CODING BUSINESS LOGIC INTO THE NETWORK THROUGH AUTOMATION

A Large Tier1 Case study

Nitin Vig & Aditya Kaul

Solution Architects – Juniper Networks

JUNIPER NETWORKS | Driven by Experience™

The autonomous network journey

- From manual to self-driven networks powered by AI

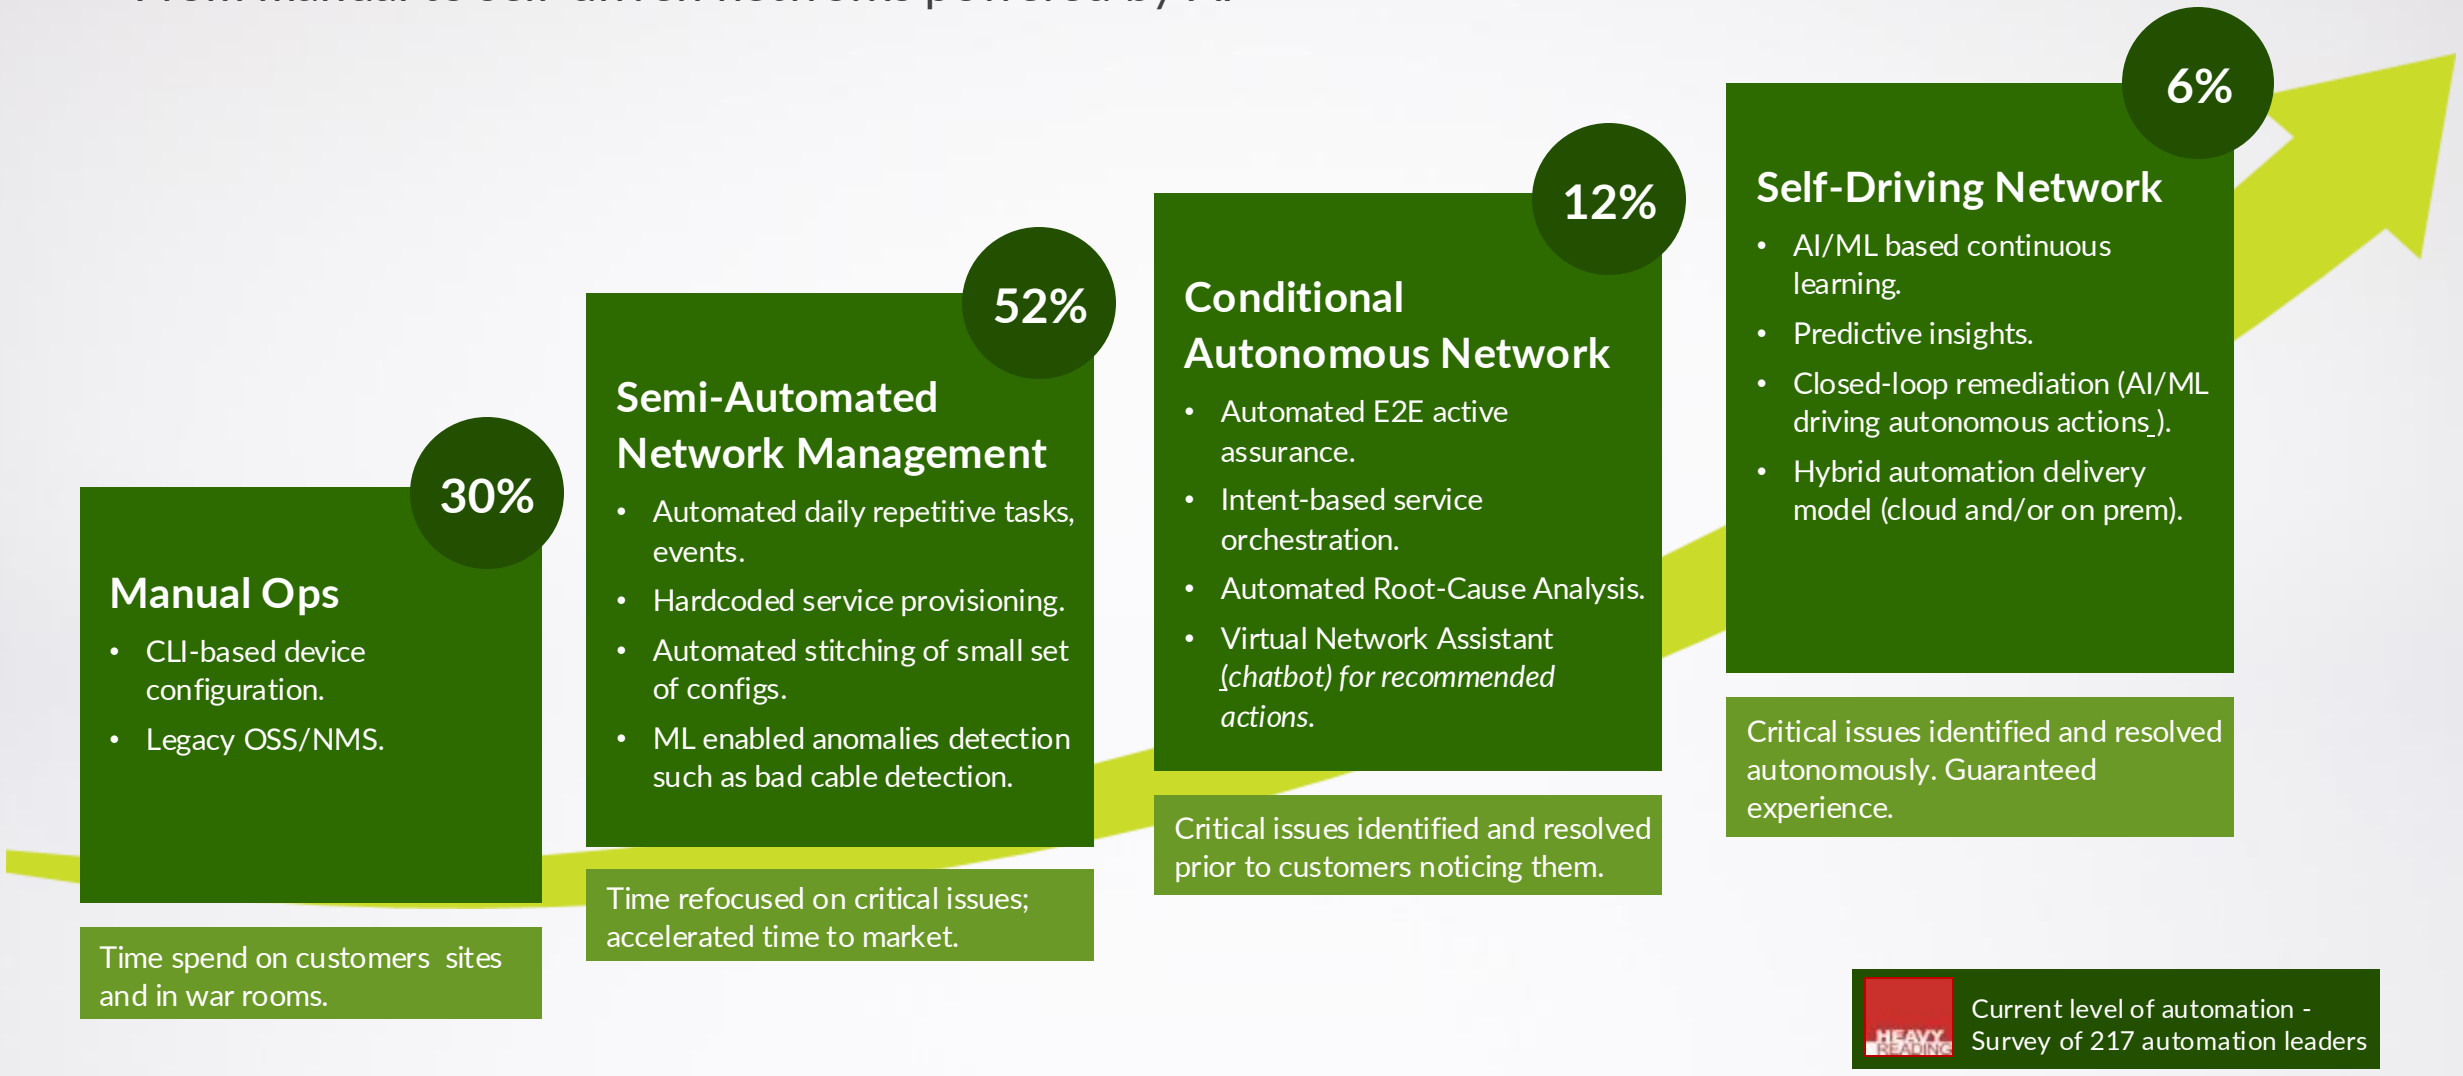

Motivation & Goals

Executive Summary

Automation is a key initiative that cuts across architecture layers in a Service Provider environment

Networks need a technology shift and investments to maintain leadership position and profitable growth

Automation

- Create operational simplicity
- Predict demand and service experience needs
- Enable greater levels of transparency and control to the network and services.
- Differentiate customer experience
- Increase speed to market
- Create market offers (Service on demand)
- Reduce the cost to serve at scale.

