SDN-based Automated Peering Optimization Challenges and Solutions

Reda Laichi: reda.laichi@nokia.com

Hamid Ould-Brahim: hamid.ould-brahim@nokia.com

NANOG 74, Vancouver, October 2018

Agenda

- **Defining the Problem**
- SDN and Automation
- Use Cases
- Summary and References

Internet traffic reality

2000 → 2018

From web browsing to social media, video streaming and online gaming

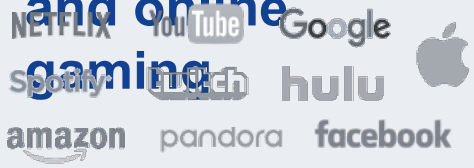

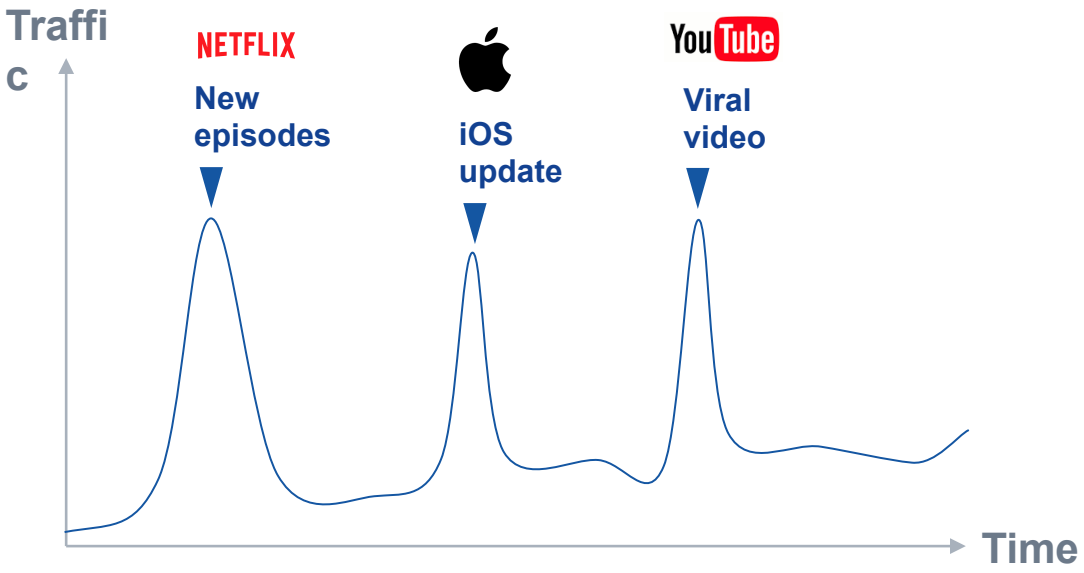

Internet traffic became much more versatile, dynamic and unpredictable

Problem Space
Objectives & applicability...

What do we want to address?

- Better bandwidth management, automatic congestion resolution, better traffic symmetry
- Better SLAs, & application Performance (Latency, Packet Loss)
- Reducing transit peering cost, addressing OPEX & CAPEX

What traffic are we interested in?

What part of the network?

Existing Limitations...

Routing mechanism - Border Gateway Protocol (BGP)

- ✖ Unaware of link capacity & real-time utilization
  - ↳ Packet loss and congestion
- ✖ No real-time path performance indication
  - ↳ High latency
- ✖ No end-to-end performance indication
  - ↳ Sub-optimum overall performance

Visibility

- ✖ Lack or limited visibility
  - ↳ Limited traffic engineering and steering

Organization and tools

- ✖ Multiple teams involved
  (network operations, peer engineers, OAM probing, EMS alarms…)
  - ↳ Slow communication
- ✖ Complex, manual processes
  - ↳ Error-prone configurations
- ✖ Reactive model
  - ↳ Inadequate for sudden real-time event changes

Agenda

- Defining the Problem
- **SDN and Automation**
- Use Cases
- Summary and References

What is being done today?
Webscale players leading the way...

