

Internet Peering Concepts and Trends

Sai Nyan Lynn Swe
CCIE # 38501 (R&S , SP and DC)
OPTIMITY Co Ltd

- The Internet for the Future
- Peering Intro and Internet Trends
- Peering Network Design
- Peering Network Telemetry
- Peering Security
- Future

The Internet for the Future

New Normals

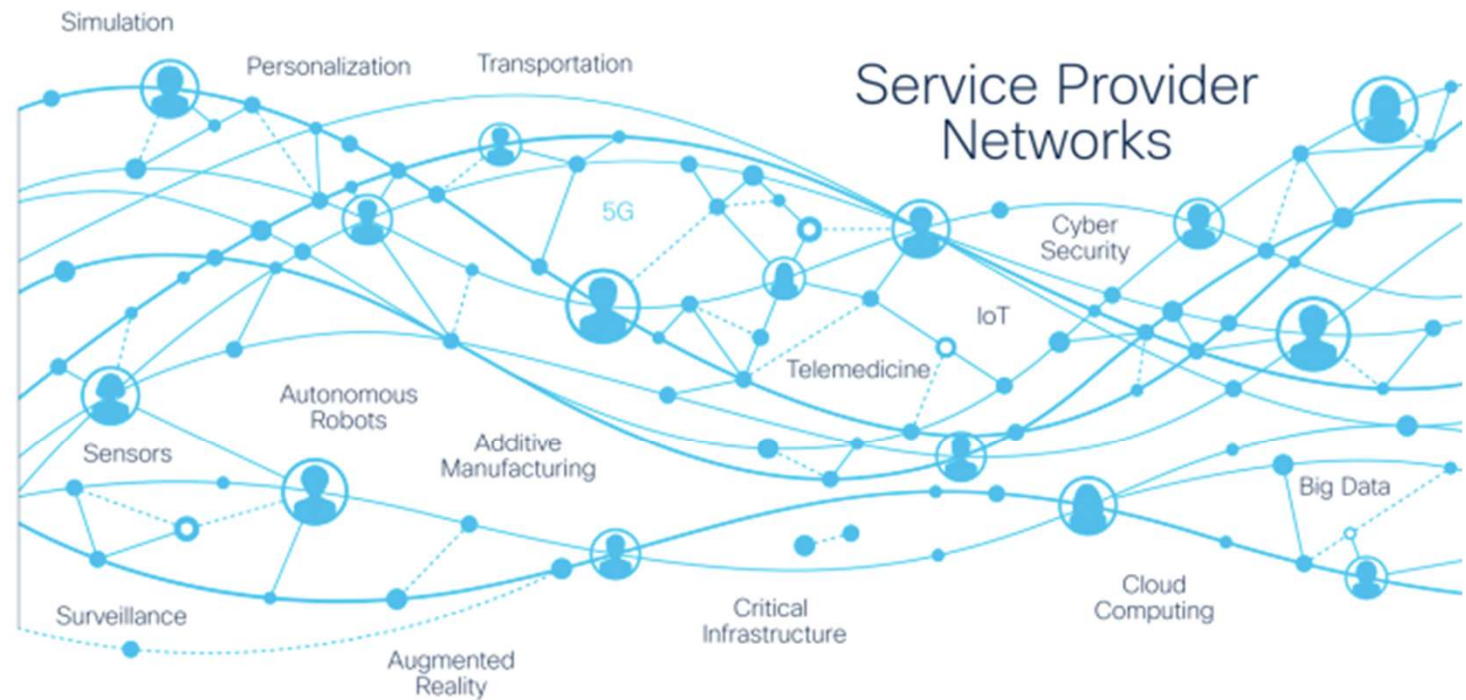
For the way we Work, Live, Play, and Learn

New Participants

Many remain unconnected and emerging IoT

New Potentials

The foundation of economies, governments, and societies




Critical Infrastructure

Requires Mass-scale, Trustworthy Networks




The Exponential Growth of the Internet

The Story of the Emperor, Inventor, and the Game of Chess



The Internet Enters the 2nd Half of the Chessboard

By 2022, Massive Scale Grows Even Larger



Akamai Internet Station

Welcome to your source for global network traffic
and security data from the world's largest edge
platform

CURRENT GLOBAL TRENDS

● JAN 06, 2023 05:11:18 GMT

3B hits/minute


Past 24 hours



GLOBAL PAGE VIEWS

7 Tbps


Past 24 hours



GLOBAL STREAMS

110M

Past 24 hours



DAILY ATTACKS



The New Era is here

>75%

of all Internet traffic will be
video
Up from 60% in 2018



66%

connected flat-panel TV
sets will be 4K
Up from 33% in 2018



3.6

Networked devices and
connections per person
Up from 2.4 in 2018



66%

Of the population will be
using internet by 2023
Up from 51% in 2018



110 Mbps

Average broadband
speed
Up from 45.9Mbps in 2018



44 Mbps

Average mobile speed
Up from 13.2Mbps in 2018



Challenges for Service Providers

Bandwidth Continues to Grow 50% Year-over-Year

The world has gone mobile

Changing Customer Expectations With AI, VR




3X Mobile Data Traffic Growth
(13-44 Mbps) From 2018-2023




Ubiquitous Access
to Apps & Services

Massive IP traffic growth, driven by video

27% CAGR
2018-2023



Connected 4K TV Sets (M)



Rise of cloud computing

Changing SP Architectures/ Service Delivery



Changing Enterprise Business Models
Efficiency & Capacity

Digitization leading to IoT

Emergence of the Internet of Things



People



Process



Data



Things

Market Dynamics

Explosive Internet Growth

2018 Cisco VNI



Economic Challenges for SPs

IHS Markit Report

0.5%

Flat Revenue Growth
(2017 - 2022 CAGR: 0.5%)

11x

\$1 of CapEx in 2020 has to
do 11X the work it did in 2012

5x

Today, operators spend \$5 of
OpEx for each \$1 of CapEx

SPs Want More for Less



Reduce Costs (CapEx, OpEx) and Latency.
Increase Capacity.



Create New Revenue.
Improve Experiences and
Time to Service



Increase Trust
and Security



HOWEVER, BUDGETS
REMAIN FLAT

The SP Market Is Subject to Five Megatrends



*Explosive growth
of video and
mobile*

- Traffic growth driven by video (>80% in 2022) and unlimited plans
- SPs' networks to evolve into cross-medium, converged delivery networks



Advent of 5G

- Service awareness and enablement
- Network assurance, policy, and SLA for Enterprises



*Changing
subscriber user
experience*

- Integration across multiple networks
- Self-service / control / immediacy



*Rise of cloud
and web scale
players*

- Majority of traffic to originate from a small number of content providers
- Emerging Distributed Cloud requirements



Cyber security

- Widening attack surfaces and increasing rate of attacks
- Multiplication of IoT endpoints / DDoS

Five Architectural Shifts Redefining SP Networks

1 Convergence

2 Subscriber Experience

3 Compute and Storage


4 Peering

5 Automation

Next Generation Requirements

- High availability (5 9s+)
- Fast converging (targeting now < .5 sec)
- Low latency (<50ms) and low jitter for real time communication services
- Unicast and multicast traffic (Layer 2 or Layer 3)
- Ultra-High Scalability (thousands to 100,000+ nodes, global scale)
- Traffic Engineering and Steering as needed
- Architectures driven by business objectives
- Fault-domain isolation and service segmentation
- Simplicity
- Greater Efficiency (higher average utilization)
- Secure and Programmable Infrastructure
- Maintenance with little to no customer impact

Next Generation Architectural Decisions



- High Capacity and Scale
- Software Defined & Controller Based
- Virtualized
- Automated
- A Combination of Hardware and Software Worlds

A New Era in Network Architectures



~2 to 5+ Year
Transition
Happening today



IP NGN Era

Designed to support a set of services

Static traffic patterns

Manual configuration (CLI)

Apps Independent of Network

Intent Driven Networks


Designed to support any kind of services

Dynamic traffic patterns

Automation (APIs, Controllers, ...)

App & Network Interaction

IP and Optical Networks Today



IP and Optical Networks Evolution

Converged SDN Transport

High Density Routers
Up to 260 Tbps

400GE ZR/ZR+

Simple Line System
Mux/ Demux/ Amplifier

Automation


Unified Capacity Planning

Unified Path Optimization

Orchestration & Assurance


Unified EMS

Single Layer
Hop-by-hop
Design



Our new reality...

