



OpenSource BGP Networking with ODL

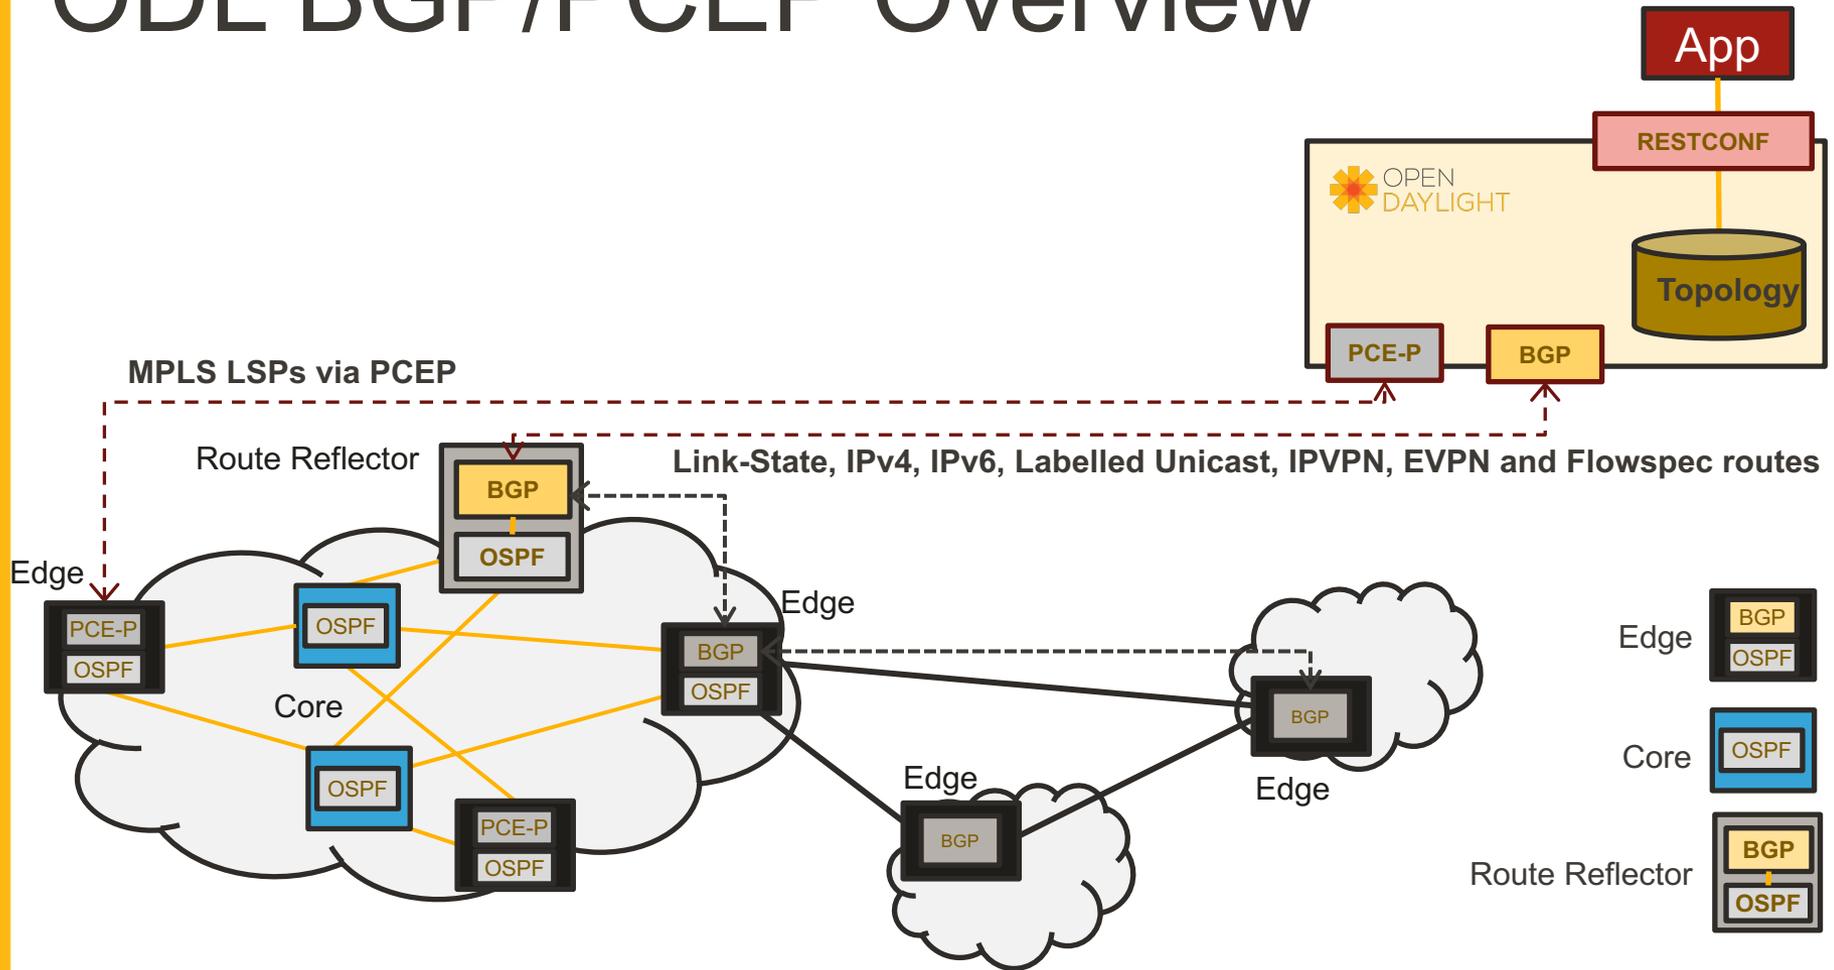
FOSDEM, Brussels - Feb 2017

Giles Heron

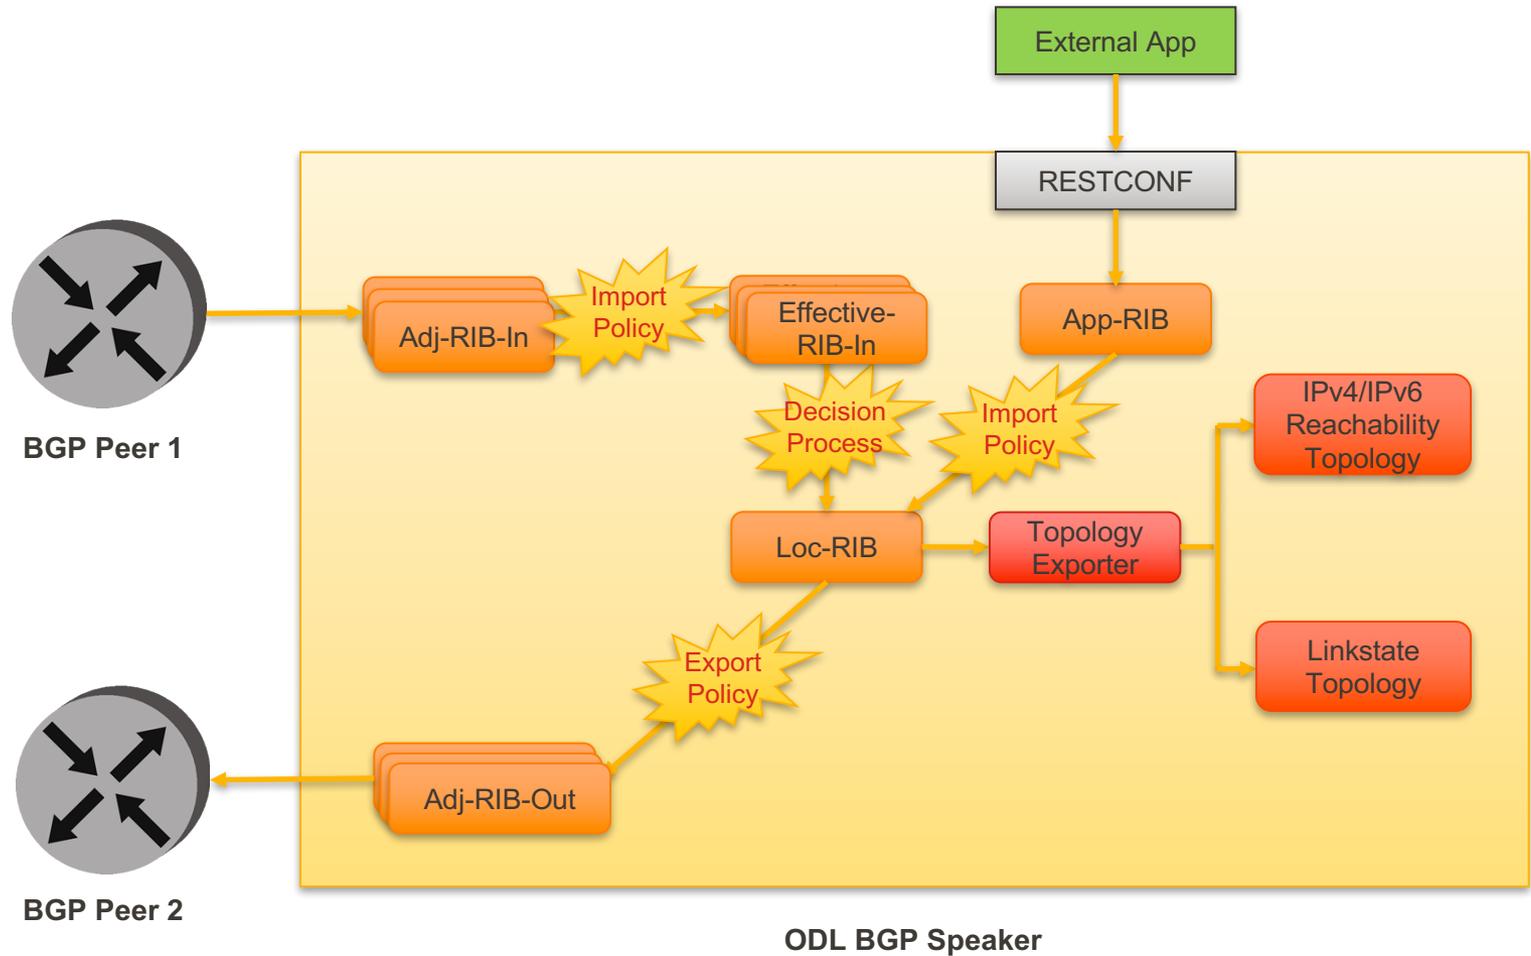
Agenda

- OpenDaylight BGP/PCEP Overview
- Use Cases
 1. Traffic Optimisation using BGP-LS and PCE-P
 2. DDoS Mitigation using BGP FlowSpec
 3. CDN Localisation using OpenConfig BGP YANG Models

