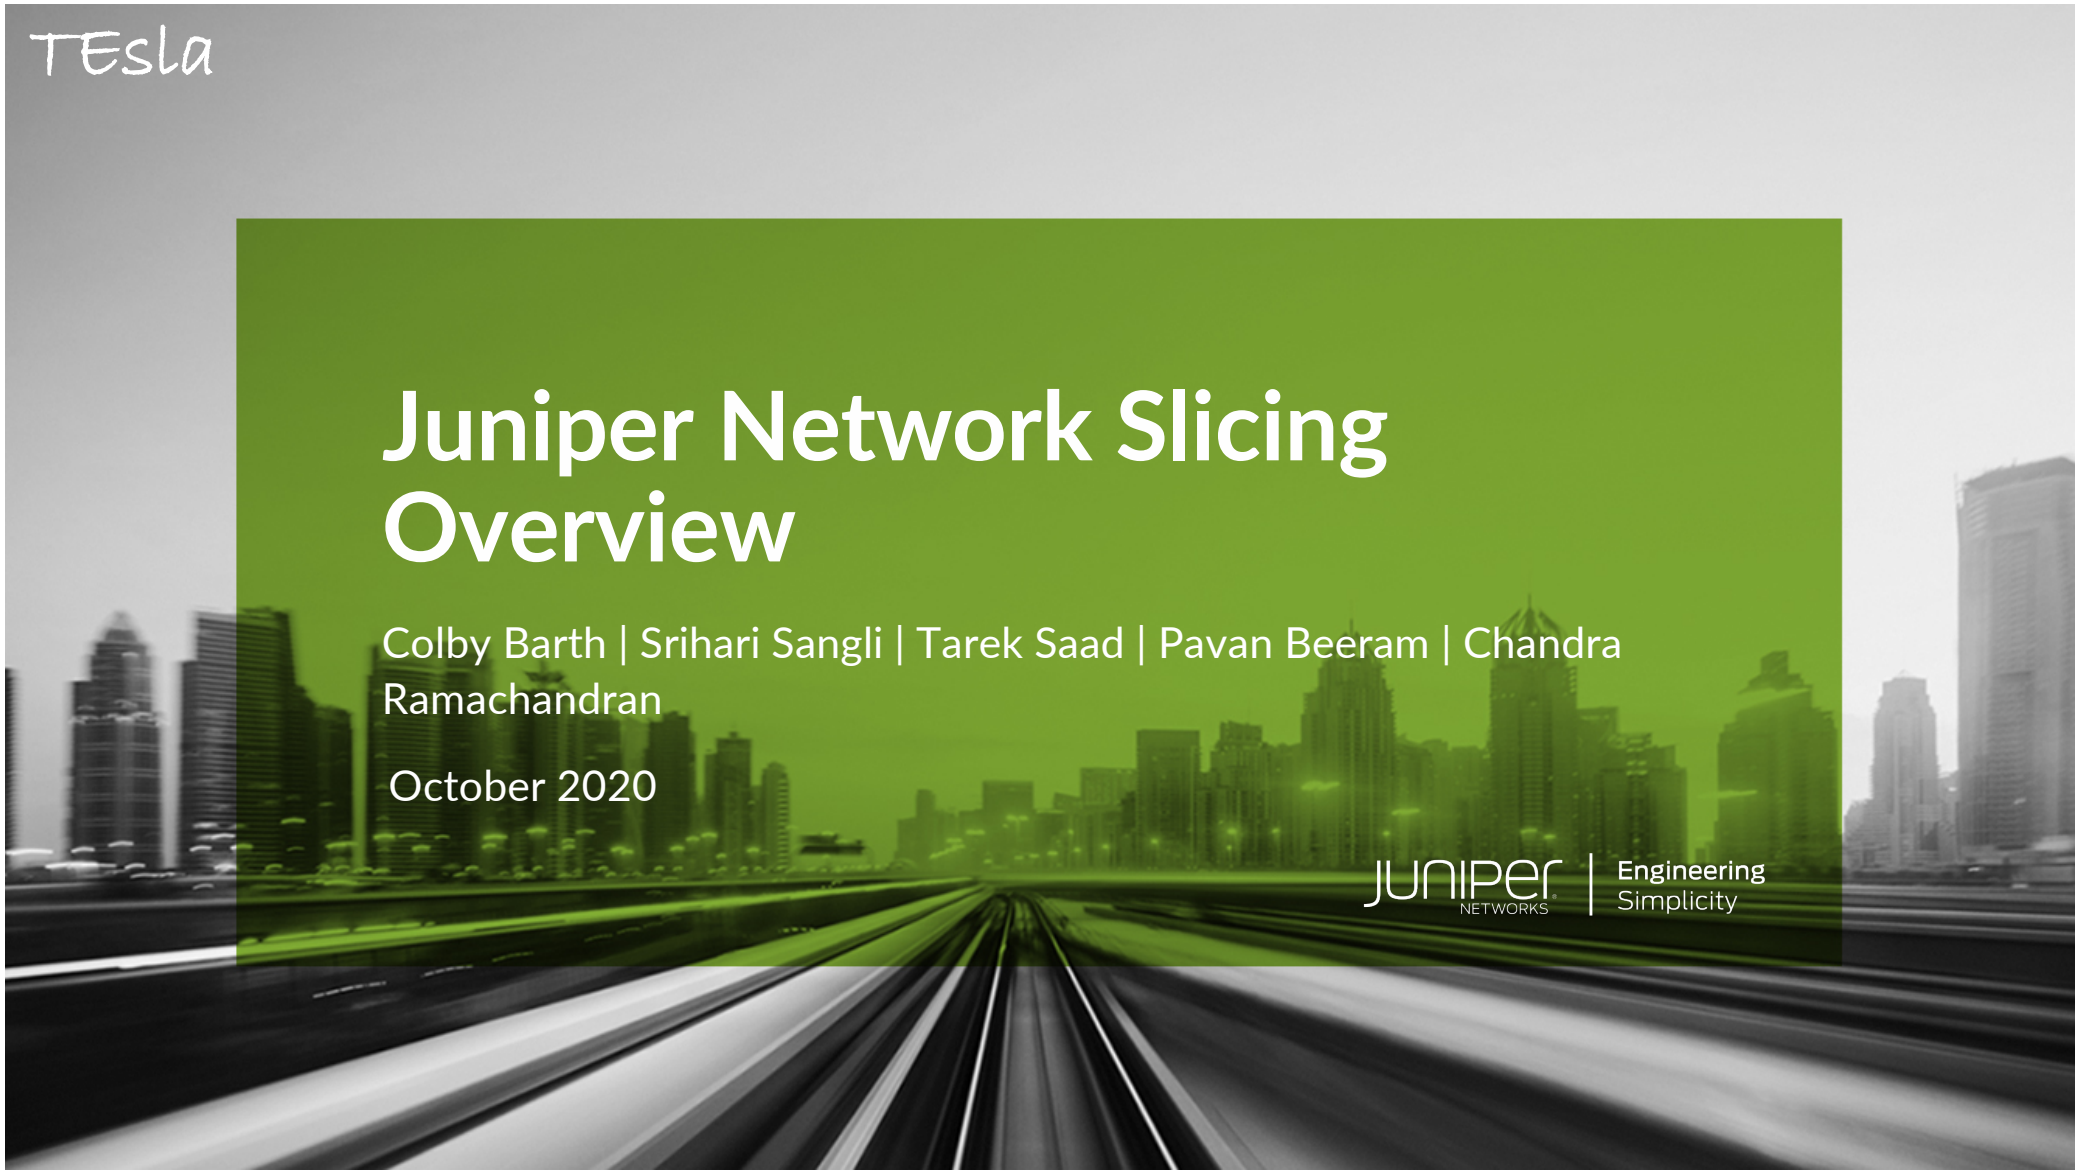

Juniper Network Slicing Overview

Colby Barth | Srihari Sangli | Tarek Saad | Pavan Beeram | Chandra Ramachandran

October 2020

Juniper Networks | Engineering Simplicity

Legal Statement

- This statement of direction sets forth Juniper Networks' current intention and is subject to change at any time without notice.
- No purchases are contingent upon Juniper Networks delivering any feature or functionality depicted in this presentation.

Team BIO's

Colby Barth is a 10+ year veteran Distinguished Engineer at Juniper Networks with over 20 years of industry experience working in the Routing Protocols group

Srihari Sangli is a Distinguished Engineer at Juniper Networks. With over 25 years of industry experience, he is in Juniper Routing team driving SR, BGP & TE solutions

Pavan Beeram is a Distinguished Engineer in Juniper's Routing team with over 20 years of experience designing and developing Traffic Engineering solutions. He currently serves as the co-chair of IETF TEAS WG.

Tarek Saad is a Principal Engineer in Juniper's Routing team driving SR & TE solutions. He has over 17 years of industry experience designing software for wide range of routing protocols. He currently serves as the co-chair of IETF MPLS WG.

