TE 2.0

Julian Lucek

Key trends in Traffic Engineering

- Advent of SDN Controllers for the WAN
  - Automation, instead of manual processes
  - Stepping stone to Self-Driving Networks
- TE to meet service requirements
  - Path Diversity for P2P and P2MP LSPs
  - Minimum Latency paths
- Streaming Telemetry for real-time visibility of “network weather”
  - Automated traffic management
- Some operators prefer RSVP for TE, while others prefer segment routing
  - So controllers need to support both

NorthStar Controller Concept

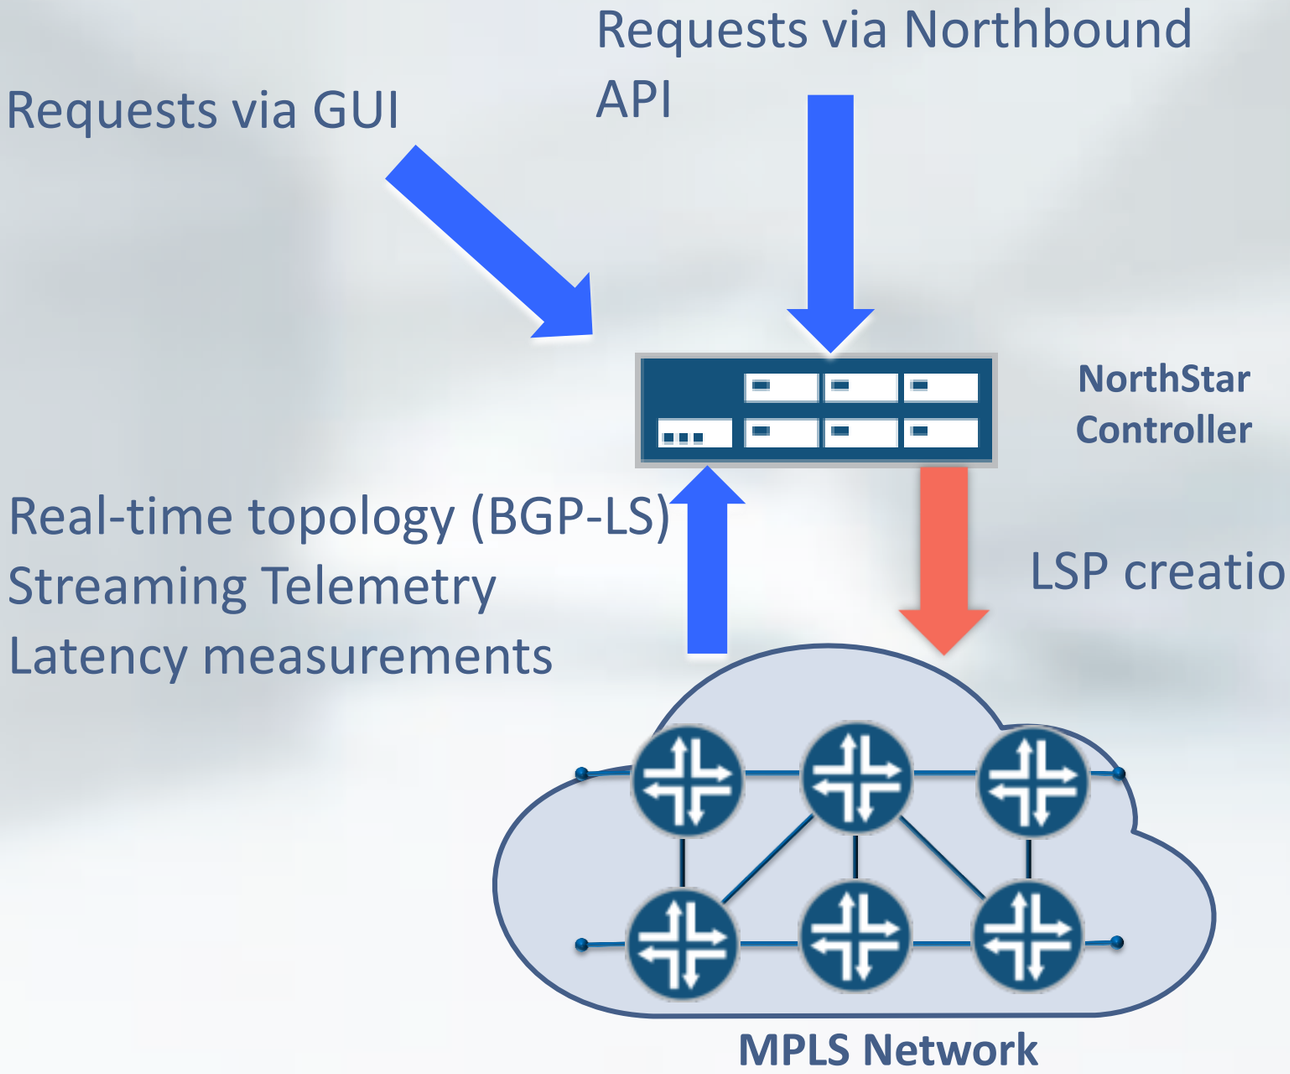

- Creation of LSPs according to service requirements.
- Placement of LSP paths according to traffic demands.
- Automation of some network operations.

BGP-LS:Discovery of topology and link parameters

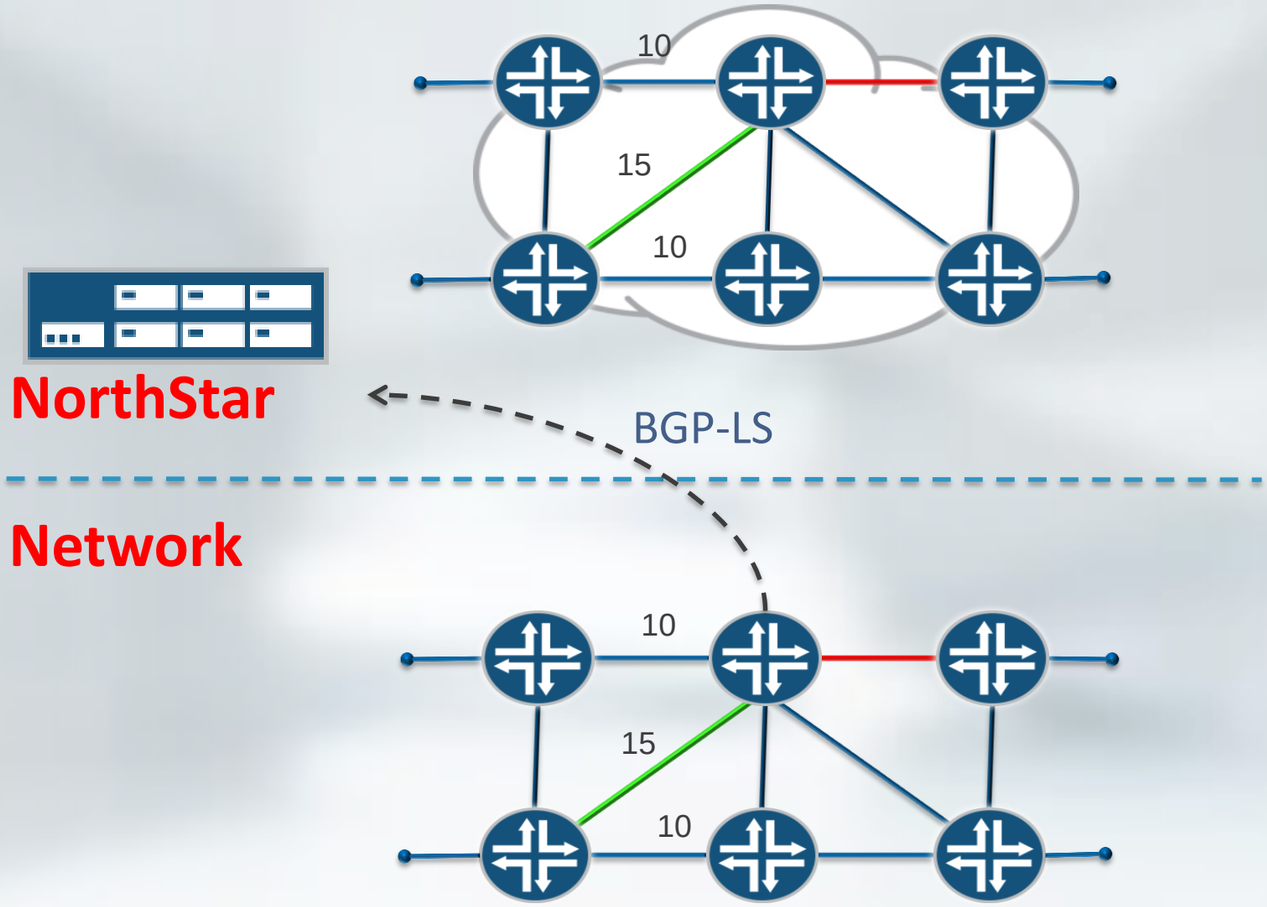

BGP-LS peering between NorthStar and at least one node in the network gives NorthStar real-time knowledge of topology and TE parameters of each link.

- IGP metric
- BW availability at each priority level
- Admin-groups (colours)
- SRLGs

PCE: A standards-based approach

What is it?

- PCE: Path Computation Element (RFC 4655)

  An entity that can calculate paths in the network

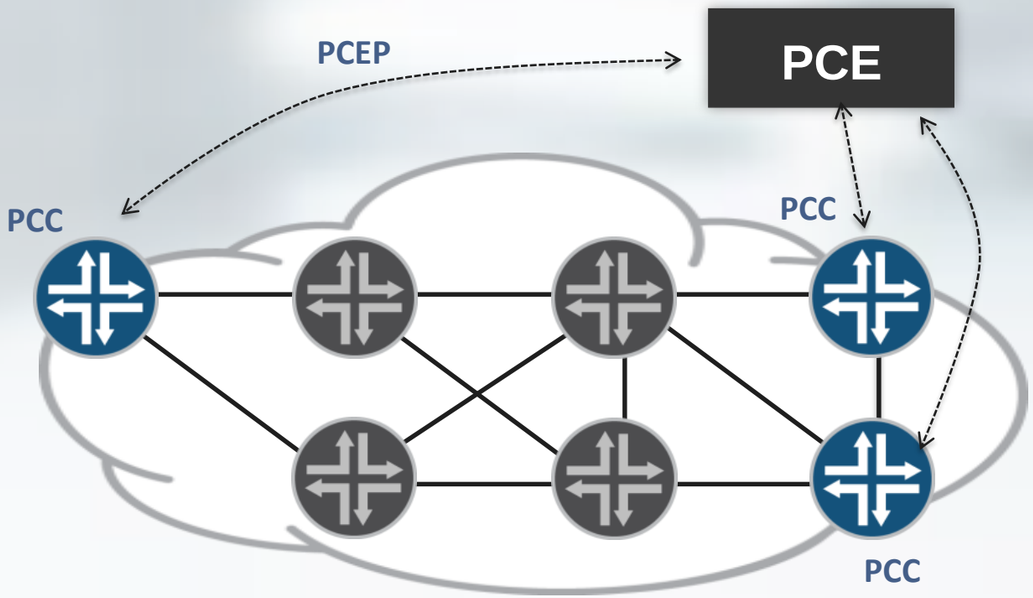

Components

- PCE: Path Computation Element

  Computes the path

- PCC: Path Computation Client

  Receives the path. Sets up LSP using RSVP signalling or Segment Routing.

- PCEP: PCE protocol (RFC 5440)

  For PCE/PCC communication

Meeting service requirements

Path Diversity

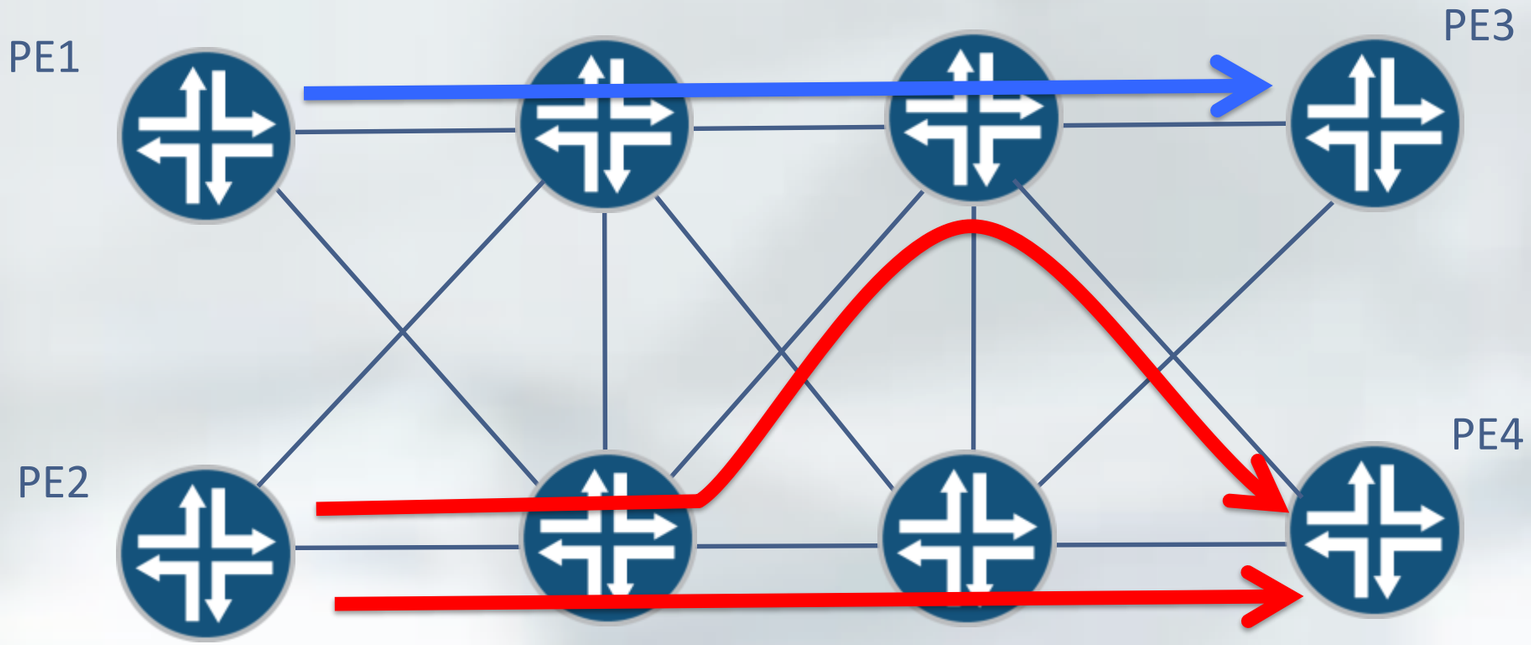

- The two paths must be engineered such that they have no nodes or links in common – including the PEs!
- Difficult to achieve path diversity when each ingress node calculates its own path – probability of no fate-sharing at all is only 1 in 16!
- A central controller can calculate both paths to ensure diversity.