ODL BGP/PCEP Overview



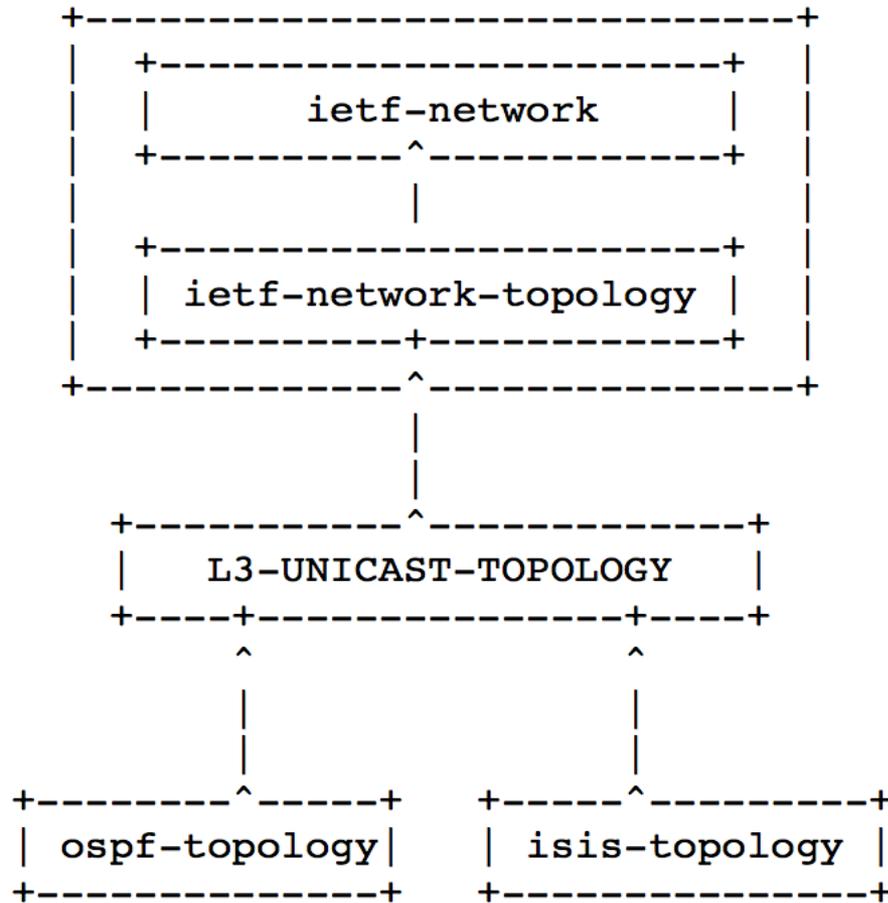
Processing of BGP Routes



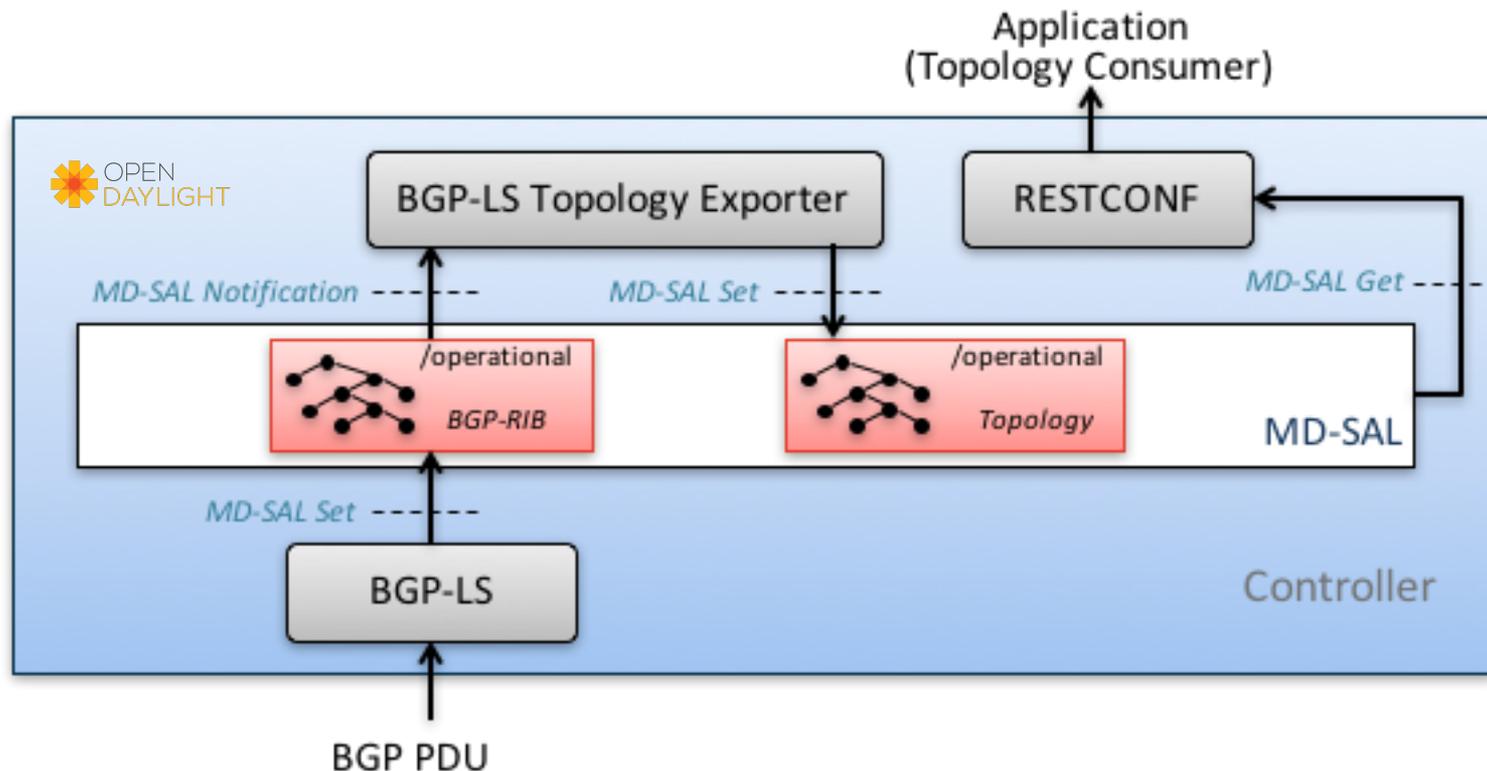
The I2RS Topology Model

- Best Approach for Network-Layer Abstraction
- Simple list of nodes / list of links
 - Finally Progressing to RFC
 - Has “termination points” to disambiguate interfaces
 - Supports topology layering
- Base model can be used for different protocols
 - Map protocol-specific info into node-id, link-id, tp-id
 - Augment models for additional protocol-specific info
 - Enables protocol-independence (as per requirements)