Business outcomes

Cost saving

- Service prioritization

Improved customer experience

- Strict Diversity guarantee

Revenue growth

- Latency-aware routing
- Smart Bandwidth

*The **cost of doing nothing** is that the business will miss market commitments and face declining revenues.*

Current Challenges vs Performance Outcomes

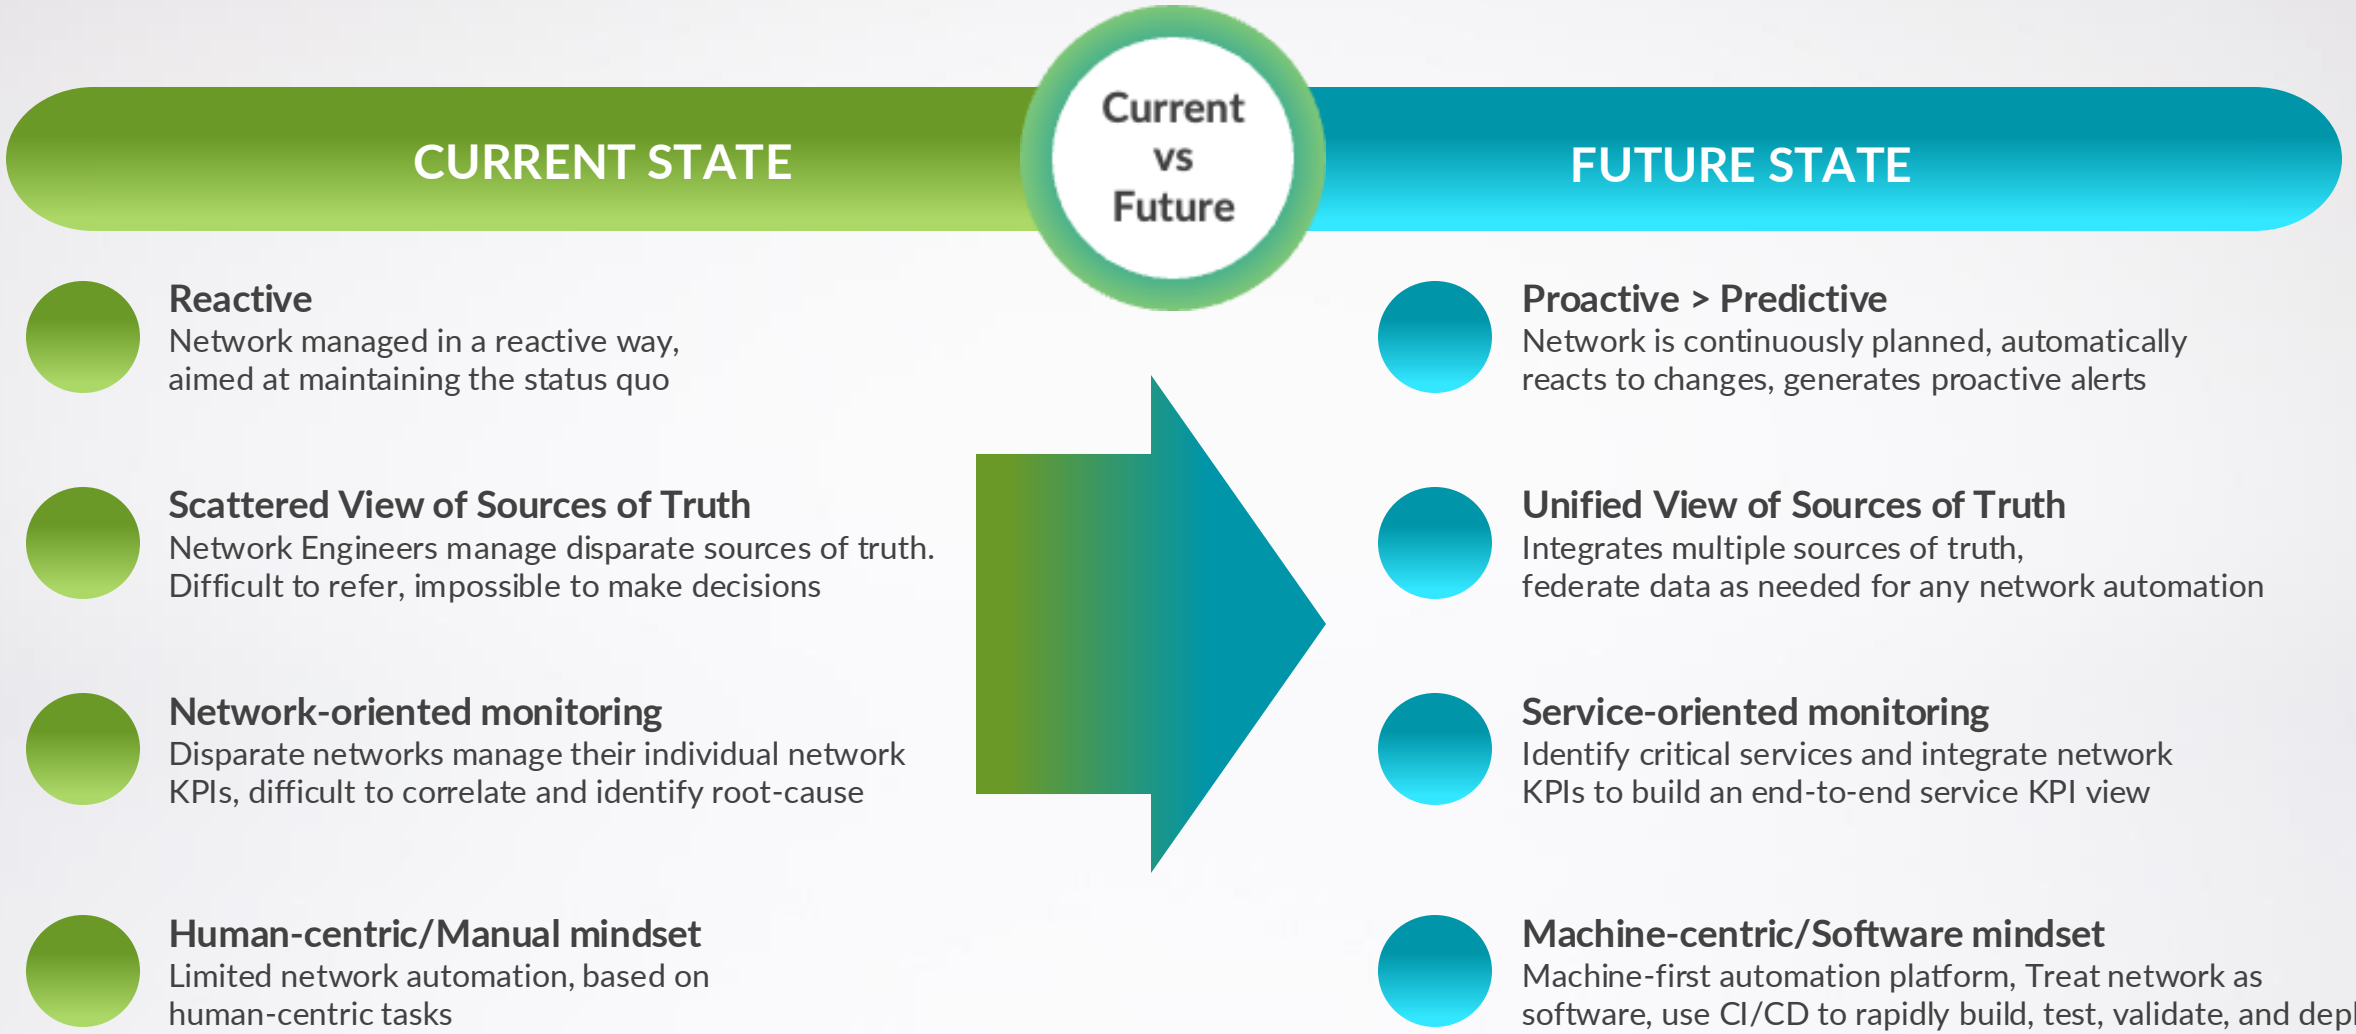

Customer goals

Use-case		Business impact	
Service Prioritization	Direct lower priority traffic to longer latency paths	• Safely run the network hotter • Delay upgrades especially on premium routes • Added capability to compute paths based on circuit cost	Cost saving
Strict Diversity Guarantee	Dynamically build 'disjoint' network paths during normal conditions & network outages	• Deliver true service separation end-to-end (a 'fully-redundant' service) • A more resilient network that automatically reacts to faults and network design changes (maintains diversity during failure)	Customer experience
Latency-aware routing	Engineer traffic for specific services based on measured or derived network latency	• Transport selected traffic flows over lower latency paths as defined by policy • Rapidly react to detected latency changes and recompute new paths. • Replace home-grown solution (Robot) by off the shelf solution	Revenue growth
Smart Bandwidth	Detect congestion and offload excess traffic to cheaper, longer latency links	• Opportunity to create new service tiers (Premium, Standard+, Standard) • Better utilization of unused network capacity • Better experience for high-profile customers	Revenue growth
Planned or 'unplanned' event automation	Automated traffic routing during a planned/unplanned event based on pre-defined rules	• Advance preparation of altered network behaviour designed to accommodate a known forthcoming event (know what to expect) • Reduce or eliminate impact to customers from known events	Customer experience