WAN: Live-Live with P2MP TE for path diversity: current deployments

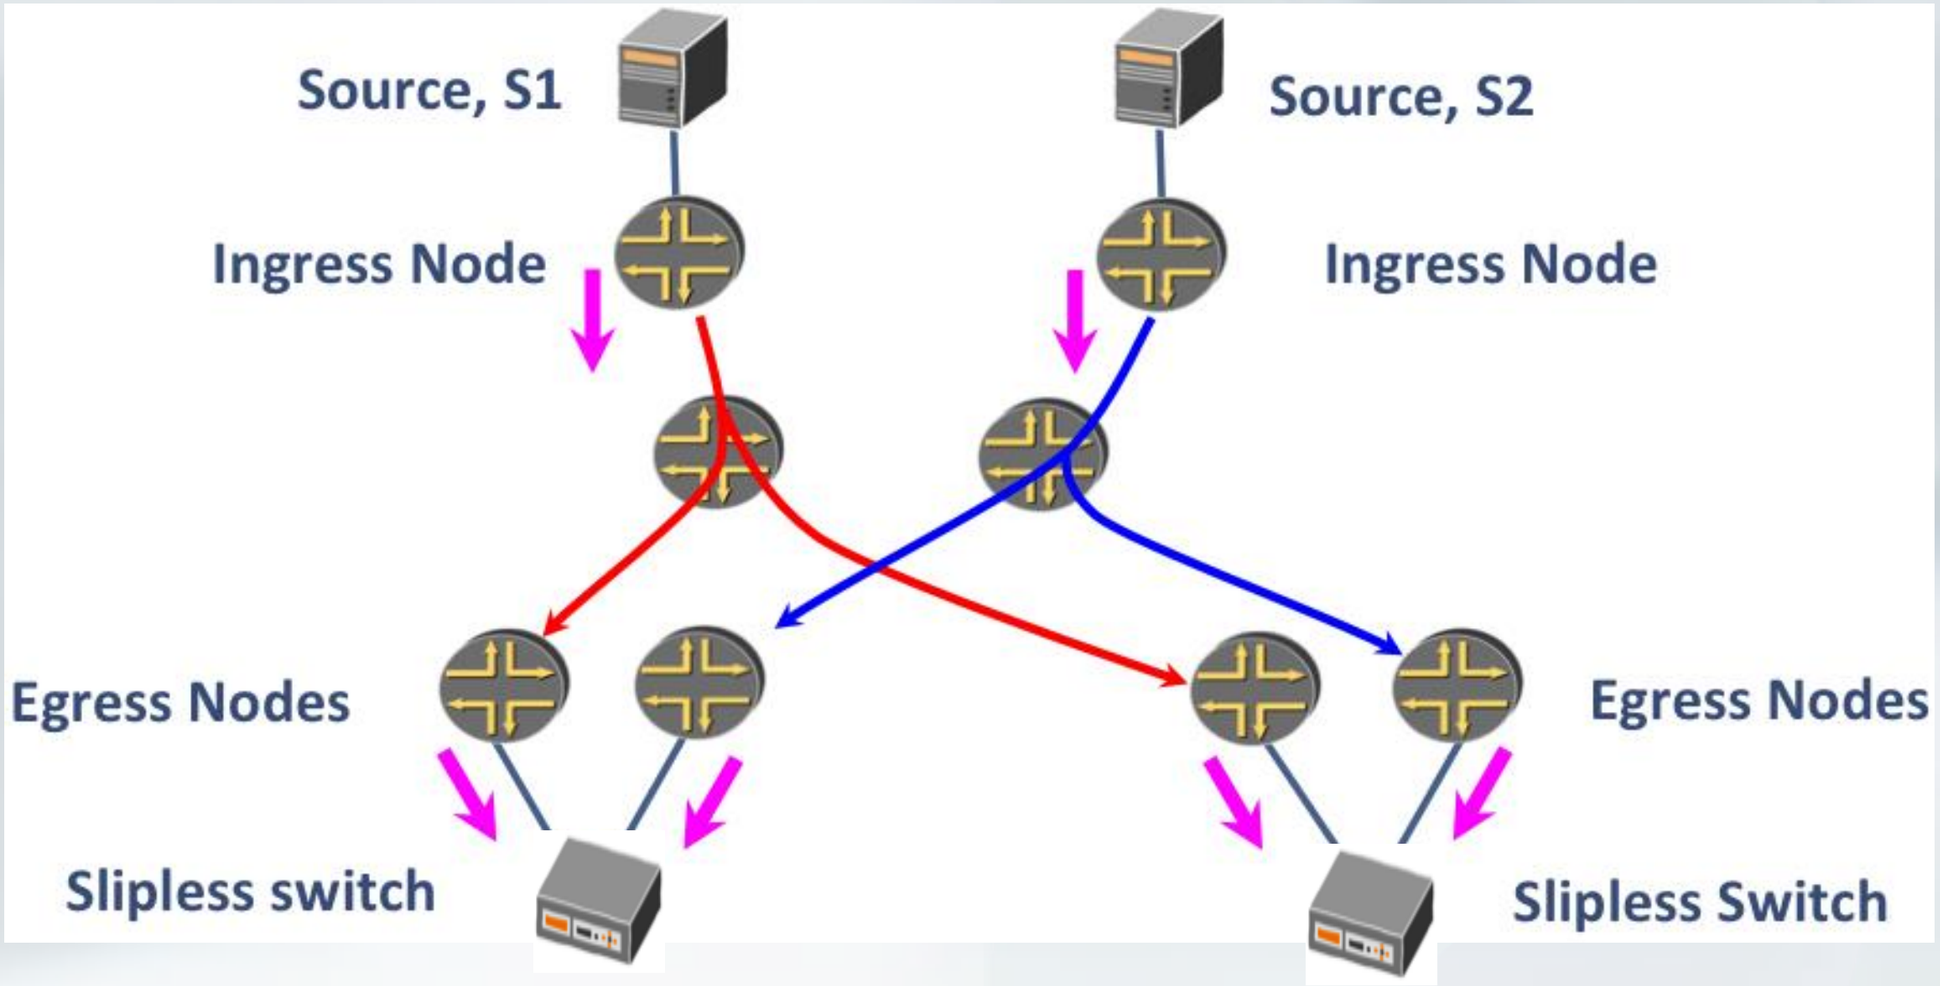

Pair of diverse P2MP LSPs: computed manually by human operator, or using offline tools. We have several such deployments e.g. BT, NTT, ORF.

WAN: Live-Live with P2MP TE for path diversity: with SDN Controller

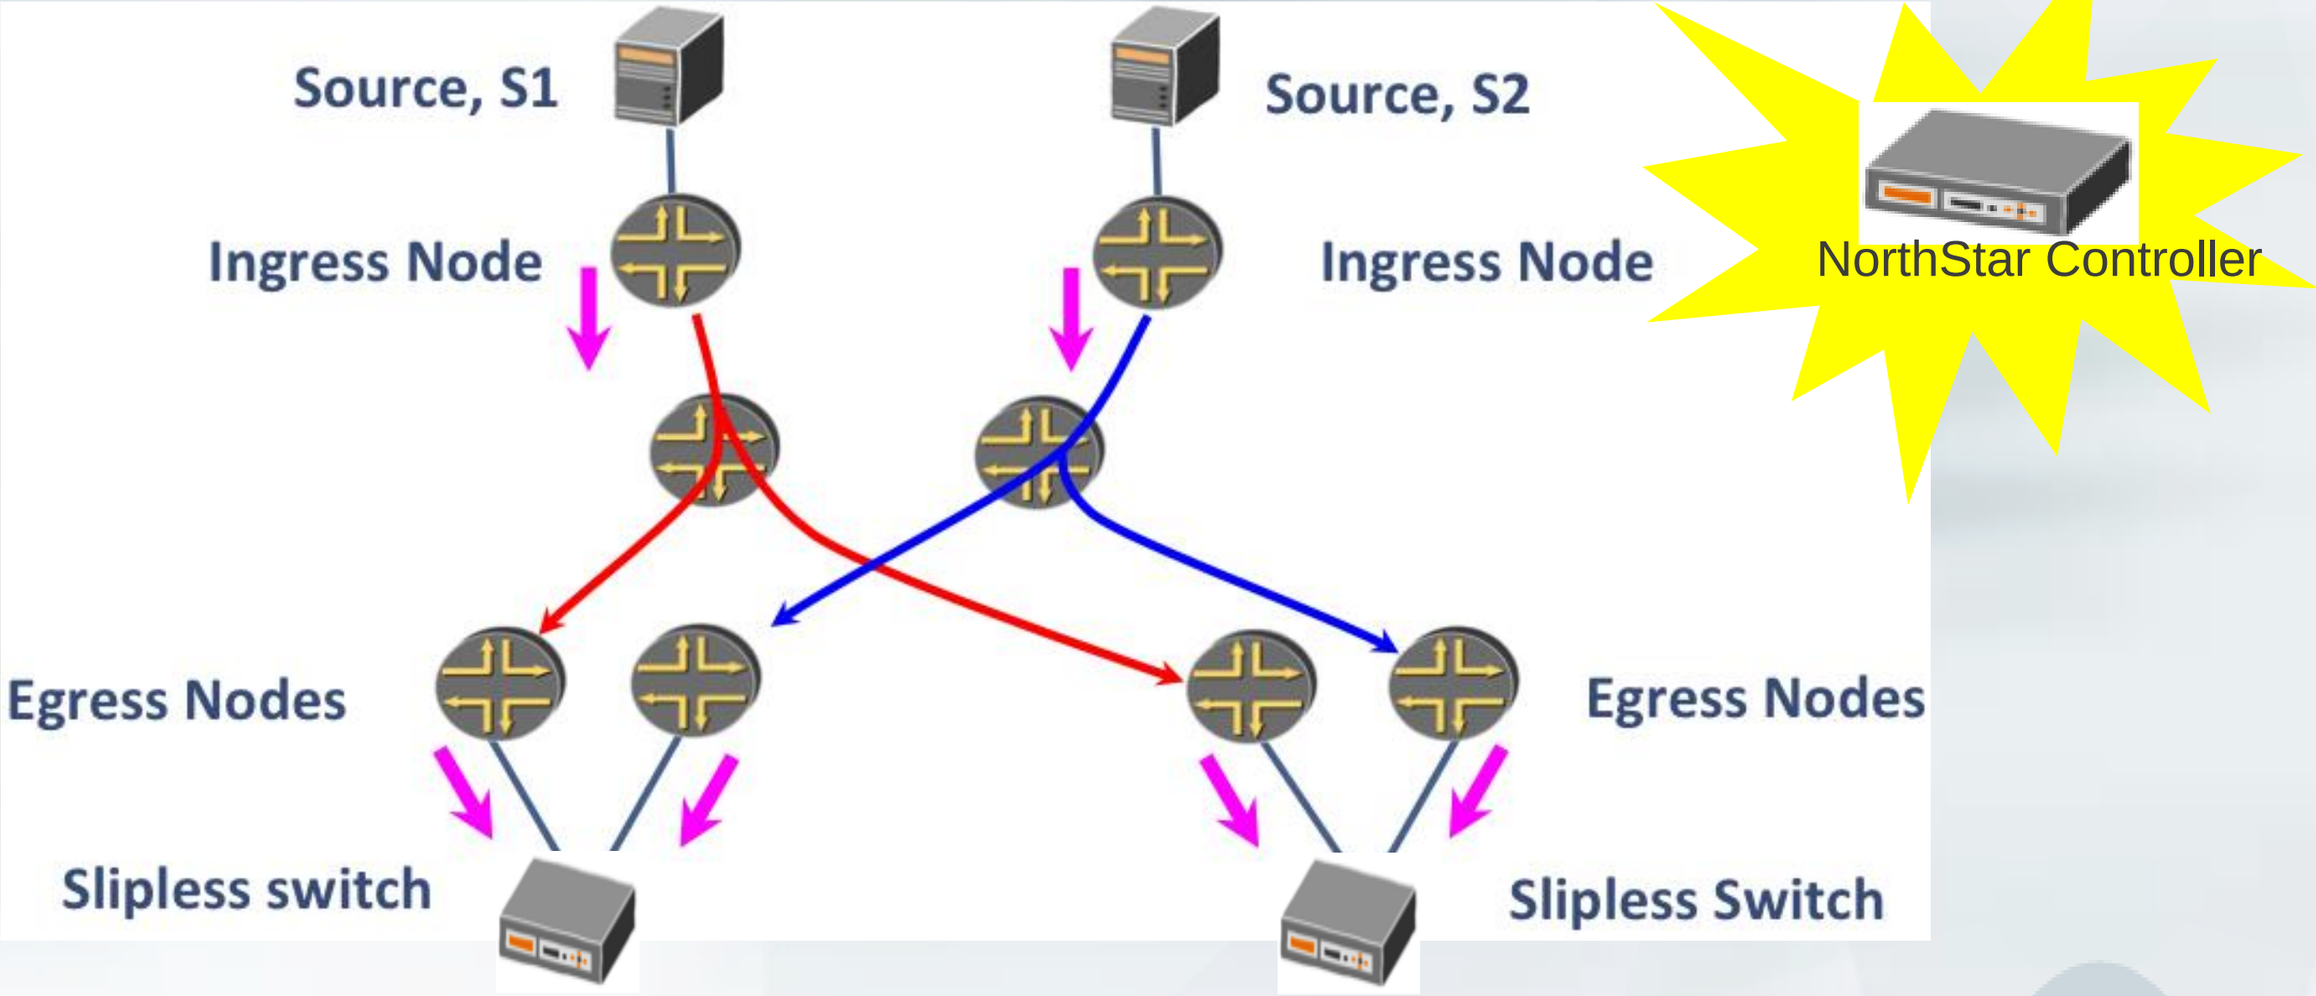

- Pair of PCE-initiated diverse P2MP LSPs, computed by NorthStar: 2H 2017
- Intelligent Bandwidth Calendaring by NorthStar: 2H 2017

