Packet Design

Making MPLS VPN Management Scalable With Network-Wide Traffic Analysis

Cengiz Alaettinoglu
cengiz@packetdesign.com

The State of IP/MPLS VPN Management

- **Virtualized nature of IP/MPLS networks creates a management challenge**
- **How do we answer:**
  - How does Customer X from host1 to host2 route?
  - Given a packet captured in the middle of the network
    - Which customer owns it, where is it going and where is it coming from?

This is what we know about traffic:

- **Per class link utilization via SNMP interface counters**

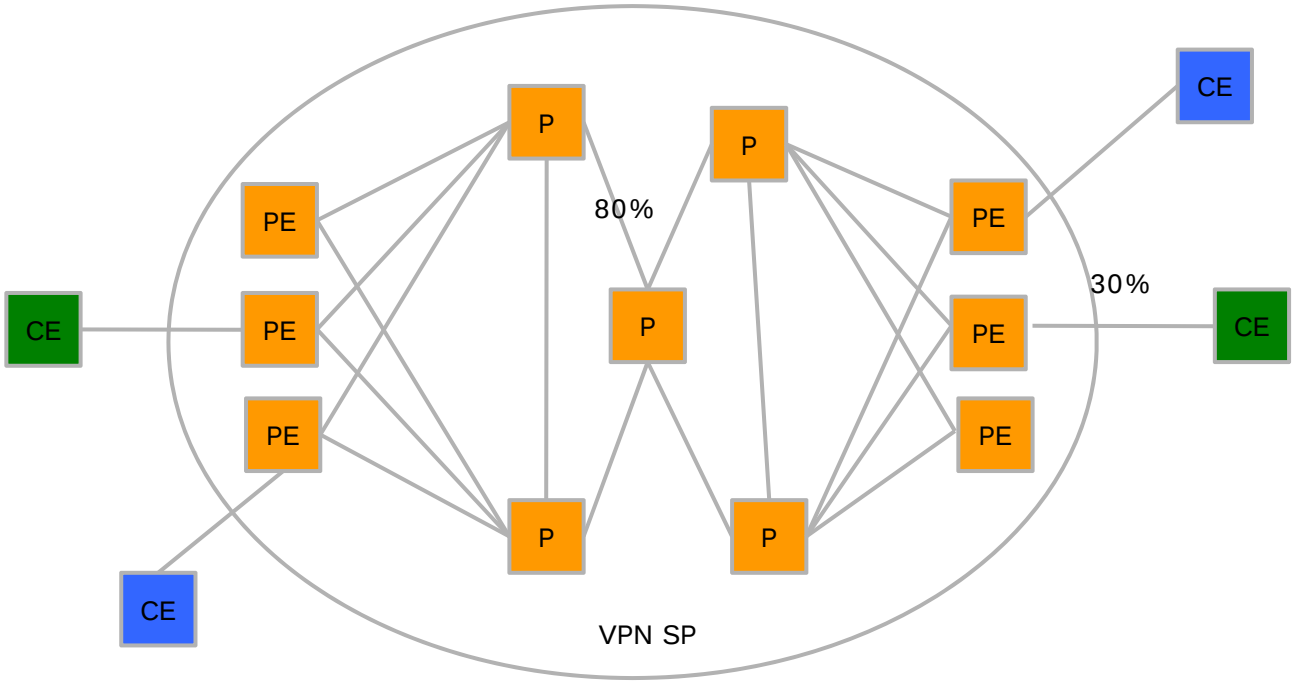

Service Assurance Challenges

- **Where is the traffic coming from where it is going?**
  - Needed for engineering the network
- **Which routers and links are carrying a customer VPN's traffic?**
  - Needed for diagnosing customer traffic issues and ensuring SLAs
- **Is CoS traffic within profile on all links in the network?**
- **How fast is the traffic growing given current trends?**
  - Needed for capacity planning
- **Will adding a new customer impact meeting SLAs of the other customers?**
- **Can a planned outage proceed without effecting SLAs?**
- **What was the state of the network before a problem occurred?**

IPFIX comes to rescue (Net-Flow v9)

- **IPFIX, an IETF standard, is a way to export traffic as unidirectional flows**
  - Not about a single packet but about a *flow* of packets

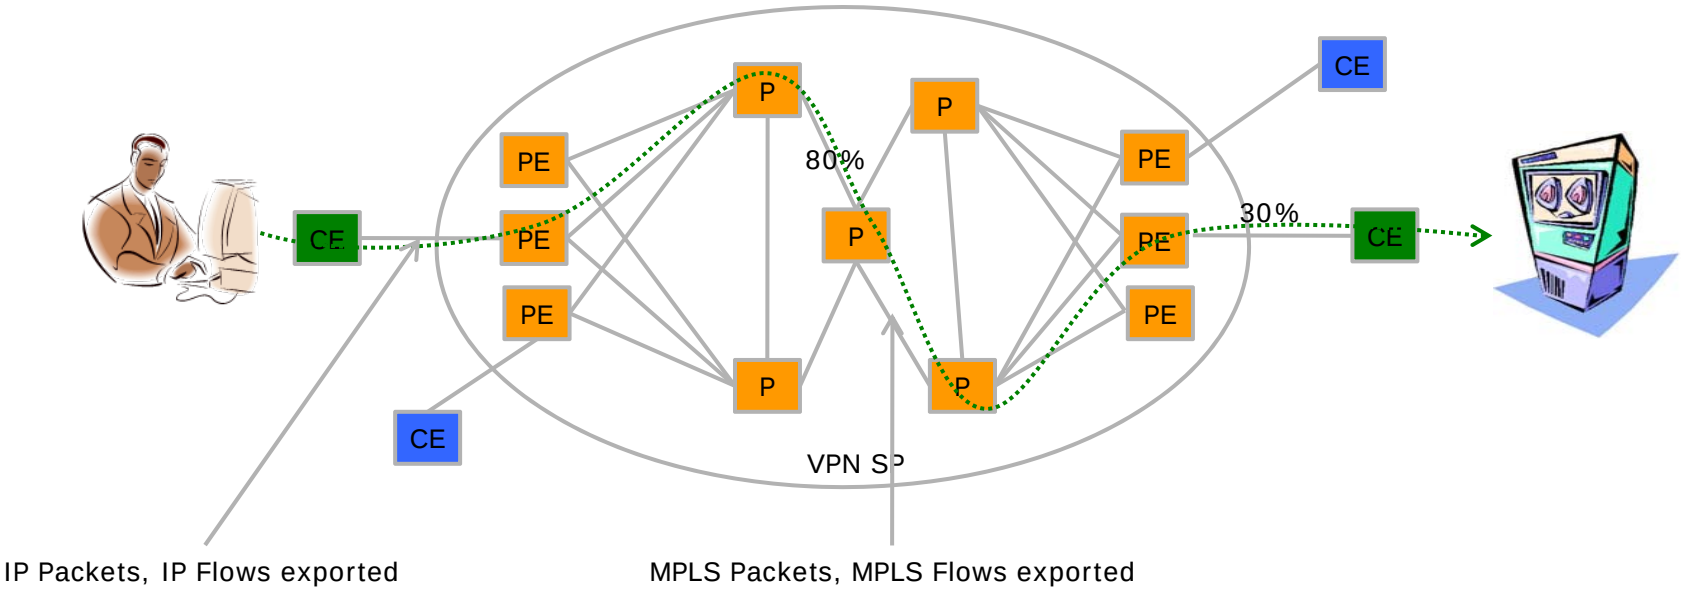

How does IPFIX work?