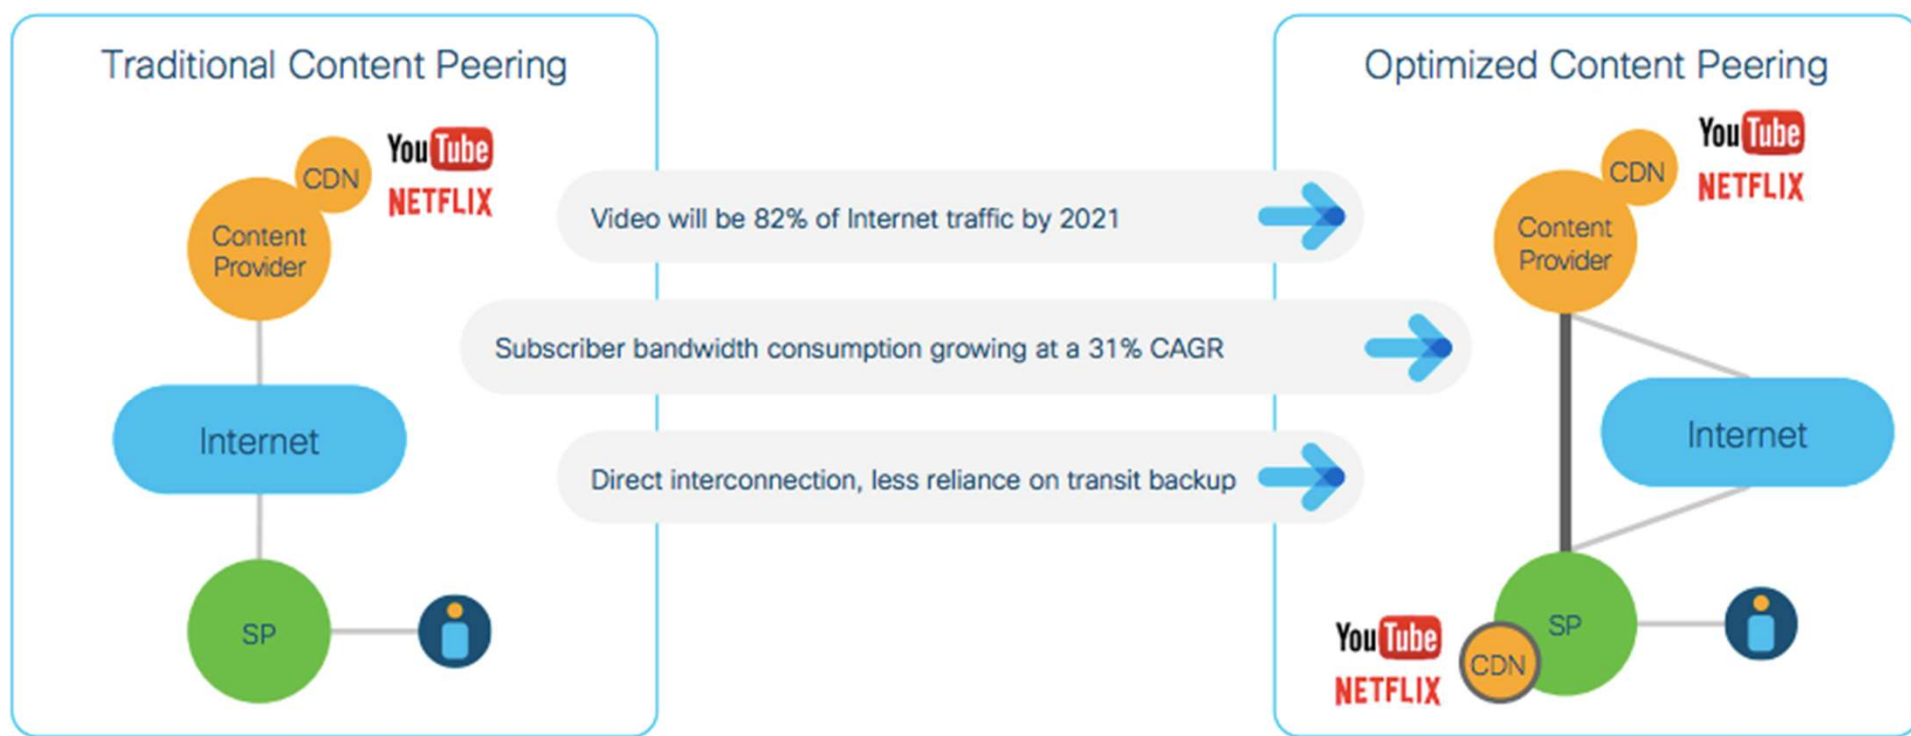
Public Peering Traffic Increases Since March 9, 2020



- “The Internet is Essential”
- Netflix Streaming effect
- Huge Spike in Webex use
- Sustained Busy Hours

How has Peering changed?

“Flattening” of the Internet powered by 2 major transitions



- The Internet for the Future
- **Peering Intro and Internet Trends**
- Peering Network Design
- Peering Network Telemetry
- Peering Security
- Future

What is Peering?

“Peering is the interconnection and exchange of IP data between two networks under different administrative control.”

Peering is the glue holding together the Internet, without it the flow of data across the Internet would not be possible.

Peering represents an important administrative, operational, and security boundary between IP networks.

“Peering” in 2020 = Interconnection covering Content Delivery, Business to Business Services, and Traditional Peering

While the fundamental role of peering hasn't changed, traffic patterns, location, operation, and security requirements have, so peering must evolve as well.

Internet Evolution

“Public” Internet circa 1995

- Low bandwidth clients, dial-up
- Many smaller regional Internet providers
- ~16M users
- Wireline only
- Static content
- More widespread content sources contributed to volume



Today's Internet

- High-speed Internet is widely available
- 100s of millions mobile users
- 4 billion+ users worldwide
- Static content replaced with video
- Traffic **volume** driven by fewer sources
- Leads to “flattening” of Internet: Direct interconnection between producer and consumer networks



Interconnection Types


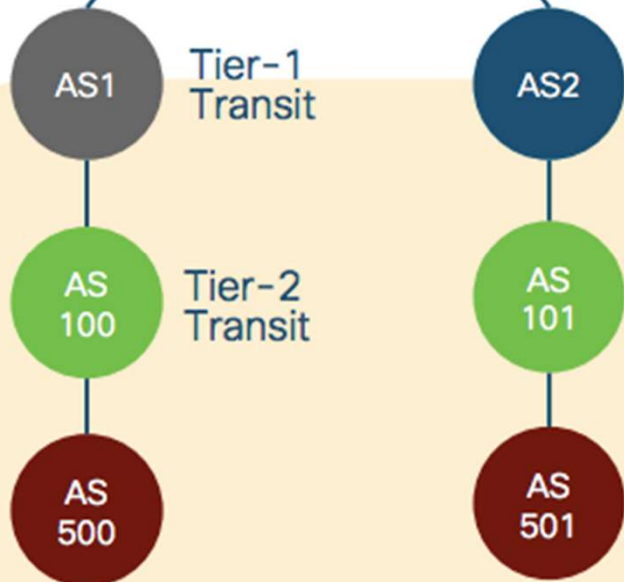


- Public or private fabrics interconnect many networks worldwide
- Highest percentage of traffic **volume** today carried over PNI
- Largest SP and content providers trending to more PNI
- CDN is a type of PNI, may or may not include BGP

“Peering” vs. Transit



- Transit providers provide reachability between their “downstream” ASNs and the rest of the global Internet
- Direct Peering “short circuits” or optimizes traffic distribution
- Expectation is peer will advertise prefixes for itself and any downstream networks (not transit) to other peers



Interconnection Growth


In 1995, ~20 IXPs, today more than 700 worldwide



What content is dominating Internet traffic?

26% CAGR
2017-2022

Exabytes
per Month




Where is traffic coming from?

CDNs will deliver 72 percent of Internet traffic by 2022

30% CAGR
2017-2022

Exabytes
per Month




■ Non-CDN Internet Traffic
(44%, 28%)

■ CDN Internet Traffic
(56%, 72%)

“Flattening” of the Internet

Average AS Path Length (excluding prepending)




- AS Path represents the number of BGP “hops” a prefix traverses
- Even with many more providers, the length has not increased
- Increased density and not increased width
- Graph shows 2012 but trend has continued relatively unchanged

Internet Global Routing Table by Numbers


67057 unique ASNs in global BGP routing table


817505 IPv4 prefixes, 80514 IPv6 prefixes

IPv4 Prefixes




IPv6 Prefixes






- **The Internet for the Future**
- **Peering Intro and Internet Trends**
- **Peering Network Design**
- Peering Network Telemetry
- Peering Security
- Future