Topology Model Augmentation - Example

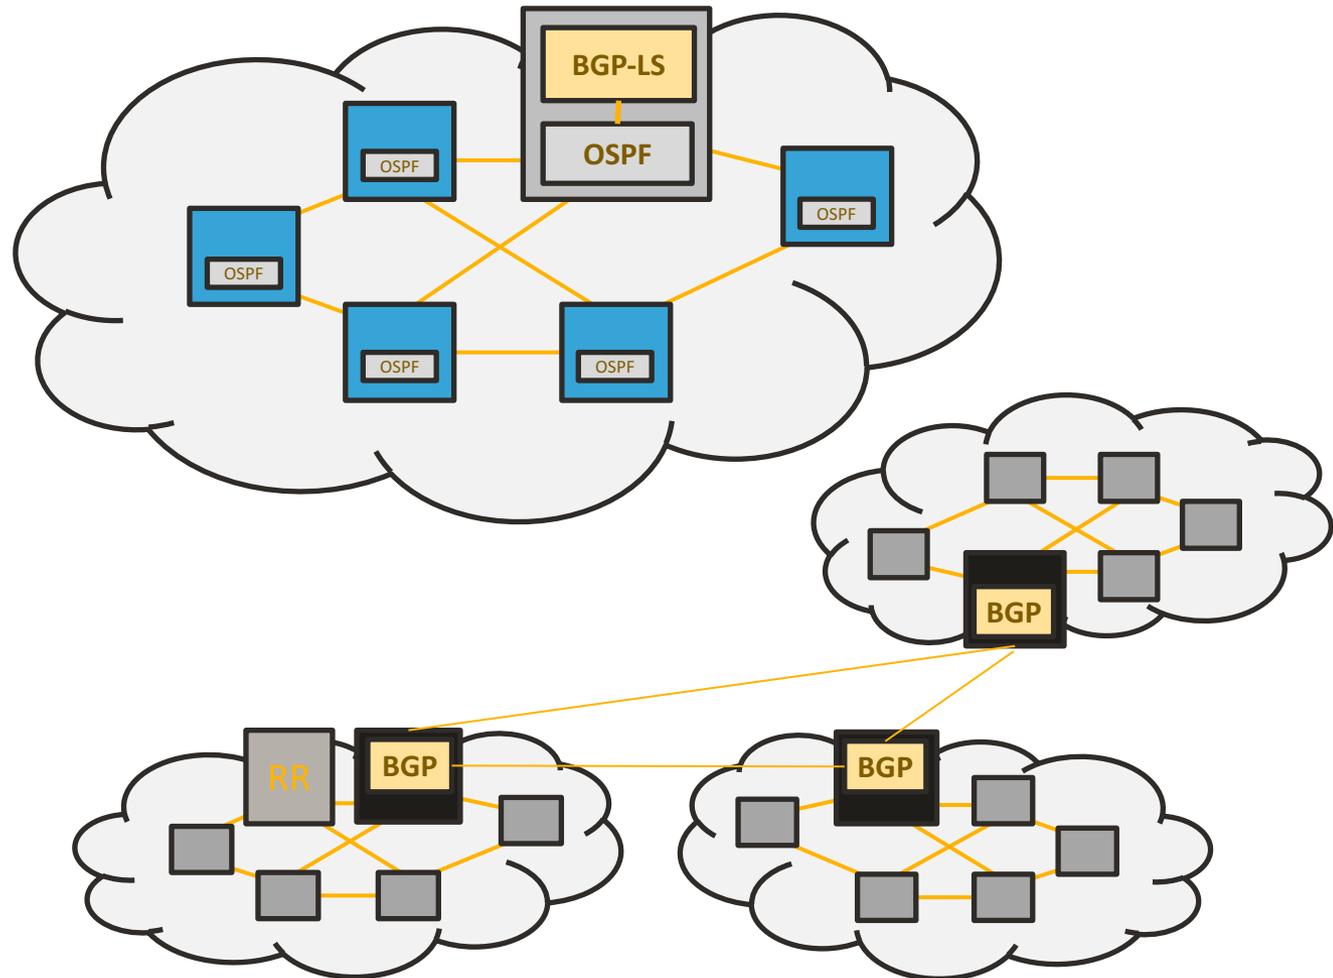


OSL BGP-LS Topology Exporter



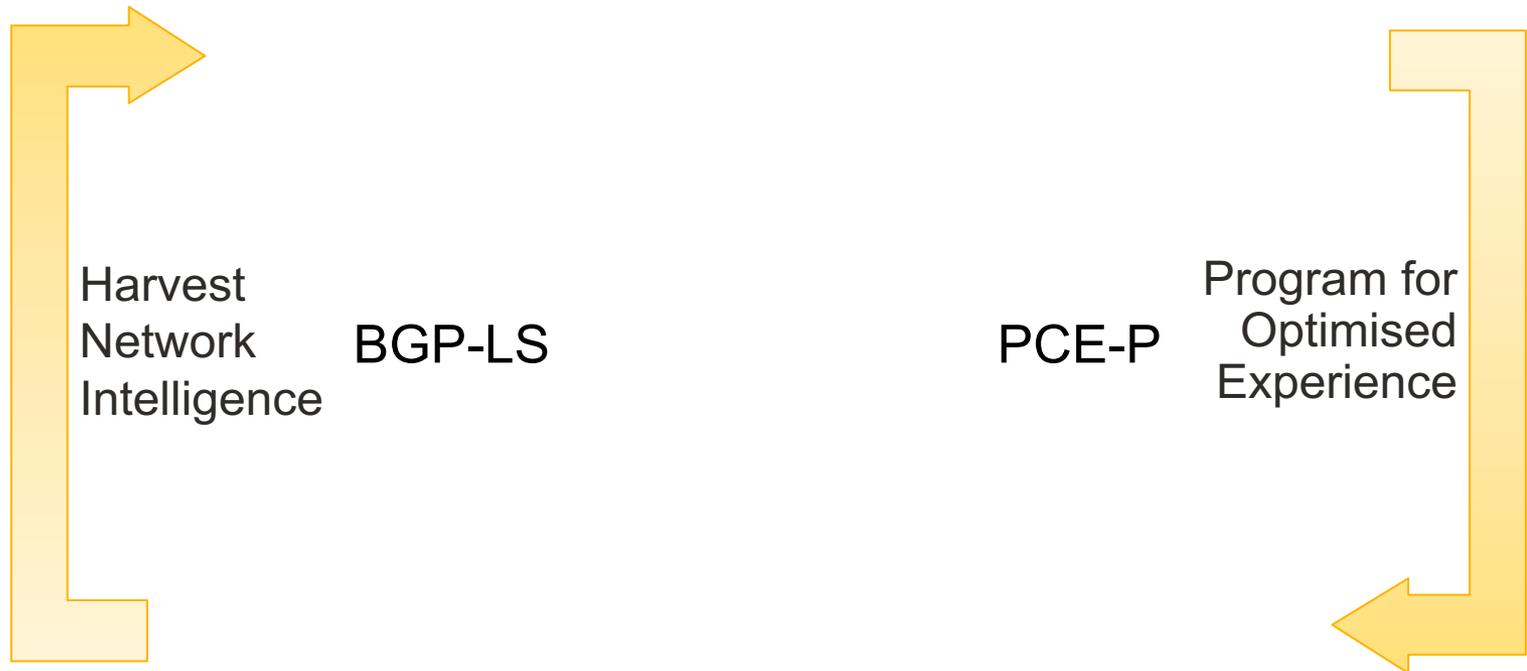
BGP/PCEP Topologies

- Link-State
- IPv4
- IPv6
- PCE-P



Traffic Optimisation using BGP-LS and PCE-P

BGP-LS and PCE-P



DDoS Mitigation using BGP FlowSpec