Webscale players

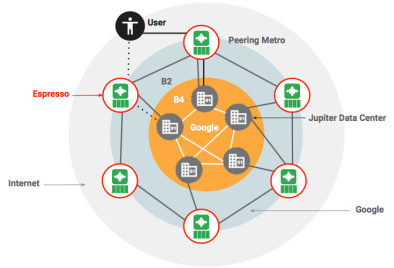

Source: delivery.acm.org, 2017, Taking the Edge off with Espresso: Scale, Reliability and Programmability for Global Internet Peering

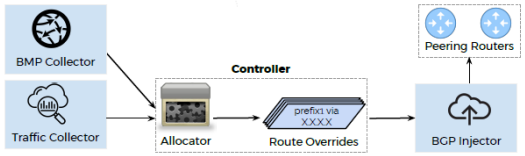

Source: research.fb.com, 2017, Engineering Egress with Edge Fabric Steering Oceans of Content to the World

Gaming companies

of Legends every day ensures that all the internet traffic from Netflix watchers and Spotify listeners doesn't interrupt their game.

The internet was not built for gaming

Source: https://qz.com/790208/how-the-company-behind-league-of-legends-rebuilt-its-own-internet-backbone-so-that-its-faster-for-gamers/

Existing route controllers/appliances

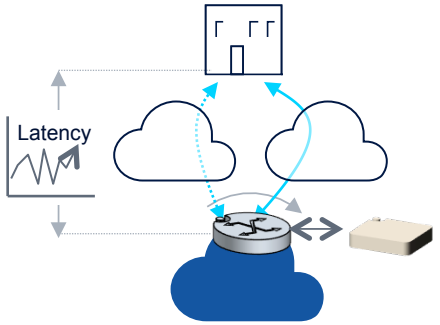

The Approach: Ten Thousand Foot View...
Closed-Loop Automation for Peering

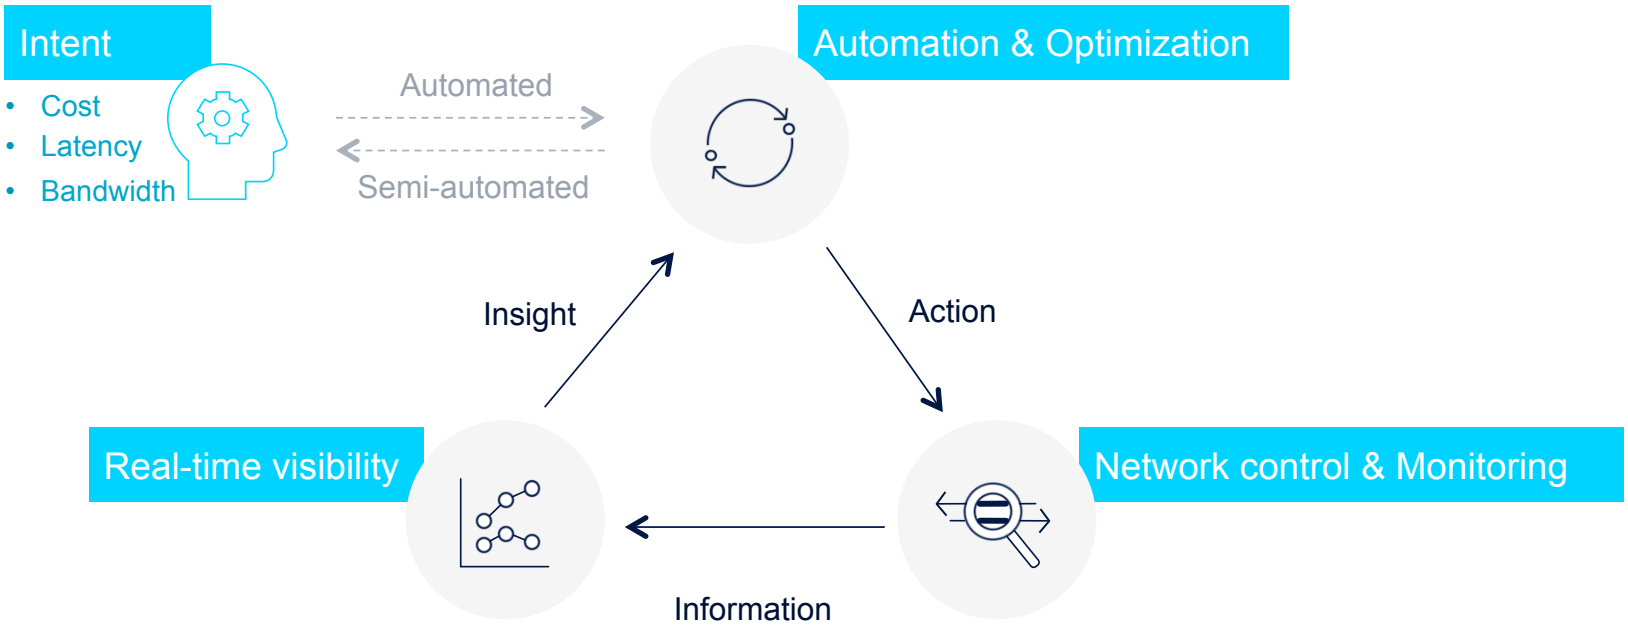

Complexity of the “automation” problem?