- **Routers/switches cache a record of active flows traversing their interfaces**
  - Includes source/destination addresses, ports, ToS bits
  - MPLS labels, exp bits
  - Timestamps
- **Stats are updated for each active flow as more packets are seen**
  - Increments bits and packets for the flow
- **When flow ends (TCP FIN/RST, idle timer, absolute timer, etc), an IPFIX record is generated summarizing the stats for the flow**
- **Flow records are put into a IPFIX packet and sent to a flow collector**

Flow Record Evolution

Netflow V5:

- IPv4 Centric
- Src/Dst IP
- Proto
- Src/Dst Port
- Src/Dst AS
- BGP NextHop
- Tos Byte

IPFIX (Netflow V9):

+

- IPv6
- MPLS Label Stacks
- FEC for top label
- Non-flow based data
- Multicast

Issues

- **Where to collect flow records?**
  - 100K+ CE-PE interfaces in large VPN providers
  - 1000s of core interfaces
  - If not collected everywhere how can we know which links carry what customer traffic?
- **IPFIX records contain MPLS labels, not the friendliest piece of information, how to map them to PEs, customers in a scalable fashion**
- **Solution: Route-Flow fusion**
  - Enables scalable IPFIX deployment
  - Maps flow to routing concepts that the SPs can control

Route Analytics Technology

- **Listens passively to routing updates**
- **Creates a real-time network map**
  - As up to date as routers
- **Analysis of current paths**
  - Paths are computed using the same procedures as routers
- **Historical view with breakdown of instability**
  - Full routing event history/forensic audit trail
  - Flapping links, prefixes
  - Ability to look at state of routing at any point in recorded history

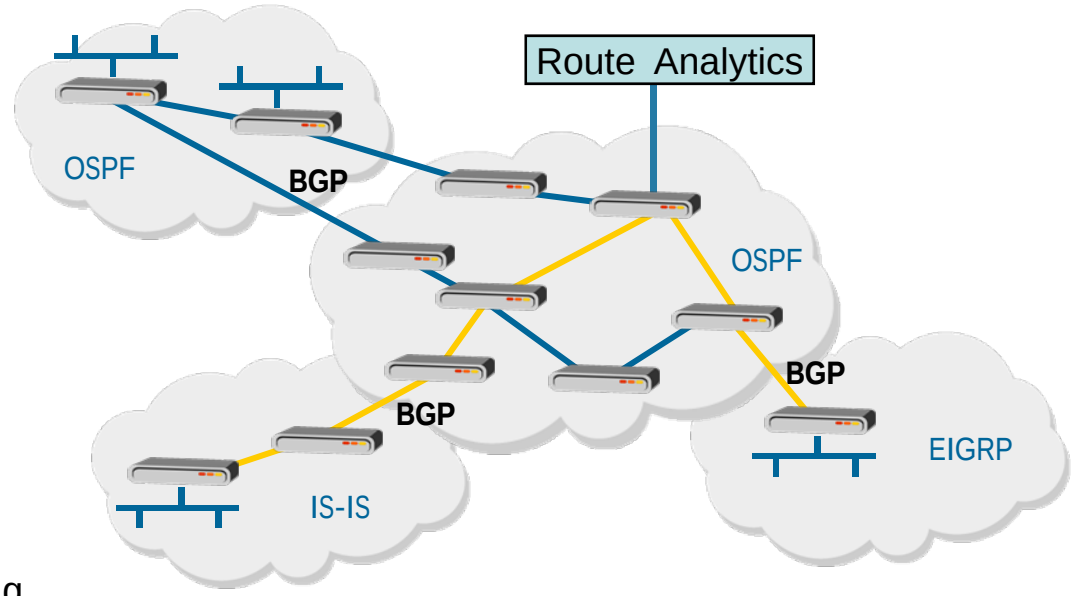

- Works across protocols (OSPF, IS-IS, BGP, EIGRP, RFC 2547bis)
- Manages and organizes the deluge of data

Route-Flow Fusion

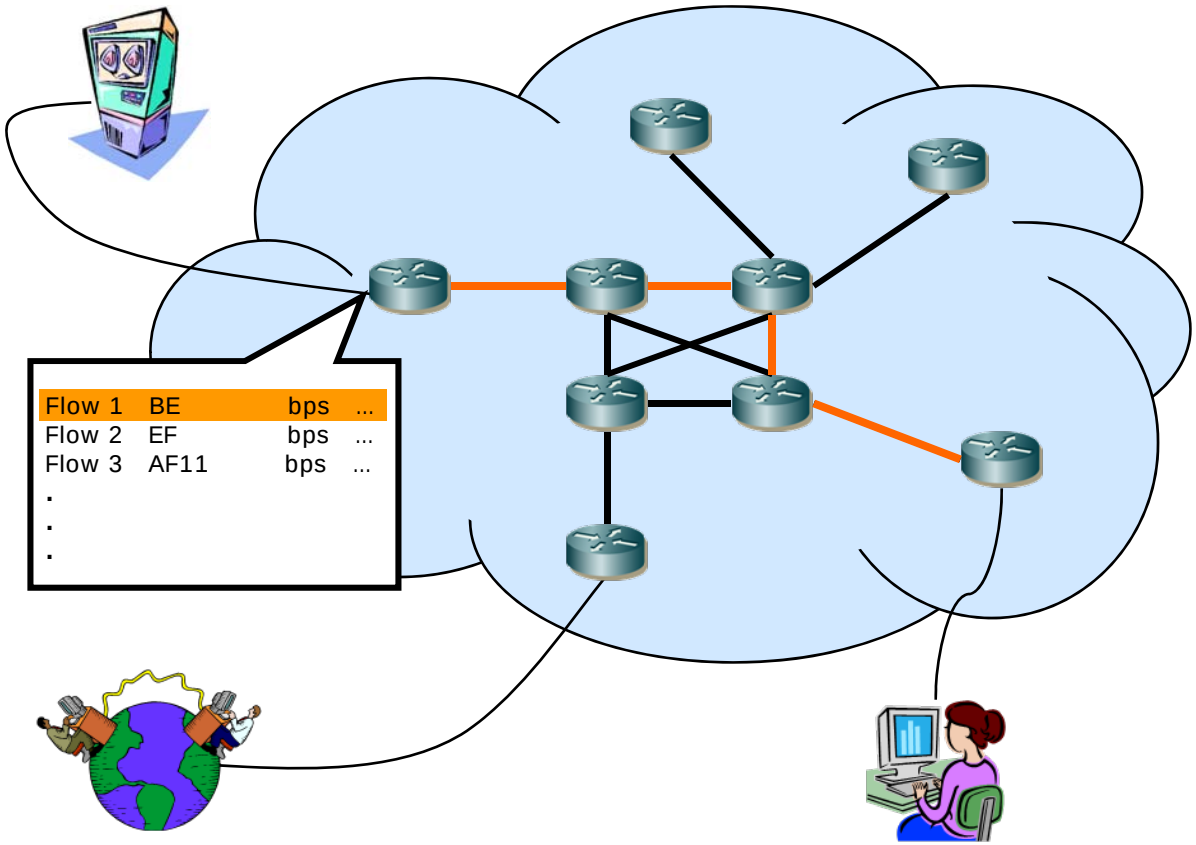

- Collects flow data at key network locations
- Computes traffic flows across topology