Towards a more resilient peering fabric




Traditional Peering



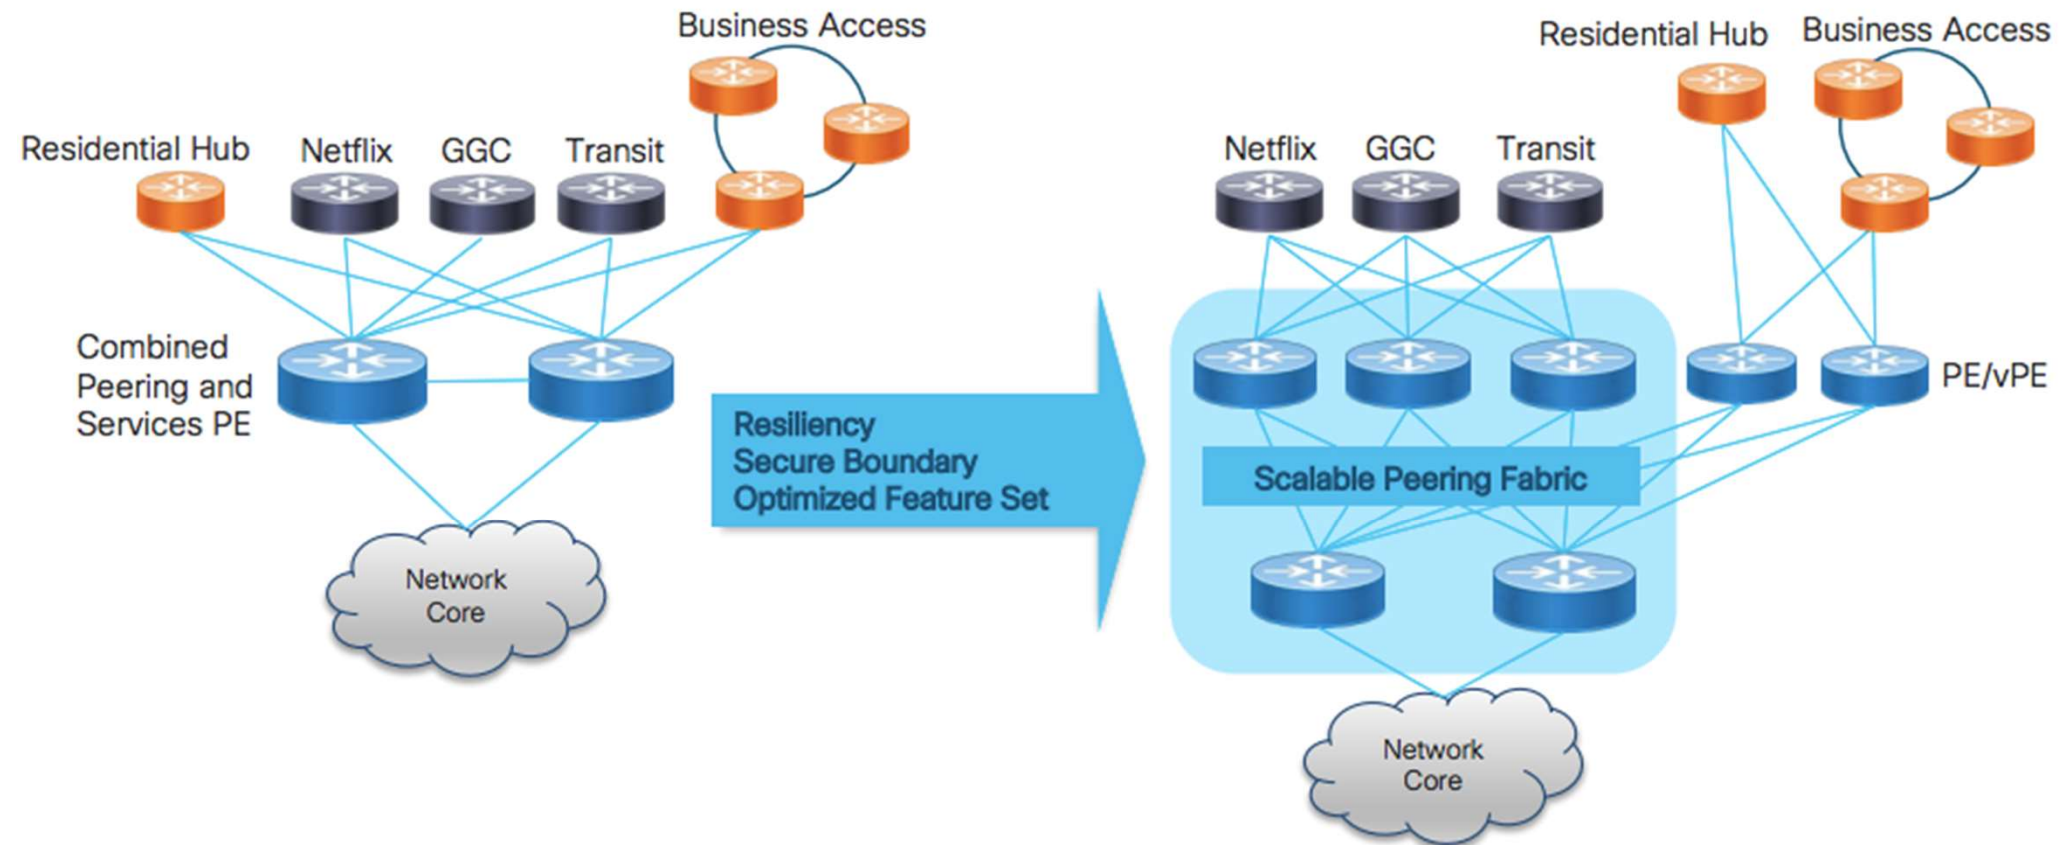
- Horizontal scaling adds resiliency
- Less reliance on long-haul backup for metro or DC Peering
- Reduced blast radius during maintenance or failure
- Simplified SR control-plane

OR




- Greater resiliency and capacity scale
- Optimized feature sets at each layer
- Optimized fabric for both ingress and egress content delivery

Network function separation




Peering Location – SP Services DC




- Localize peering close to user service termination
- Requires flexibility to connect both traditional BGP peering and content caches

Content Provider Cache Aggregation

Peering Facility




Distributed Cache in Agg/Access



How do I influence peering traffic patterns?

Inbound Traffic	Detail
Prefix Advertisement	Suppression, longer prefixes
MED (multi-homed peer)	Some peers (transit) will listen to MEDs and carry traffic over their network to reach yours Typically set to IGP metric
AS Path Length	AS_PATH length influences peer route selection, prepending used for ingress TE
Outbound Traffic	Detail
Local Preference	Highest priority BGP attribute used for path selection
MED	"Metric" attribute also used in outbound path selection
TE Methods (SR-TE, RSVP-TE, EPE)	Steer traffic to specific location or peer using TE overlay methods

Hot potato vs. cold potato routing



- Hot potato (red) has routing policy to always route 10.0.0.0/24 to closest AS1→AS2 interconnect
- Cold potato (blue) carries traffic across AS1 network to AS1→AS2 interconnect point closest to final AS2 destination
- Transit providers (paid) will typically use cold potato, peers will be use hot potato