Chandra Ramachandran is a s/w developer in Juniper's MPLS Routing Team with close to 20 years of experience in designing & developing Routing Protocols & MPLS Traffic Engineering applications.

Agenda

Brief Review/Overview

Deployment modes discussion

Conclusion

Overview: goals, requirements & intent

Why are we here?

- How do 5G network requirements translate to transport network SLAs?
- What enhancements are needed for an IP transport network to meet new 5G requirements?
- What new differentiation in service offerings can network providers offer on top of their existing network infrastructure?

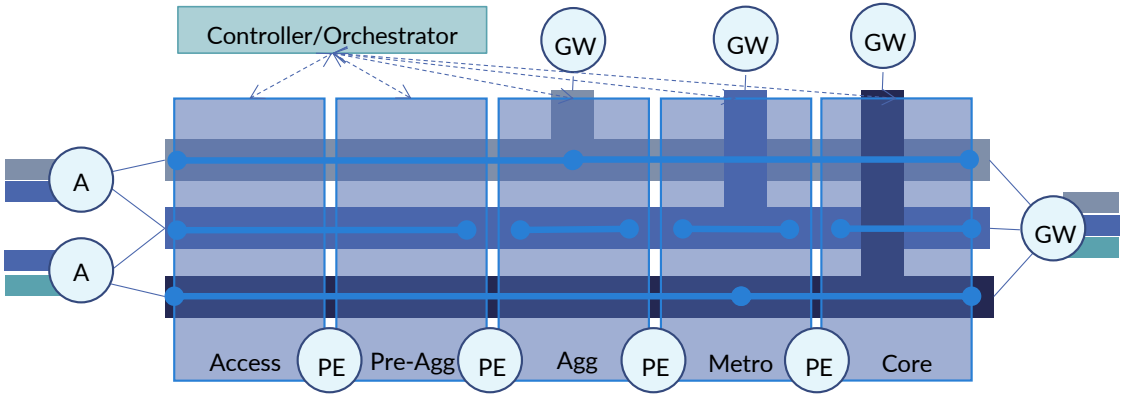

Multi-Domain(AS/Area/Level) End-to-End Network

A model driven approach

SDN is about automation/programming via centralized “Controllers”

- Requires standardized models for describing Network Services to facilitate such automation/programming

These models, in turn, MUST cater to various deployment scenarios

- Model is ‘compiled’ into configuration(NETCONF) instructions and/or protocol extensions(PCEP, BGP, …)
- Facilitates the Slice specific SLO’s within the network e.g. end-to-end BW and latency guarantees