Intelligent Bandwidth Calendaring

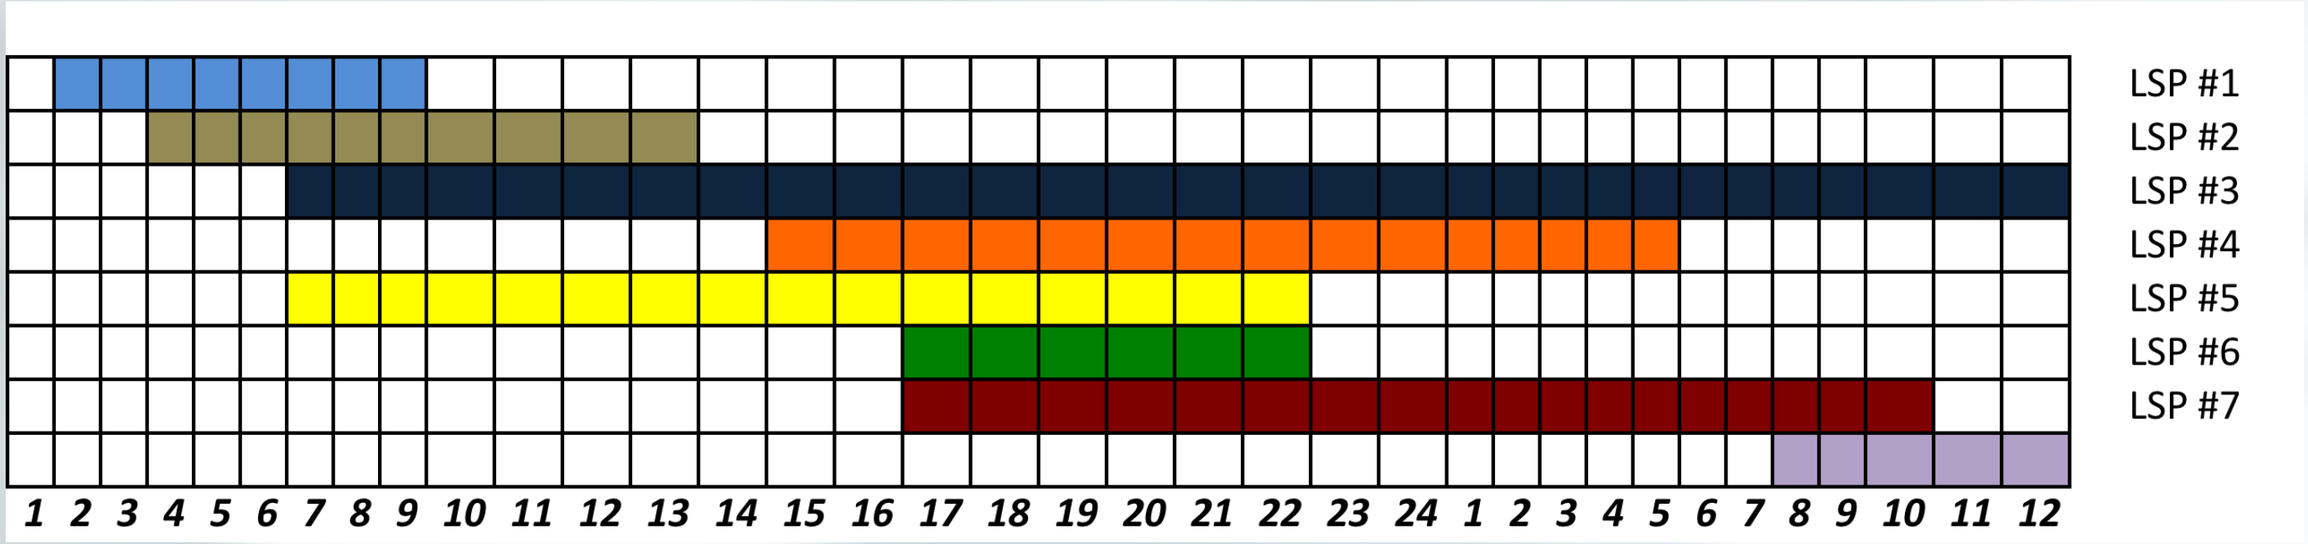

- When new LSP (P2P or P2MP) is requested, NorthStar performs path computation taking into account previous bookings that overlap in time and resources.

Programmable cost function

Default behaviour is to compute a path with lowest IGP metric. Numbers in blue show the IGP metric.

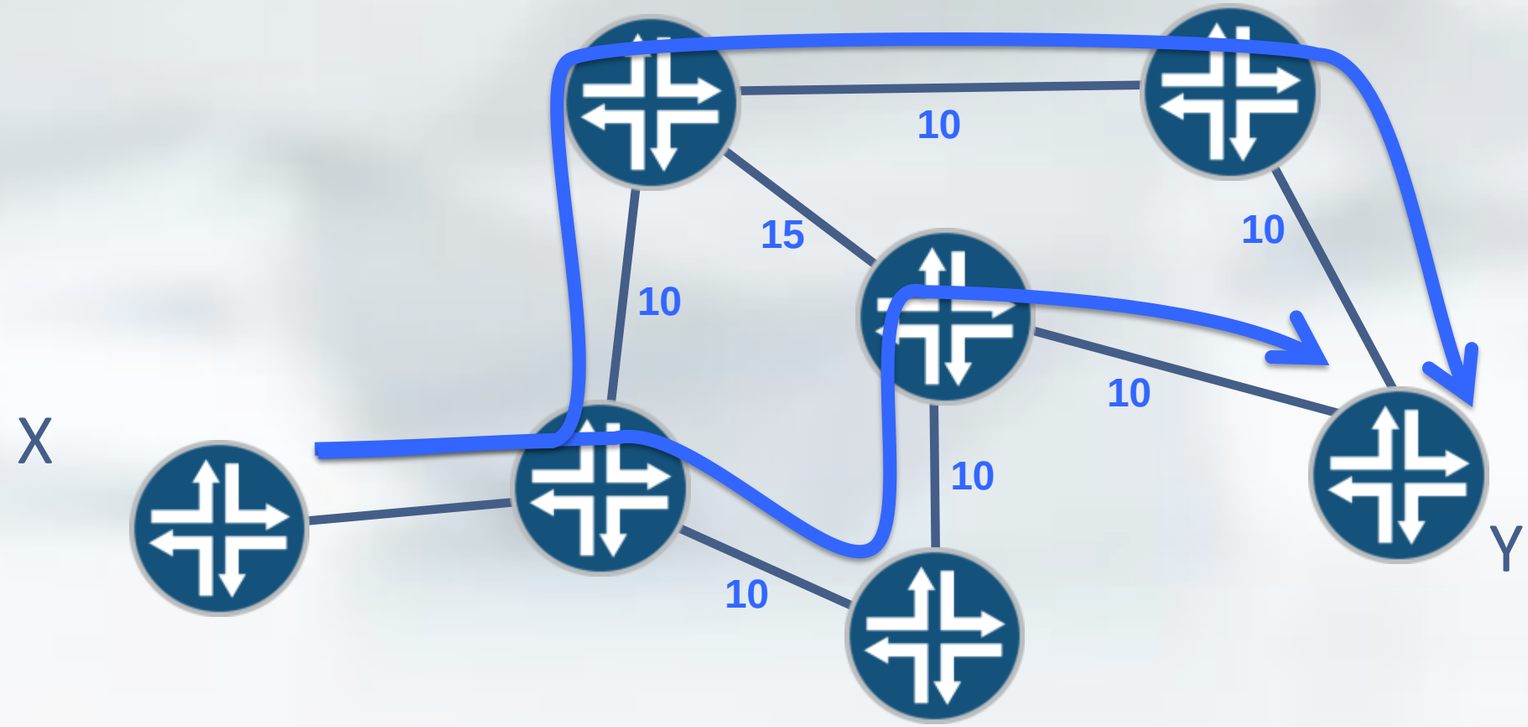

Lowest IGP metric paths →

Blue numbers show IGP metric

Programmable cost function

Optionally, can specify *latency* or *hop-count* as the criterion for LSP path computation

Can choose cost function on a per-LSP basis.

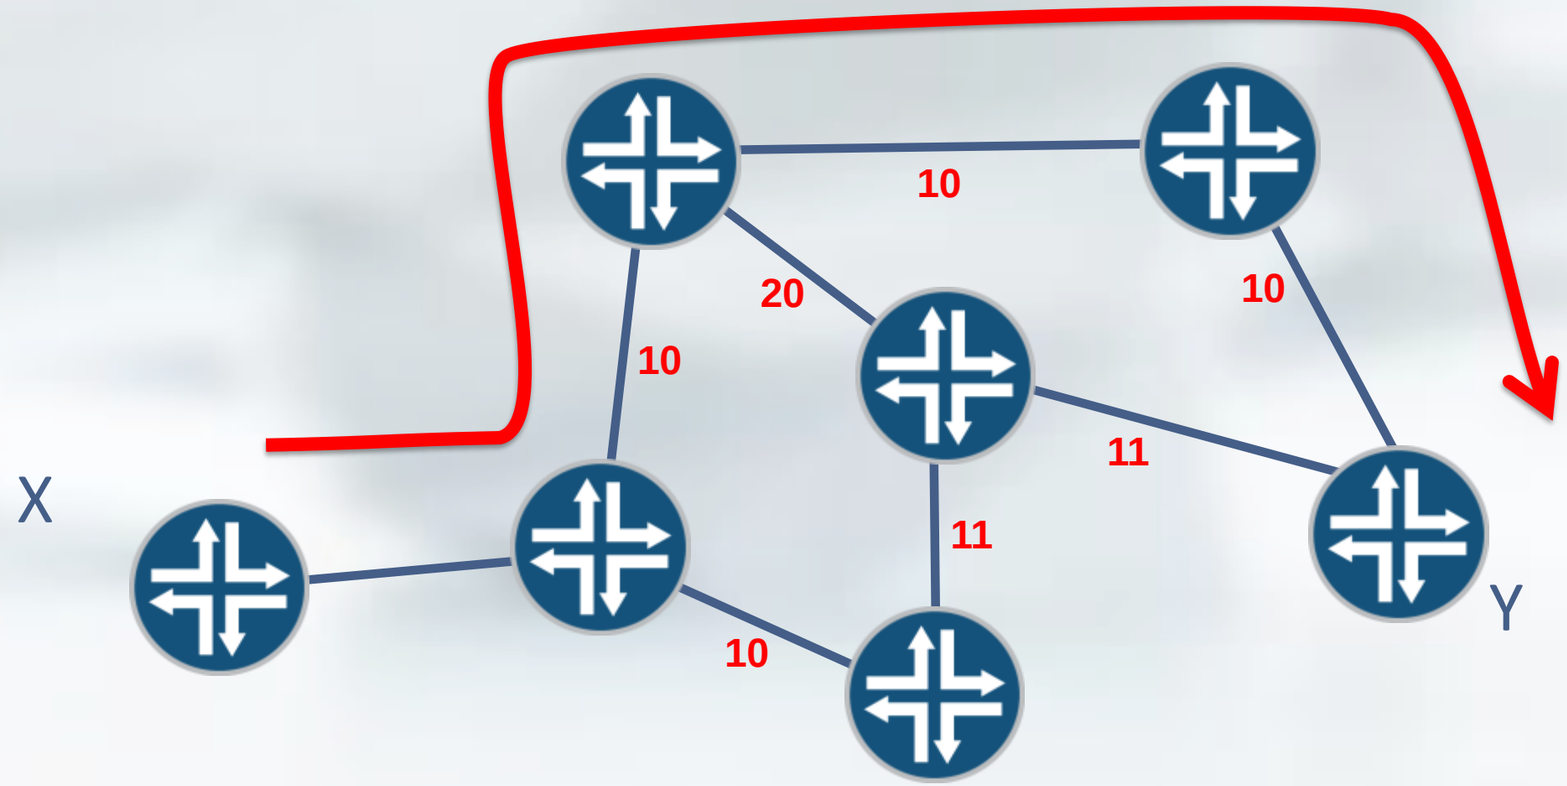

Lowest latency path →

Red numbers show delay

Acquiring latency information

- NorthStar supports the following methods for acquiring link latency information
  - input via GUI
  - input via northbound REST API
  - learnt from an optical controller (draft-ietf-teas-yang-te-topo-06)
  - measured by routers using time-stamped probe packets (RPM) and reported to NorthStar via Streaming Telemetry

