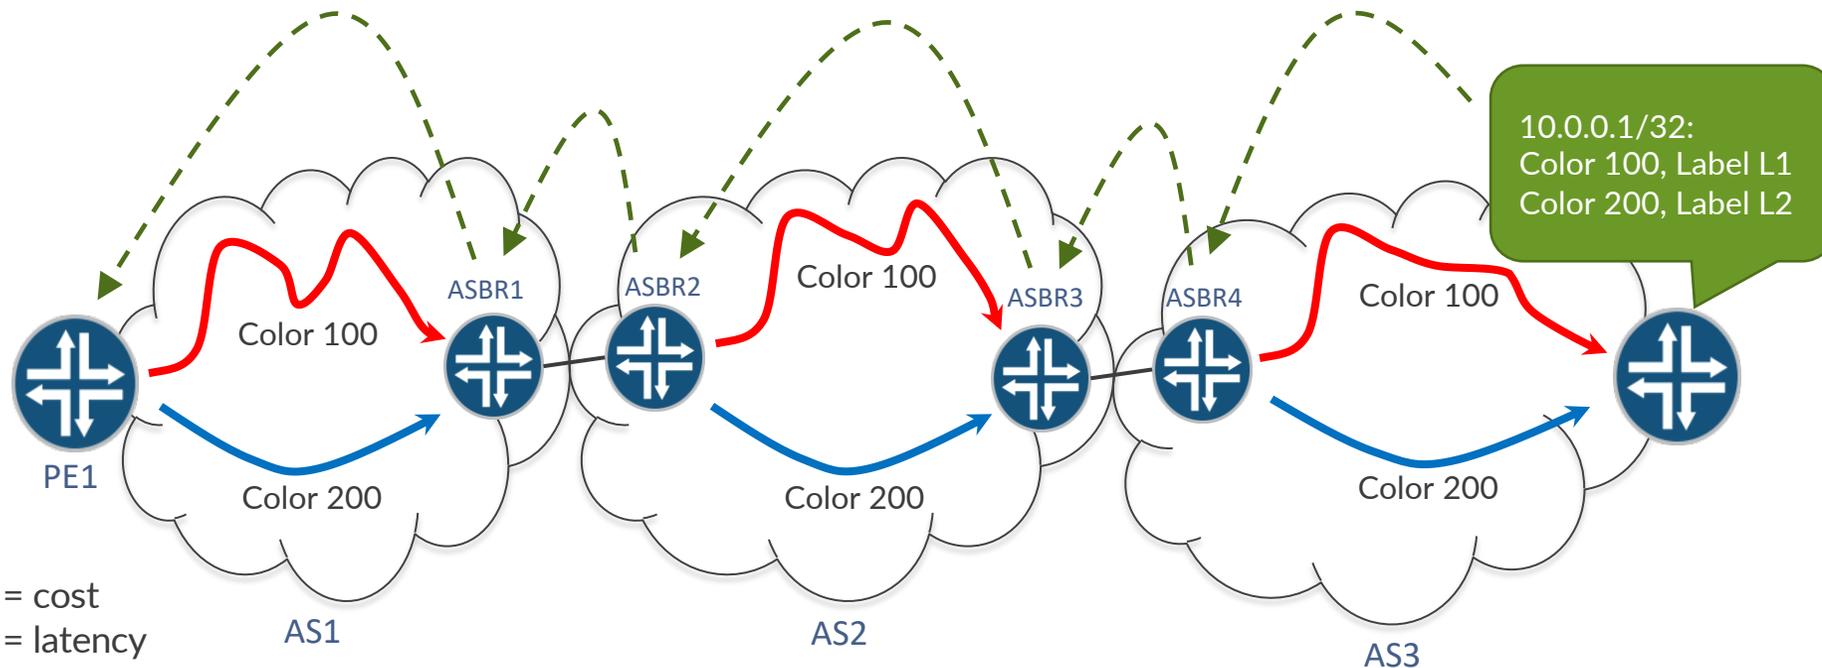
Mapping traffic to the preferred transport

- Colors provide a convenient way to auto-map prefixes/services to a particular transport, e.g., mapping a particular VPN to a particular type of TE-tunnel or to a particular Flex-Algo.
- A prefix (plain IP or VPN) with a color community is automatically mapped onto Flex-Algo with matching color. By default, BGP color community is the same as flexible algorithm ID.
- For inter-domain cases, BGP-CT could be used to map service prefixes into colored tunnels



BGP Classful Transport (BGP-CT)

- A mechanism for extending color-mapping across multiple ASes or domains.
- No need to expose internal topology of a domain to any other domain.
- Each domain can make its own choice of transport technology independently of what other domains are using, e.g. RSVP-TE/SR-TE, with or without controller, or Flex-Algo. BGP-CT acts as the “glue” between domains



Summary

- Many networks require only course-grained TE (as use cases mentioned earlier)
- Flex-Algo is a powerful tool that removes complexity and satisfies lightweight TE requirements
- Flex-Algo can co-exist with LDP and RSVP
- Every node that needs to be part of Flex-Algo need to support Flex-Algo extensions
- Benefits of deploying Flex-Algo:
 - Each SR path is reduced to a MPLS single label
 - No need to specify TE policy on a controller
 - Greater scalability because no need to configure per LSP static SR-TE on each node
 - Operational simplicity
- Juniper blog about Flex-Algo: <https://blogs.juniper.net/en-us/industry-solutions-and-trends/igp-flexible-algorithms-flex-algo>



THANK YOU

Eduardo Haro – Sr. Sales Engineer

eharo@juniper.net

<https://www.linkedin.com/in/haroeduardo/>

JUNIPER
NETWORKS

Driven by
Experience™