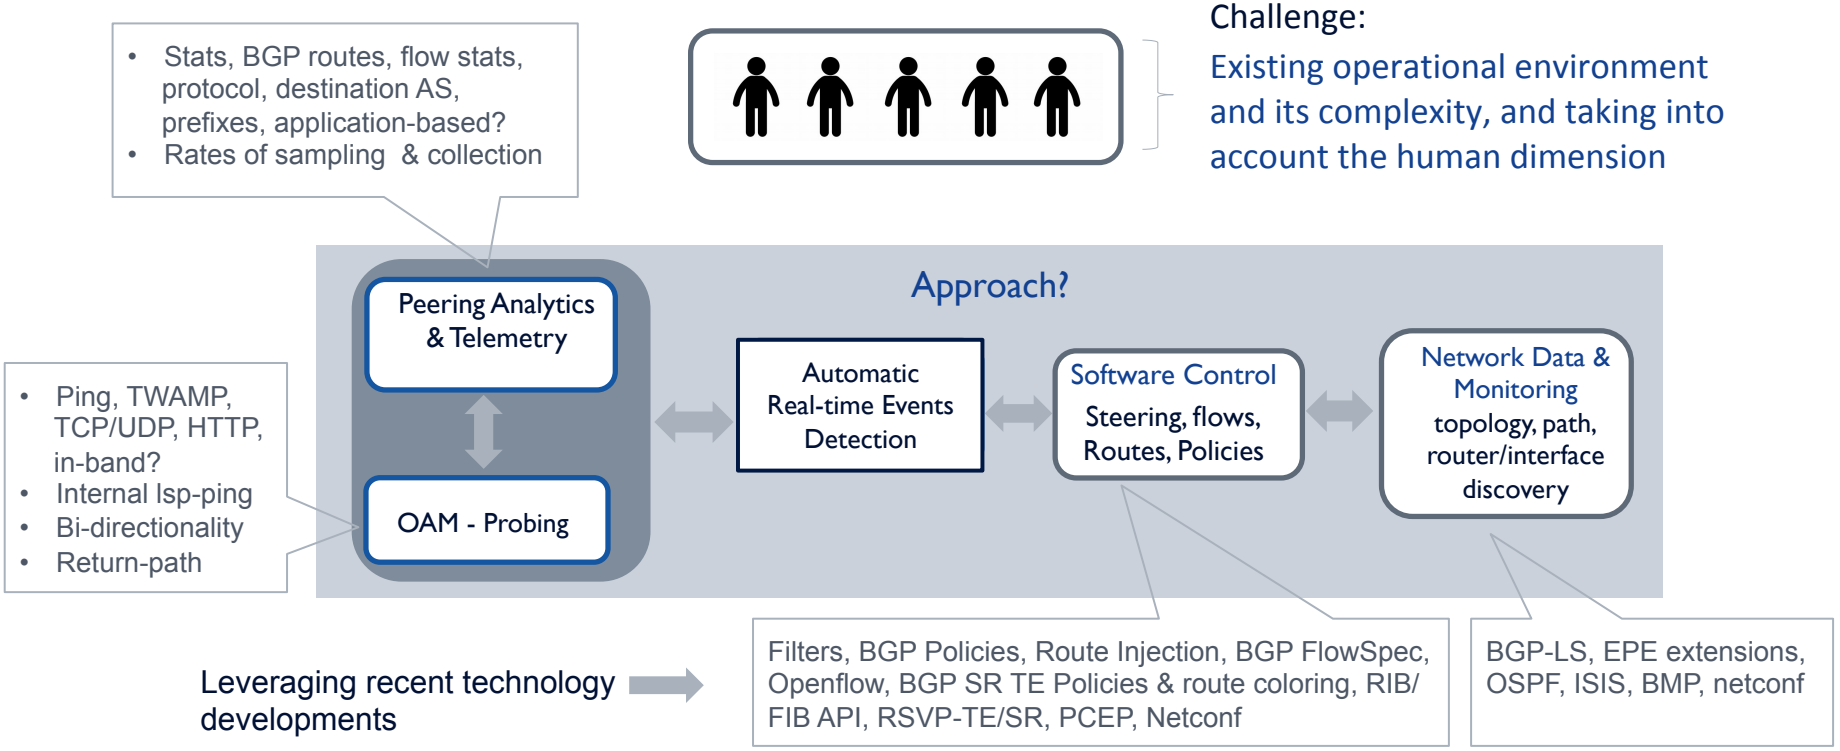

The other challenge is in system integration, open interfaces and multi-vendor...at scale

Agenda

- Defining the Problem
- SDN and Automation
- **Use cases**
- Summary and References

Use Case 1: Local Peer Engineering
Congestion example for outbound traffic

Problem:

100s of peering partners, which link is best and is **least congested**?