“Gluing” together disparate features does not facilitate such flexibility

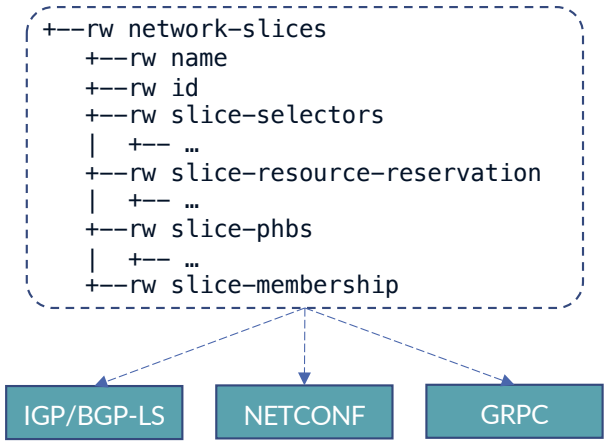

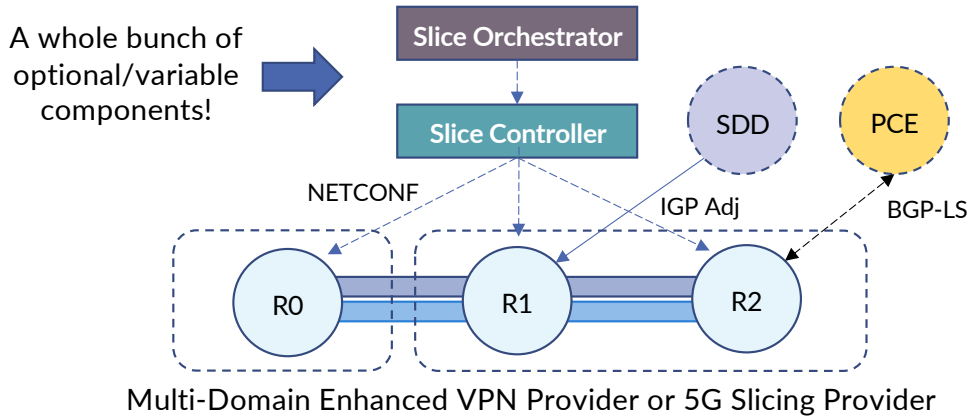

Multi-Domain Enhanced VPN Provider or 5G Slicing Provider

Underlay transport agnostic

Packets carry a Slice identifier (additional label(SR, RSVP, LDP), static service label, specific forwarding label set(Flex-algo, pop-n-forward), …

- Step1 to realizing Network Slicing can not be 1st be to migrate to Segment-routing!

The slice identifier is associated with a Slice specific PHB that is divided into child-queues to instantiate per Slice hierarchical CoS for any underlay transport technology

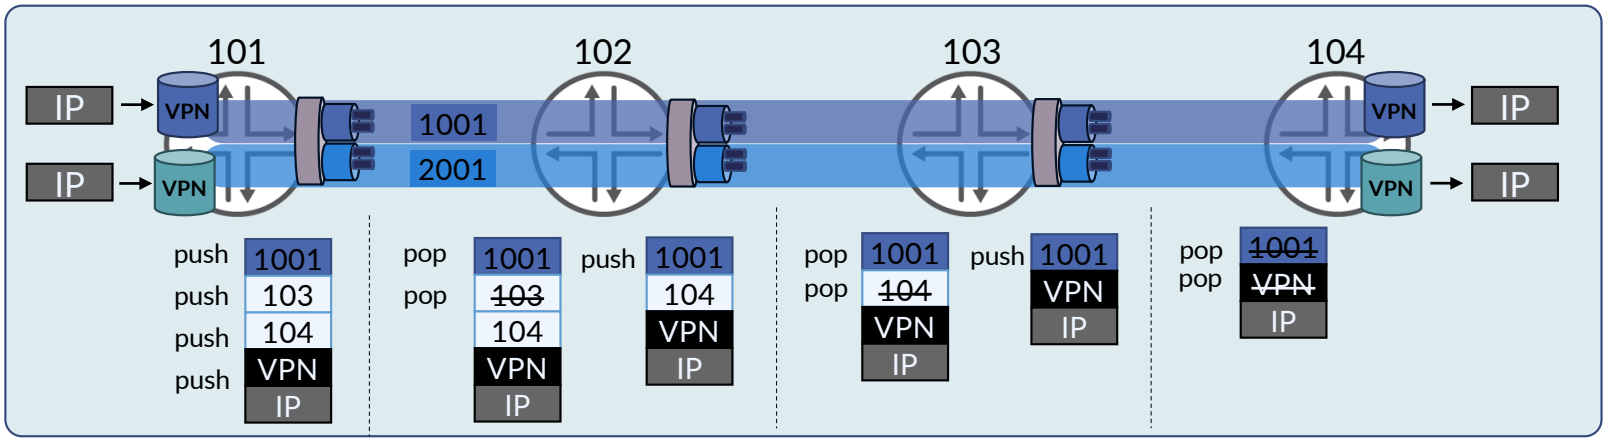

SR-TE, PHP transport LSP with VPN service example

Slice aware path placement

Lowest delay with per Slice/Q granularity example

Slice Unaware TE (Relaxed mode)

- Path placement may include TE attributes, e.g. delay metrics, link admin-groups, but not Slice specific
- E.g. bounded delay path including any purple or green links

Flex-algo as a TE solution (Relaxed mode)

- Flex-algo includes TE attributes, e.g. delay metric, link admin-groups, but is not Slice specific
- E.g. Least delay path including any purple or green links

Slice Aware TE (Elastic/Guaranteed mode)

- Path placement includes Slice specific resources
- Per slice Traffic Engineering Database (TED)
- E.g. bounded delay path within a Slice considering per Slice available BW