Inter-AS Ingress Peer Traffic Engineering

Traditional BGP Methods

- Disaggregation
- Advertisement Suppression
- AS-Prepend
- **No guarantees**

Content Provider Peering

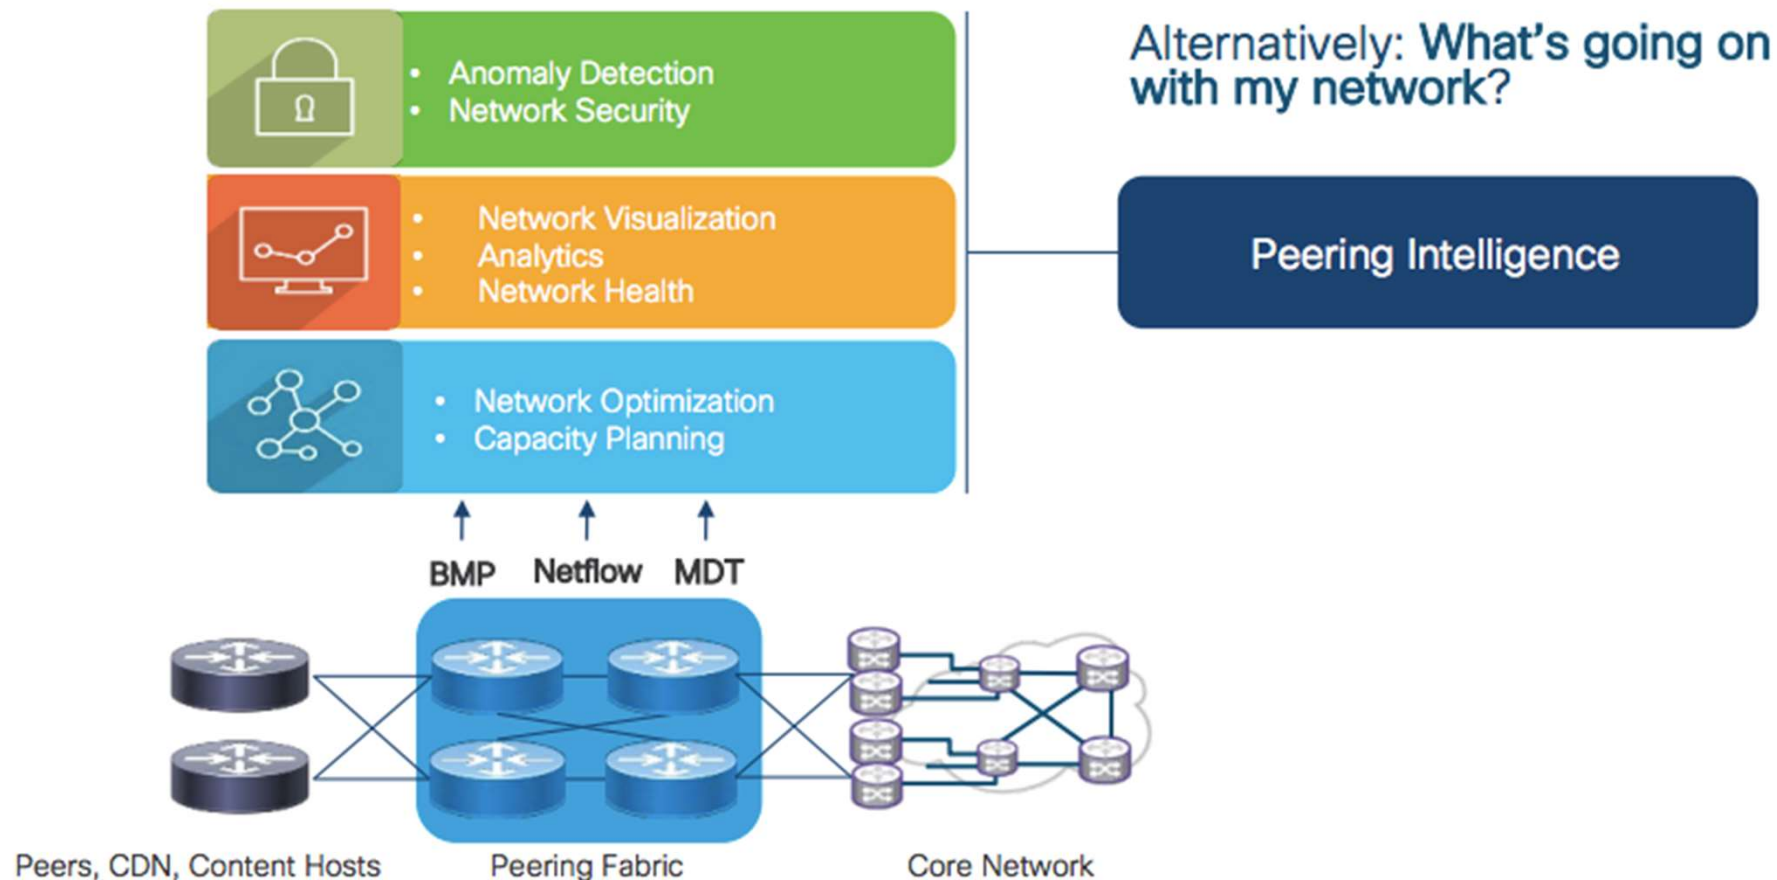
- Most allow influence by community

SR-TE Egress Peer Engineering – Traffic Steering

- Static routing
 - Route IPv4/IPv6 routes into a defined SR Policy
- BGP Flowspec
 - Use BGP-FS rule to direct traffic into a SR Policy or VRF using SR for reachability to egress nodes
- Per-Flow Traffic Steering
 - Utilize CoS markings to map inbound traffic to specific SR Policies
 - Can use QPPB to mark traffic based on destination BGP attributes prior to steering

- **The Internet for the Future**
- **Peering Intro and Internet Trends**
- **Peering Network Design**
- **Peering Network Telemetry**
- Peering Security
- Future

Peering Data Provides Network Insights for Planning, Policy and Control



Periodic Streaming Telemetry

- Data is collected on node, “pushed” to collection entity at periodic intervals
- Cisco calls this model-driven telemetry (MDT)
- Best suited for time-series data, EG: interface statistics, router CPU
- Can also apply to network topology, EG: delay measurement between nodes
- Optimized data collection and optimized transport
- NETCONF/RESTCONF subscriptions can also be considered “streaming telemetry”


Event Driven Telemetry

- Data is pushed asynchronously from node based on state change or monitored event
- SNMP Traps, Syslog, Cisco EEM, Junos event scripts, and RMON are examples of existing event driven telemetry
- Modern approaches use YANG models and same structured encoding as periodic streaming telemetry
- BGP Monitoring Protocol (BMP) can also be thought of as event-driven telemetry

Model-Driven Telemetry for Peering

Higher Resolution Metric Data

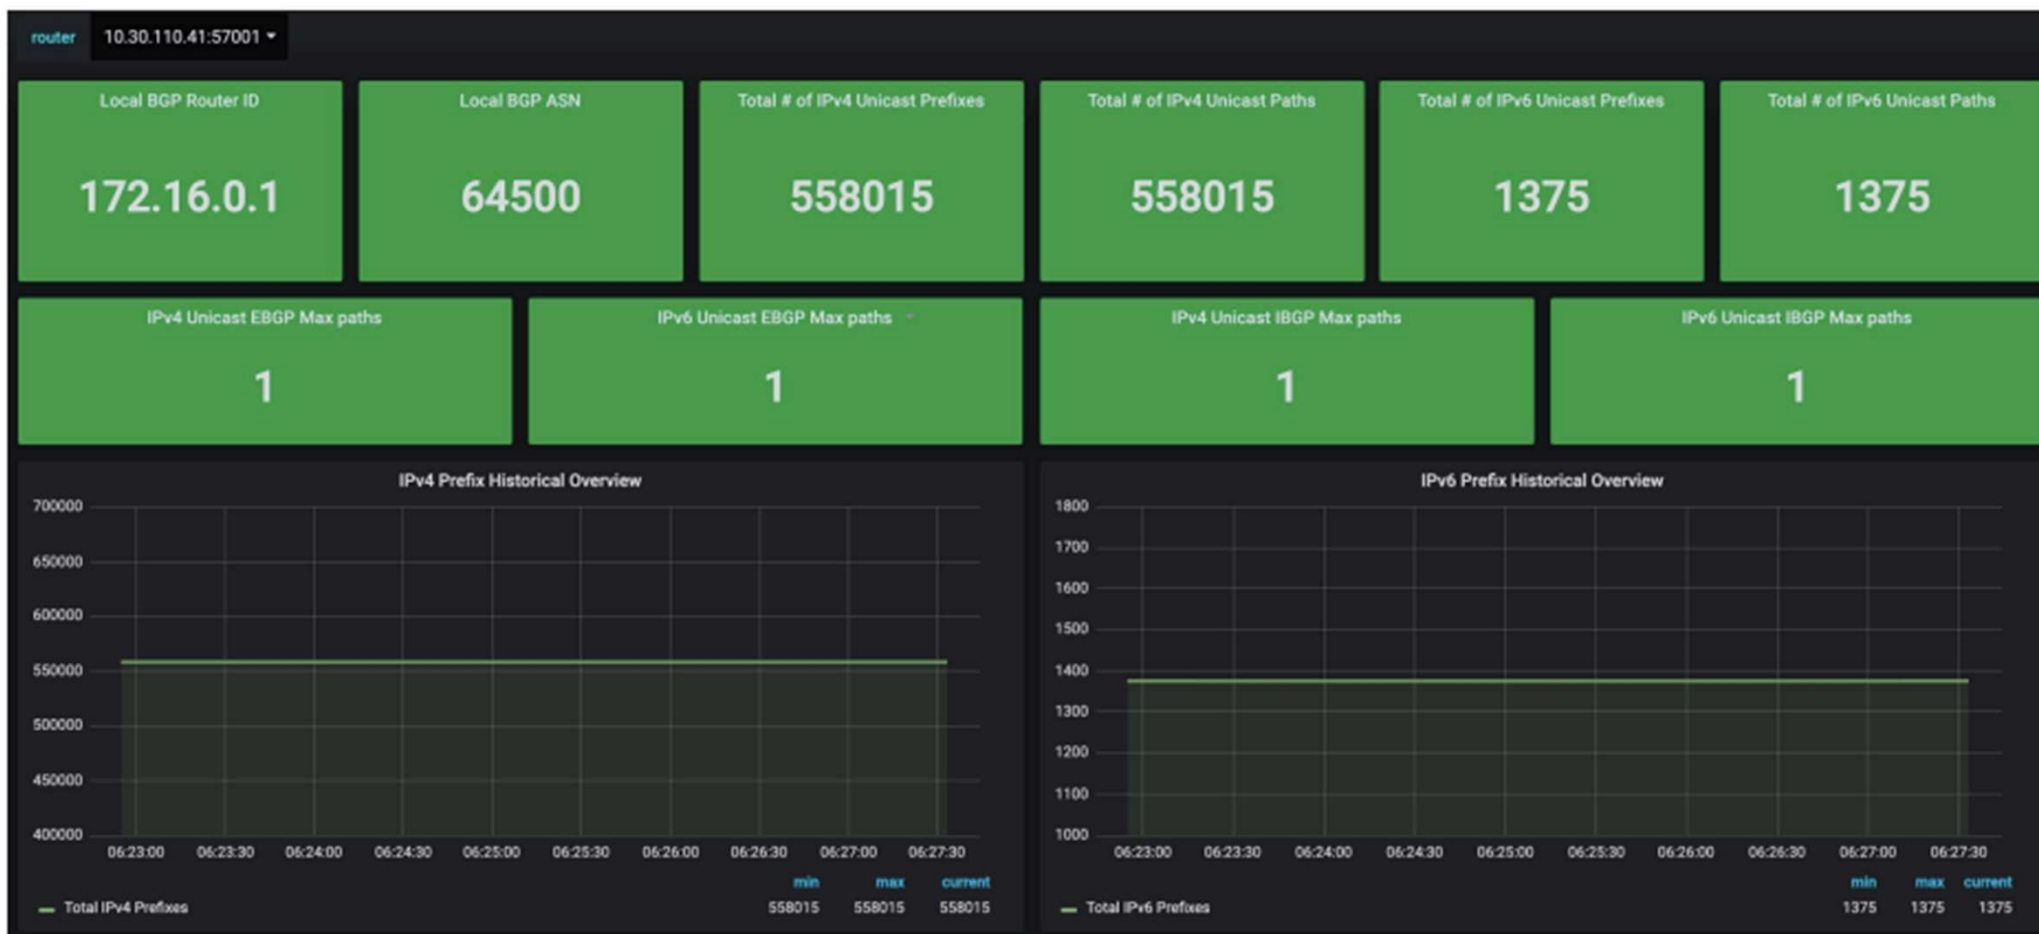
- Quickly detect anomalies when coupled with thresholds or machine learning
- Increased visibility into traffic patterns
- Expose hidden oscillations
- See instant impact of network changes or maintenance events