Route-Flow Fusion

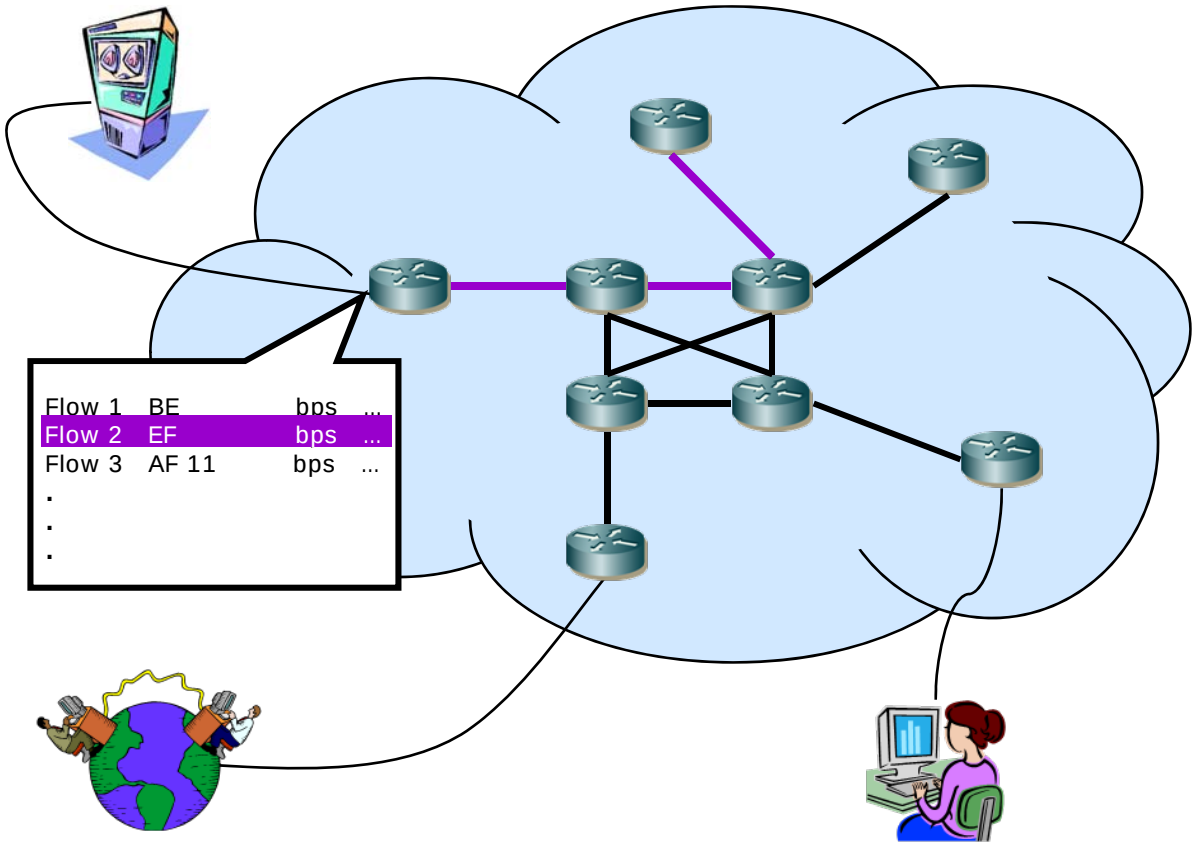

- Collects flow data at key network locations
- Computes traffic flows across topology

Route-Flow Fusion

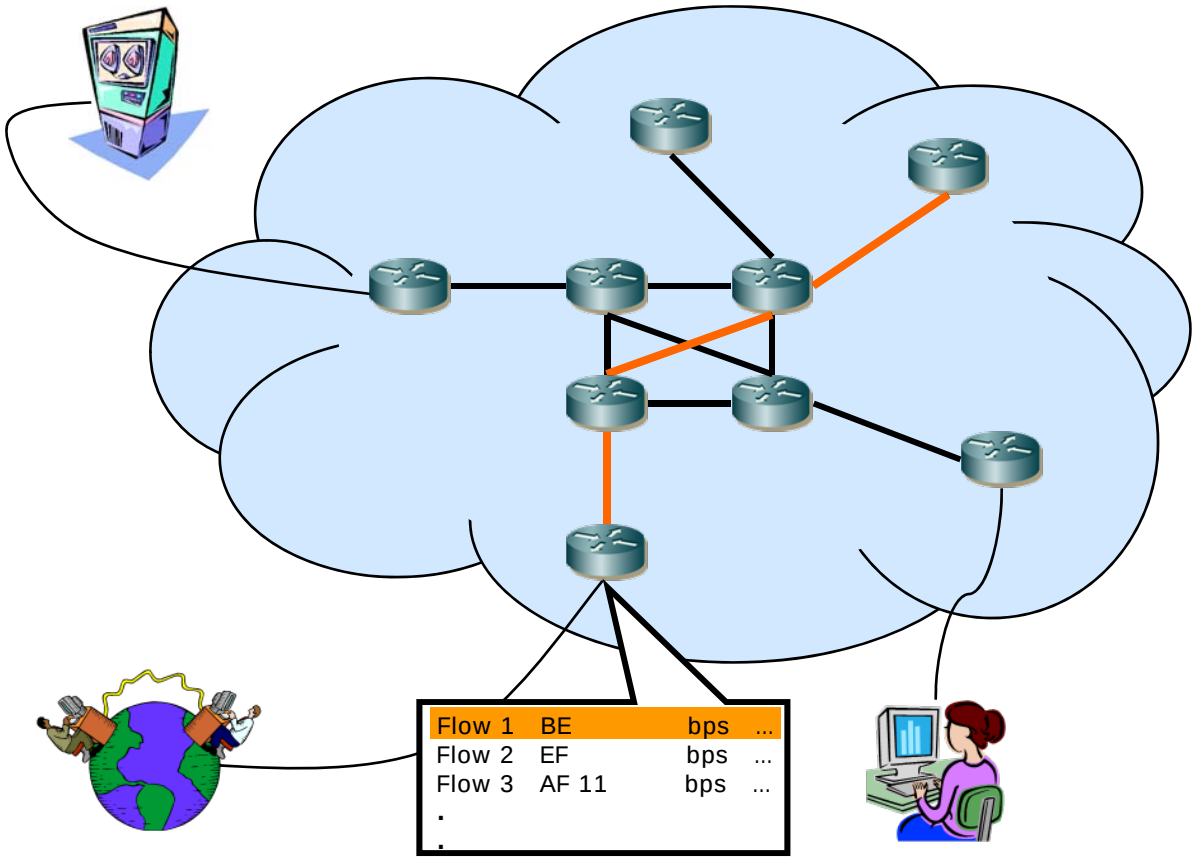

- Collects flow data at key network locations
- Computes traffic flows across topology