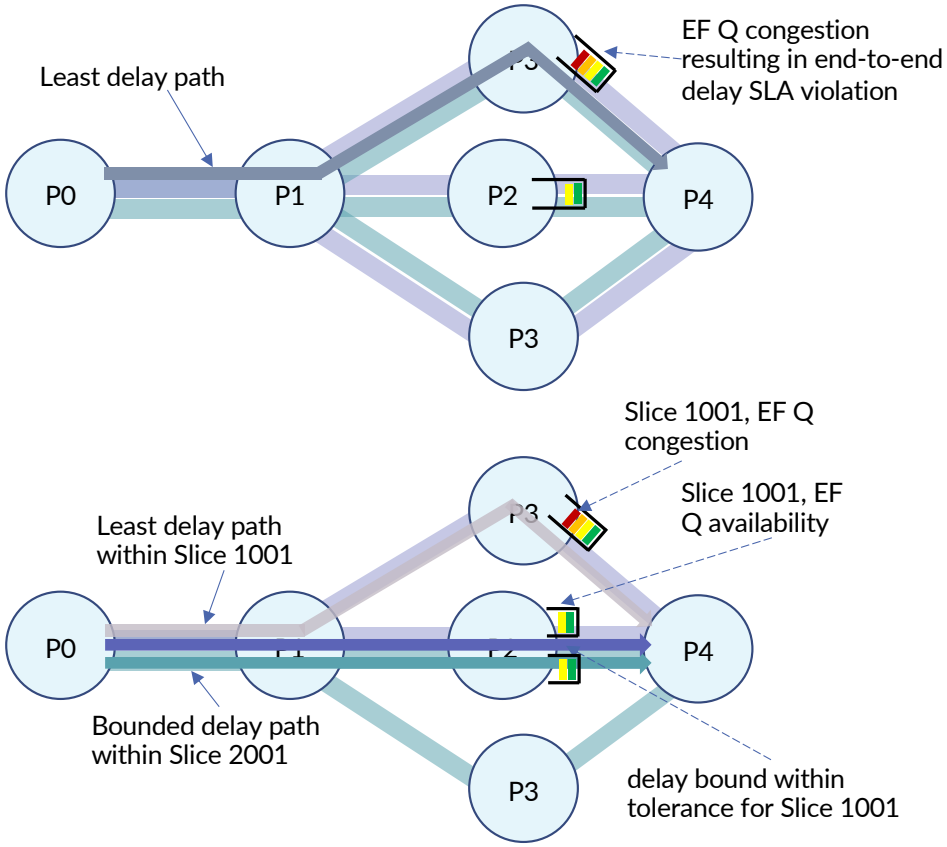

3 modes of operation

Relaxed, Elastic, and Guaranteed

Relaxed, Elastic, & Guaranteed Modes

Hybrid Scenarios are also permitted

Relaxed Mode

Data-plane only Slicing solution

Slice Unaware TE

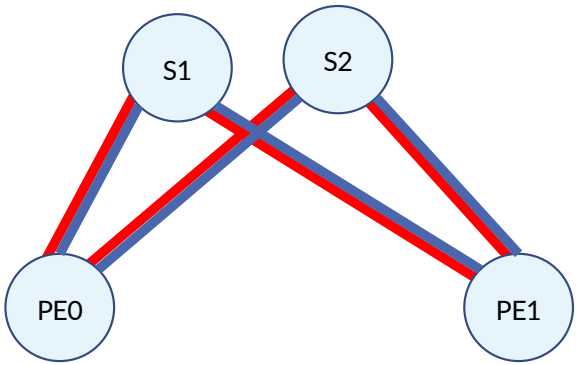

Slice specific data-plane resource allocation

Elastic Mode

Control-plane only Slicing solution

Slice aware TE

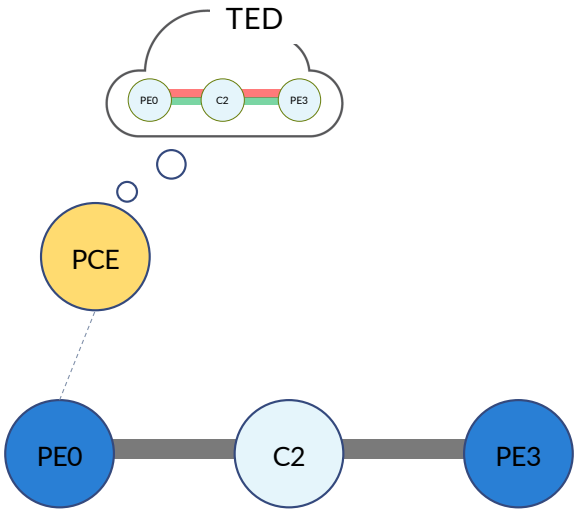

Slice aware TE
Simple, common, 4-8Q EXP based PHBs may still be deployed

Guaranteed Mode

Control & Data-plane Slicing solution

Slice aware TE

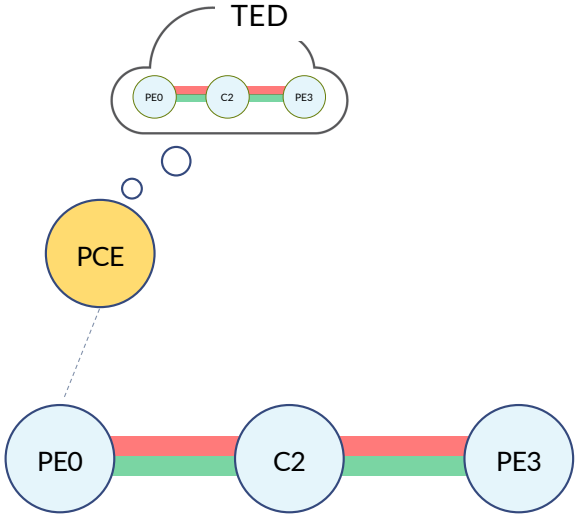

Slice aware TE & Per Slice H-QoS, with guaranteed control-plane & data-plane resources

Data plane only network slicing
Relaxed Mode

Slicing uses transport or Slice label inferred PHB

Slice data plane

- Slice indicator is present
- Per slice CoS profile is applied on participating links and nodes
- Transit nodes classify incoming traffic (e.g. using Slice label) and apply per slice scheduling