Legacy Network

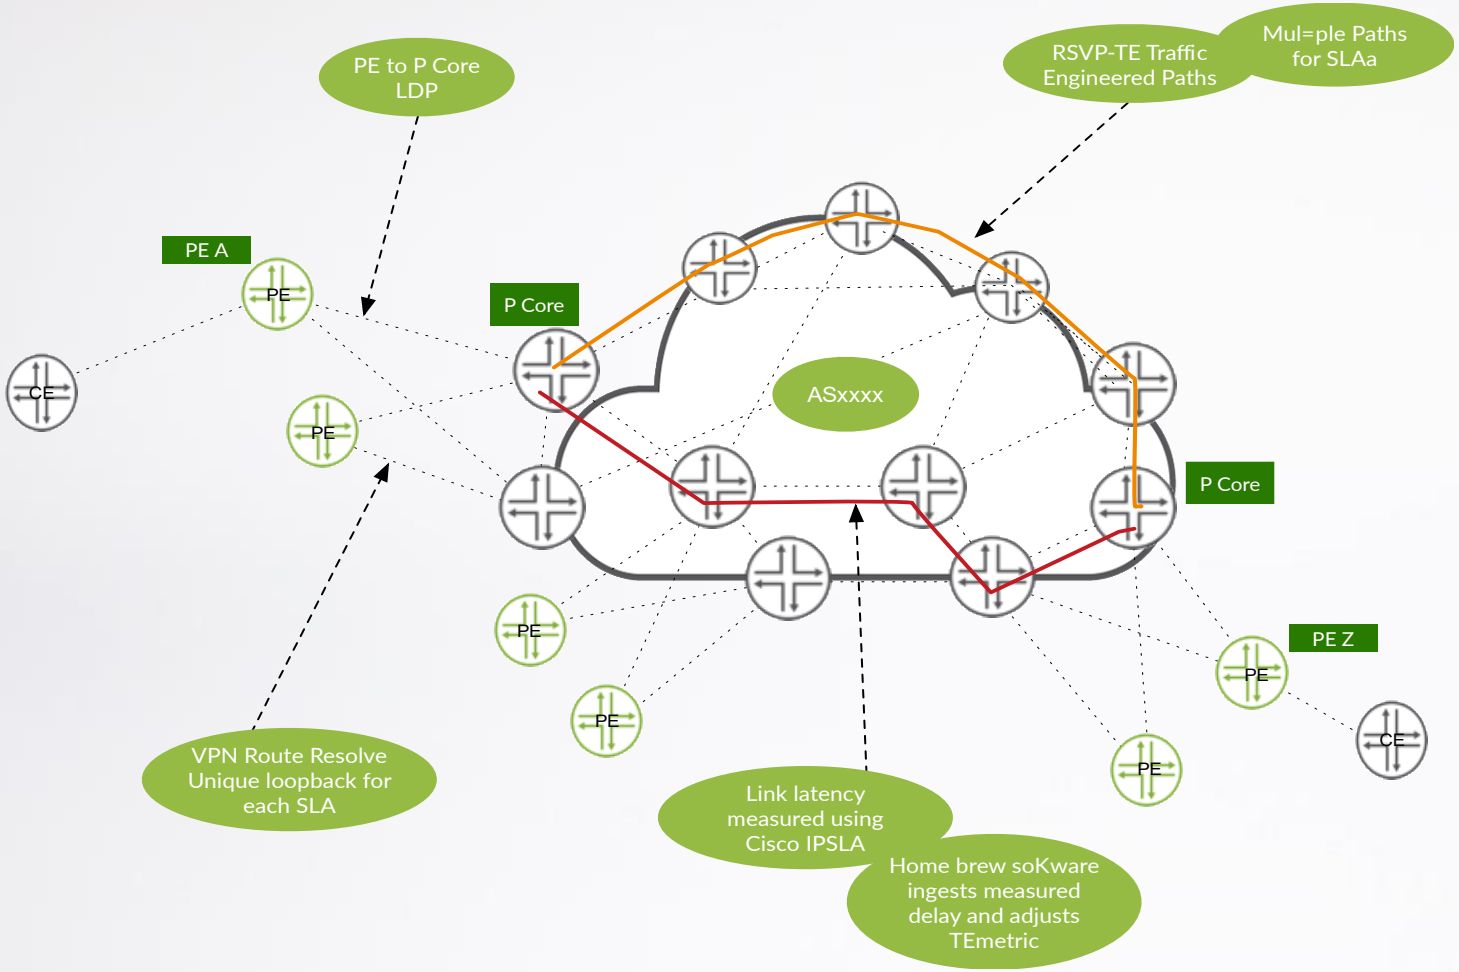

Network Status Quo

Automation done in parts, but not as a holistic solution aligned to business objectives

- Home-grown automation tool to update network metrics; Dependent on SME, difficult to scale
- Multi-vendor networks with no or limited end-to-end automation capabilities
- Multiple provisioning, workflow and operation tools with no integration.
- Limited service-differentiation capabilities
- Demand prediction using excel sheets and personnel insights.

Goal Network

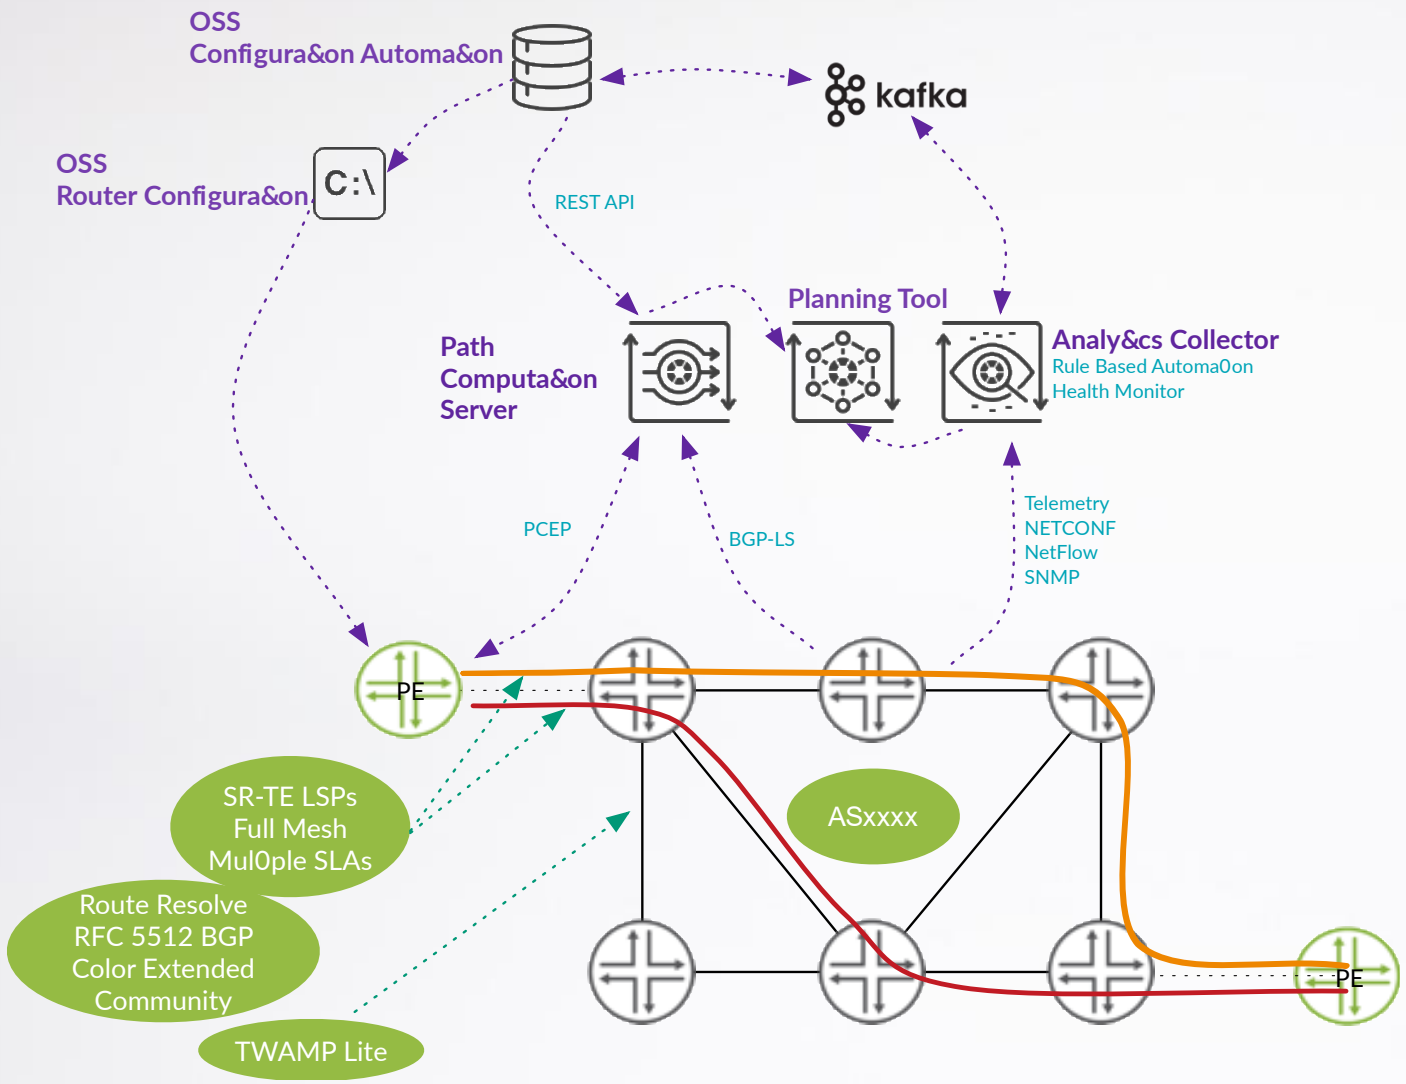

Network Future State

Automation taken up as key initiative that integrates multiple vendors and tools to achieve business objectives.

- Focus on specific use-cases that either improve customer experience or generate new revenue or both.
- Dual Vendor Network , MPLS/SR network based on Industry Standards
- Multiple automation vendors platforms: Integrated Service provisioning, Workflow automation, PCE controller, Analytics engine and Network planner to achieve the specific use-cases
- Summary of the outcomes:
  - Greater transparency and control on the network and services running on top.
  - Automation Framework based on IETF standards using multi vendor platforms
  - Create new services on demand with differentiated customer experience
  - Increase speed to market & Reduce the cost to serve at scale.

Foundational Components

Foundational capabilities for Automation

Industry-standard

Open-source

Traffic engineering

- **ISIS**: Traffic Engineering Database
- **LSPs for traffic engineering (RSVP/SR)**:
  - Traffic-engineered paths (Static & Dynamic)
  - Color-based forwarding
  - Baseline SLA
- **BGP-LS**:
  - Network topology discovery
  - IGP/TE metrics, Link delay, admin-groups
- **PCEP**: Dynamic path computation
- **PCE Controller:** IETF Compliant
  - Path Computation Element (PCE)
  - TE topology management
  - SLA-based routing methods (IGP metric, TE metric, delay, measuredDelay, userCost)
  - Value-added SLA on top of Baseline

Network planning

- **Offline planning tool:**
  - Network capacity planning

Telemetry & Analytics

- **Metric collection & reporting:**
  - Metric reporting: IETF Open Config Telemetry, SNMP, Netconf, NetFlow
  - Metric collection: TWAMP, IPSLA, RPM
- **Collector & Analytics tool:**
  - Multiple ingest from devices (Telemetry, SNMP. Netconf, NetFlow, Syslog)
  - Data analytics using Time-series DB, Embedded rules and machine learning
  - Closed-loop automation: Value-add SLA automation
  - Share data/alerts with 3rd party systems from different vendors for reporting

Workflow automation

- **Workflow Automation Tool:**
  - Automate existing manual/semi-automated operational workflows (MoPs)
  - Act as a glue amongst various tools (Analytics, Configuration, Ticketing, Inventory, OSS)

Network source of truth

- **Multi vendor Configuration management & Service provisioning tool:**
  - Act as a single source of truth for all network configuration
  - Provisioning devices (baseline/golden config) and services (L2/L3 VPN)
  - Managing device configuration files

Additional capabilities

- **Multi Vendor OSS:** Operator front door
  - Router/Link onboarding, Add/remove/changes
  - REST API integration to other systems such as SDN controller to add/update LSPs
  - Network inventory management
- **Message bus (Kafka):**
  - Exchange data/alerts across multiple systems
- **Integration between multiple systems using REST APIs**

Automation Strategy

Operate

- Schedule Maintenance events
- Add Links and routers
- Monitor Health
- Add more high value services!