Route-Flow Fusion

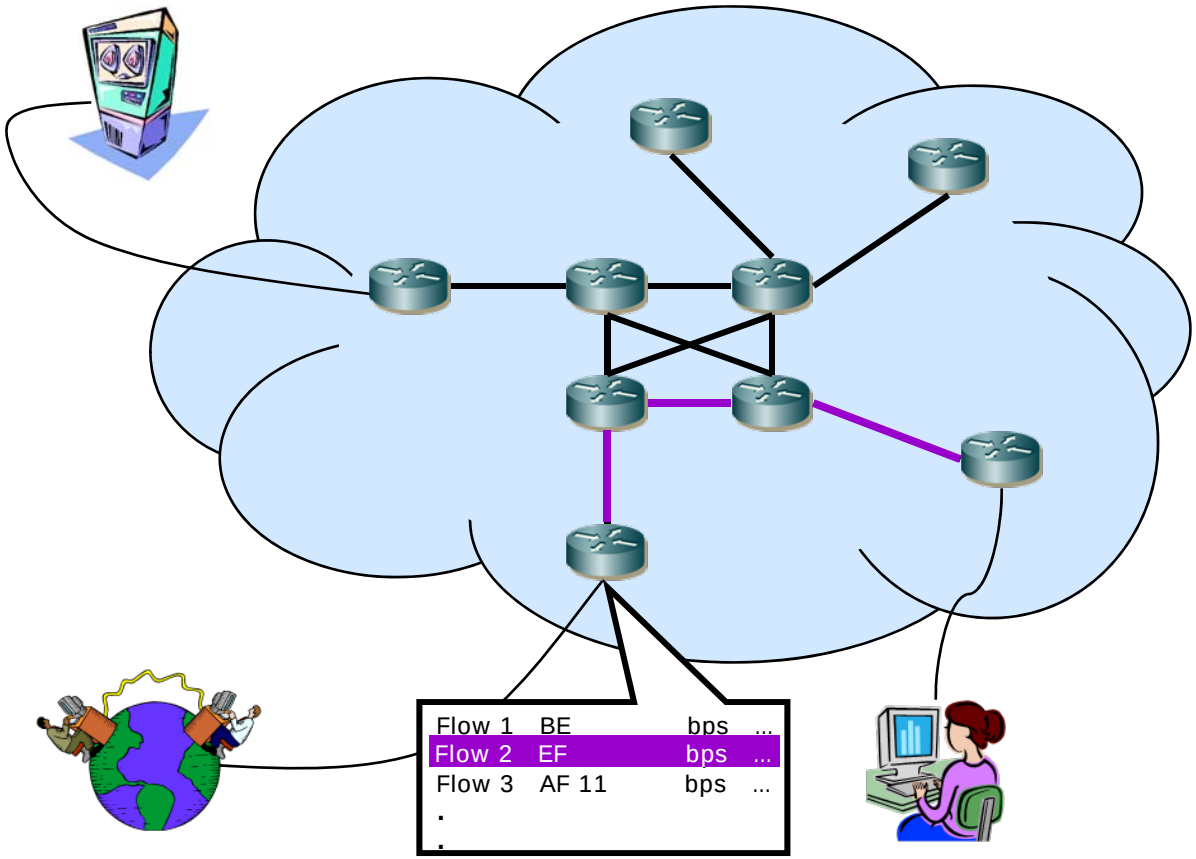

- Collects flow data at key network locations
- Computes traffic flows across topology

Route-Flow Fusion

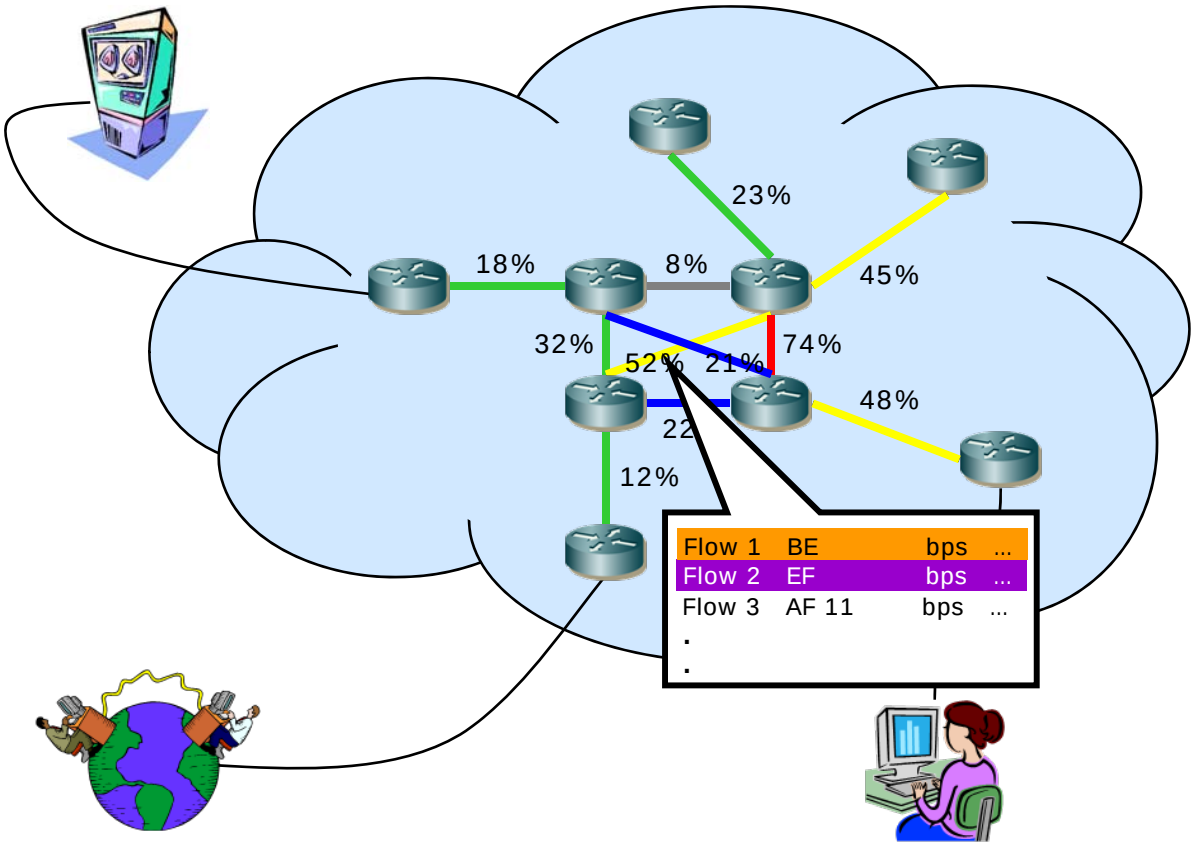

- Collects flow data at ***key*** network locations
- Computes traffic flows across topology
- Computes per CoS traffic levels for ***every link***
- Stores complete traffic and routing history for analysis, diagnostics, and planning

Route-Flow Fusion for MPLS/VPNs

- **Key locations are PE-P interfaces**
- **Use routing (IGP/BGP/LDP) to complete the path of the flow**