Slice control plane

- No control plane awareness of slice resource(s)
- No slice-aware path placement/TE

Use-case

- Suitable when no BW engineering is required and ECMP is leveraged between endpoints (e.g. Spine/Leaf deployment)
- Does not address all network slice SLOs being standardized at IETF

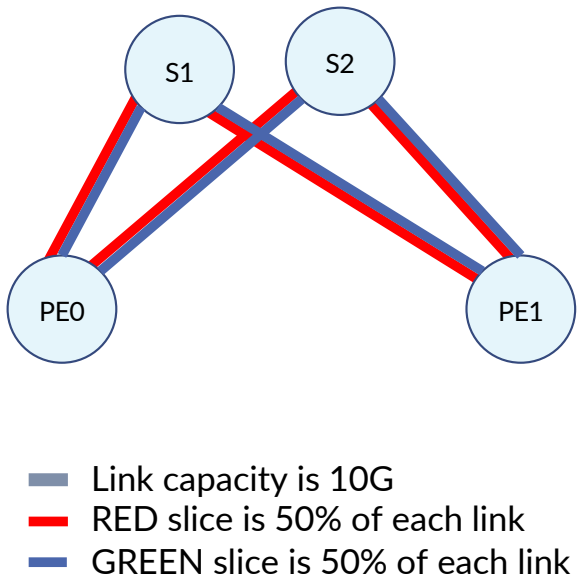

Slicing example – Relaxed mode

RSVP or SR-TE transport network(s), Slice Indicator = MPLS top-most label

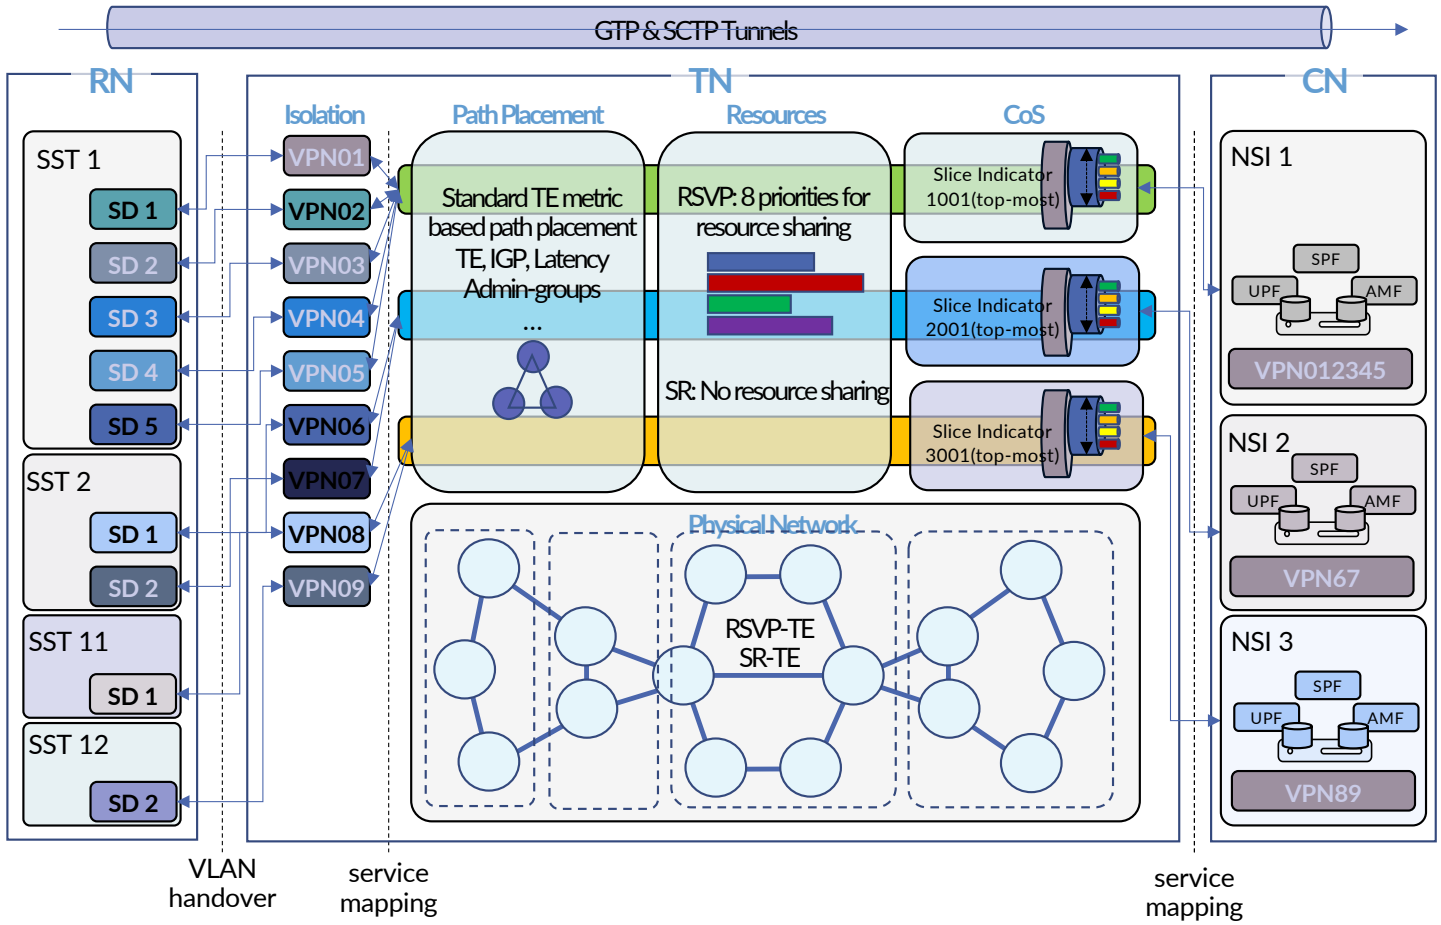

Slicing example – Relaxed mode

Flex-algo transport network(s), Slice Indicator = MPLS top-most label

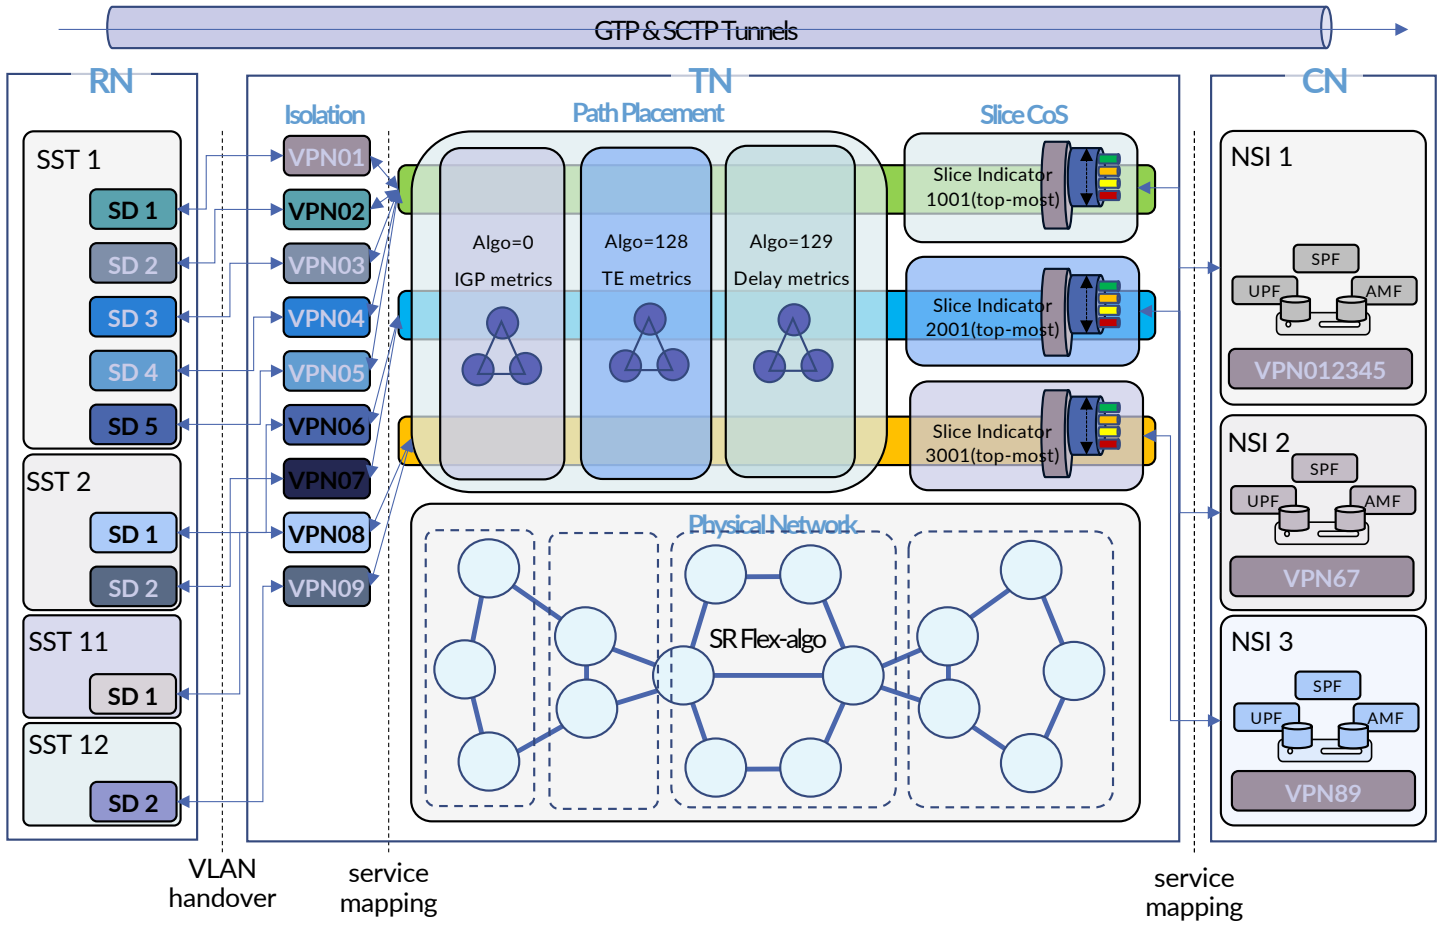

Control plane only network slicing
Elastic Mode

Compute & Resource management are Slice aware

Control plane only network slicing