Operator
Planner
Engineering

Value Use Cases - Add value to the Baseline Foundation

- Planned Event Driven Automation
- Delay Based Routing
- Strict Diversity
- Service Prioritization
- Smart Bandwidth

PCEP and BGP-LS

JTI, gRPC, SNMP, NETCONF, NetFlow

REST API for adds/moves/changes

Playbooks, User Functions, automation Rules

Kafka Publish and Consume automated updates

Foundation Use Case

- Baseline network; known routing state
- Baseline network with known SLAs

Upgrade IGP
- TED

Add TWAMP

Underlay Baseline
- Deploy RSVP/SR
- Migrate services

PCS and Analytics Infrastructure
- Automation tool
  - Path Computation Element (PCE) controller
  - Analytics Engine

Foundation Network – Topology Acquisition

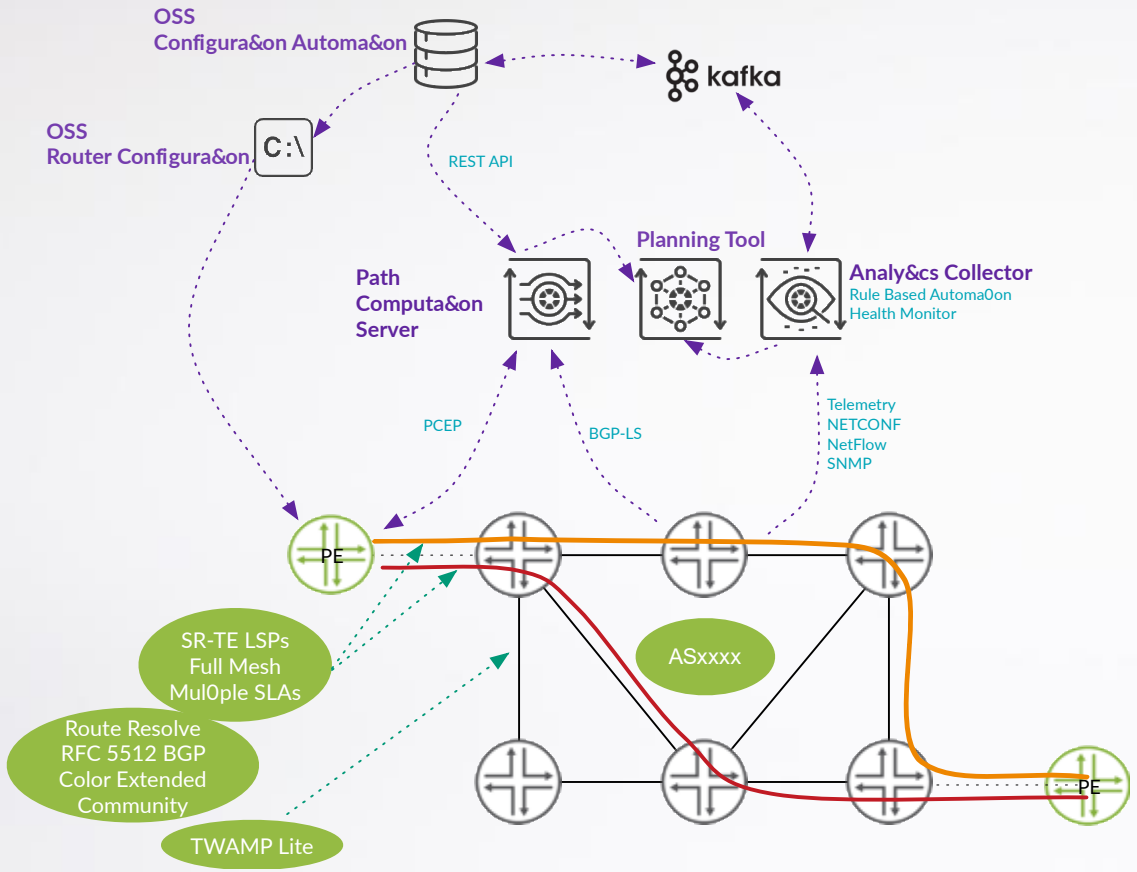

- RSVP/SR-TE LSPs configured on the PCC are delegated to **PCE** using PCEP.
  - Delegated RSVP/SR-TE LSPs are added to the *PCE database*.
- Links and Nodes (i.e., routers) are learned via BGP-LS on **PCE**
  - Links and Nodes are added to the *PCE database.*

BGP-LS: RFC7752 Update Message

- ❑ MP_REACH_NLRI
  - ❑ Link
    - ❑ IPv4 interface address
    - ❑ IPv4 neighbor interface address
- ❑ BGP-LS Type Code 29
  - ❑ admin-groups/affinity-bits
  - ❑ TE metric/IGP Metric
  - ❑ SR Adjacency SIDs
  - ❑ Unidirectional Link Delay

Path Computation Element Protocol (PCEP):
RFC5540/RFC8664

- ❑ PCC PCRpt
  - ❑ SR-ERO sub object
  - ❑ Record Route Object (RRO)
  - ❑ LSP Attributes
    - ❑ admin-groups include-all, include-any, exclude-any
  - ❑ LSP Object flags Delegate bit
- ❑ PCE PCUpd
  - ❑ Explicit Route Object (ERO)

Foundation Network – Use Case Design Principals

- PCE RSVP/SR-TE LSP *Value added* Path Optimization

1. Prune Links based on constraints
2. Select paths that satisfy Advance Properties
3. Routing Method determines which metric to use (see next slide)
    a. Calculate lowest cost path.

Constraints

- admin-groups
    - *include-all, exclude, include-any*
- Bandwidth Sizing
    - *Link Utilization Threshold* constraint
- Explicit Path

Advanced Properties

- Symmetric Pair Group
    - *RSVP/SR-TE LSP A-B and B-A*
- Diversity Group/Diversity Level
    - Two(2) RSVP/SR-TE LSPs require diverse paths
- SliceID
    - *include-all, exclude, include-any*

Routing Method

- routeByDevice
- default
- ISIS/OSPF
- delay
- adminWeight

Foundation Network – Use Case Design Principals

- PCE Link Properties

metrics	IGP metric	routingMethod = ISIS/OSPF
	TEmetric	routingMethod = default
	measuredDelay	routingMethod = delay, RFC 8570 IS-IS ISIS extended IS reachability , Unidirectional Link Delay, TWAMP Lite.
	userCost	routingMethod = adminWeight, user defined metric representing monetary cost.
	delay	routingMethod = delay, configured PCE controller value overriding *measuredDelay*
Advanced Properties	admin-groups	Router configured admin-groups;RFC5305 Sub TLV 3
	Link Utilization Threshold	Specify the threshold value for link utilization when traffic on a link exceeds this value, PCE controller triggers re-routing for label switched paths (LSPs).
	Packet Loss Threshold	Specify the threshold value for link utilization when traffic on a link exceeds this value, PCE controller triggers re-routing for label switched paths (LSPs).
	slices	Decimal value that represents a logical network on a physical network.

Router objects advertised via BGP-LS

PCE value add database configured objects

Use-cases

Use Case – Delay Based Routing

Outcome	Revenue growth

Delay Measurement	Ingest Delay Measurement - Cisco IPSLA, JUNOS RPM, IETF TWAMP - TWAMP direct to PCE controller via BGP-LS - Multiple protocols supported for migration/upgrade
Analytics Engine Calculation	Analytics Engine Playbook - Calculate *delay* with any user defined *overrides.* - *E.g., artificially increase/decrease as required based on design* - Control if Link pruned, measured, maintenance.
Update Link Delay	Analytics Engine Playbook - REST API PATCH PCE controller Link delay - REST API PATCH PCE controller Link in/out service.
PCE Optimization	PCE Optimization - PCE optimizes RSVP/SR-TE LSPs - RSVP/SR-TE LSPs have guaranteed lowest delay SLA.