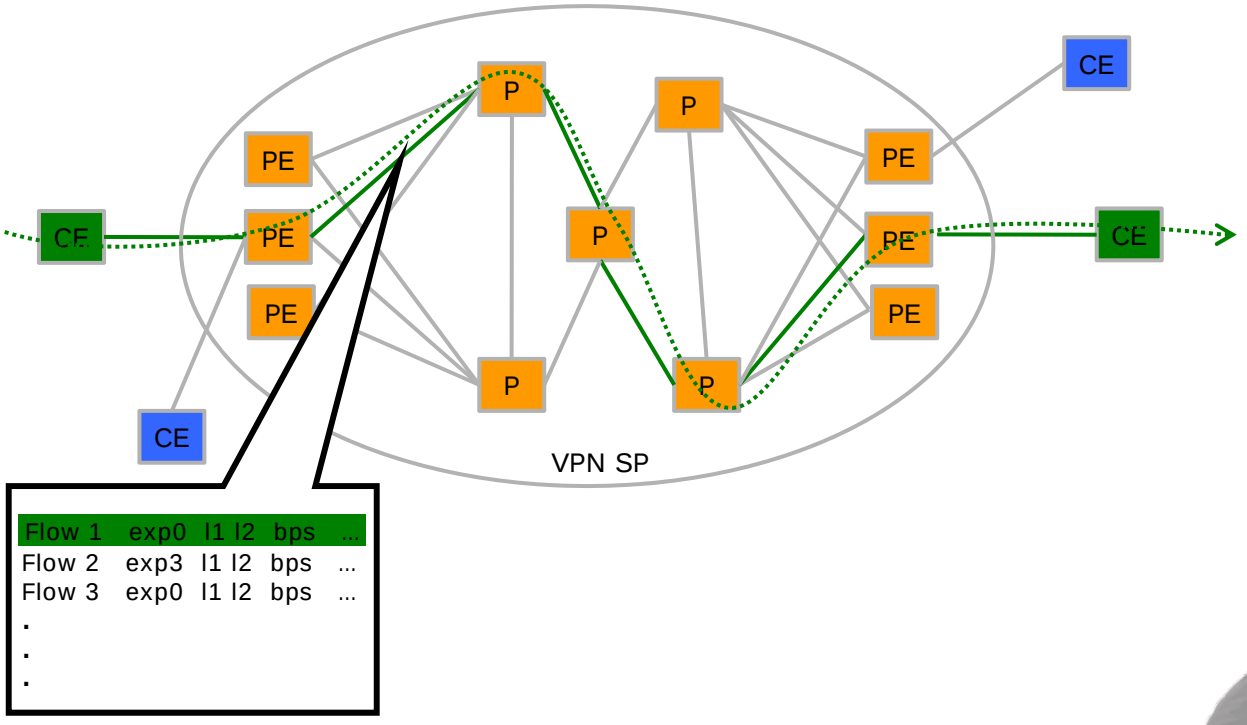

Route-Flow Fusion Details for MPLS VPNs

- **IPFIX flows**
  - MPLS labels, customer private source and destination addresses, traffic volume and timing, exp bits (CoS)
  - Deployed at P routers on PE facing interfaces
- **Label to FEC bindings via LDP from P routers**
  - Fuses LDP label in flow record to a egress PE
  - We now have flow + egress PE

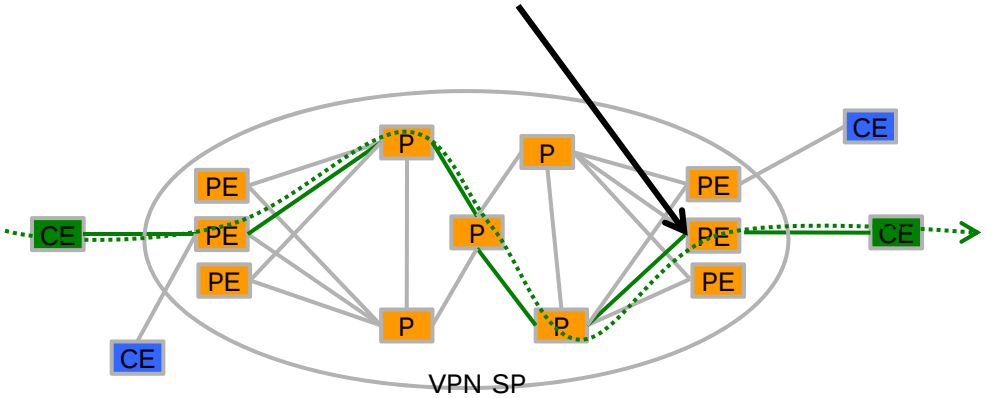

Route-Flow Fusion for MPLS VPNs

- **IGP routing**
  - We now have flow + egress PE + IGP path

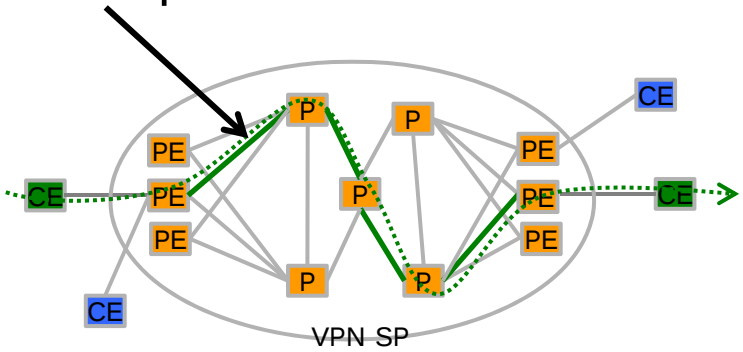

- **MP-BGP routing**
  - MP-BGP routes distribute the per VRF MPLS label which is fused with 2nd MPLS label in flow record
  - Prefix in that VRF gives us route targets, i.e. the customer
  - We now have ingress PE + egress PE + IGP path + VRF + customer + CoS + …

Resulting Fusion Creates Many Dimensional Flow Space

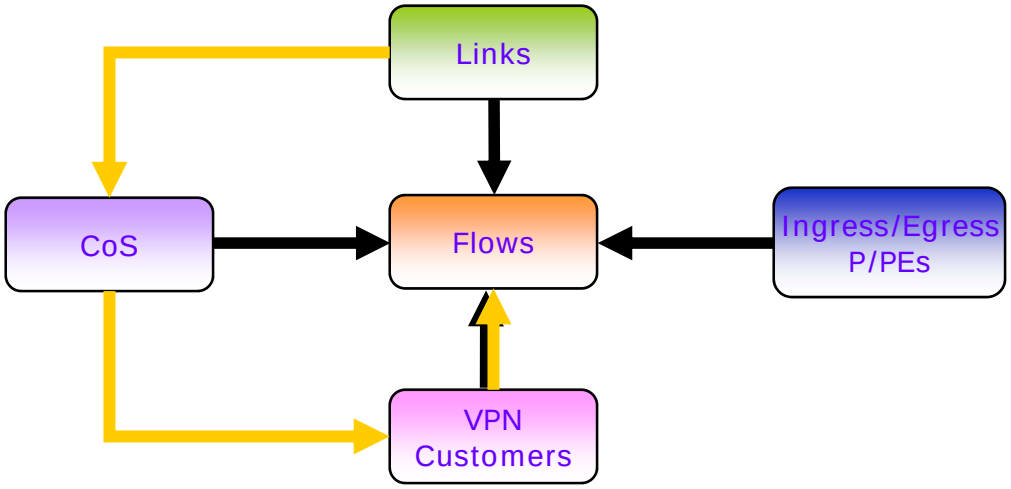

- **Further, fusion enables projection**
  - Pick one dimension and then drill into other dimensions