- Link resources are sliced in control plane
  - Per slice link maximum and available BW
- Ingress routers/PCE forms per slice TED using the link-state
  - Slice-aware path computation and path placement
  - CSPF uses the slice aware TED to select optimal TE path
- Control plane preemption in case of contention on a specific link resource
  - In case of degradation of AE link, control plane can preempt LSP(s) to avoid congestion

Data plane

- No per slice classification of traffic or per slice PHB on transit routers
- Policing can happen on incoming links

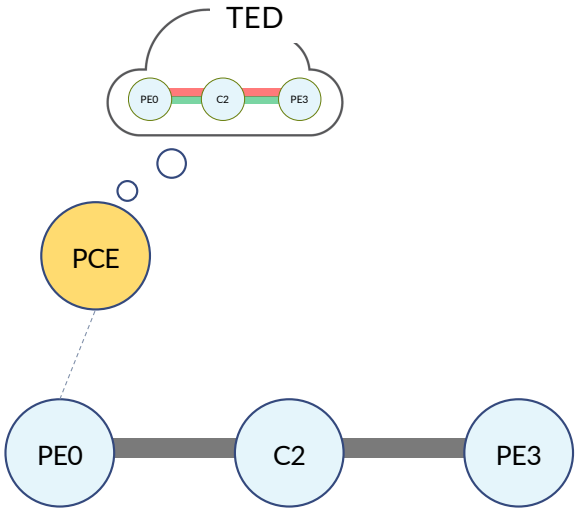

Simple, consistent, 4-8Q
PHBs may still be deployed

Slicing example – Elastic mode

TE transport network(s), Slice aware path placement, Slice Indicator = MPLS top-most label