Approach:

1. Automatic link congestion detection

   Real-time stat/data collection and correlation (gRPC, Netconf/SNMP) Interface & Flow Stats (IPFix)

2. Determine optimal alternate peer based on bandwidth availability

   Match selected traffic flows at local ASBR based on IP flow/route/AS/5 tuple, or application match (requires transaction-based steering)

   Redirect to next hop using FlowSpec, openflow, filters/Netconf, RIB/FIB API, BGP policies)

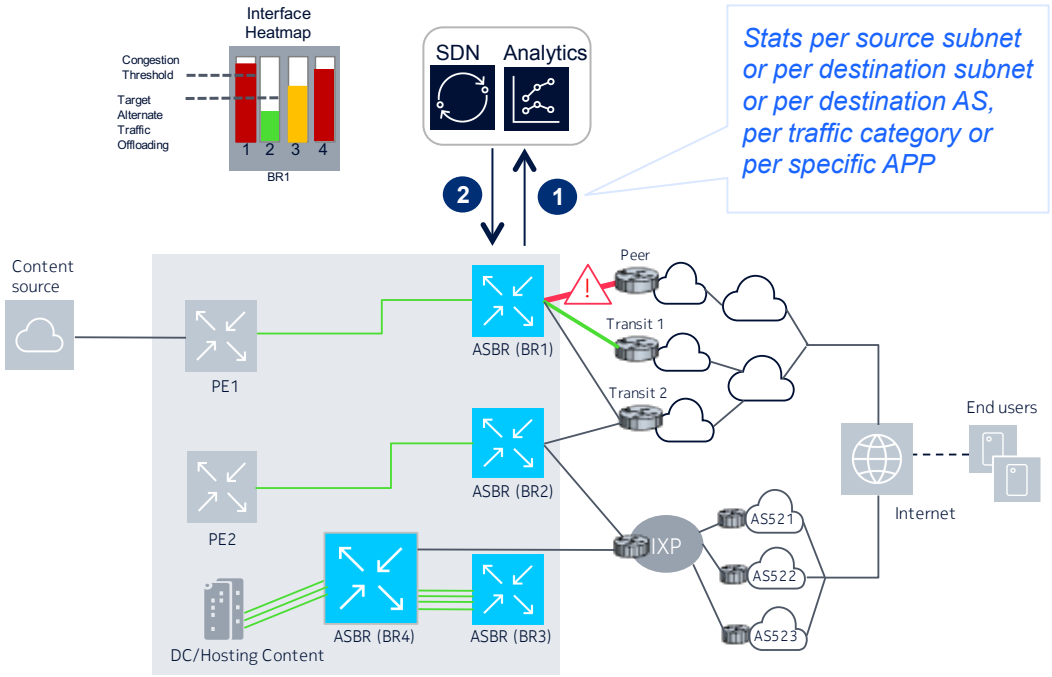

Use Case 2: Egress peer engineering

Problem:

100s of peering partners, which **peer and egress link** is best and is **least congested**?

Approach:

1. Automatic link congestion detection

   Real-time stat/data collection and correlation (gRPC, Netconf/SNMP) Interface & Flow Stats (IPFix)

2. Determine optimal alternate ASBR & alternate peer based on topology and bandwidth availability.

   Auto-create/use existing tunnel/path to alternate ASBR (can be PCE initiated)

3. Steer selected IP flows at the edge of the network across the newly/existing tunnel and encode the egress peer link label/segment ID.