Use Case – Delay Based Routing

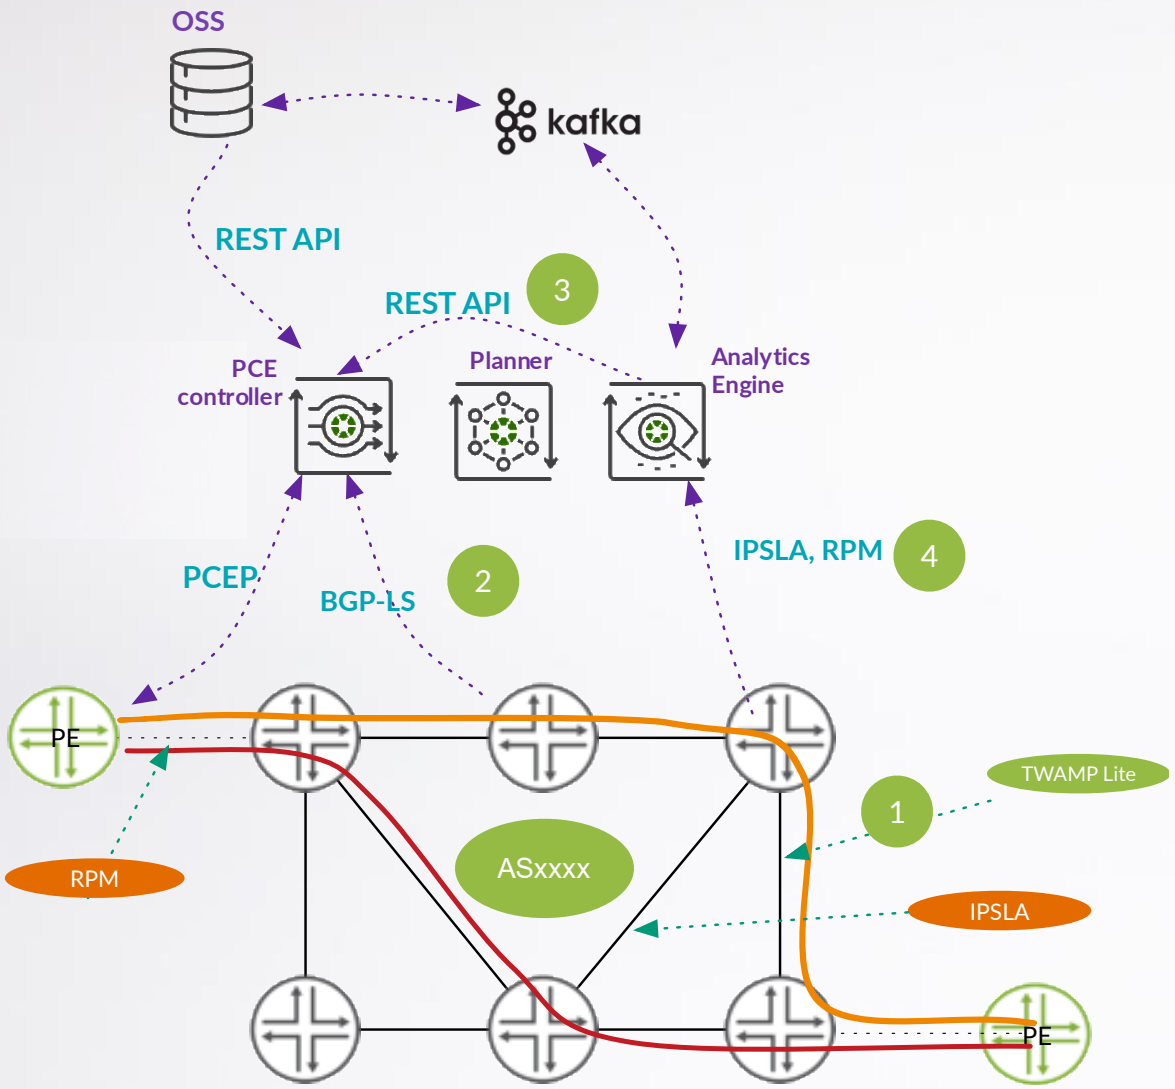

1. Link Delay Measurement Protocol enabled on all network links
2. TWAMP is sent to PCE controller in BGP-LS

   Analytics engine needs measure delay to calculate a usable delay
3. Analytics Engine REST API GET for TWAMP Delay
4. Analytics Engine playbooks ingest legacy protocols:
   - JUNOS RPM & Cisco IPSLA

Use Case – Delay Based Routing

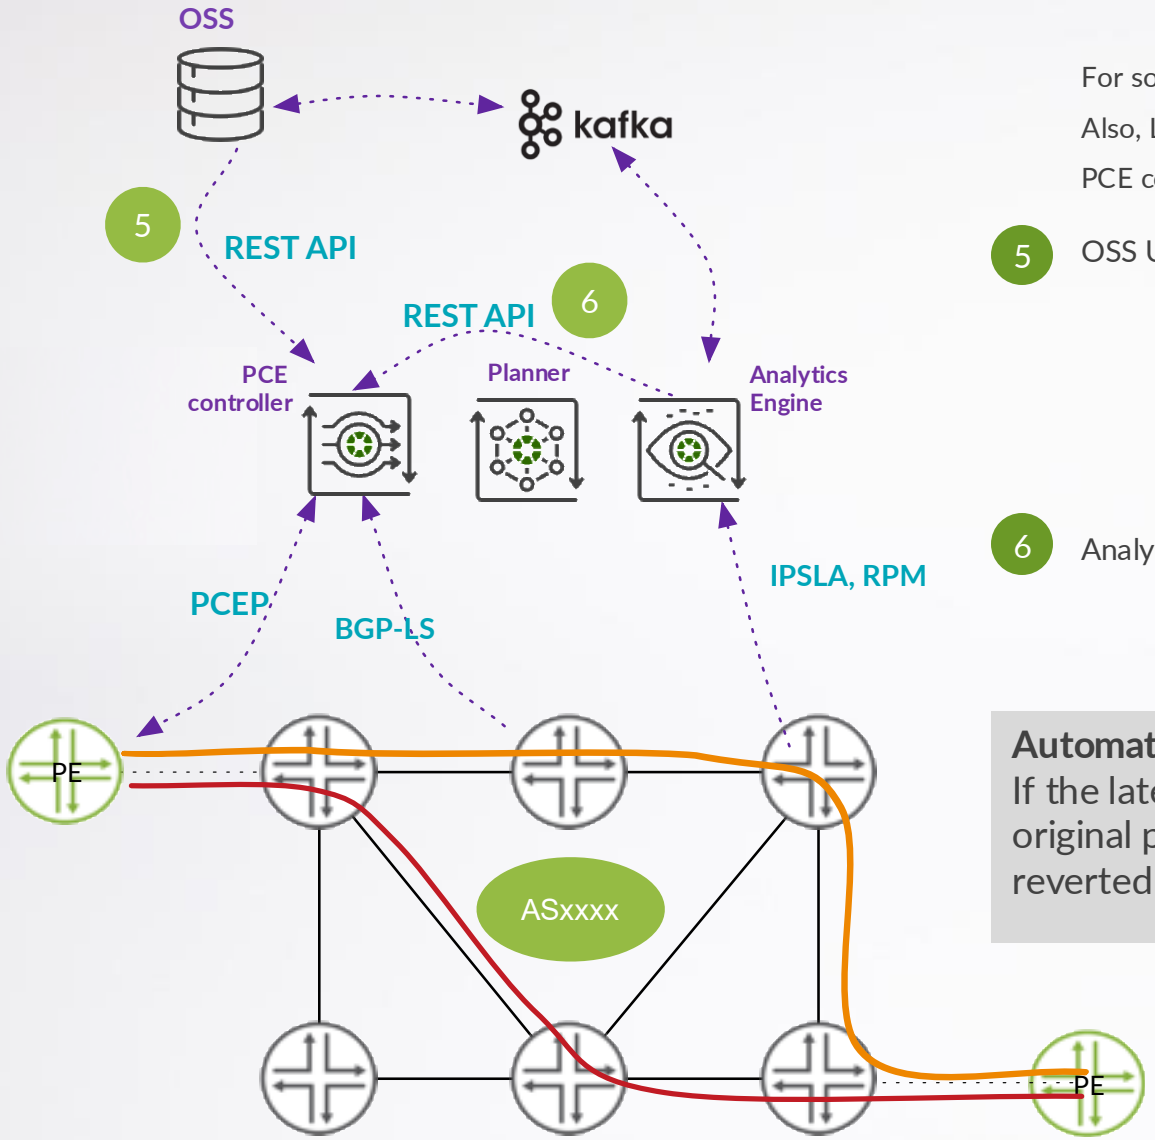

For some Links, the measured delay needs to be overridden

Also, Links need to be taken out of service for routingMethod = delay, but not other routingMethods

PCE controller Link Custom Attributes are used as Operator defined variables

5 OSS Updates REST API Updates PCE controller Link Custom Attributes

- delaymax
- delaymin
- measurementprotocol
- managed
- maintenance
- Override

6 Analytics Engine REST API GET

- TWAMP Delay
- Link Custom Attributes

Automated reversion:
If the latency improves on the original path, the traffic will be reverted to that path.

PCE controller Database Link JSON Example

```
"endA": {

    "userProperties": {
        "delaymax": "500",
        "delaymin": "65",
        "delayexpect": "7"5,
        "maintenance": "false",
        "managed": "true",
        "measurement": "twamp",
        "override": "900"
    },
    "utilization": {

        "measuredDelay": 57.568,

    }
},

"id": "L192.168.127.5_10.0.57.0_192.168.127.7_10.0.57.1",
"linkIndex": 27,

"slices": [
    {
        "sliceId": 4294967295
    }
],
"topoObjectType": "link",
"topologyIndex": 1
```

Use Case – Strict Diversity Guarantee

Outcome	Customer experience

RSVP/SR-TE LSP	Custom RSVP/SR-TE LSPs for high value customers - Head end diversity. - Link diversity - Optimized based on delay.
BGP Color Community	BGP Color Community - Customer specific BGP Color community - Measures prefix bps - Compares prefix bps to % of Inter-AS link - Triggers alarm if in excess.
OSS update RSVP/SR-TE LSP	Analytics Engine Playbook - Publish Kafka Notification OSS - Consume Kafka Notification
Monitor	Analytics Engine - Monitor Path Diversity of RSVP/SR-TE LSPs - Alarm is Diversity is not satisfied.