BGP FlowSpec - Overview

- RFC 5575
- Similar to OpenFlow but uses BGP to distribute match/action rules

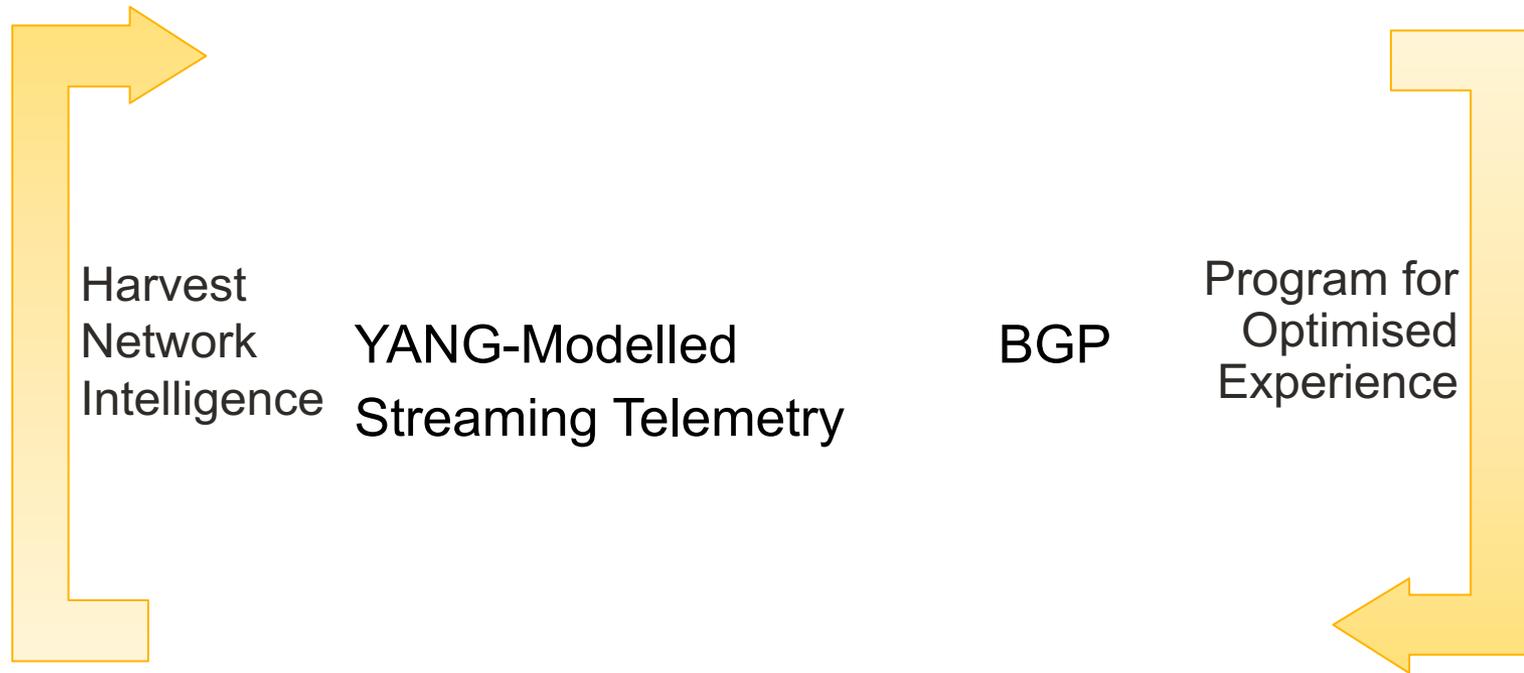
Matches:

- Source / Destination IP prefix
- IP Protocol
- Source / Destination TCP/UDP port
- ICMP Type / Code
- TCP Flags
- Packet Length
- DSCP Field
- Fragment (DF, IsF, FF, LF)