MPLS VPN Traffic on Core Links Visualized

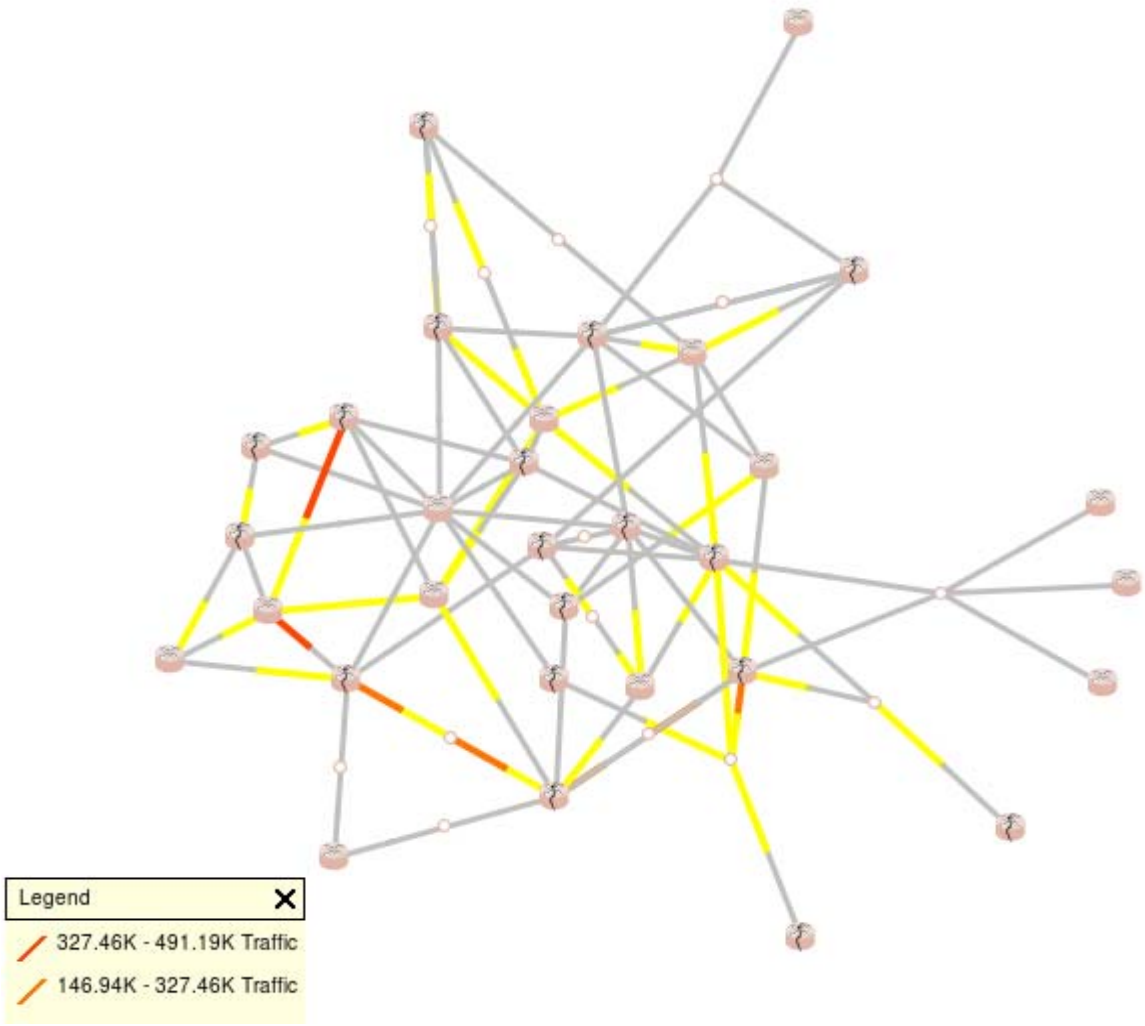

MPLS VPN Traffic on Core Links

Link

5 Minute | Compare | Drill Down

Source	Destination	5 Minute Averag	-1 day Average	-1 hour Average	-5 minutes Average
[redacted]	[redacted]	2.79M	1.37M	1.59M	2.42M
[redacted]	[redacted]	2.35M	1.97M	2.15M	1.95M
[redacted]	[redacted]	2.20M	1.81M	2.03M	1.81M
[redacted]	[redacted]	1.56M	302.21K	361.15K	534.27K
[redacted]	[redacted]	1.56M	302.21K	361.15K	534.27K
[redacted]	[redacted]	1.37M	1.17M	1.12M	1.24M
[redacted]	[redacted]	1.37M	1.17M	1.12M	1.24M
[redacted]	[redacted]	1.16M	857.69K	857.47K	1.05M
[redacted]	[redacted]	1.16M	857.69K	857.47K	1.05M
[redacted]	[redacted]	1.06M	853.34K	1.08M	1.71M
[redacted]	[redacted]	959.96K	1.15M	1.62M	1.13M
[redacted]	[redacted]	959.96K	1.15M	1.62M	1.13M
[redacted]	[redacted]	840.06K	554.78K	575.71K	737.59K
[redacted]	[redacted]	817.87K	614.26K	882.43K	561.39K
[redacted]	[redacted]	817.87K	614.26K	882.43K	561.39K
[redacted]	[redacted]	710.62K	514.46K	693.69K	769.63K
[redacted]	[redacted]	546.78K	467.42K	486.41K	472.12K
[redacted]	[redacted]	437.62K	174.98K	509.74K	547.14K
[redacted]	[redacted]	422.63K	353.00K	389.53K	414.00K
[redacted]	[redacted]	422.63K	353.00K	389.53K	414.00K