   Note: Colored BGP route to BGP-SR TE Policy tunnel can be used for steering

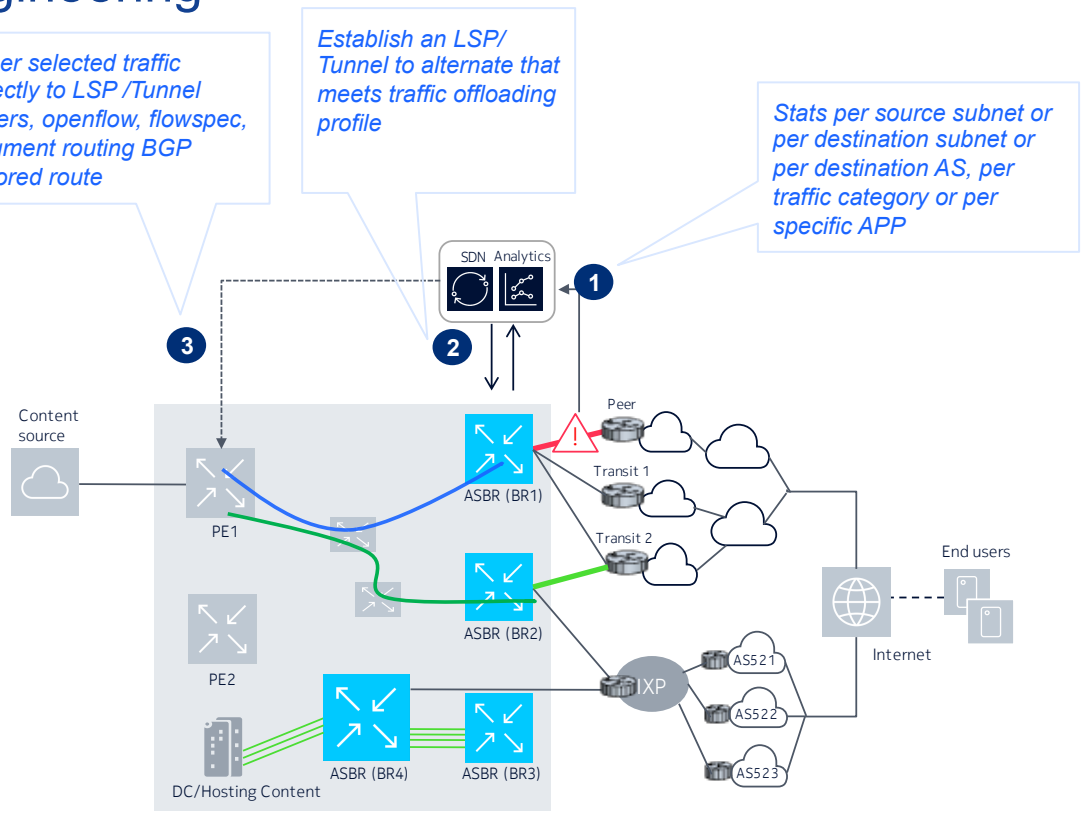

Use Case 3: Locale and EPE with Latency-based Steering
Performance-based Optimization End-to-End

Problem:

100s of peering partners, BGP best path on default peer causing high latency for top traffic or selected application/branch destination traffic.

Approach:

1. Real-time probing from each source PE and each alternate ASBRs of discovered top/VIP destinations
2. Determine optimal alternate ASBR based on end-to-end latency/performance data using path computation based on abstract topology Auto-create the LSP/ tunnel to alternate ASBR
3. Auto-steer selected traffic to the LSP/ tunnel to alternate ASBR/ASBR+Egress peer link