NorthStar and Segment Routing

NorthStar and SPRING/Segment Routing

- If bandwidth “reservations” are required, Segment Routing has no mechanism to reserve bandwidth at the network layer, so a controller is mandatory to perform admission control and keep track of bandwidth assignments.
- Path computation on controller is agnostic about how path is instantiated (RSVP or SR)
- PCEP extensions for SR-TE LSP

Spring and PCEP

Global topology, demand, and LSP view

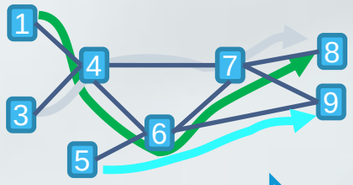

Centralized path computation

Path setup via PCEP

PCEP Message (SPRING flavour):
To make traffic take path R1-R4-R6-R7-R8, send it to R4 with label stack = {**406,607,708**}

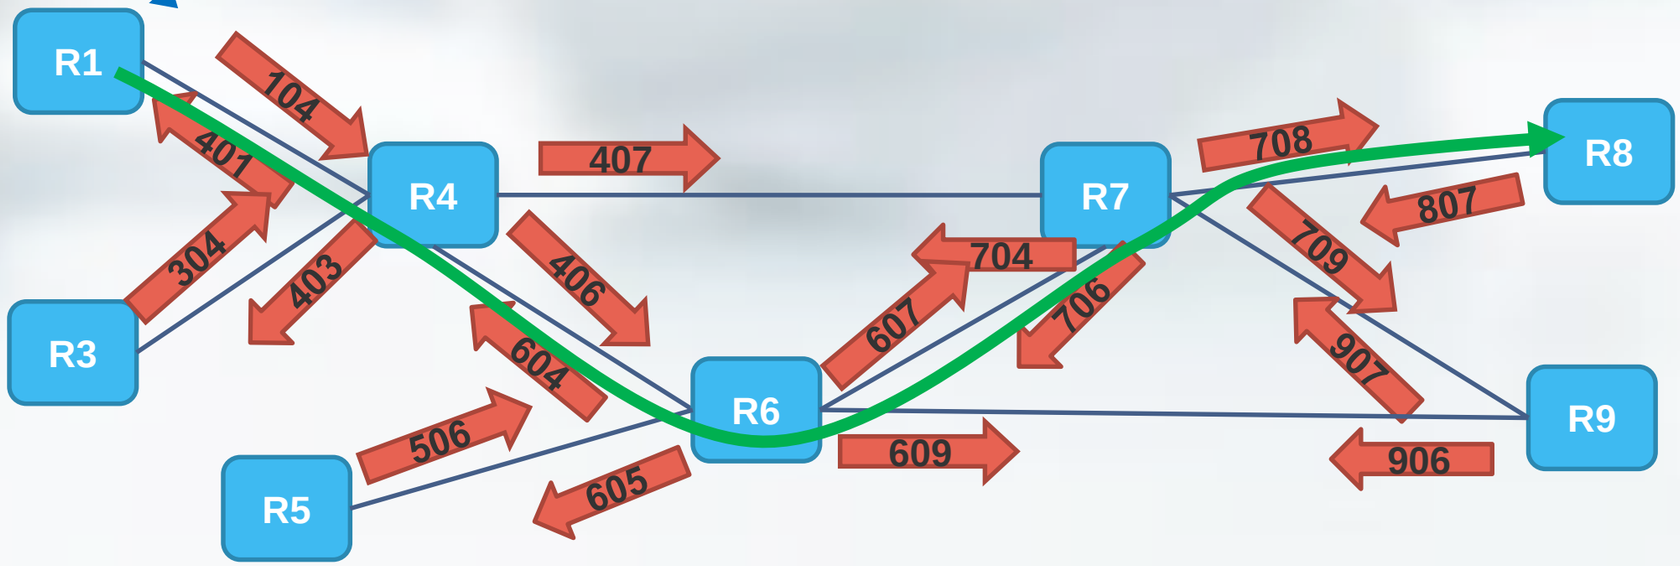

NorthStar: SPRING-TE overview

- Adjacency SID & node SID learning via ISIS or BGP-LS
  - BGP LS Extensions (draft-gredler-idr-bgp-ls-segment-routing-ext-03)
- New PCEP capability, ERO sub-object and TLVs
  - draft-ietf-pce-segment-routing
- SPRING-TE LSP creation, visualization & optimization