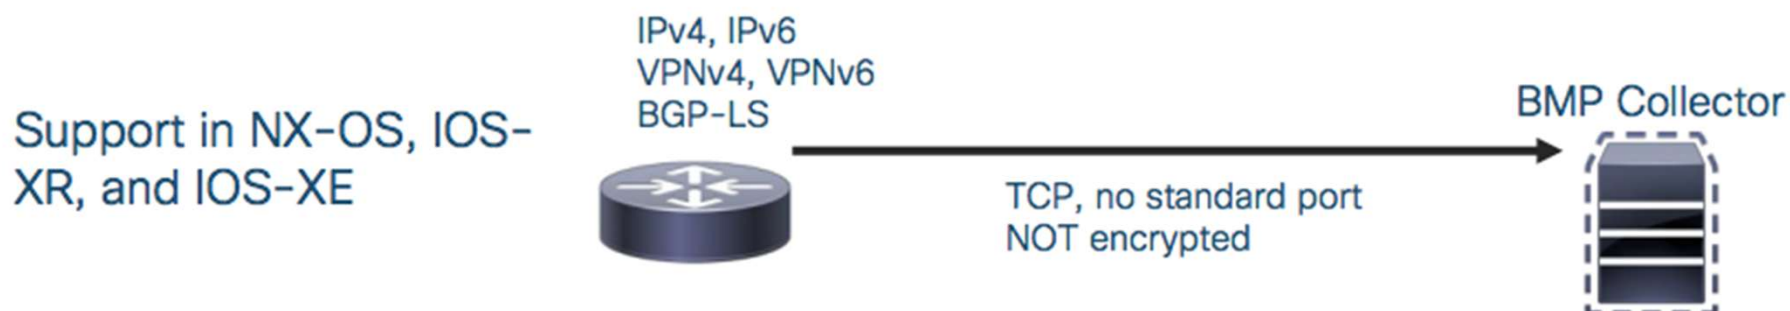
Network and Device Health Monitoring

- Monitoring queuing resources, can be important across peering or fabric where ingress/egress interfaces are the same speed. Similar in concept to datacenter microburst detection
- Monitor hardware FIB capacity and RIB memory

Easily Build Peering Dashboards

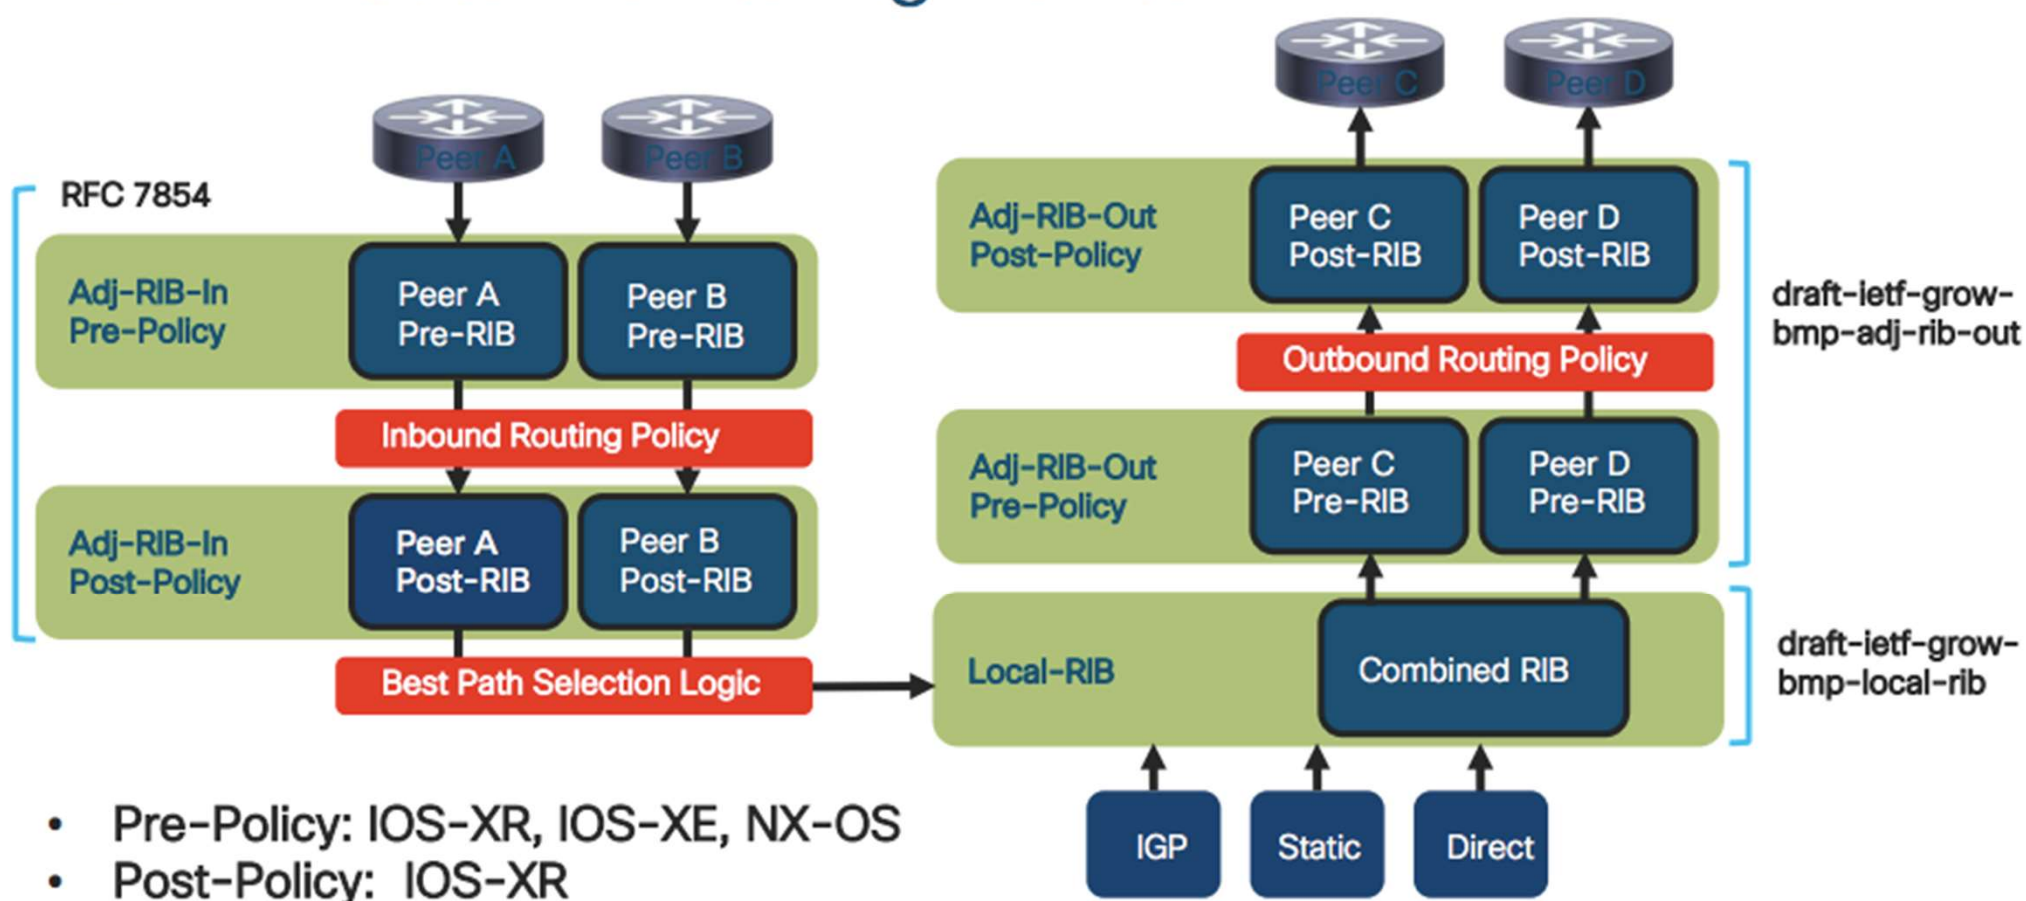


BGP Monitoring Protocol



BMP Message Type	Data
Route Monitoring	Per-peer NLRI and ongoing NLRI updates
Statistics Report	14 periodic stats values, EG: denied prefixes, RIB counts
Peer Down Notification	Peer down, includes local/remote notification msg
Peer Up Notification	Peer in Established state, includes open msg
Initiation Message	sysName, sysDescr, additional info
Termination Message	Termination reason, additional info
Route Mirroring	Exact copy of BGP message and context