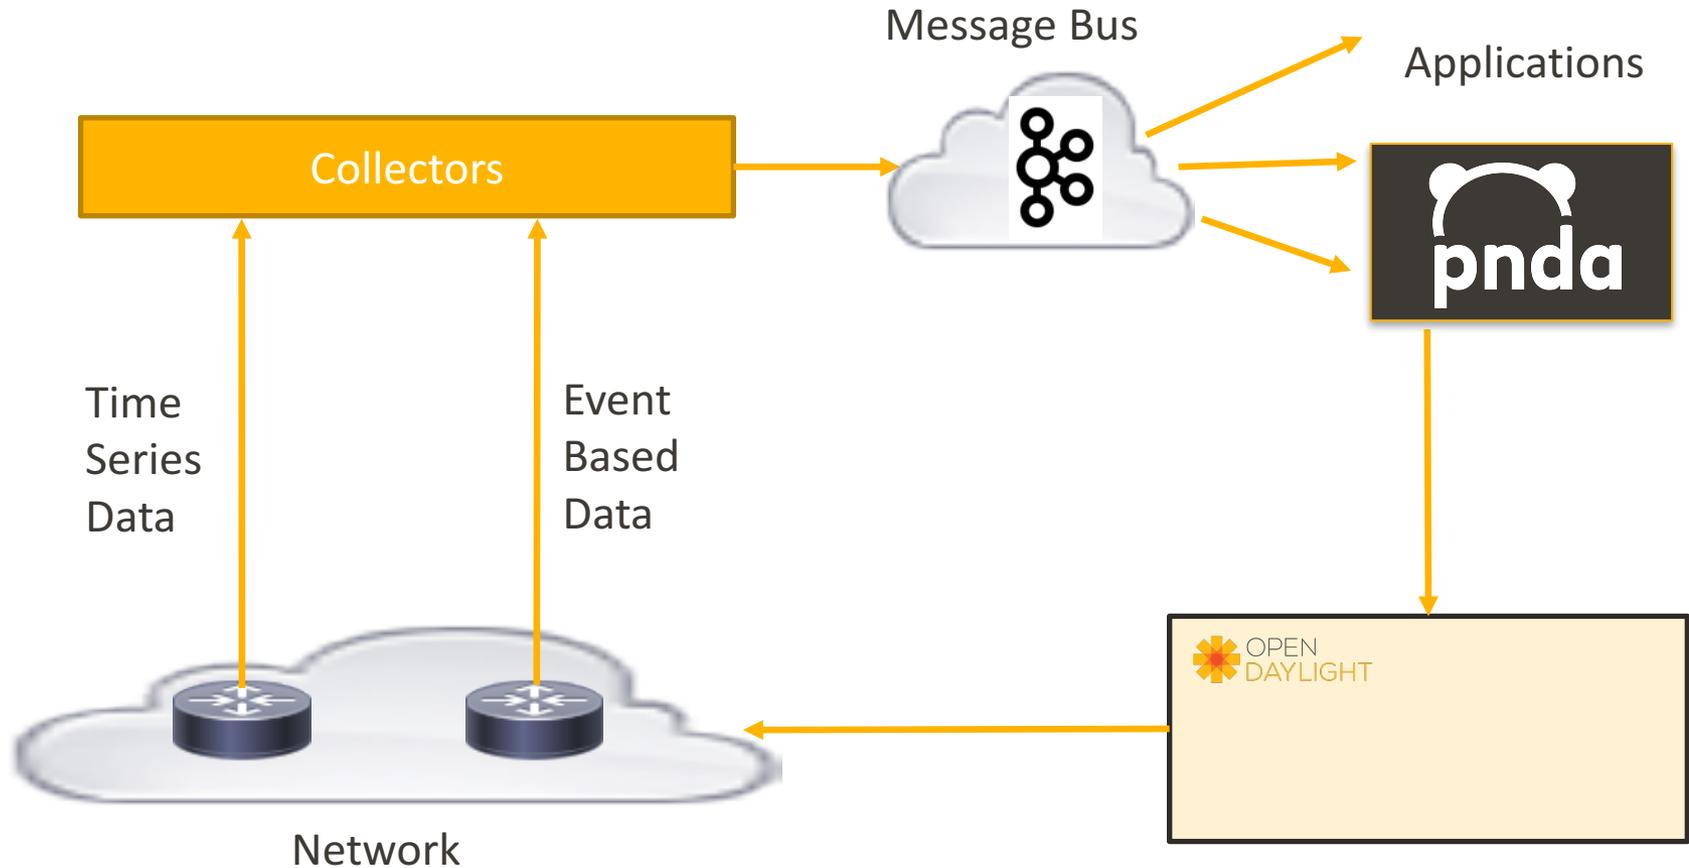
Actions:

- Rate limit
- Traffic sampling
- Redirection
- Traffic marking (DSCP)
- And more... (optional)