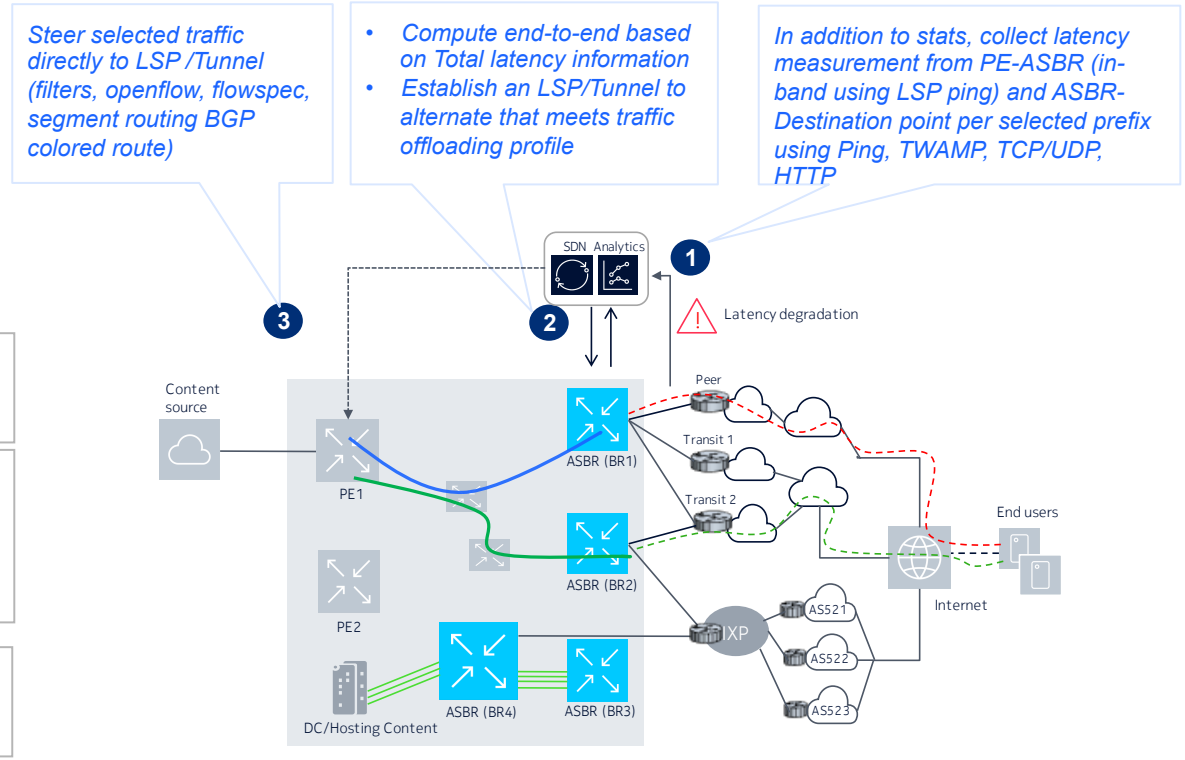

Use Case 4: Controlling Inbound Traffic
The example of Auto-pilot mode

- **Problem:** The links from transit providers are congested due to high incoming traffic. Automate traffic shifting per BGP communities to alternate TP.
- **Solution/Approach:** Integrates and automate analytics data showing top BGP community traffic.
- Monitors bandwidth availability down to Customer devices and performs BGP route extraction and analysis (using BMP) – Extract topology information, LSPs, telemetry
- Using BGP Policies or route injection, selectively shift customer incoming traffic to alternate peer links or to even to an alternate transit provider
- Optionally steer selected customer prefixes to specific LSPs

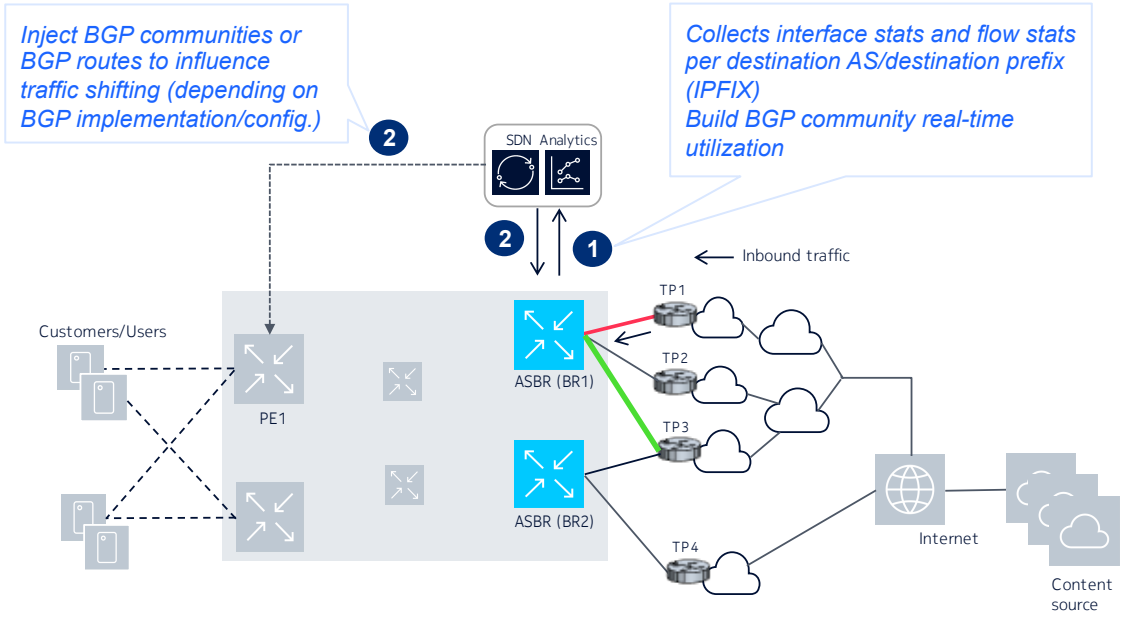

- Automation hierarchical policies: Automate within the same router, other routers same site, completely different region (slice)
- Need to factor in the BGP convergence time

Use Case 4: Controlling Inbound Traffic
Example of Operator Triggered Change

Problem: The links from transit providers are congested due to high incoming traffic.