BMP Route Monitoring Points



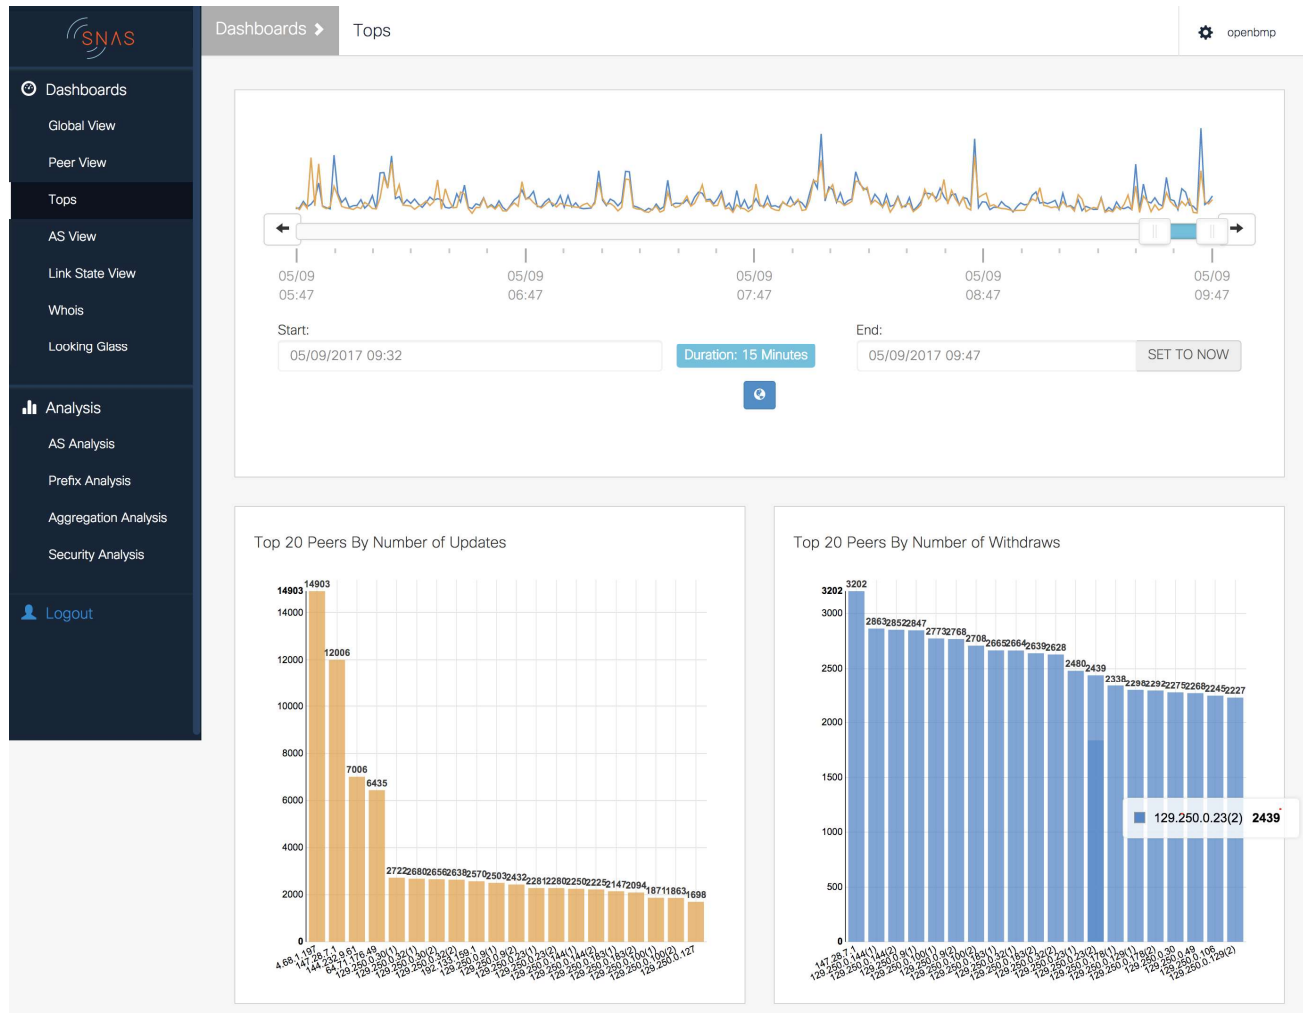
BMP Security Use Cases and Resources

- Monitor peers and prefixes for instability
- Monitor peers for “bad” attributes such as invalid/private ASNs, long ASN lengths, internal communities, bogon prefixes etc.
- Forensic analysis of routing events, having a historical log of routing changes can be invaluable in root cause analysis
- Use diff from pre-policy to easily detect specific rejected prefixes



- SNAS, formerly OpenBMP, available at <https://snas.io>
- All-in-one Docker container easy to spin up
- Alternatives are PMACCT,

Streaming Network Analytics System Sample Report



Netflow / IPFIX

- Has been around for many years
- Cisco Netflow v9 latest Netflow version
- IPFIX – IETF standard flow export
- Peering BGP data must be associated with flow information to be the most meaningful
bgp attribute-download in XR
- Modern traffic rates require sampling.
1:4000 is sufficient for accurate traffic modeling, but dimension for your network
- Application-level visibility is becoming more difficult with encrypted traffic increasing, but peering data is only reliant on SRC/DST IP and still valid

Capacity planning use cases

- “Who should I peer with?”
- “Where should I peer with X,Y,Z?”
- “Should I build local peering or add caching to optimize my network?”
- “Should I change my network topology?”



Peering Capacity Planning

1. Derive traffic matrix
 - SR Traffic Matrix
 - RSVP-TE tunnels
 - Netflow flow source router/interface to egress interface
2. Develop network growth model
 - Use historical data to grow interfaces and links at realistic rates, not the same rate across all links
 - Machine learning, or humans, can add intelligence to the model over time. Filter anomalies and predict seasonal changes
3. Simulate network failures
 - Balance cost vs. consumer experience and SLAs

Resources for Finding Peers

- Peering DB
 - www.peeringdb.net
 - Database of peering locations, who is peering at those locations, and what their peering policies are
- Networking and Peering Conferences
 - NANOG, RIPE, APRICOT, etc.
 - Meet other providers and IXP organizers
 - Negotiate peering terms and interconnection cost
- Content cache providers
 - Netflix OpenConnect
 - Google Global Cache
 - Akamai
 - Apple

Cisco Peering Telemetry Open Source

Application	Collection Method	Use Cases
 Streaming Network Analytics System	BMP (BGP Monitoring Protocol) BGP-LS	BGP performance monitoring BGP security monitoring BGP routing anomalies IGP topology BGP looking glass
 telegraf	Model-Driven Telemetry Open-source collector with Cisco MDT plugins in mainline release	Collect, process, and output router telemetry GNMi or static configuration Input gRPC,JSON telemetry data Output to Telegraf supported streams (Kafka, InfluxDB, etc.)

- **The Internet for the Future**
- **Peering Intro and Internet Trends**
- **Peering Network Design**
- **Peering Network Telemetry**
- **Peering Security**
- **Future**

Peering Edge Security Threats

Leading Threat Concerns*

DDoS Attack (88%)

Infrastructure
Security (55%)

BGP Route Hijacking
(25%)

Description


Distributed Denial of Service
Volumetric traffic to overwhelm network and hosts

Compromise of network control-plane
Compromise of network devices

Man-in-the-middle attack
ASN hijacking has also been an issue

Peering DDoS Mitigation - RTBH

- Remote Triggered Black Hole
 - Applicable for content, SP, enterprise
 - Black hole could be sinkhole, honey pot
- S/RTBH
 - Drop based on source address and not destination
 - Uses Unicast RPF with BGP NH set to /32 with static route to Null0
- Upstream providers will often match specific community to allow customers to trigger RTBH (see resources for more info)
- Cymru has UTRS, global RTBH network



Peering DDoS Mitigation – BGP Flowspec

- New AFI/SAFI NLRI, IPv4 defined in RFC5575, IPv6 nearing RFC status
- Distribute data-plane ACLs using MP-BGP
- Match on packet criteria then drop, police, redirect, or mark matched traffic
- Foundation for scalable distributed DDoS protection