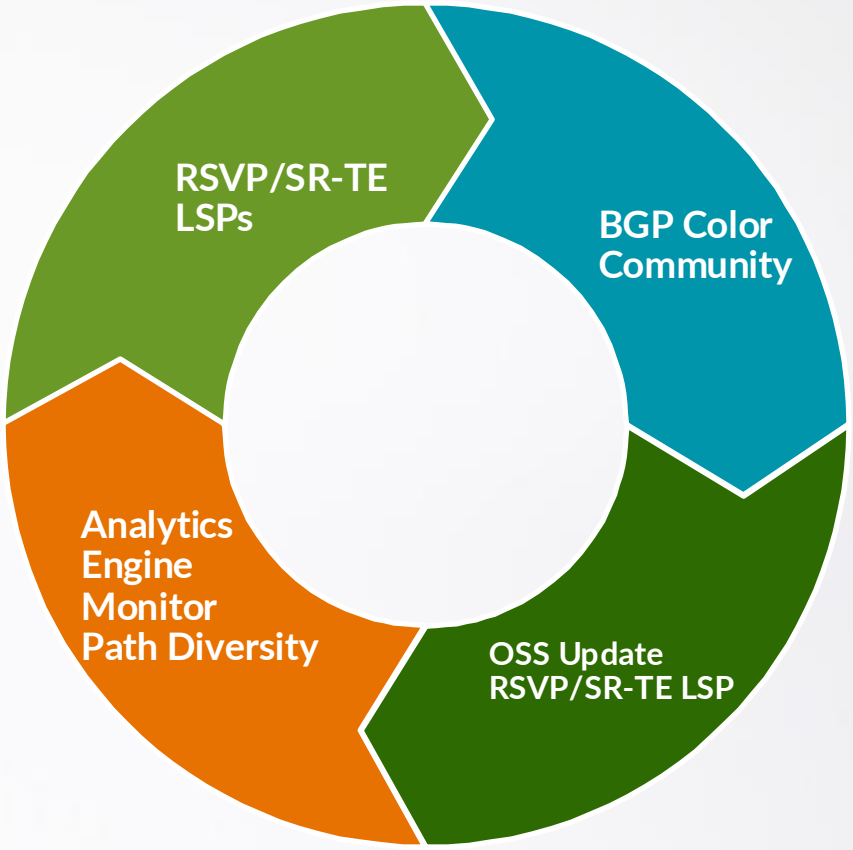

Use Case – Strict Diversity Guarantee

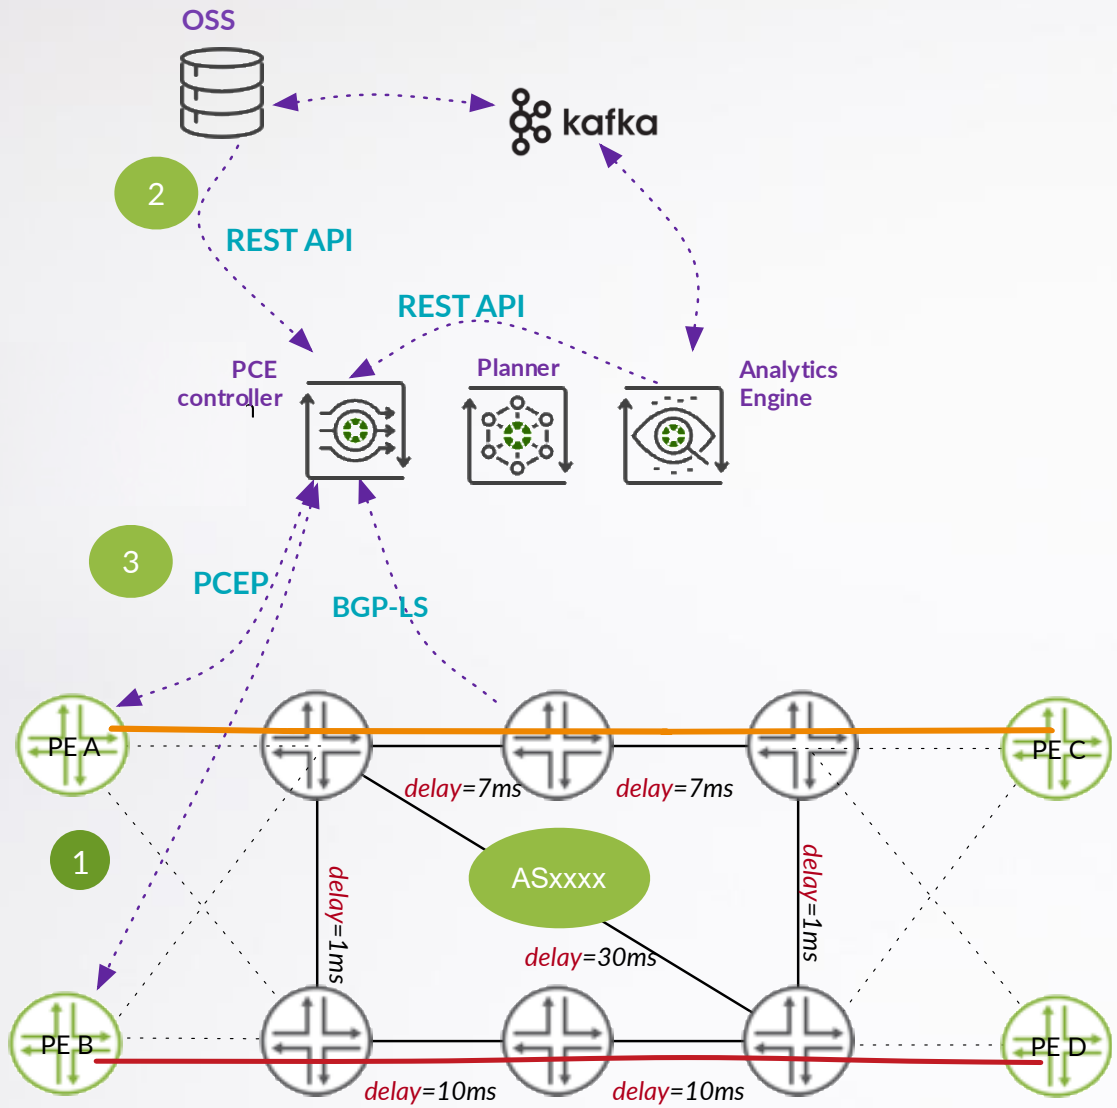

1. OSS configures two (2) RSVP/SR-TE LSP in PE A and PE B
   - Customer-specific BGP color community
   - PCC delegated to PCE controller

2. OSS REST API PATCH RSVP/SR-TE LSP
   - routingMethod = delay
   - PCE controller Diversity parameters
     - diversityGroup
     - diversityLevel

```
REST API PATCH /traffic-engineering/api/topology/v2/1/te-lsps/<index>

{
    "plannedProperties": {
        "design": {
            "routingMethod": “delay”,
            "adminGroups": {
                "attributeIncludeAll": 4096,
                "attributeExclude": 3224371456
            }
            "diversityGroup": "PEA-PEB-MASTERGIP9999999",
            "diversityLevel": "site"
        }
    },
    "pathType": "primary",
    "lspIndex": <index>,
    "provisioningType": "SR"
}
```

3. PCE controller optimizes network using PCEP
   - diversityGroup RSVP/SR-TE LSPs are optimized together.

Use Case – Service Prioritization

Outcome	Cost Saving

Ingest NetFlow	Analytics Engine NetFlow from Inter-AS Link - destinationIPv4Prefix(Length) - ingressInterface - octetDeltaCount
Analytics Engine Calculation	Analytics Engine Playbook - Resolve interface name - Measures prefix bps - Compares prefix bps to % of Inter-AS link - Triggers alarm if in excess.
Publish Result	Analytics Engine Playbook - Publish Kafka Notification OSS - Consume Kafka Notification
Update Router	OSS - Update router policy - add/remove $LOWCOST *color* to prefix

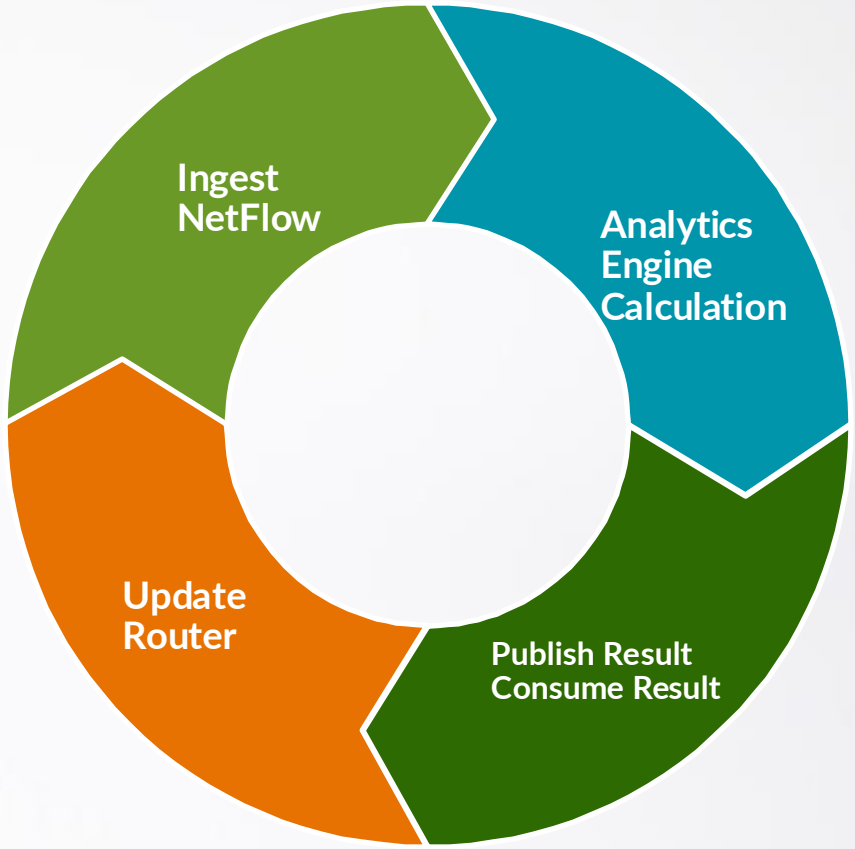

Use Case – Service Prioritization

- By default traffic uses the PREMIUM path
- Regrade traffic to use the $LOWCOST path