482 entries

History of MPLS VPN Traffic on a Link

Breakdown to Customers

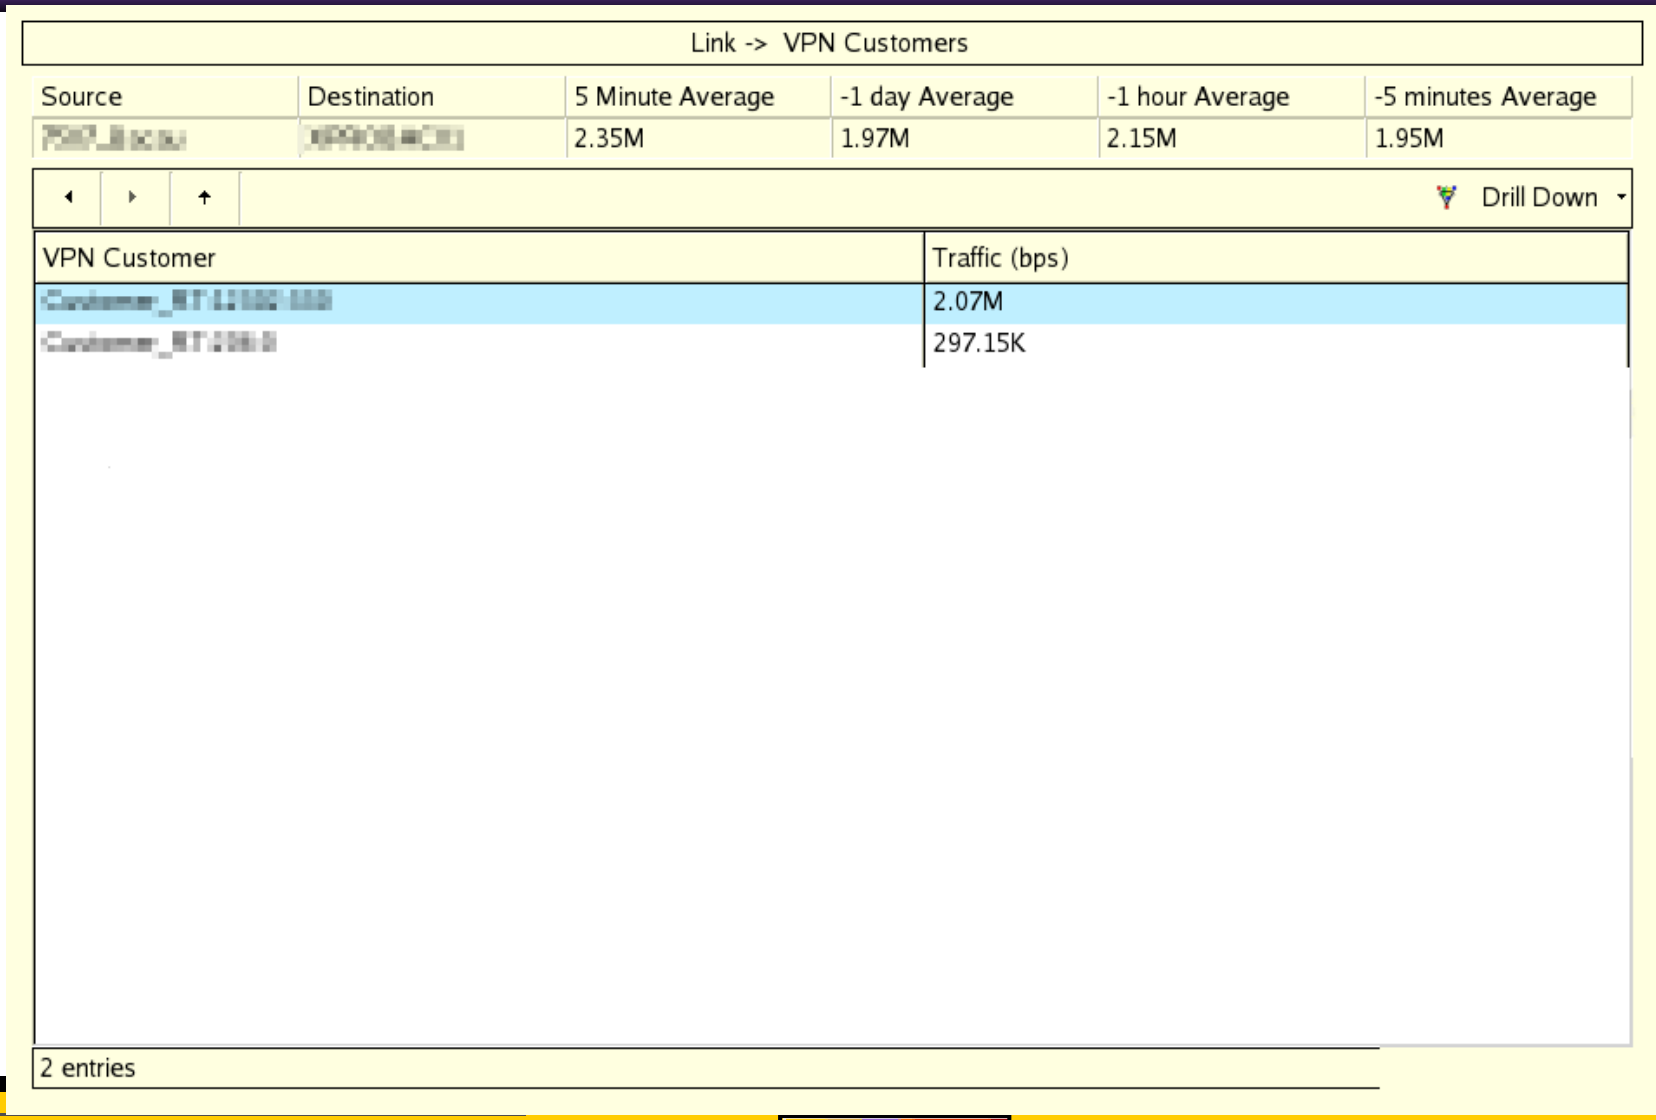

Breakdown to CoS

Link -> Customer -> Cos

Source	Destination	5 Minute Average	-1 day Average	-1 hour Average	-5 minutes Average
[redacted]	[redacted]	2.35M	1.97M	2.15M	1.95M

VPN Customer	Traffic (bps)
[redacted]	2.07M

Drill Down

CoS	Traffic (bps)
3	1.38M
0	361.50K
1	1.35K
4	64.79K
6	1.29K
5	30.23K
2	247.40K

7 entries

Modeling Changes with Route Analytics

- **Route Analytics simplifies modeling**
  - Works on routing model of the network
    - Processing routing messages changes the model
  - Mock-up a routing message to model a change
    - Computes new paths
    - Re-maps the flows to new links, customers, ASes, etc
  - Similarly, mock-up flow records to change traffic loads

Online Network Modeling

- **Make edits to ‘as-running’ network and ‘actual’ traffic loads**
  - Eliminate the need to build the model ground-up
  - Add new customers, new sites or new routing policy
  - Edit existing Customer profiles or PE-PE traffic matrix
- **Analyze the effect network-wide**
  - ‘Before and After’ comparison reports quickly highlight the effect of changes
  - Understand impact before maintenance, validate operation afterwards
  - Reduce configuration errors and avoid implementation surprises