Server Config

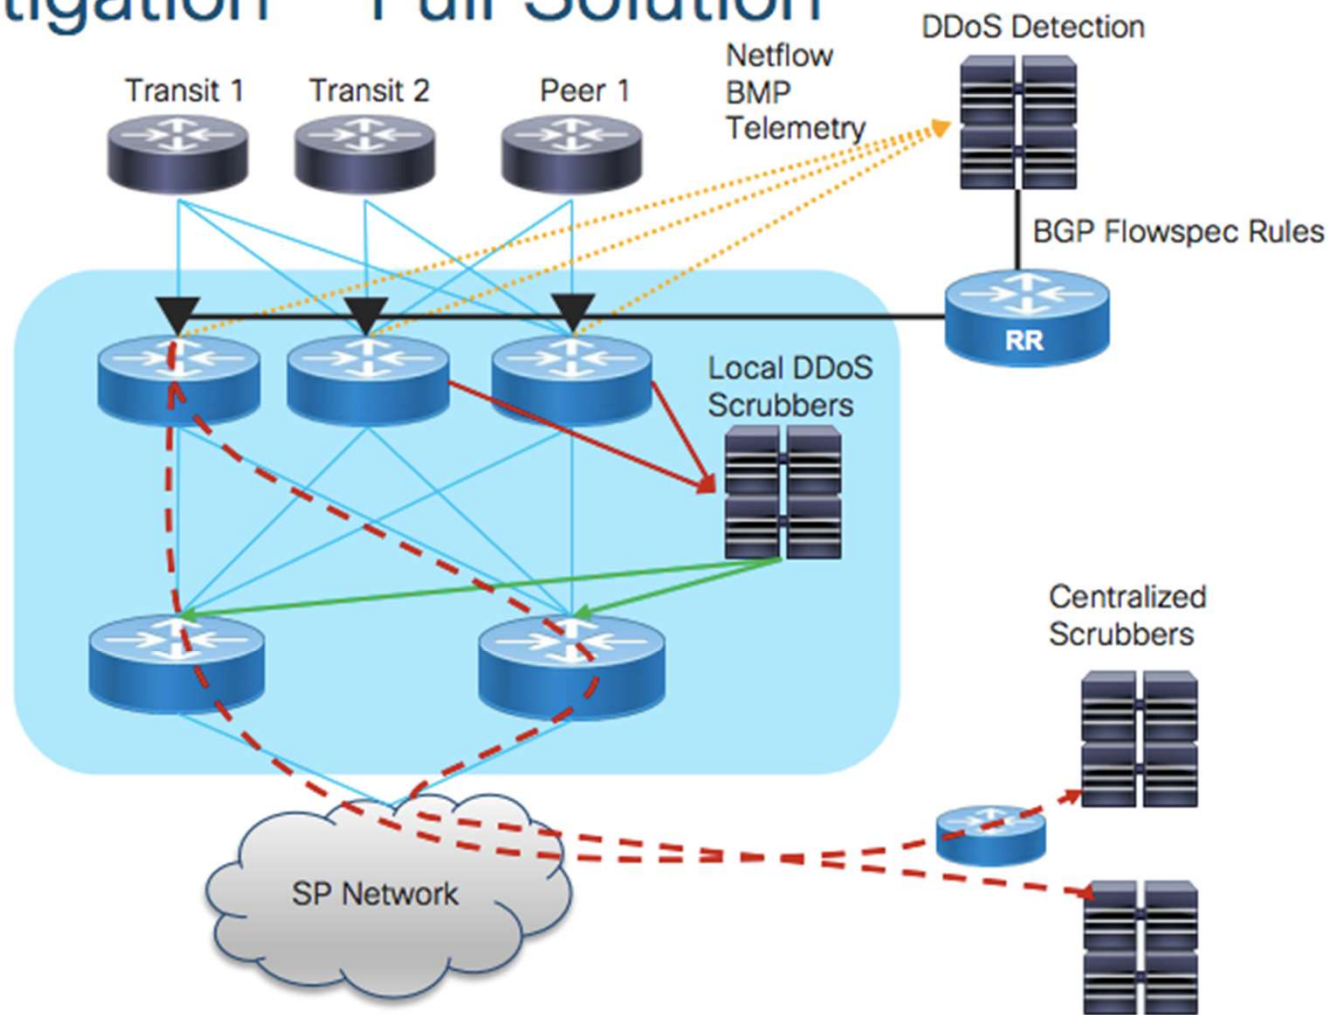
```
class-map type traffic match-all memcached
  match destination-port 11211
  match protocol udp tcp
end-class-map
!
policy-map type pbr drop-memcached
  class type traffic memcached
    drop
  !
  class type traffic class-default
  !
end-policy-map
!
flowspec
  address-family ipv4
    service-policy type pbr drop-memcached
```

Client Config

```
flowspec
  address-family ipv4
    local-install interface-all
```


Peering DDoS Mitigation – Full Solution

- Granular Netflow helps identify attacks faster
- High scale 5-tuple ACLs block high-volume identified traffic
- BGP Flowspec automates traffic filtering, policing, and redirection
- SR-TE to steer and load-balance traffic to end scrubbers



DDoS Traffic Steering using SR-TE and BGP-FS

- BGP Flowspec redirects traffic to SR Policy
- SR-TE to steer and load-balance traffic to end scrubber/DPI
- Manually defined EPE SID in XR 7.1.1



DDoS Traffic Steering using SR-TE and BGP-FS

Head-end Configuration

```
segment-routing
traffic-eng
  segment-list pr1-ddos-1
    index 1 mpls label 16441
    index 2 mpls label 28000
  segment-list pr1-ddos-2
    index 1 mpls label 16441
    index 2 mpls label 28001
  policy pr1_ddos1_epe
    color 999 end-point ipv4
192.168.14.4
  candidate-paths
    preference 500
    explicit segment-list pr1-ddos-1
    !
    explicit segment-list pr1-ddos-2
    weight 100
```


**Per-flow load balances across
equal weight paths**

SR-TE policy database

```
Color: 999, End-point: 192.168.14.4
Name: srte_c_999_ep_192.168.14.4
Status:
  Admin: up Operational: up for 00:17:25
Candidate-paths:
  Preference: 500 (configuration) (active)
  Name: pr1_ddos1_epe
  Requested BSID: dynamic
  PCC info:
    Symbolic name: cfg_pr1_ddos1_epe_discr_500
    PLSP-ID: 517
  Explicit: segment-list pr1-ddos-1 (valid)
    Weight: 100, Metric Type: TE
    16441
    28000
  Explicit: segment-list pr1-ddos-2 (valid)
    Weight: 100, Metric Type: TE
    16441
    28001
Attributes:
  Binding SID: 25384
  Forward Class: 0
  Steering BGP disabled: no
  IPv6 caps enable: yes
```

Increasing BGP Session Security with TCP-AO

- Session threats
 - TCP RST attacks
 - Snooping
 - SYN flooding
 - Peering is being used for more critical applications than just best-effort Internet
- Question: When was TCP MD5 authentication obsoleted?
 - Answer: Obsoleted in 2010
- TCP-AO – TCP Authentication Option – RFC 5925
 - Use HMAC-SHA2-256 hash at minimum
 - Protects BGP TCP connection by authenticating TCP segments
 - Does NOT provide session encryption
 - Supported in IOS-XR in 6.5.3, IOS-XE in 16.12
 - Recommended in RFC 7454 (2015)



TCP-AO IOS-XR Configuration

Key chain and TCP AO Config

```
tcp ao
  keychain TCP-AO-KEY
    key 1 SendID 100 ReceiveID 100
  !
!
key chain TCP-AO-KEY
  key 1
    accept-lifetime 00:00:00 january 01 2018 infinite
    key-string password 0204034B0A131B29
    send-lifetime 00:00:00 january 01 2018 infinite
    cryptographic-algorithm AES-128-CMAC-96
```