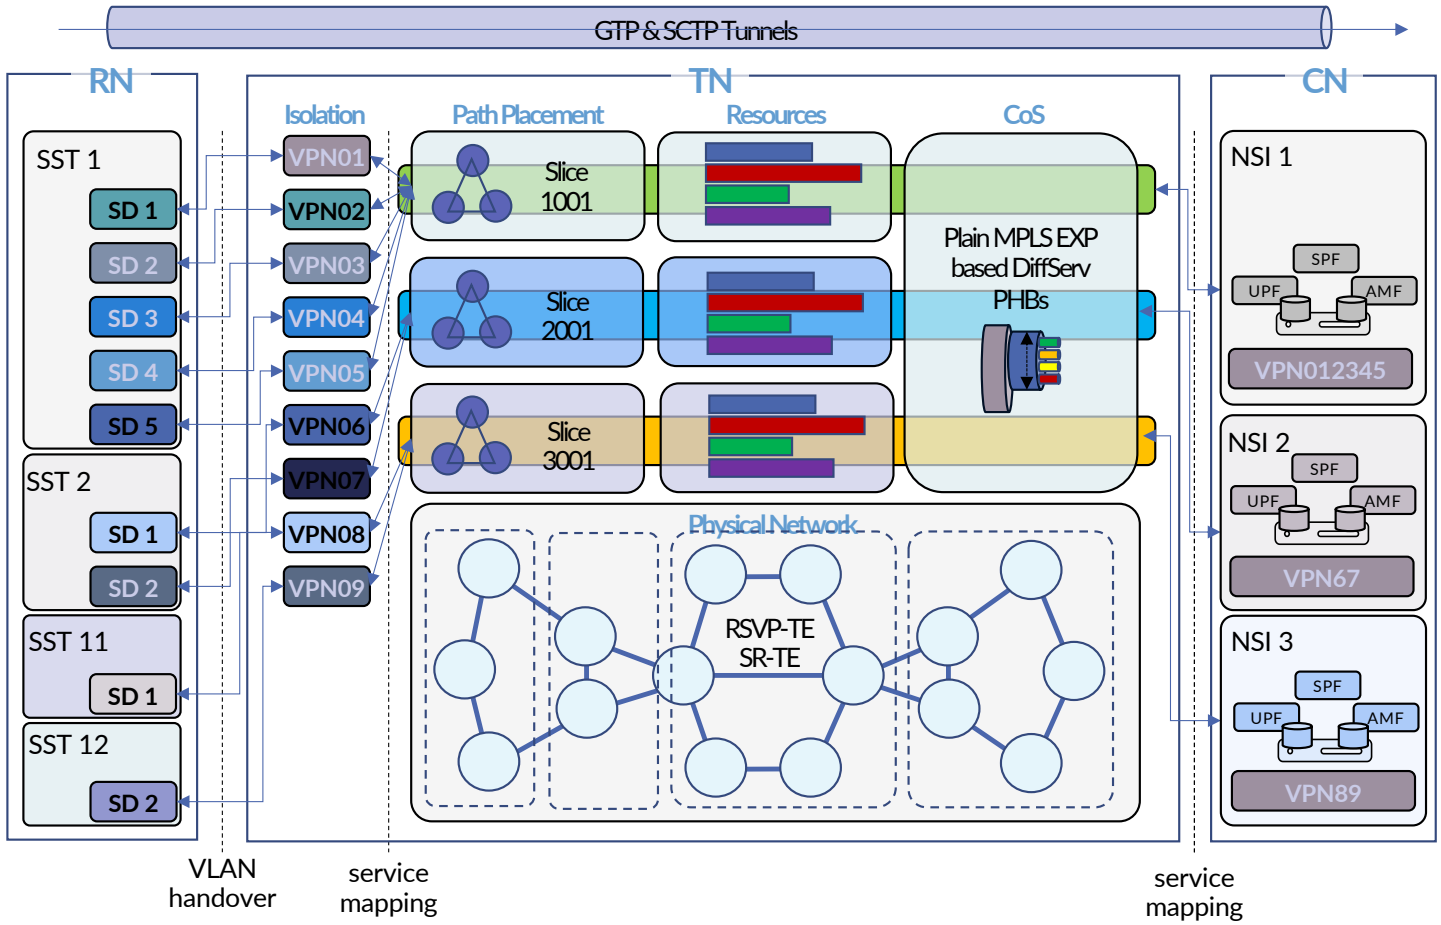

Control plane and data plane network slicing
Guaranteed Mode

Slice label inferred PHB
Compute & resource management are Slice aware

Control plane and data plane network slicing

- Combination of the previous two
- Control plane TED helps ingress router and/or PCE do proper placement of LSPs based on per slice link available BW
- Data-plane CoS on transit nodes provide guarantees in case of congestion on a link
  - Can occur when CP is slow to update reservation
  - Can occur during degradation of link while control plane preemptions are in-progress
- Covers strict and shared resource slice isolation requirements