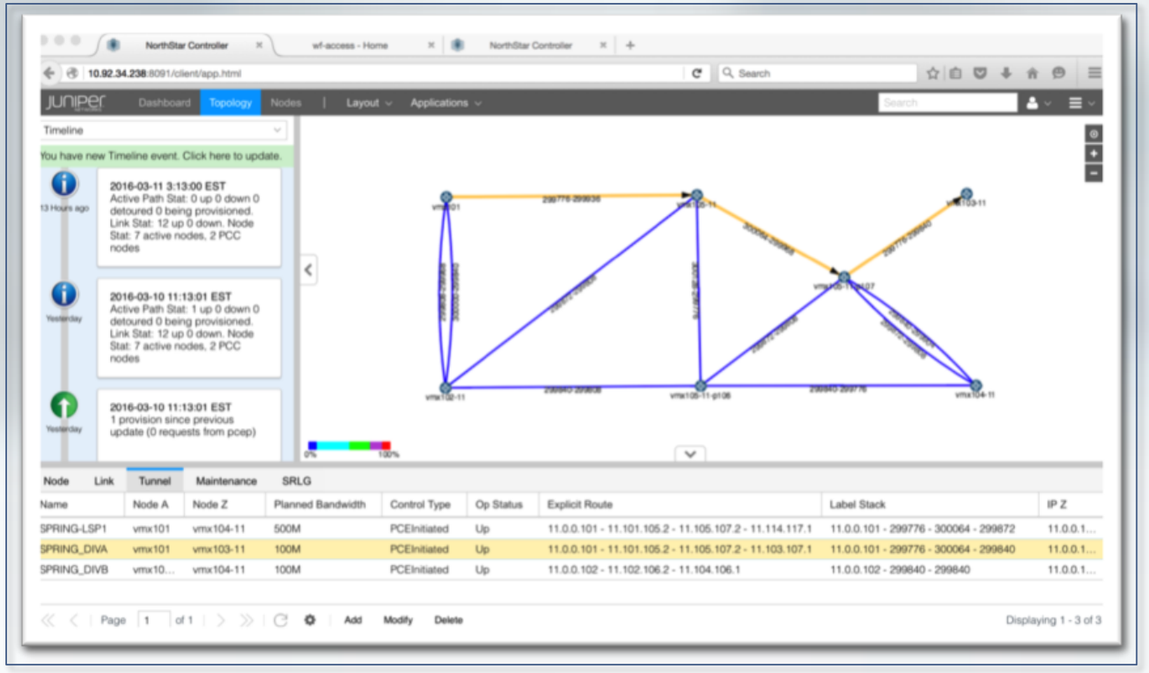

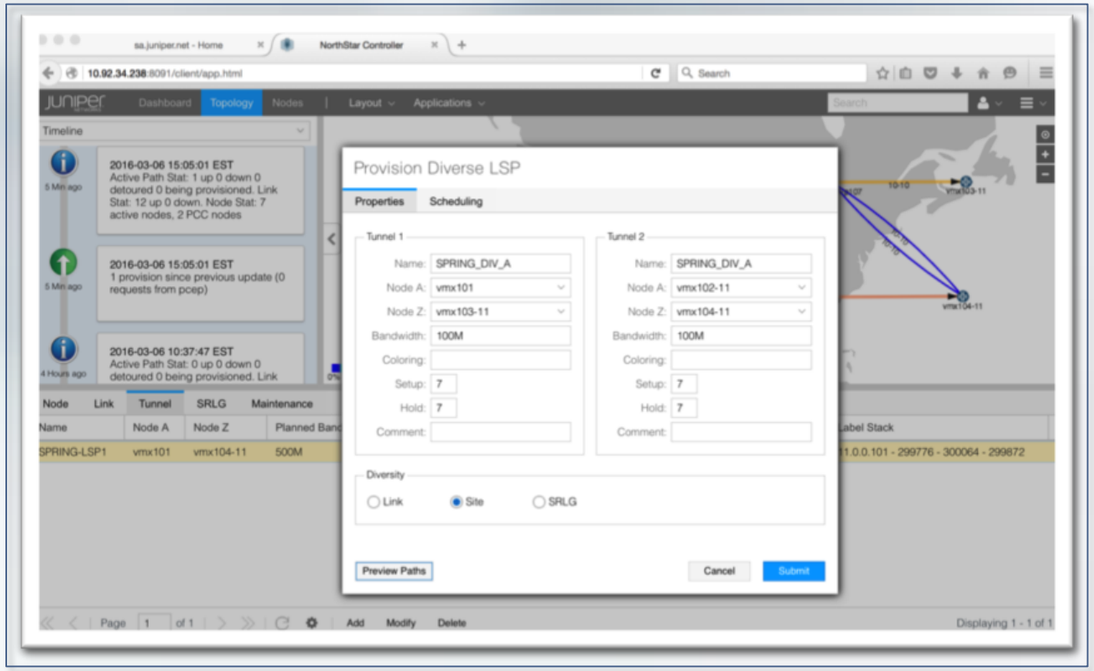

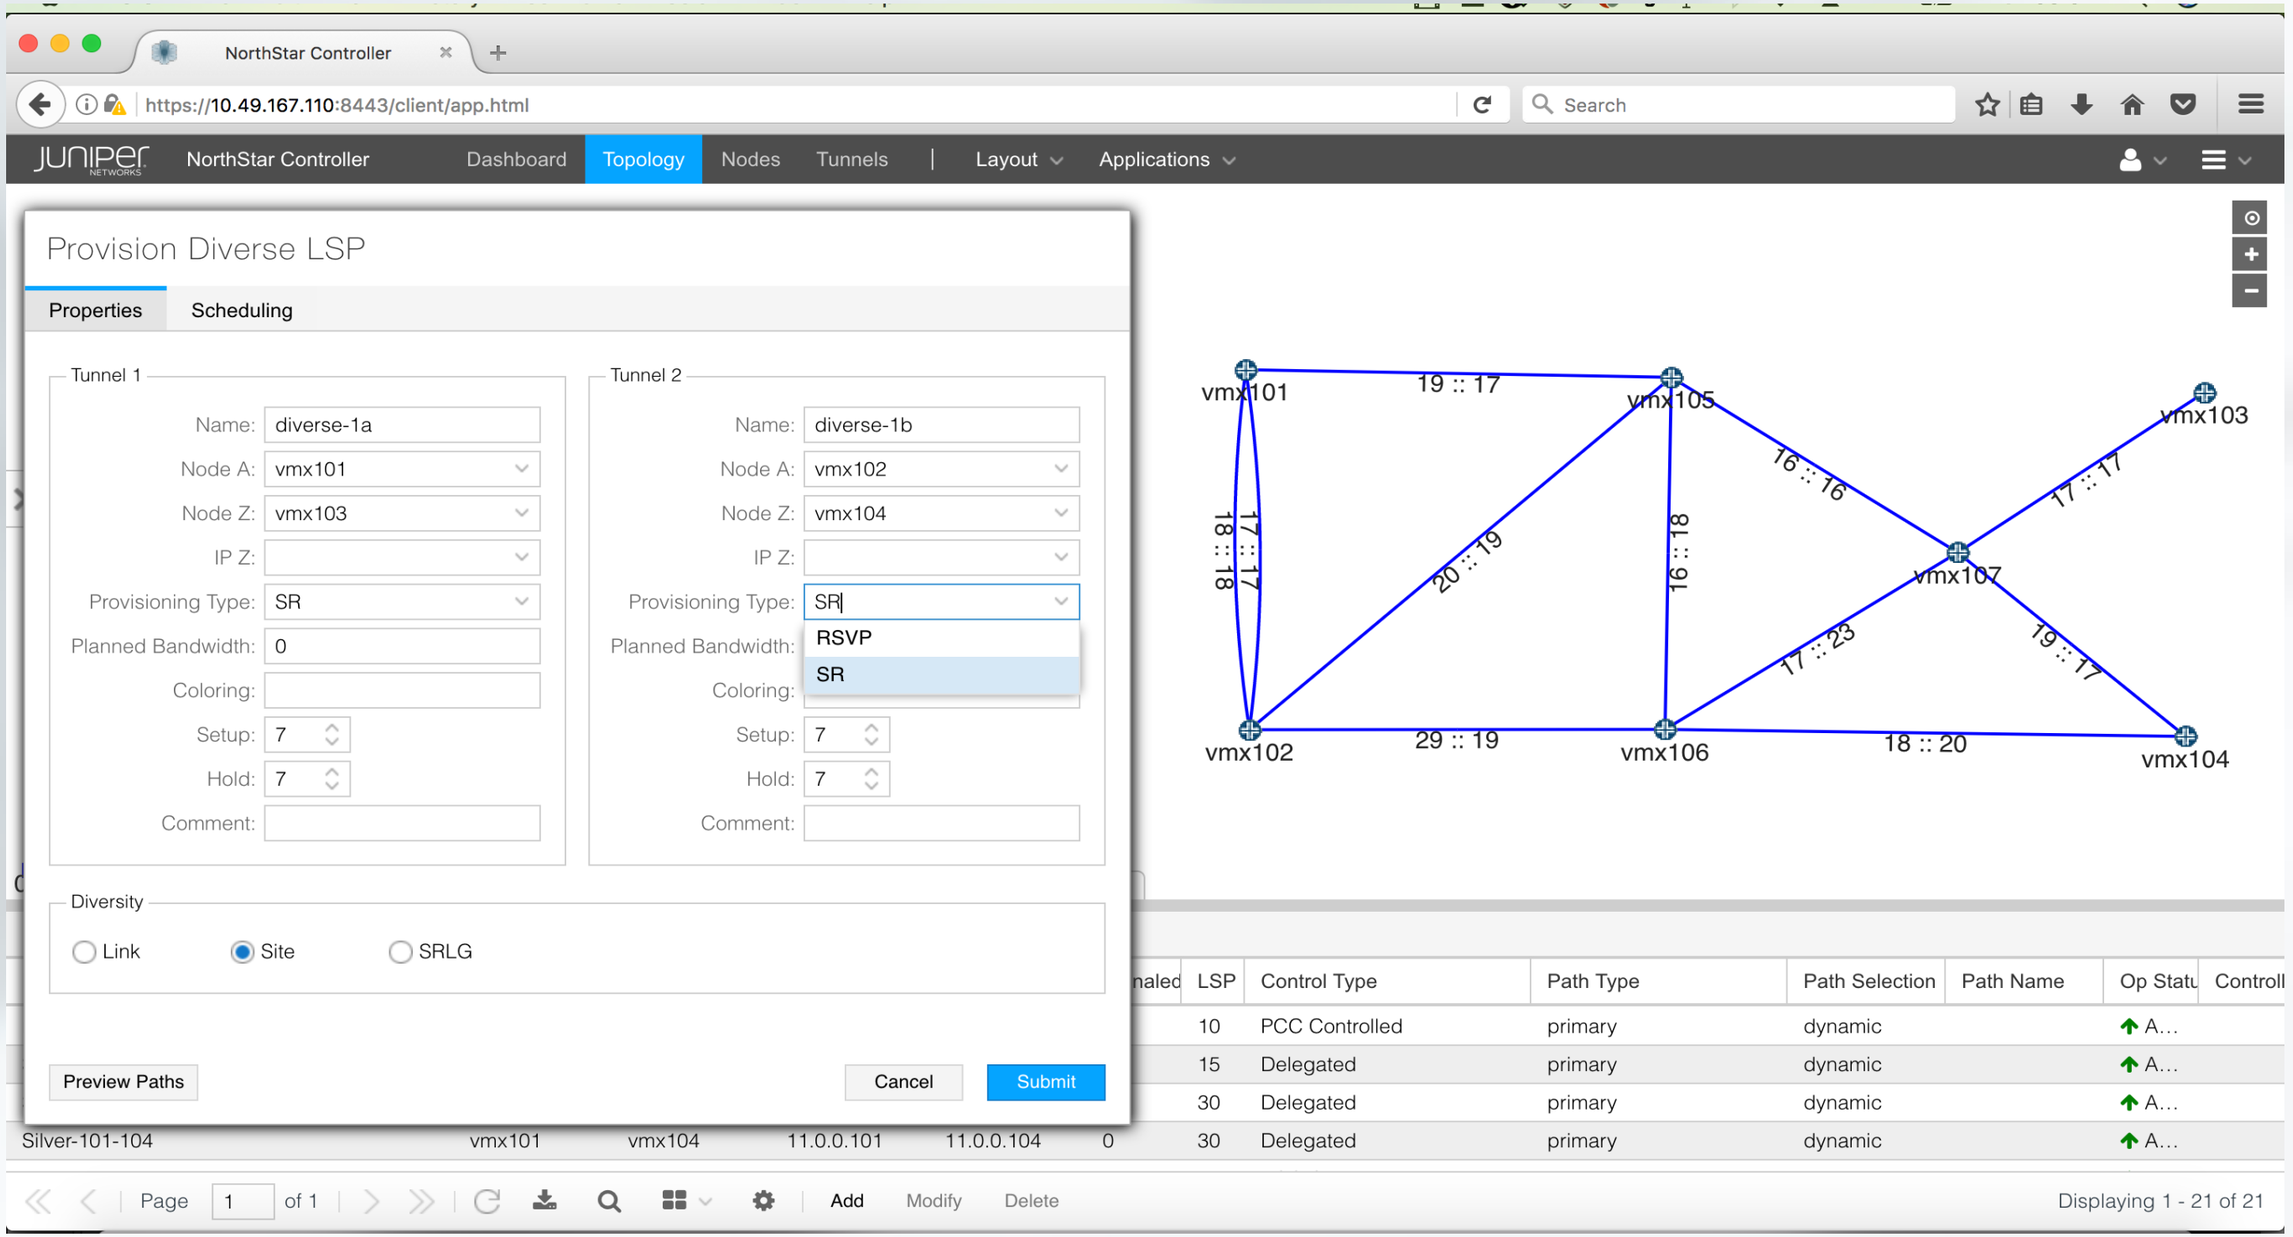

NorthStar Analytics

“Big data” telemetry

Pressure, temperature, vibration readings etc

Interface statistics, queue high-water mark, LSP statistics, memory utilisation, FIB utilisation, RSVP protocol statistics, optics diagnostics...

Streaming Telemetry

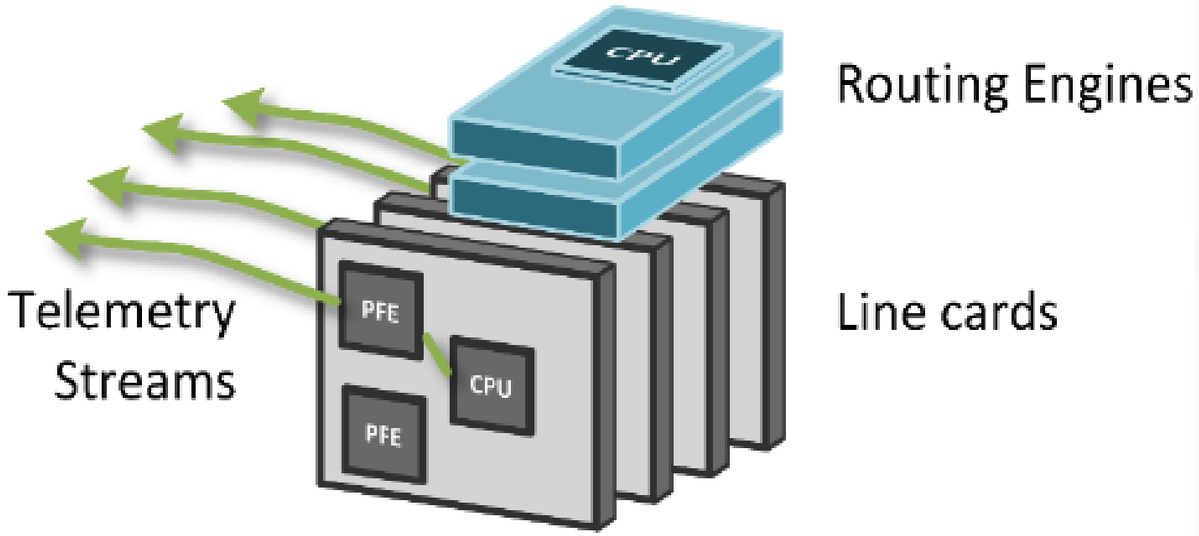

Traditionally, SNMP used to retrieve statistics

- Quite heavy-weight: poll/respond paradigm, via control processor on network device

Streaming Telemetry

- Statistics sent at high frequency (e.g. every few seconds) without any polling needed from collector
- Statistics collected on line-cards (e.g. interface/queue stats, LSP stats) sent directly from line-cards to collector
    - no bottleneck from control processor