Peering Policy Intent: Shift top traffic per destination AS/ prefix to South/East region and pick and choose which prefix to steer to alternate tunnel once traffic is shifted.

Solution/Approach:

- Slice the peering/network in terms of automation/ optimization zones.
- Performs BGP route extraction and analysis (using BMP) – Extract topology information (BGP-LS, IGP), discovers LSPs.
- Using BGP Policies or route injection, selectively shift customer incoming traffic to alternate transit provider by injecting the change to the PEs.
- Optionally steer selected customer prefixes to specific LSPs

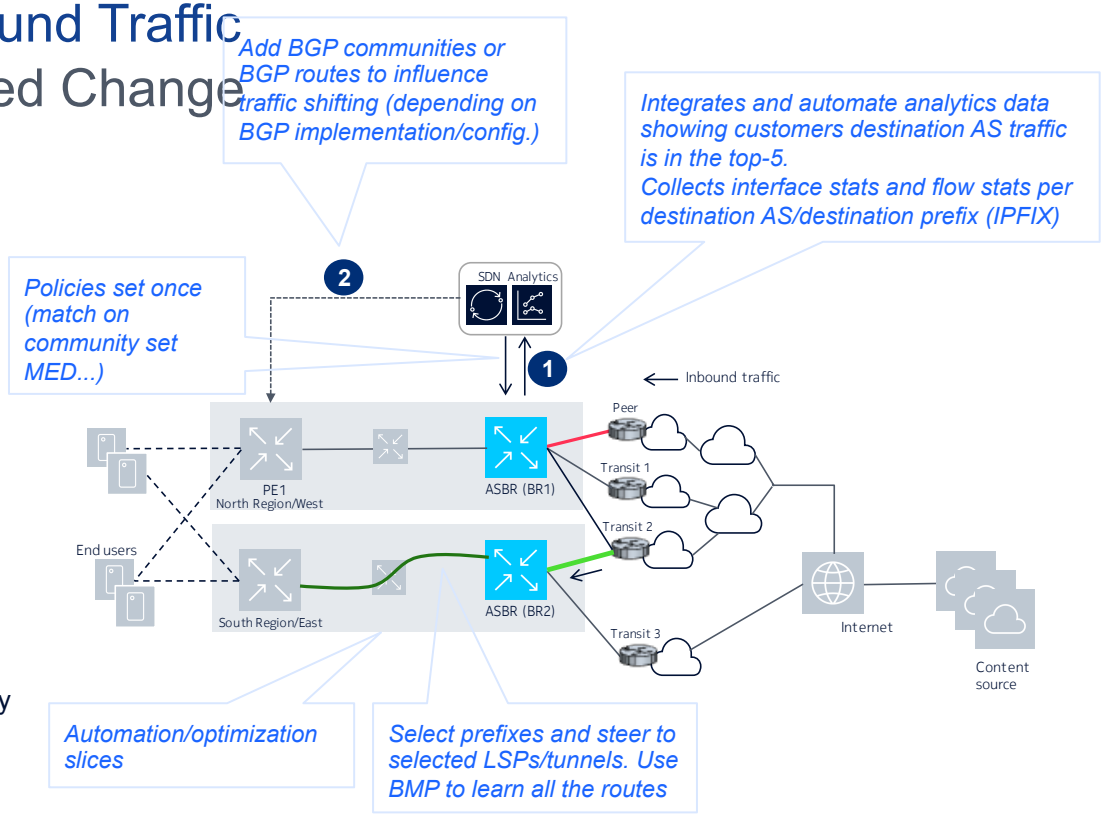

Use Case 4: Controlling Inbound Traffic – Behind The Scene

Example of Operator Triggered Change

From Transit provider 2/Slice6011 to transit provider 2/ Slice6012

```
policy-statement eBGP {
    ...
        from {
            protocol bgp;
            as-path CUSTOMER1-ASN;
        }
        then {
            community add TP2_SLICE6011_SUPPRESSED;
            community add TP1_SUPPRESSED;
            community add TP2_SLICE6012_USED;
            community add TP2_USED;
            accept;
        }
community TP1_SLICE6011_SUPPRESSED members 1000:12227;
community TP1_SUPPRESSED members [ 1000:12201 1000:12222 ];
community TP2_SLICE6012_USED members 1000:12176;
community TP2_USED members [ 1000:12151 1000:12171 ];
```