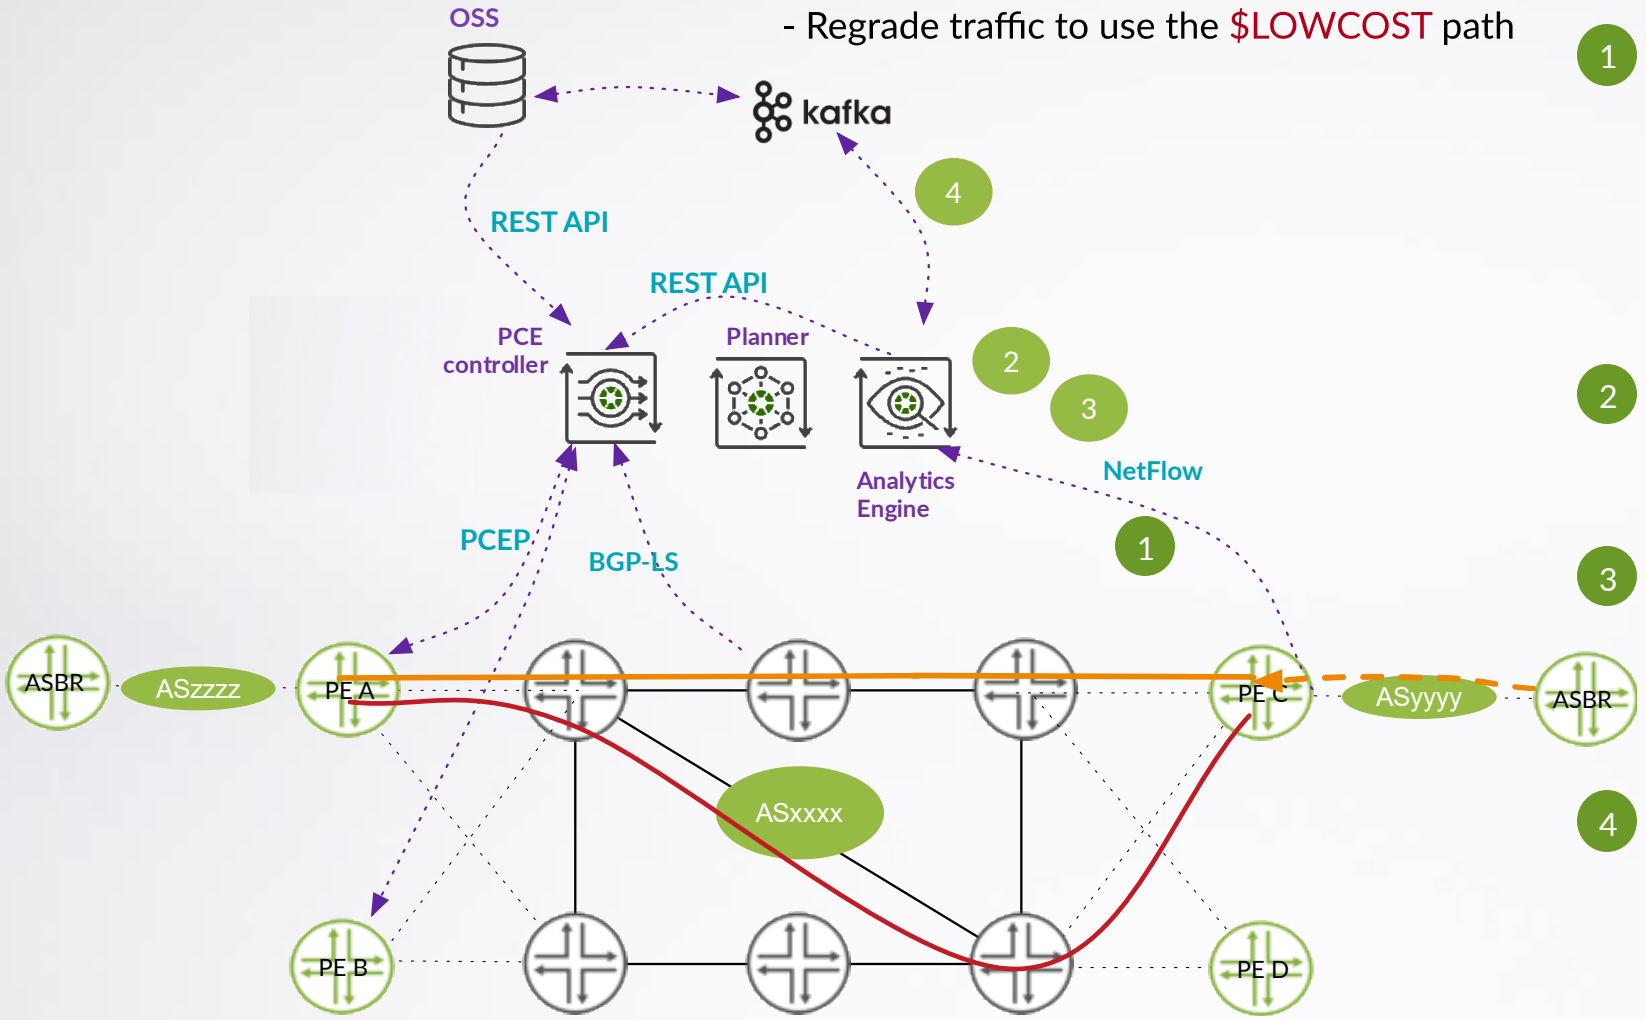

1. Analytics Engine ingests Netflow from ASBR
   - bgpNextHopIPv4Address
   - bgpDestinationAsNumber
   - destinationIPv4Address
   - ingressInterface
   - octetDeltaCount
   - sourceIPv4Address
   - destinationIPv4PrefixLength
2. Analytics Engine collects:
   - ASBR interface Bandwidth
   - Operator configured threshold % ASBR interface bandwidth
3. Analytics Engine calculates
   - Total traffic to destination Autonomous system (AS)
   - If traffic exceeds operator configured threshold %
     - then Kafka publish
4. Analytics Engine Trigger
   - If traffic exceeds operator configured threshold &
     - then Kafka publish alarm

Use Case – Smart Bandwidth

Outcome	Revenue growth

Ingest NetFlow	Ingest RSVP/ SR-TE LSP on congested Links - REST API GET LSPs on or thru Link Ingest NetFlow from Inter-AS Link - bgpNextHopIPv4address
Analytics Engine Calculation	Analytics Engine Playbook - Correlate RSVP/SR-TE LSP and NetFlow records - RSVP/SR-TE LSP Ingress and Egress PE with NetFlow source and bgpNexyHopIPv4address Calculate
Publish Result	Analytics Engine Playbook - Publish Kafka Notification OSS - Consume Kafka Notification
Update Router	OSS - Update router policy - add/remove BURST *color* to prefix

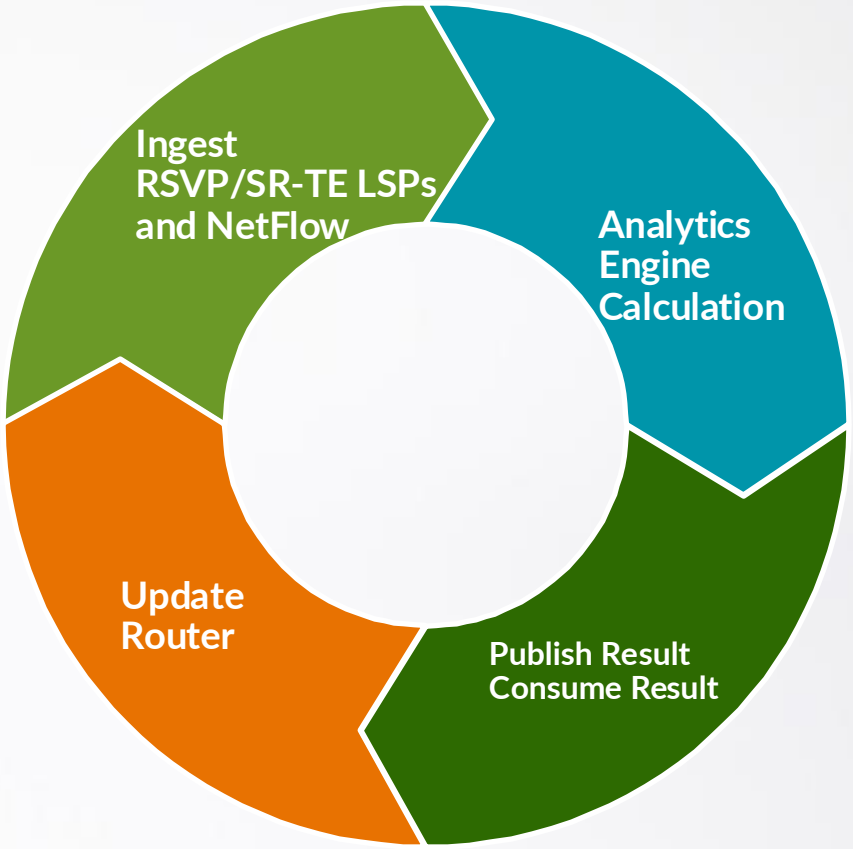

Use Case – Smart Bandwidth

- Internet traffic has four(4) Service Offerings routed over three(3) RSVP/SR-TE LSPs

1. **PREMIUM**
    a. Always routes over the most optimal low latency path
2. **STANDARD+**
    a. Always routes over the most optimal IGP path
3. **STANDARD**
    a. Routes over the optimal IGP path
    b. During congestion prefixes are *downgraded to BURST* based on measured traffic volume and business rules.
4. **BURST**
    a. Routes over the optimal IGP path
    b. During overall network congestion routes over links *not configured with Link Utilization Threshold* (and therefore $cheaper longer latency links).
    c. Routes over links that may undergo congestion.