High-Resolution Interface Statistics

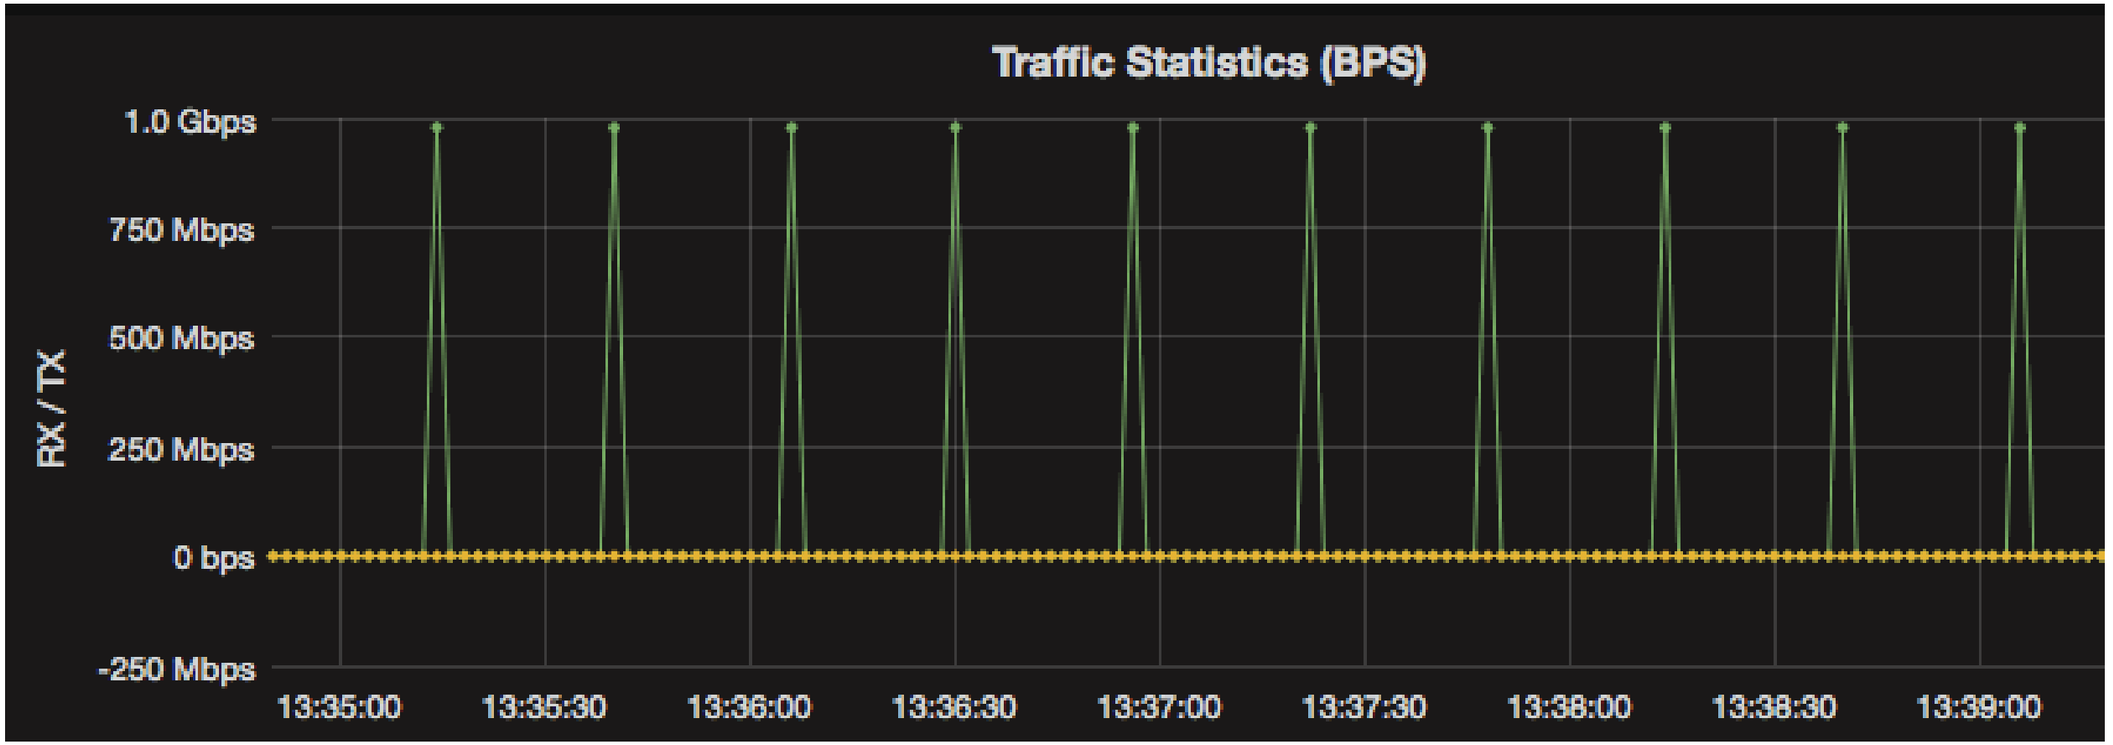

30 seconds

High-Resolution LSP Statistics

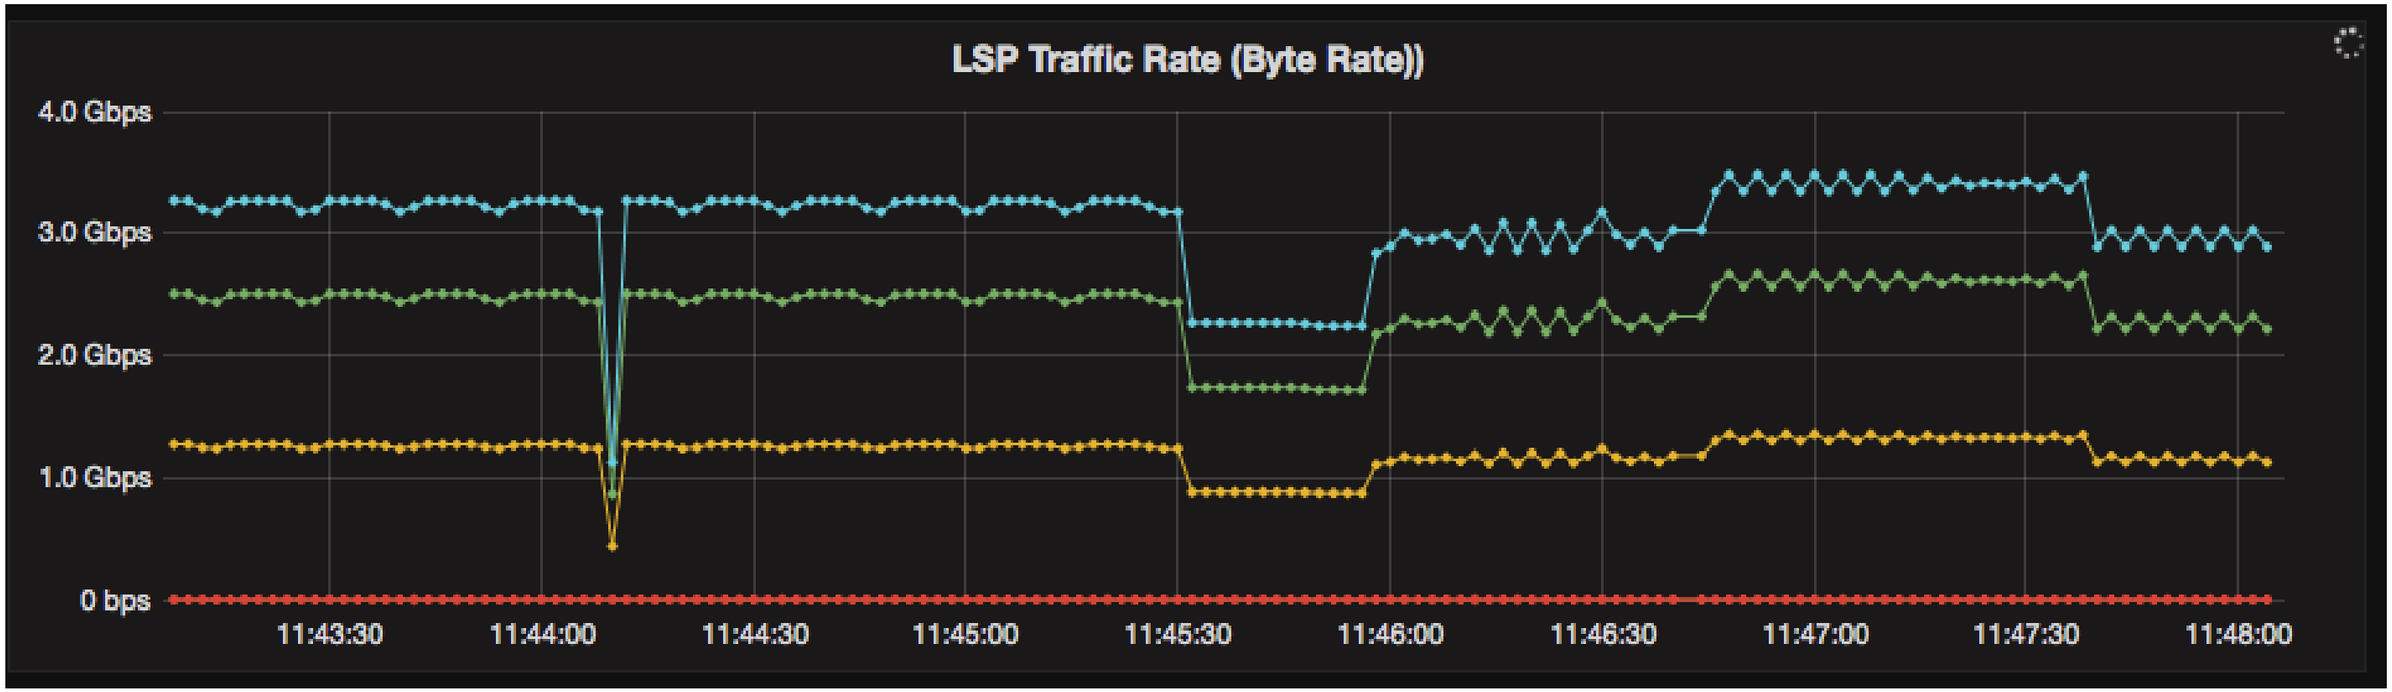

30 seconds

NorthStar Controller with Telemetry

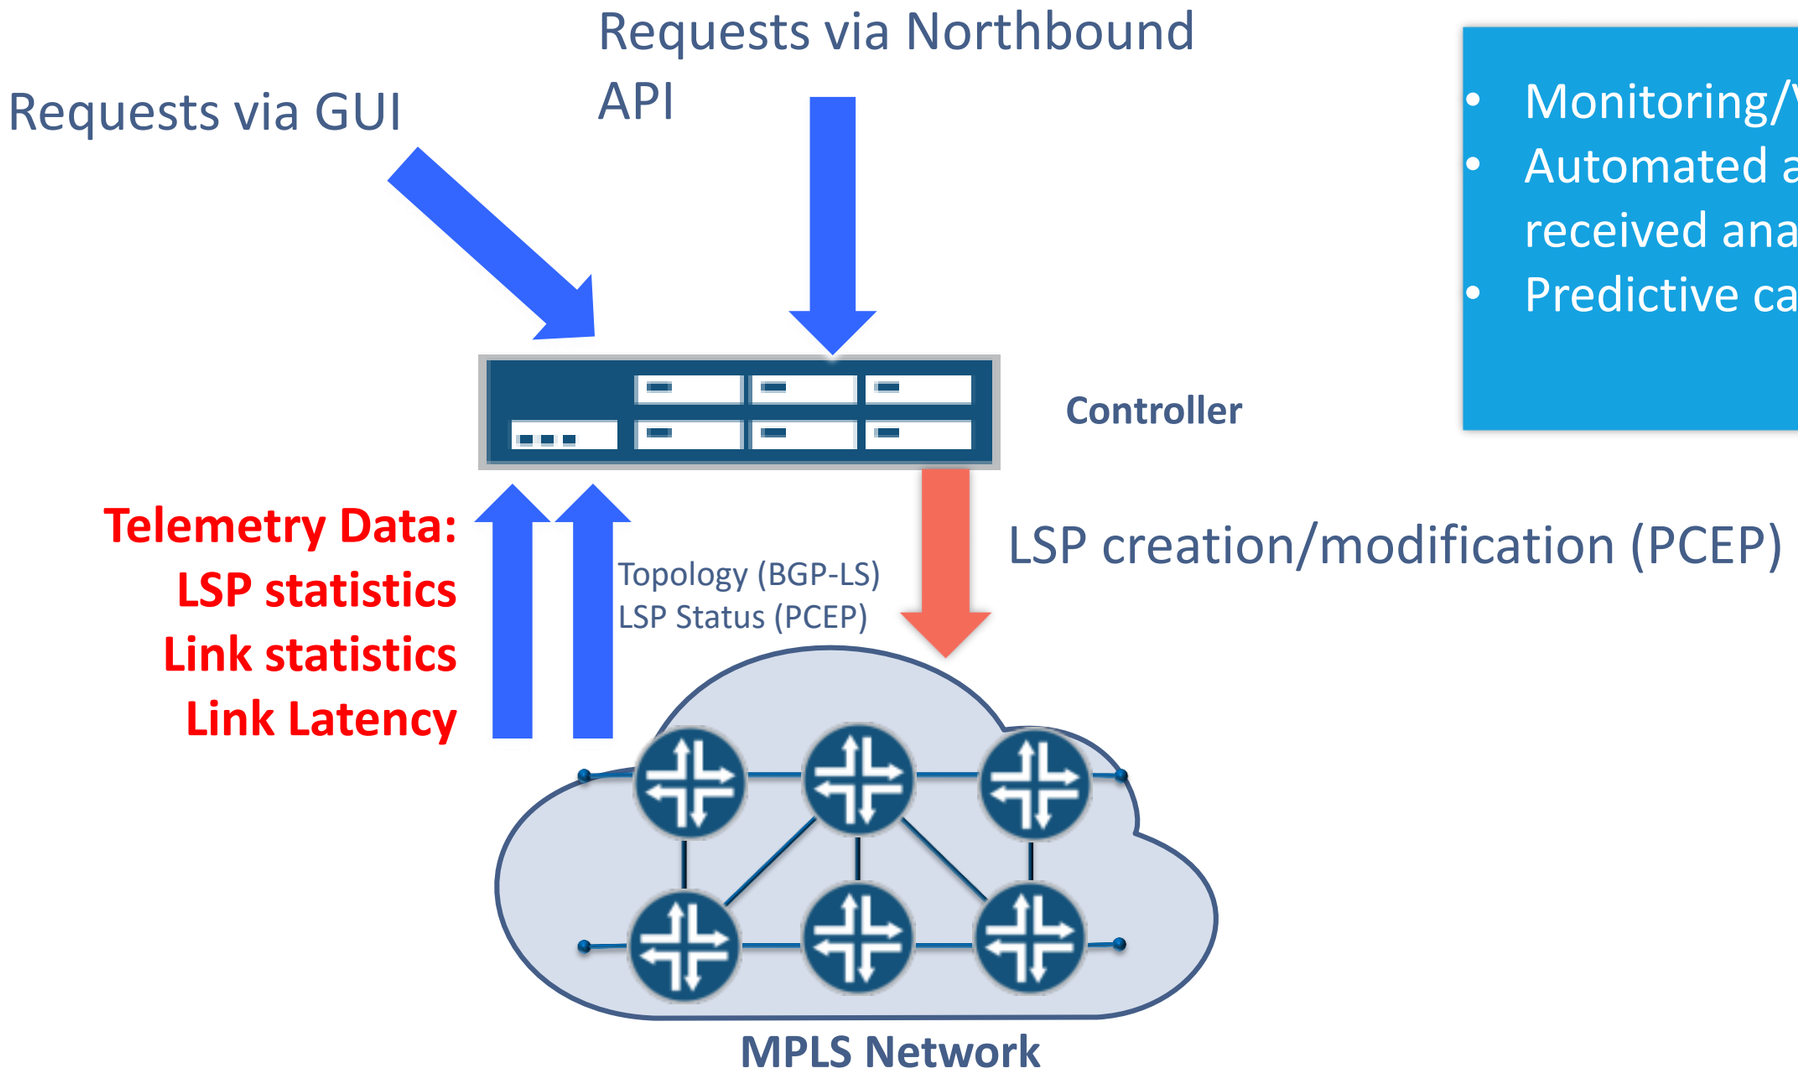

- Monitoring/Visualization
- Automated actions, based on received analytics
- Predictive capacity planning

NorthStar: Live link utilisation heatmap

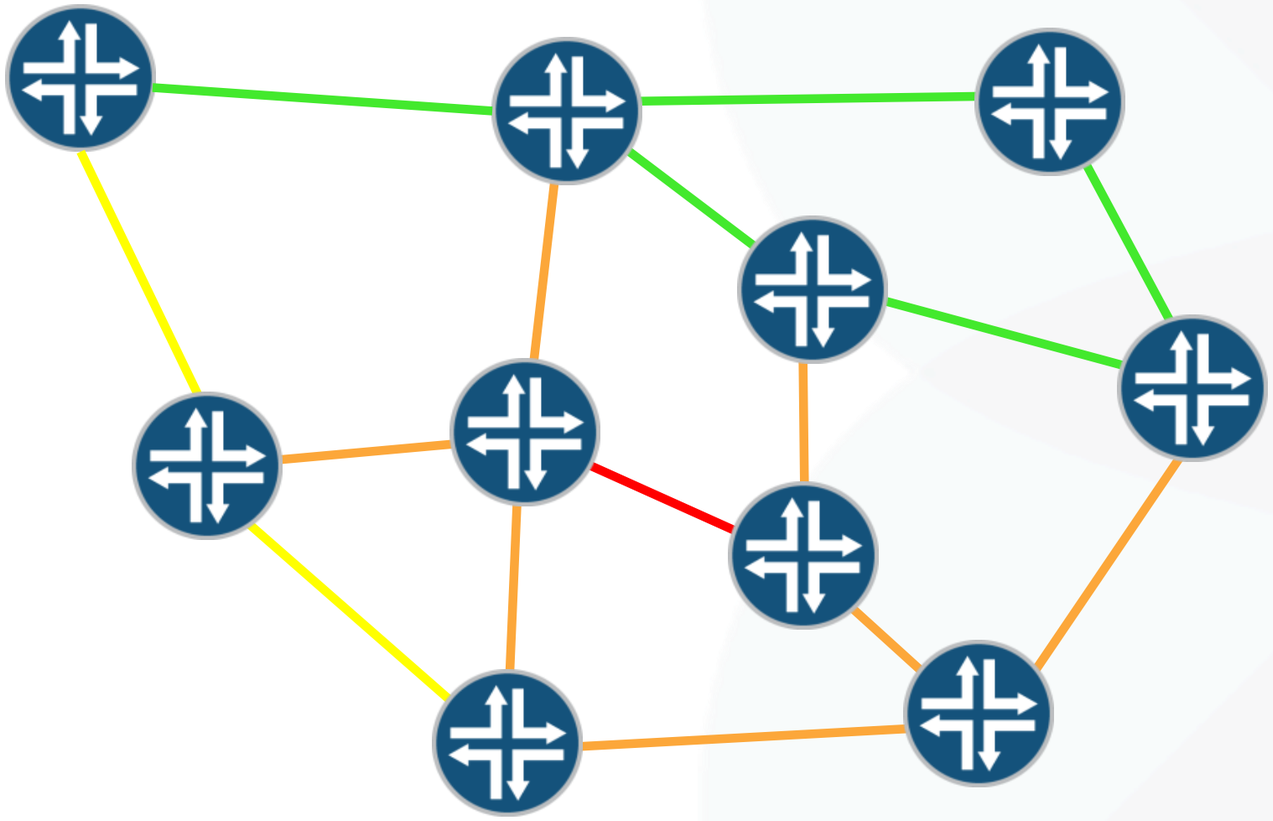

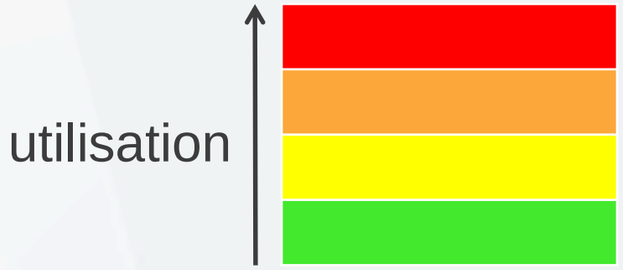

NorthStar: Per-LSP statistics

Name	Node A	Node Z	Capacity	Utilization	Delay	Packet Loss	Metric	Control Type	Path Type	Op Status	Setup	Hold	Record Route	Schedule	Last Update
LSP_PE1_PE2-EF	11.0.0.1	11.0.0.2	1G	23M	5.4ms	0%	100	PCC Initiated	Primary	Up	07	07	11.0.0.7-11.0.0.2		2016-03-01 18:00:00
LSP_PE1_PE3-EF	11.0.0.1	11.0.0.3	2G	1.1G	24ms	0%	200	PCC Initiated	Primary	Up	07	07	11.0.0.7-11.0.0.3		2016-03-01 19:00:00
LSP_PE1_PE4-EF	11.0.0.1	11.0.0.4	3G	5.25G	1.3ms	0%	300	PCC Initiated	Primary	Up	07	07	11.0.0.7-11.0.0.4		2016-03-01 20:00:00
LSP_PE1_PE5-EF	11.0.0.1	11.0.0.5	5G	3.14M	7.7ms	0%	400	PCC Initiated	Primary	Up	07	07	11.0.0.7-11.0.0.5		2016-03-01 21:00:00