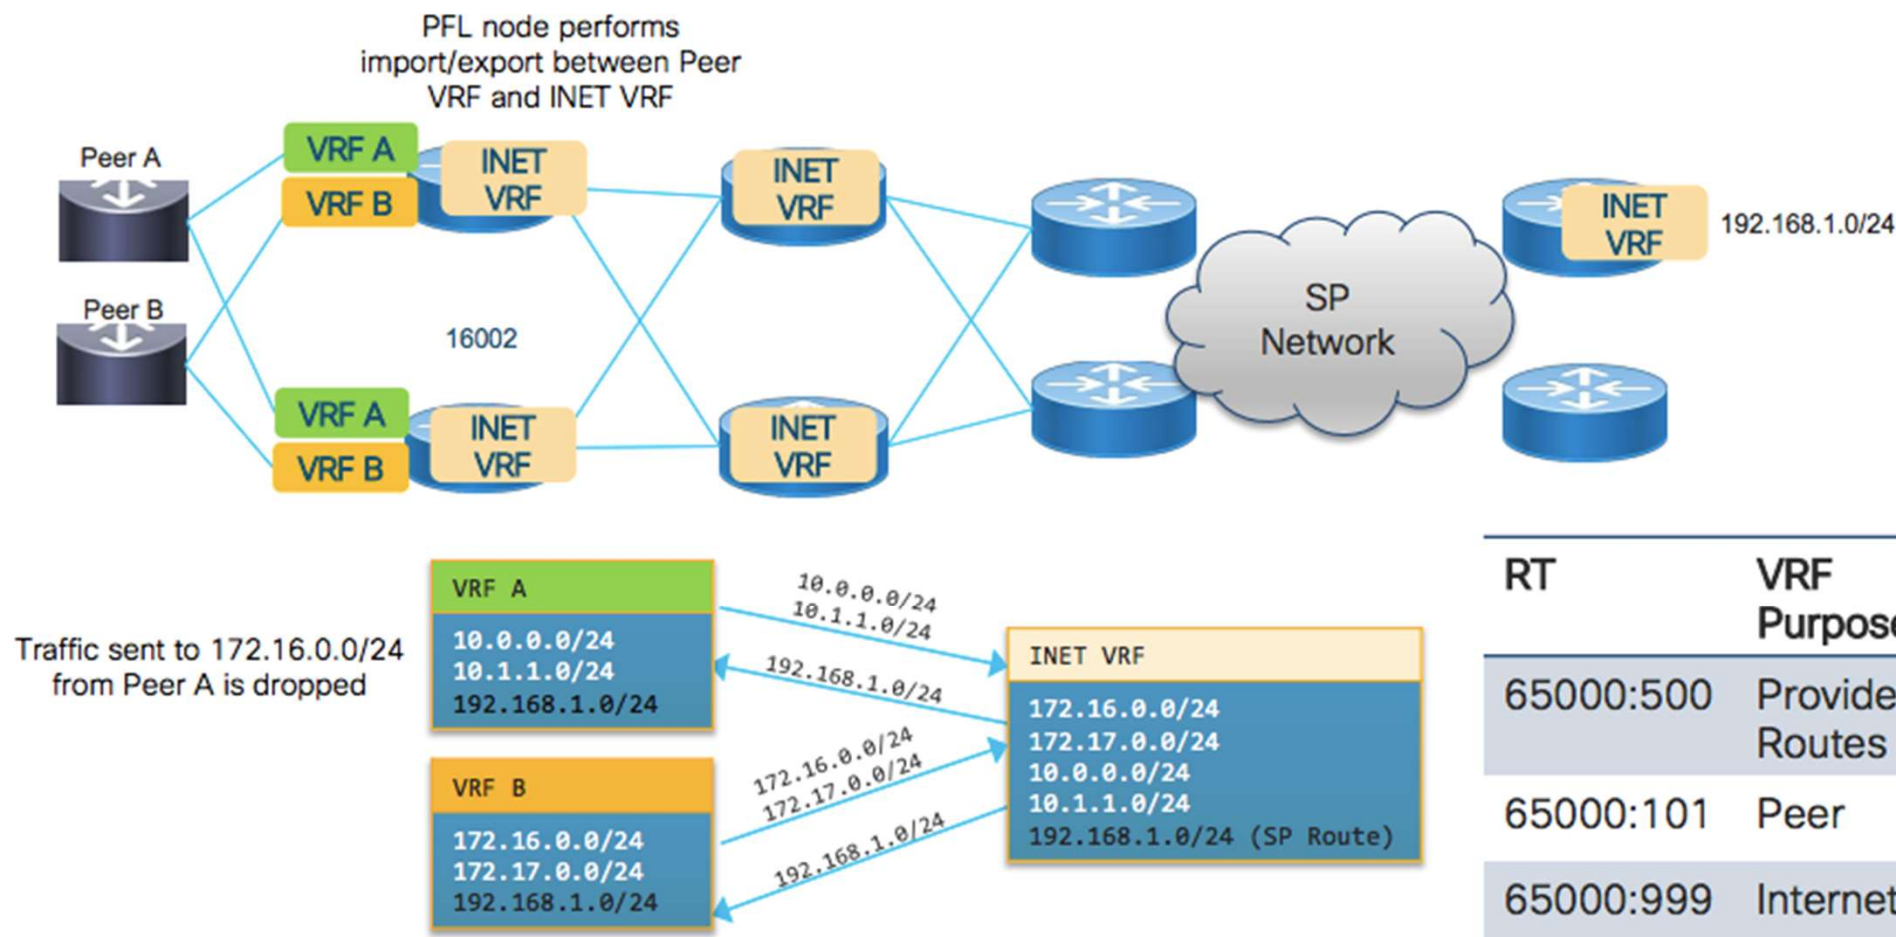
BGP Configuration

```
router bgp 100
  neighbor 1.2.3.4
    remote-as 101
    ao TCP-AO-KEY include-tcp-options enable
```


Infrastructure Security using Peering in a VRF

Peering in VRFs GRT in default VRF	All "Internet" in VRF	Peering in VRFs GRT in "Internet" VRF
<ul style="list-style-type: none">Isolate peers for data-plane securityImport only specific peer routes into customer VRFs	<ul style="list-style-type: none">Base infrastructure isolated from user and service trafficBGP diversity using RD instead of add-path extensionsEBGP peers share same Internet VRF	<ul style="list-style-type: none">Isolate peers for data-plane securityBase infrastructure isolated from user and service trafficControl inbound/outbound routing using RT

Internet in a VRF with Peer Isolation



Data Plane Boundary Concerns

- Scanning vulnerability probes and botnet C&C
- Volumetric and application-layer DoS
- CoS value retention
- Spoofed traffic
- Infrastructure attack traffic to peering edge, DNS, and other critical services

Ingress Traffic

- What should I do at the edge?
 - Filter control-plane traffic to internal infrastructure
 - Filter well-known bad traffic that won't cause user issues (chargen, etc.)
 - Fragments? Source of many attacks but may not be feasible
 - Explicitly reset CoS values on ingress
 - Monitor everything, characterize steady-state and rate-limit if you can
 - Follow security alerts from US-CERT (<https://www.us-cert.gov/ncas/alerts>), CVE feeds and other security organizations
- CDN is still an unsecured edge device
- Use BGP-FS for transient dynamic events, use stateless ACLs for well-defined long-term filters
- Route dark (unused) space to honeypot servers for threat inspection and research

Egress Traffic Filtering – Much the same as ingress

Follow BCP 38 for ingress filtering on downstream connections ☺

- Use strict filtering based on well-maintained data

Known bad protocols with no current legitimate Internet use

Automation is key to deploying filters quickly so your customers are not actors in attacks

Best Practices Summary

- TCP-AO session authentication with strong encryption (AES)
 - TCP-AO available in IOS-XR 6.5.1 w/stronger crypto algorithms
 - MD5 as a lowest common denominator
- Control-plane policing per-peer, default in IOS-XR
- Reset IPP, DSCP, EXP on inbound peering traffic
- Delete inbound communities, especially if doing VRF peering, some vendors may accept routes with an RT set from an EBGP neighbor
- Limit BGP control-plane to only configured peers
- Data-plane filters inbound and outbound
 - If feasible whitelist your own IP space at edge
 - Automation is key in maintaining accuracy
- Review BCP 84, 194, and BCP 38 if you are providing Internet service

Summary

Peering Infrastructure

- High frequency **Netflow**, **BMP**, and **Model-Driven Telemetry** Export
- **Control-Plane Policers** Per-Peer
- BGP MD5, GTSM Support

BGP Prefix Security

- Powerful **IOS-XR Routing Policy** Language (RPL)
- **BMP** ADJ-RIB-In Pre and Post Policy
- **RPKI** ROA Support, RFC8212 Default Deny

SP Network

- Peer and Internet isolation using **validated Peering and Internet in a VRF**
- IPv4 and IPv6 **BGP Flowspec**
- Integration with Leading **DDoS Detection and Mitigation Platforms**

- **The Internet for the Future**
- **Peering Intro and Internet Trends**
- **Peering Network Design**
- **Peering Network Telemetry**
- **Peering Security**
- **Future**

Today's trends continue

- Requires flexible hardware with low power footprint
- In SP networks we will continue to see peering and CDN distributed deeper in the network close to users
- Continued "compartmentalization" of the Internet as long-haul traffic levels drop over time
- Continued focus on security and peering operations
- Continued enhancements in BGP Flowspec
 - PCEP drafts on BGP-FS via PCEP
 - Flexible FS redirect based on defined Segment Routing SID list

Additional Peering Resources

- Cisco Peering Fabric HLD
 - <https://xrdocs.io/design/blogs/latest-peering-fabric-hld>
 - Details on best practices, validated model driven telemetry
- <https://github.com/cisco-ie/anx> to explore NETCONF and telemetry paths
- <http://www.team-cymru.com/>
 - Resource for security best practices, BOGON API feed
- <https://onestep.net/communities/>
 - List of communities supported by SPs to trigger route behavior
- IETF working groups
 - IDR (Inter-Domain Routing)
 - SIDR (Secure Inter-Domain Routing, now closed)
 - SIDROPS (Secure Inter-Domain Routing Ops)
 - GROW (Global Routing Operations)

Reference

SP360: Service Provider- <https://blogs.cisco.com/sp>
(<https://www.caida.org/projects/as-core>
<https://www.cidr-report.org/as2.0/>
<https://atlas.ripe.net/results/maps/>
<https://www.akamai.com/internet-station>