Automatic Injection

Deployed on Border Router:

```
community ASPATH_PREPEND_1_TP1_NORTH members 1000:12201;
term 40 {
        from {
            protocol bgp;
            community ASPATH_PREPEND_1_TP1_SOUTH;
        }
        then {
            as-path-prepend 1000;
```

Use Case 5: The Case of Destination AS Optimization

Problem:

Hosting provider with congestion uplinks and available alternate network (metro) but limited bandwidth used congestion with no BGP reconfiguration

Approach:

1. Real-time traffic utilization monitoring per destination AS.
2. Determine best next hop and steer selected AS to that next hop.

Requires next hop tracking capabilities

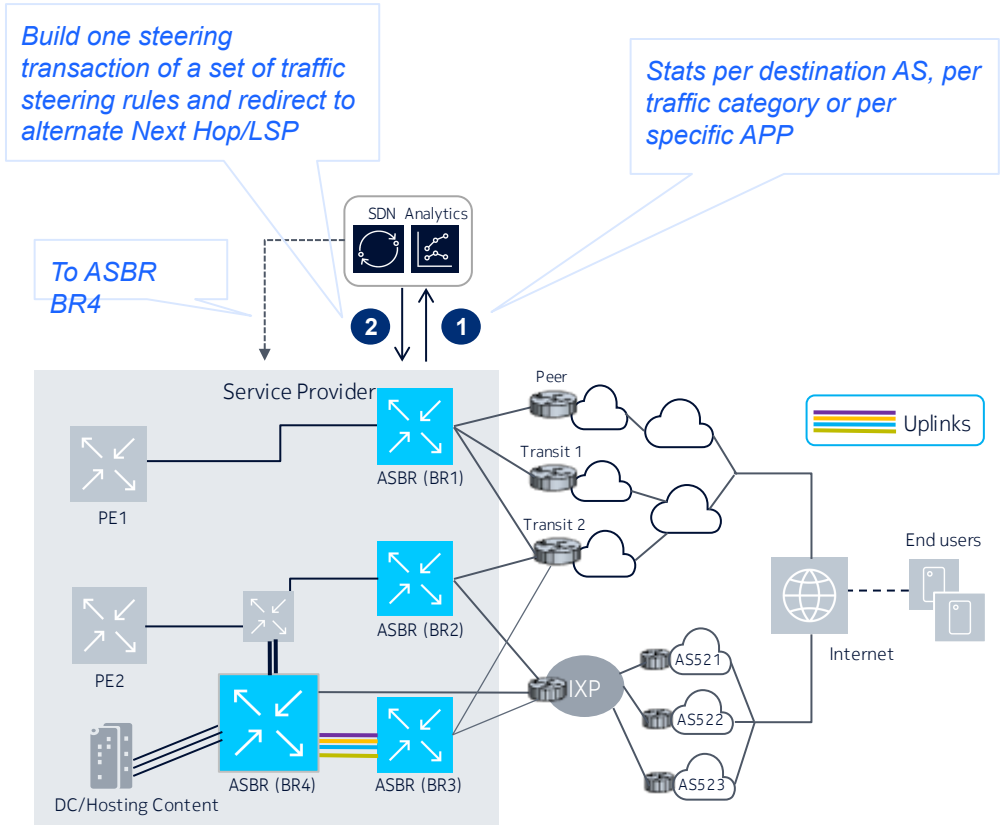

Summary and Other Considerations

- Tracking changes (uniform data retention consideration)
- Culture and operational change considerations
- Think about “Revertive” actions (Undo button)!
- Peering simulation & predictive analysis on multi-dimensional data-set
- Traffic engineering/steering on multi-dimensional data-set
- Network control changes at large scale and real-time or near real-time telemetry at high frequency.
- Keep what BGP is best at and complement existing functionality

References

- BGP-LS
  - https://tools.ietf.org/html/rfc7752
- PCEP
  - https://tools.ietf.org/html/rfc5440
- Segment Routing Architecture
  - https://tools.ietf.org/html/rfc8402
- Insight Driven Automated Network white paper
  - https://pages.nokia.com/12018.Insight-driven-automated-networking.html

NOKIA

Questions?

Reda Laichi: reda.laichi@nokia.com
Hamid Ould-Brahim: hamid.ould-brahim@nokia.com