Link latency heatmap

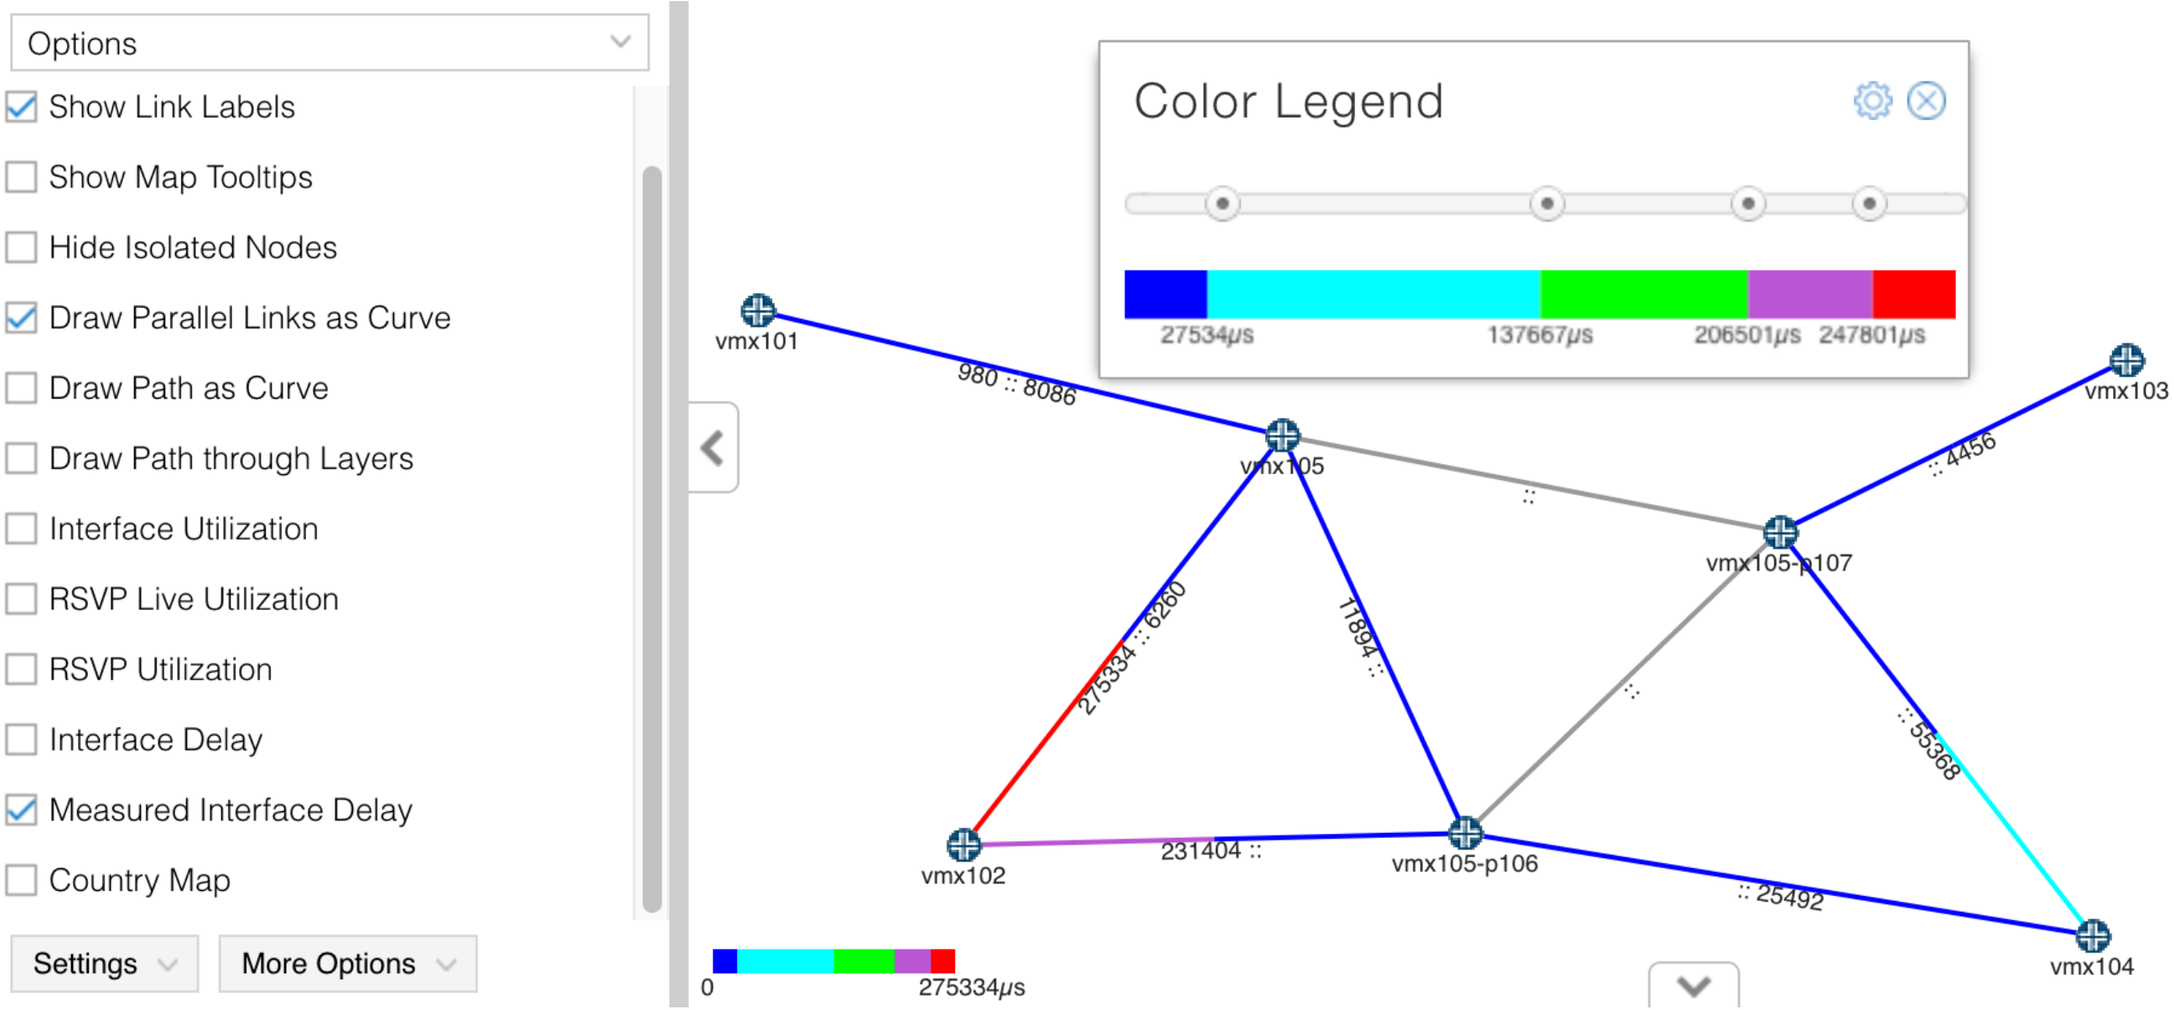

Automated actions, based on link utilisation (1)

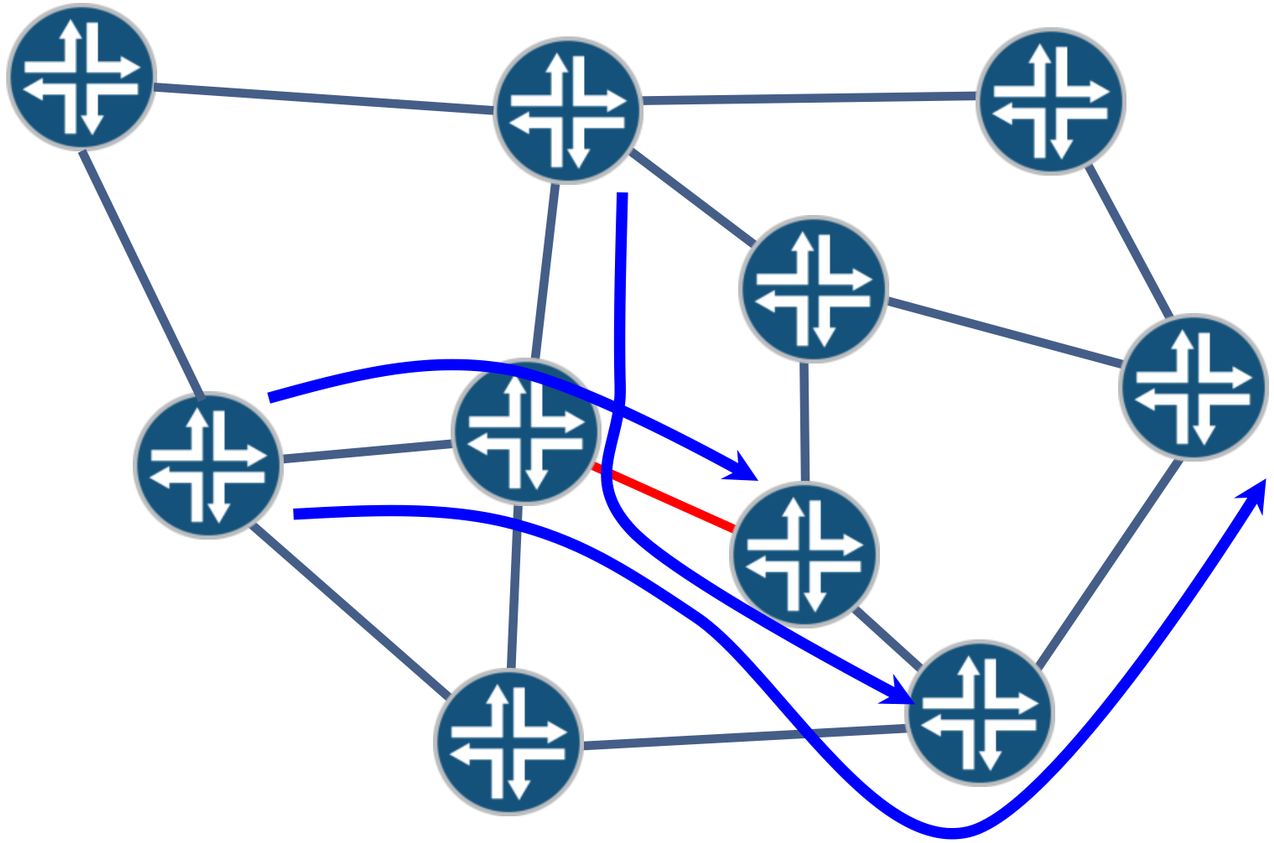

Red link is currently experiencing high utilisation

Or is *anticipated* to experience high utilisation in near future

Automated actions, based on link utilisation (2)