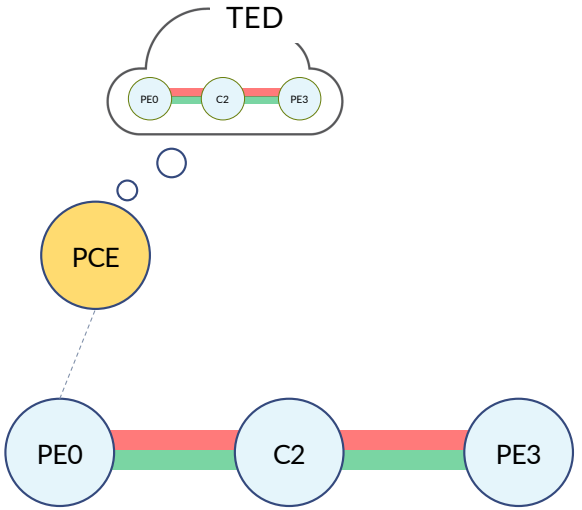

Slicing example – Guaranteed mode

TE transport network(s), Slice aware path placement, Slice Indicator = MPLS top-most label

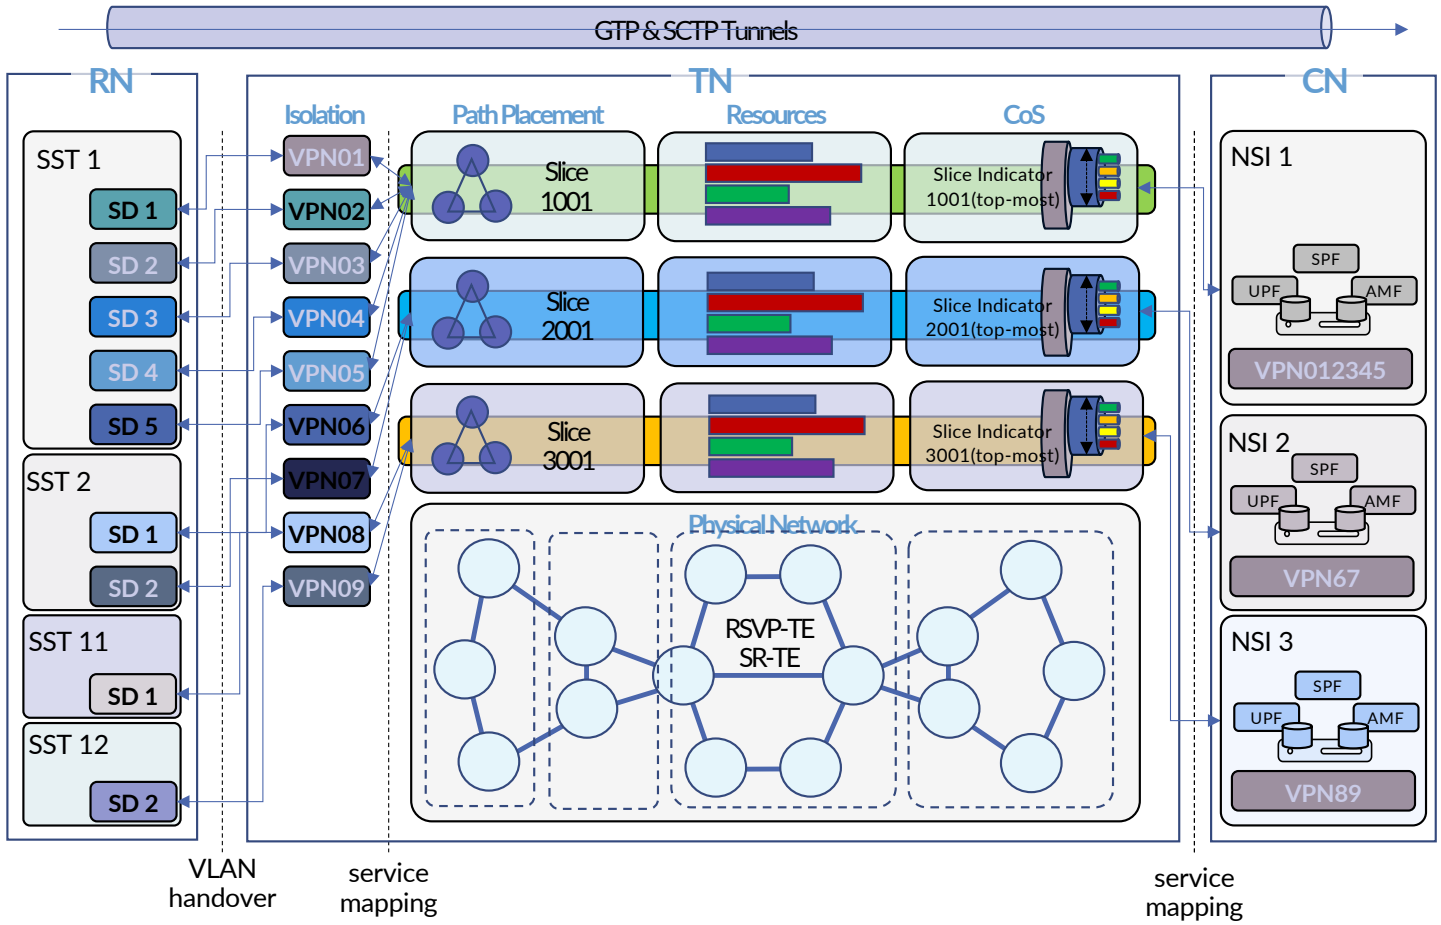

References

Realizing Network Slices in IP/MPLS Networks - https://tools.ietf.org/html/draft-bestbar-teas-ns-packet-00

YANG Data Model for Network Slice Per-Hop Definition - https://tools.ietf.org/html/draft-bestbar-teas-yang-ns-phd-00

Conclusion

Attributes of Juniper's Network Slicing solution

- The Slice definition is independent of the under-lay technology
- Any underlay technology can be used for Slice specific path placement
- Decoupled data-plane & Control-plane for highly scalable deployments
- A flexible data-plane identifier is used for Slice specific forwarding
- A Slice may include a customized topological set of nodes and links
- Integrates easily with TE Controllers/PCEs
- Easy to deploy using any of 3 modes: Relaxed, Elastic, & Guaranteed to suit the range of customer requirements


Thank you
JUNIPER NETWORKS
Engineering
Simplicity