Modeling a New Customer

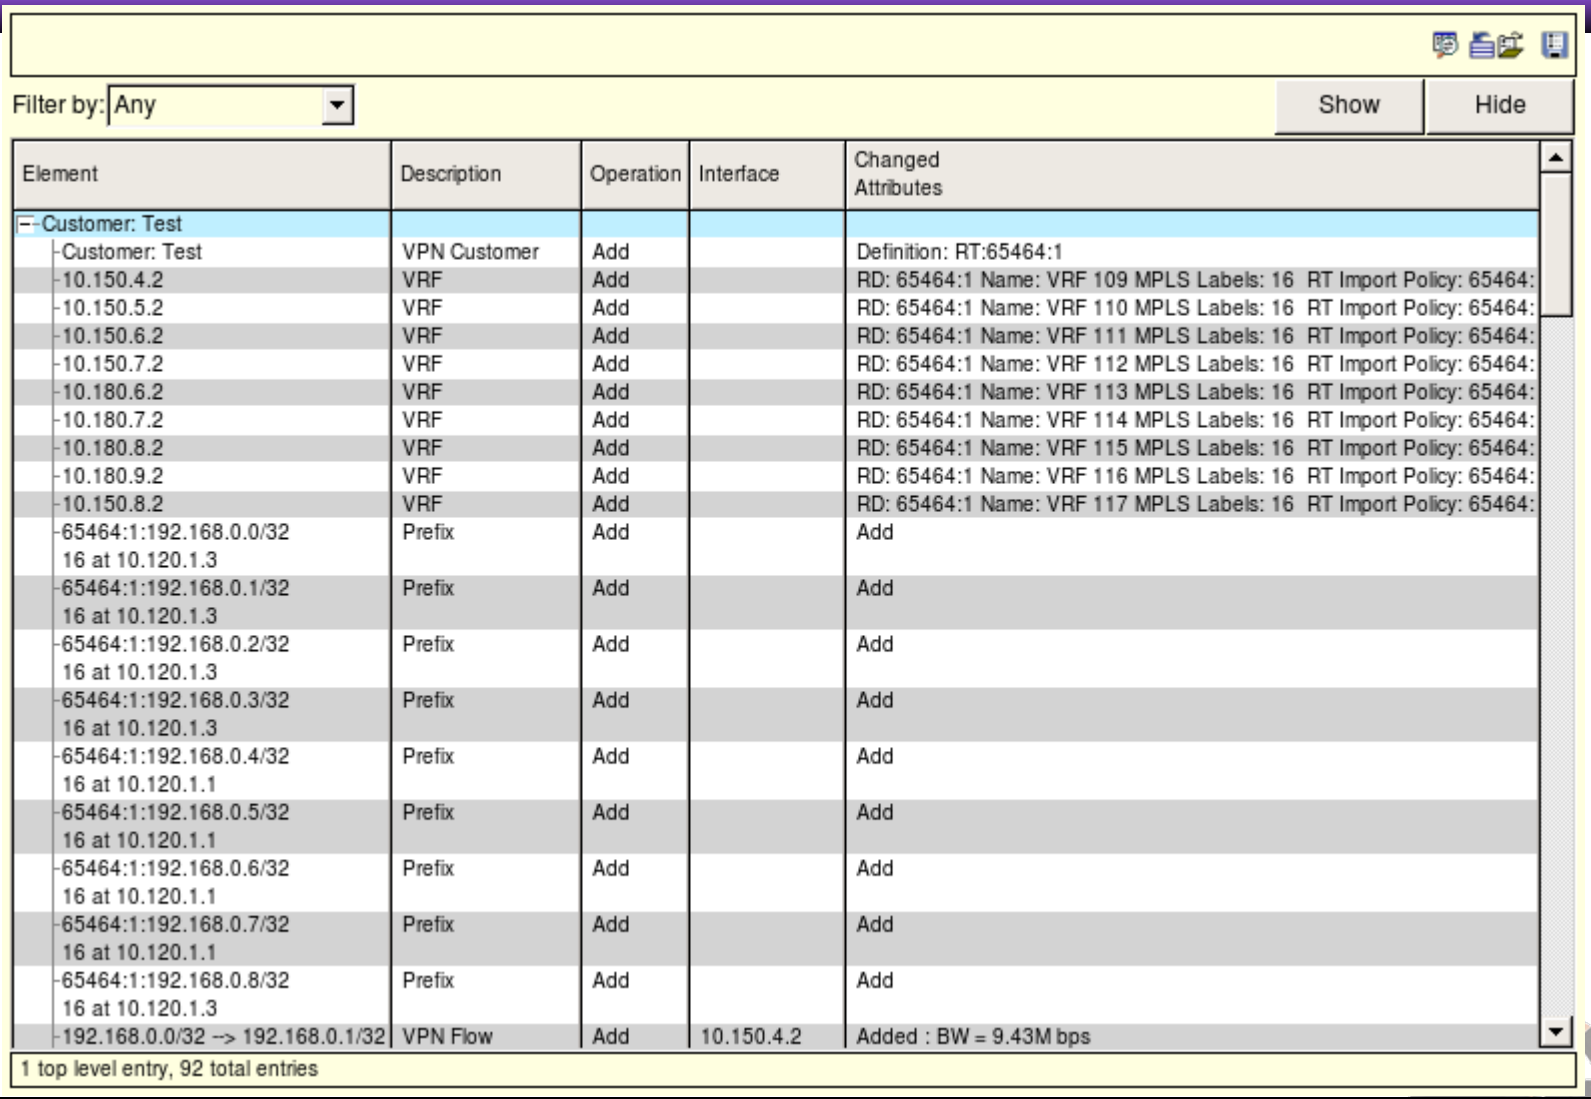

Filter by: Any

Show | Hide

Element	Description	Operation	Interface	Changed Attributes
Customer: Test				
Customer: Test	VPN Customer	Add		Definition: RT:65464:1
10.150.4.2	VRF	Add		RD: 65464:1 Name: VRF 109 MPLS Labels: 16 RT Import Policy: 65464:
10.150.5.2	VRF	Add		RD: 65464:1 Name: VRF 110 MPLS Labels: 16 RT Import Policy: 65464:
10.150.6.2	VRF	Add		RD: 65464:1 Name: VRF 111 MPLS Labels: 16 RT Import Policy: 65464:
10.150.7.2	VRF	Add		RD: 65464:1 Name: VRF 112 MPLS Labels: 16 RT Import Policy: 65464:
10.180.6.2	VRF	Add		RD: 65464:1 Name: VRF 113 MPLS Labels: 16 RT Import Policy: 65464:
10.180.7.2	VRF	Add		RD: 65464:1 Name: VRF 114 MPLS Labels: 16 RT Import Policy: 65464:
10.180.8.2	VRF	Add		RD: 65464:1 Name: VRF 115 MPLS Labels: 16 RT Import Policy: 65464:
10.180.9.2	VRF	Add		RD: 65464:1 Name: VRF 116 MPLS Labels: 16 RT Import Policy: 65464:
10.150.8.2	VRF	Add		RD: 65464:1 Name: VRF 117 MPLS Labels: 16 RT Import Policy: 65464:
65464:1:192.168.0.0/32 16 at 10.120.1.3	Prefix	Add		Add
65464:1:192.168.0.1/32 16 at 10.120.1.3	Prefix	Add		Add
65464:1:192.168.0.2/32 16 at 10.120.1.3	Prefix	Add		Add
65464:1:192.168.0.3/32 16 at 10.120.1.3	Prefix	Add		Add
65464:1:192.168.0.4/32 16 at 10.120.1.1	Prefix	Add		Add
65464:1:192.168.0.5/32 16 at 10.120.1.1	Prefix	Add		Add
65464:1:192.168.0.6/32 16 at 10.120.1.1	Prefix	Add		Add
65464:1:192.168.0.7/32 16 at 10.120.1.1	Prefix	Add		Add
65464:1:192.168.0.8/32 16 at 10.120.1.3	Prefix	Add		Add
192.168.0.0/32 --> 192.168.0.1/32	VPN Flow	Add	10.150.4.2	Added : BW = 9.43M bps

1 top level entry, 92 total entries

Effects as Before and After Reports

Link

Drill Down

Source	Destination	Capacity	Traffic Before Edit (bps)	Traffic After Edit (bps)	Traffic Change (bps)
DC-CORE2-ROUTER4	10.64.7.0/24	...	34.82M	322.14M	287.32M
10.64.7.0/24	LA-CORE-RTR5	...	34.82M	322.14M	287.32M
10.64.12.0/24	DC-CORE2-ROUTER4	...	34.82M	322.14M	287.32M
LA-CORE-RTR5	10.64.3.0/24	...	34.82M	322.14M	287.32M
10.64.3.0/24	SF-CORE-ROUTER2	...	34.82M	322.14M	287.32M
10.150.11.1	10.64.12.0/24	...	3.27M	290.59M	287.32M
SF-CORE-ROUTER2	10.64.14.0/24	...	3.27M	290.59M	287.32M
10.64.14.0/24	10.180.11.1	...	3.27M	290.59M	287.32M
10.150.4.2	10.150.11.1	...	331.72K	201.84M	201.51M
10.180.11.1	10.180.9.2	...	330.61K	190.84M	190.51M
10.180.11.1	10.64.14.0/24	...	3.10M	164.35M	161.25M
10.64.14.0/24	SF-CORE-ROUTER2	...	3.10M	164.35M	161.25M
10.64.12.0/24	10.150.11.1	...	3.10M	164.35M	161.25M
SF-CORE-ROUTER2	10.64.3.0/24	...	27.25M	188.50M	161.25M
10.64.3.0/24	LA-CORE-RTR5	...	27.25M	188.50M	161.25M
LA-CORE-RTR5	10.64.7.0/24	...	27.25M	188.50M	161.25M
10.64.7.0/24	DC-CORE2-ROUTER4	...	27.25M	188.50M	161.25M
DC-CORE2-ROUTER4	10.64.12.0/24	...	19.75M	181.00M	161.25M
10.150.8.2	10.150.11.1	...	333.79K	111.74M	111.40M
10.150.11.1	10.150.6.2	...	315.21K	96.01M	95.69M
10.150.11.1	10.150.4.2	...	314.38K	92.48M	92.16M
10.150.5.2	10.150.11.1	...	332.30K	74.10M	73.77M
10.150.11.1	10.150.7.2	...	314.04K	70.35M	70.04M
10.180.11.1	10.180.6.2	...	331.20K	67.98M	67.65M
10.180.9.2	10.180.11.1	...	301.05K	67.29M	66.99M
10.150.11.1	10.150.8.2	...	301.82K	60.68M	60.38M
10.150.6.2	10.150.11.1	...	331.20K	60.15M	59.82M
10.180.8.2	10.180.11.1	...	301.82K	59.81M	59.51M
10.180.11.1	10.180.8.2	...	333.79K	58.50M	58.17M
10.180.6.2	10.180.11.1	...	315.21K	53.76M	53.44M

86 entries

Summary

- **Route analytics provides a novel approach to MPLS VPN management and service assurance**
  - Leverages rich data in routing
- **Network-Wide MPLS-VPN aware visibility into SP's core**
  - SP's aggregate & CoS view of the network
  - VPN customer view of the network
  - Detailed drilldown reports to troubleshoot a problem and find the root cause
- **Allow SP to proactively plan for capacity and design their network**
  - Real network, real traffic and real impact!
  - Model impact of local changes in traffic and routing on global network performance