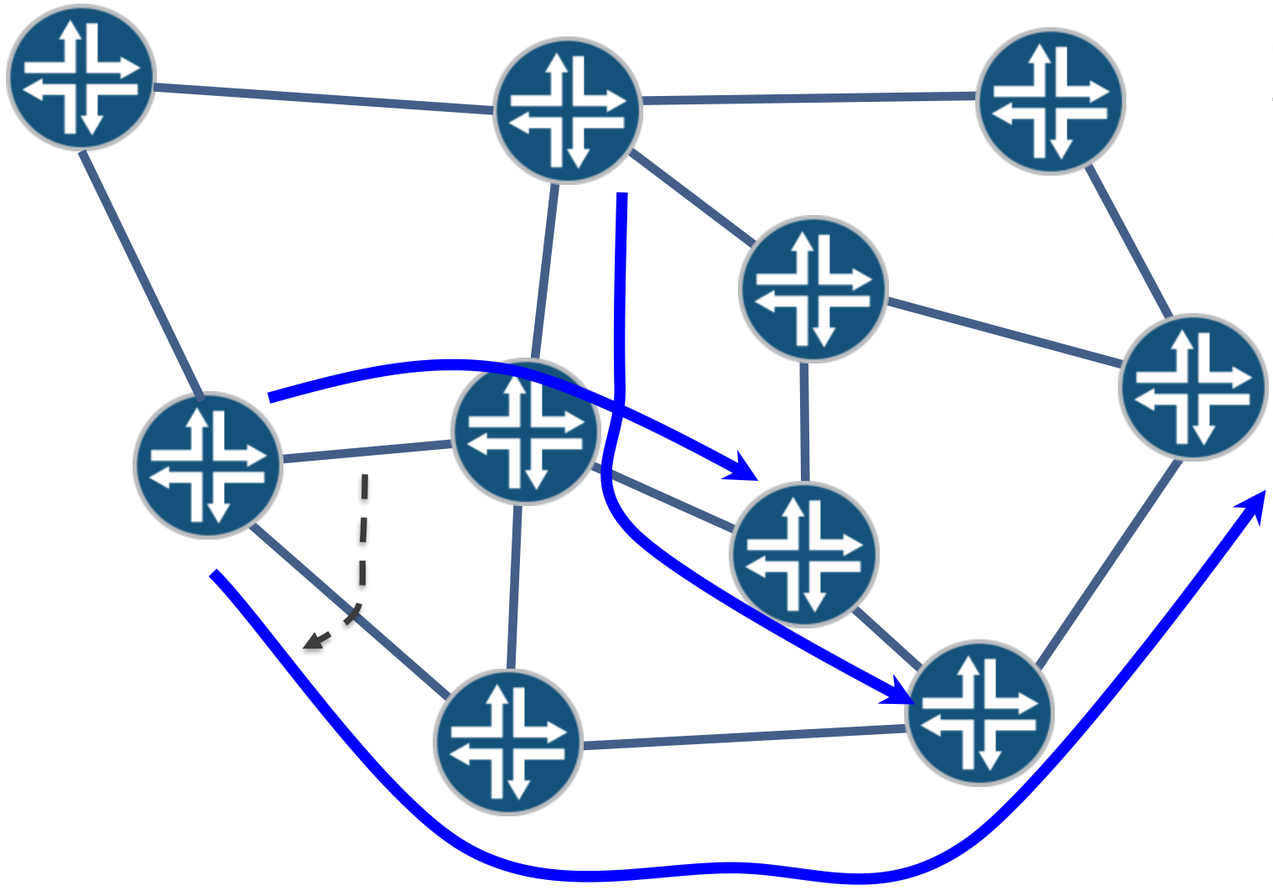

Controller knows how much traffic is travelling on each LSP

So can automatically move away some LSPs from congested link

Automated actions, based on delay measurements (1)

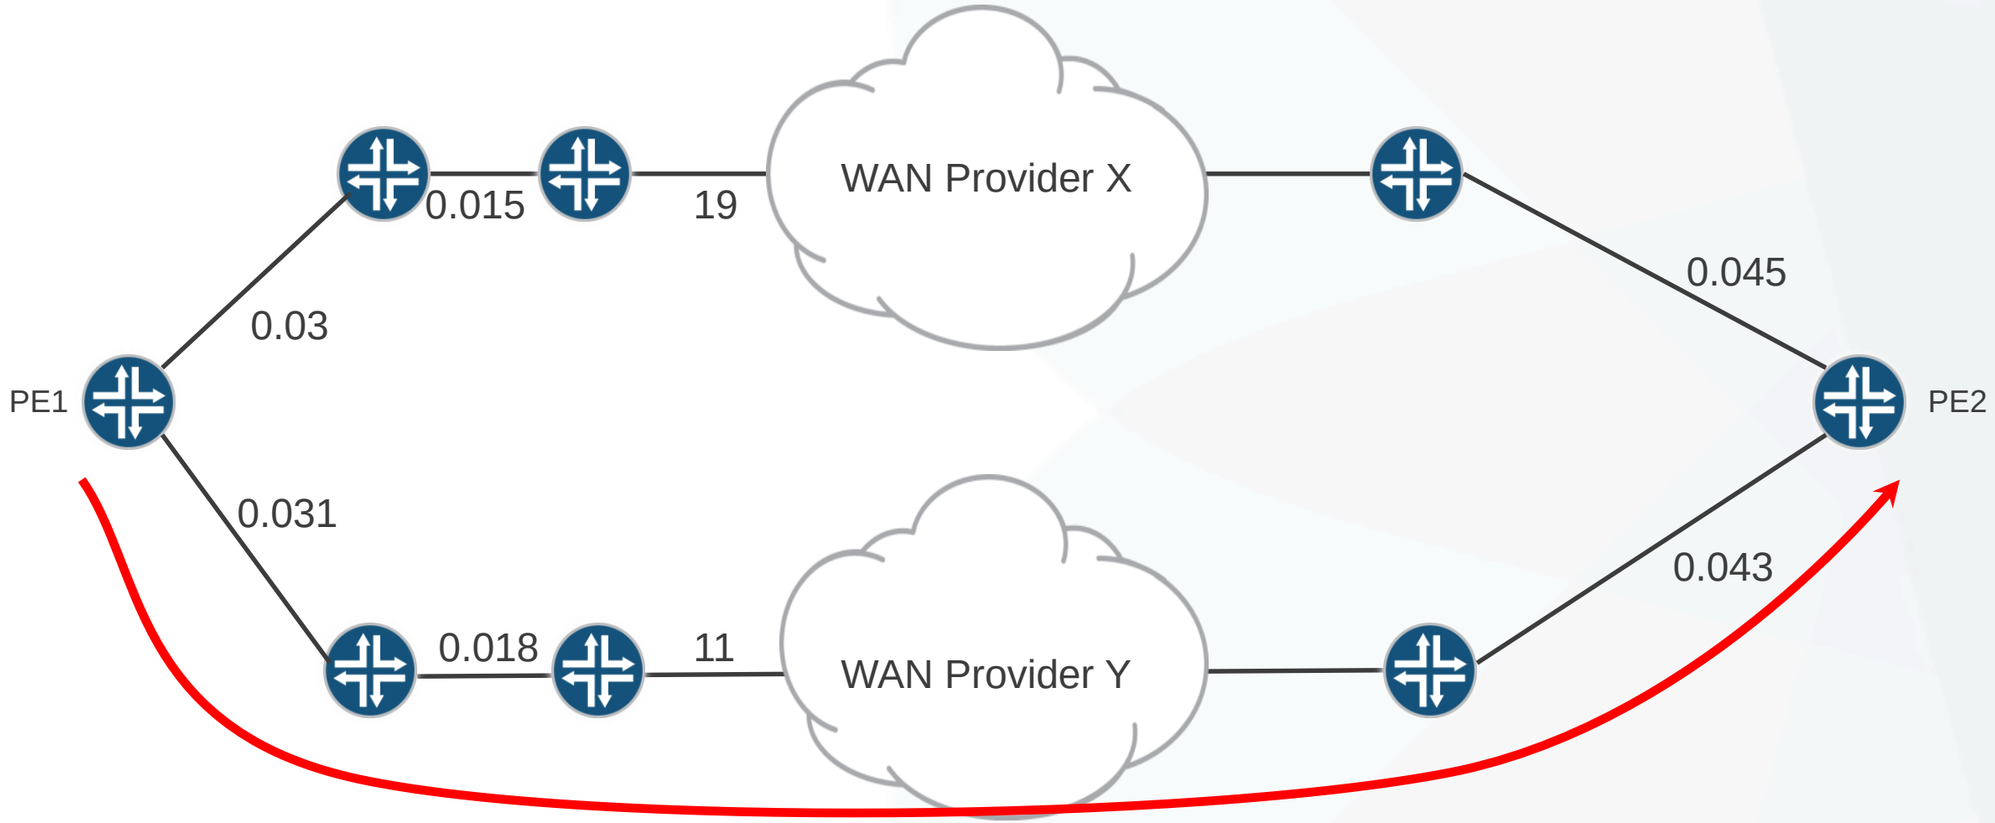

Link delay values are measured by routers and reported to Controller periodically.
Red LSP carries delay-sensitive traffic so Controller uses delay as the cost function when computing the path.

Automated actions, based on delay measurements (2)

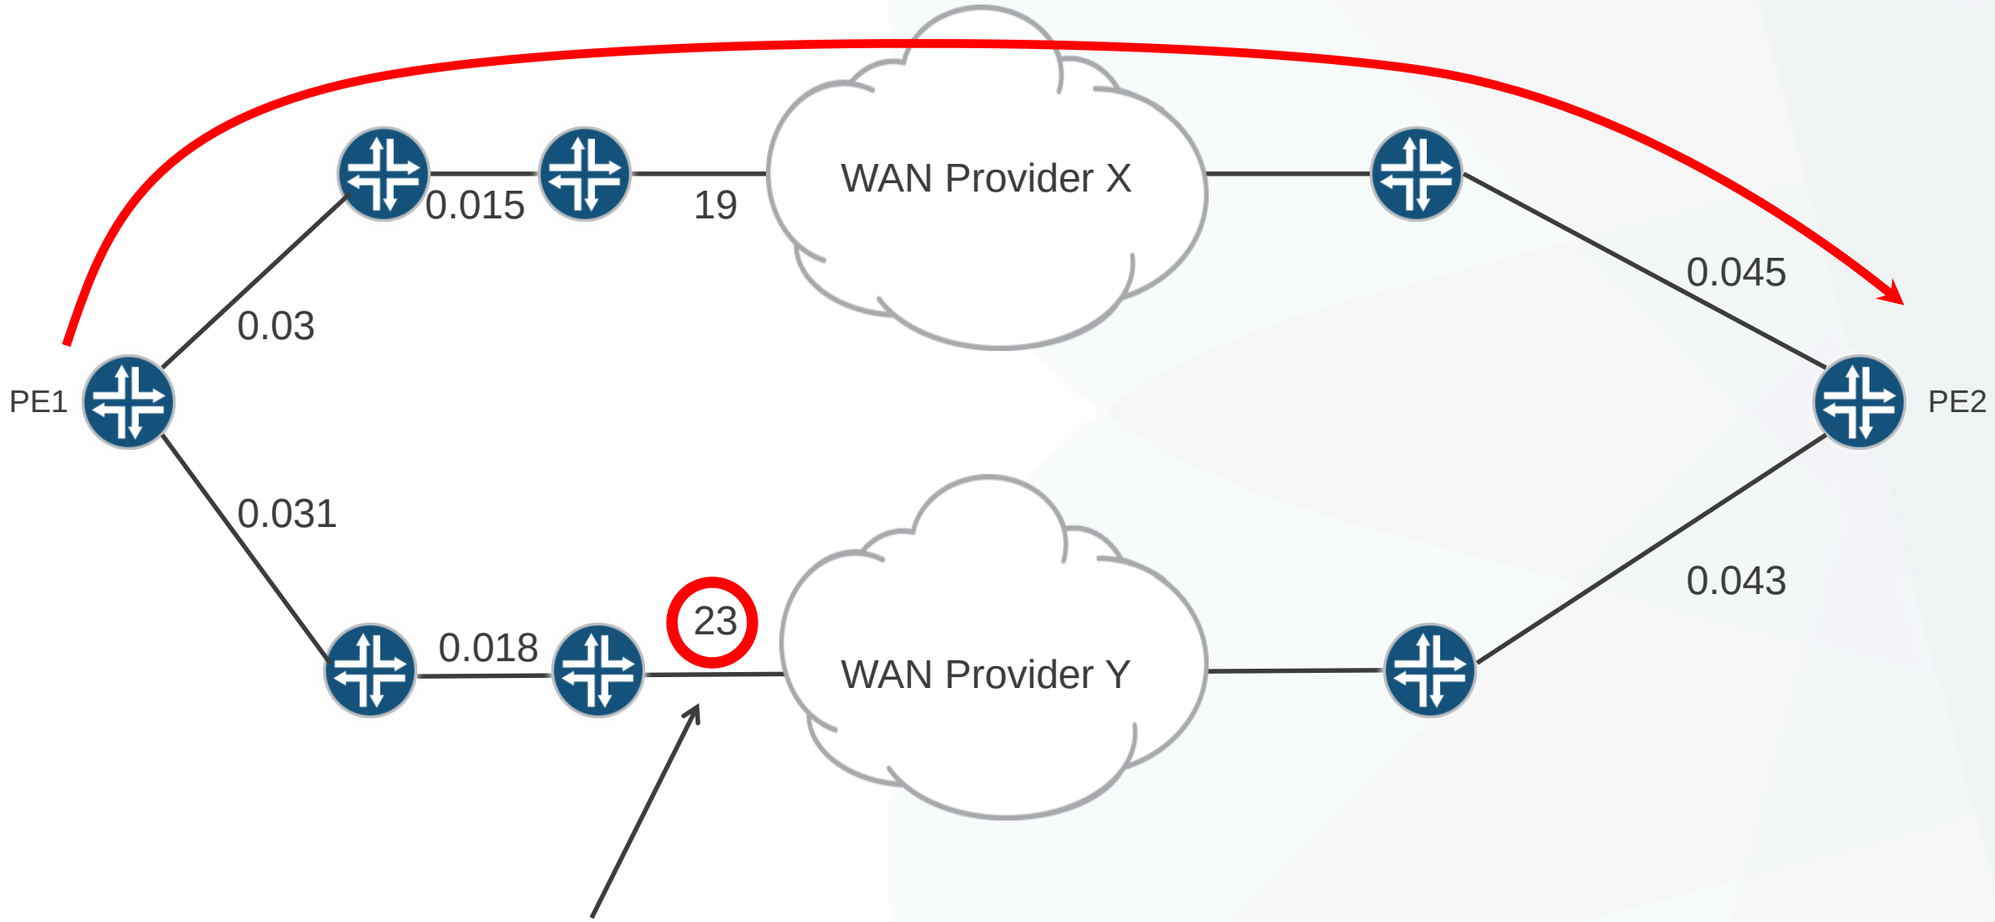

Step change in delay via Provider Y (due to reroute or protection event within Provider Y)
Controller reroutes red LSP via Provider X.

Predictive Capacity Planning