Use Case – Smart Bandwidth

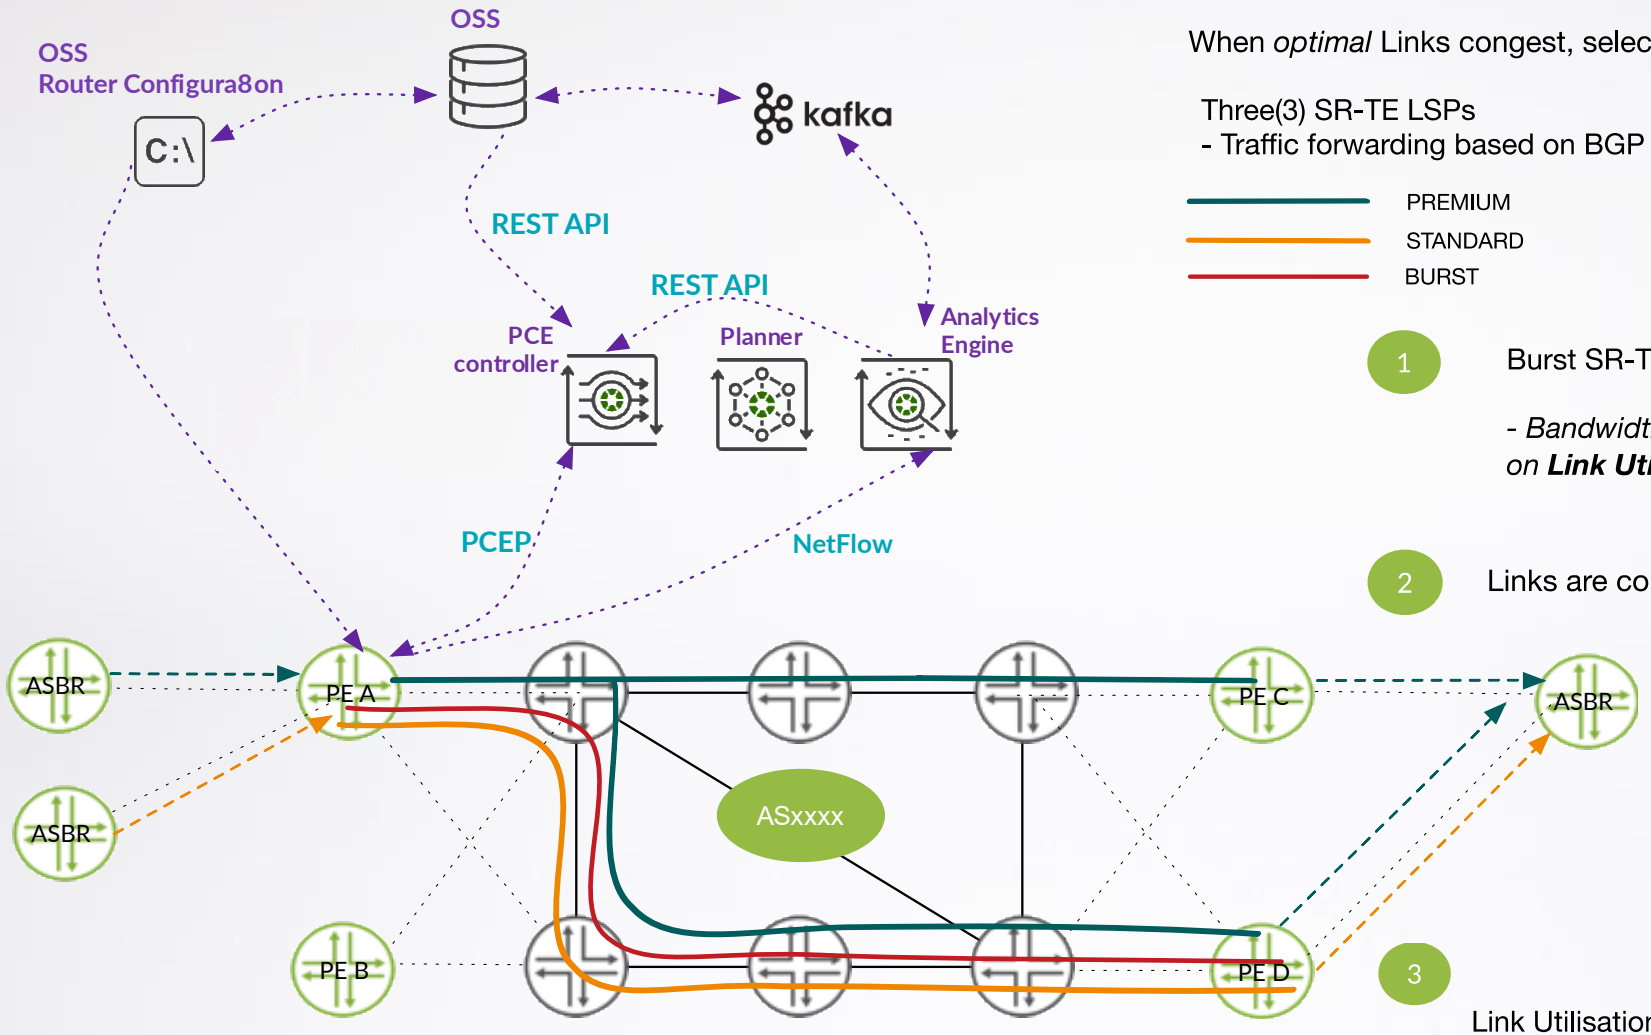

During times on no congestion all traffic can use the most *optimal* Links

When *optimal* Links congest, selectively move traffic at ASBR Inter-Links

Three(3) SR-TE LSPs
- Traffic forwarding based on BGP Community *Color*

PREMIUM
STANDARD
BURST

1 Burst SR-TE LSP is configured with *Bandwidth Sizing*

*- Bandwidth Sizing will optimize the SR-TE LSP based on **Link Utilisation Threshold***

2 Links are configured with a ***Link Utilisation Threshold***

3 BURST

Link Utilisation Threshold is reached

BURST SR-TE LSP is optimised to avoid link at 90%

Use Case – Smart Bandwidth

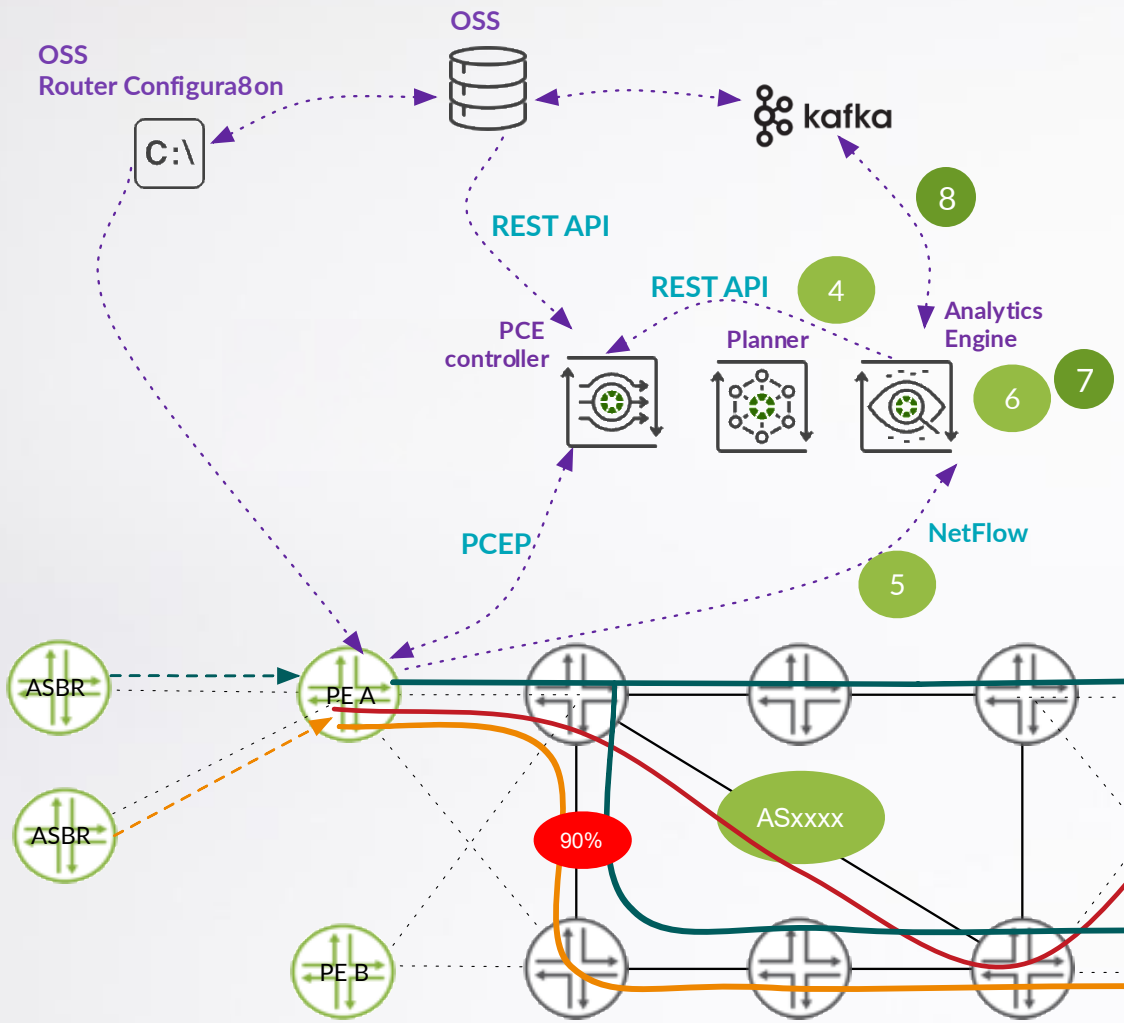

4 Analytics Engine REST API GET all RSVP/SR-TE LSPs that are **on or thru the congested link**

- This provides the Ingress and Egress PE that are congesting the traffic

5

Analytics Engine ingests Netflow from ASBR

- bgpNextHopIPv4Address
- bgpSourceAsNumber
- destinationIPv4Address
- ingressInterface
- octetDeltaCount
- sourceIPv4Address
- destinationIPv4PrefixLength

6

Analytics Engine correlates the ingested data

- RSVP/SR-TE LSP Ingress and Egress VPN PE
- Netflow PE source and bgpNextHopIPv4Address

7 Analytics Engine calculates:

- Total traffic to destination AS
- If traffic exceeds operator configured threshold %
  - then Kafka publish

8 Analytics Engine trigger

- If traffic exceeds operator configured threshold %
  - then Kafka publish alarm
  - Downgrade prefixes to *BURST*

THANK YOU

JUNIPER NETWORKS | Driven by Experience™