DDoS Mitigation

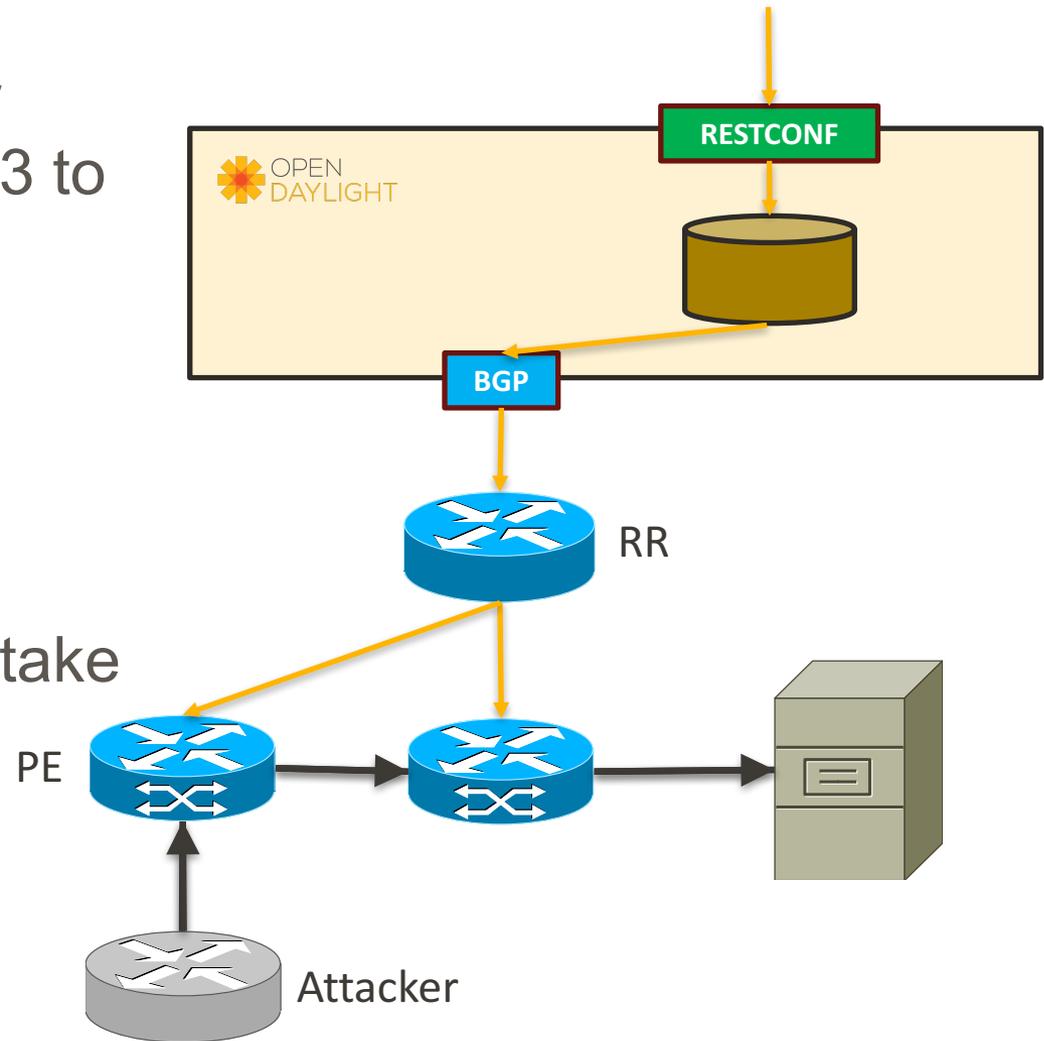


Telemetry – Close the Loop with ODL



DDoS Mitigation using FlowSpec

- Detect attacking flow (e.g. DDoS on port 53 to DNS from certain prefixes or ASNs)
- Advertise FlowSpec route via RR
- All routers learn FlowSpec route and take action



CDN Localisation using BMP/BGP with OpenConfig YANG models

What Problem are we Solving (overview)

- Broadband subscribers watch a lot of "Over-The Top" video
(In US typically 1/3 of traffic towards the subscriber is Netflix, 1/3 is 'other' OTTs, and 1/3 is everything else)
- The cost of carrying all this video across the network is prohibitive to the ISPs
- If video content can be cached closer to the user costs will drop
- But how do we figure out which video cache is closest to the user?

What Problem are we Solving (detail)

- OTT players (e.g. Netflix, Akamai, Google) have built out global CDNs
 - Connect at IXPs
 - Also deploying caches into ISP networks
- The CDN needs to be able to direct a user to the nearest available cache containing the relevant content
- Problem can be broken down into (roughly)
 1. Mapping a user to a prefix
 2. Identifying the correct cache to serve the prefix (combination of content availability, cache load, and **proximity of prefix to cache**)
- The goal of our solution will be to assist SPs/OTTs in solving problem 2.

How to Address the Problem

1. Mapping User to a Prefix

- Trivial for Netflix as their portal knows the user IP address
- For e.g. Akamai use DNS Client Subnet
- (this is out of scope for the solution we're building here)

2. Determining proximity of the Prefix to the Cache

- BGP communities (peering can be centralised or distributed)
- BGP MEDs (e.g. Netflix – requires distributed peering)
- Other?

	Communities (centralised BGP)	Communities (distributed BGP)	MEDs Distributed BGP
Portal			NETFLIX
DNS Client Subnet			

Determining proximity of a prefix to a cache: BGP Communities vs MEDs

Cache List:

Cache B - Site 2001

Cache C - Site 1001

Cache E - Site 101

User List:

Pool A'/25 - Site 3001, 2001, 2002, 1001, 1002, 1003, 101, 102

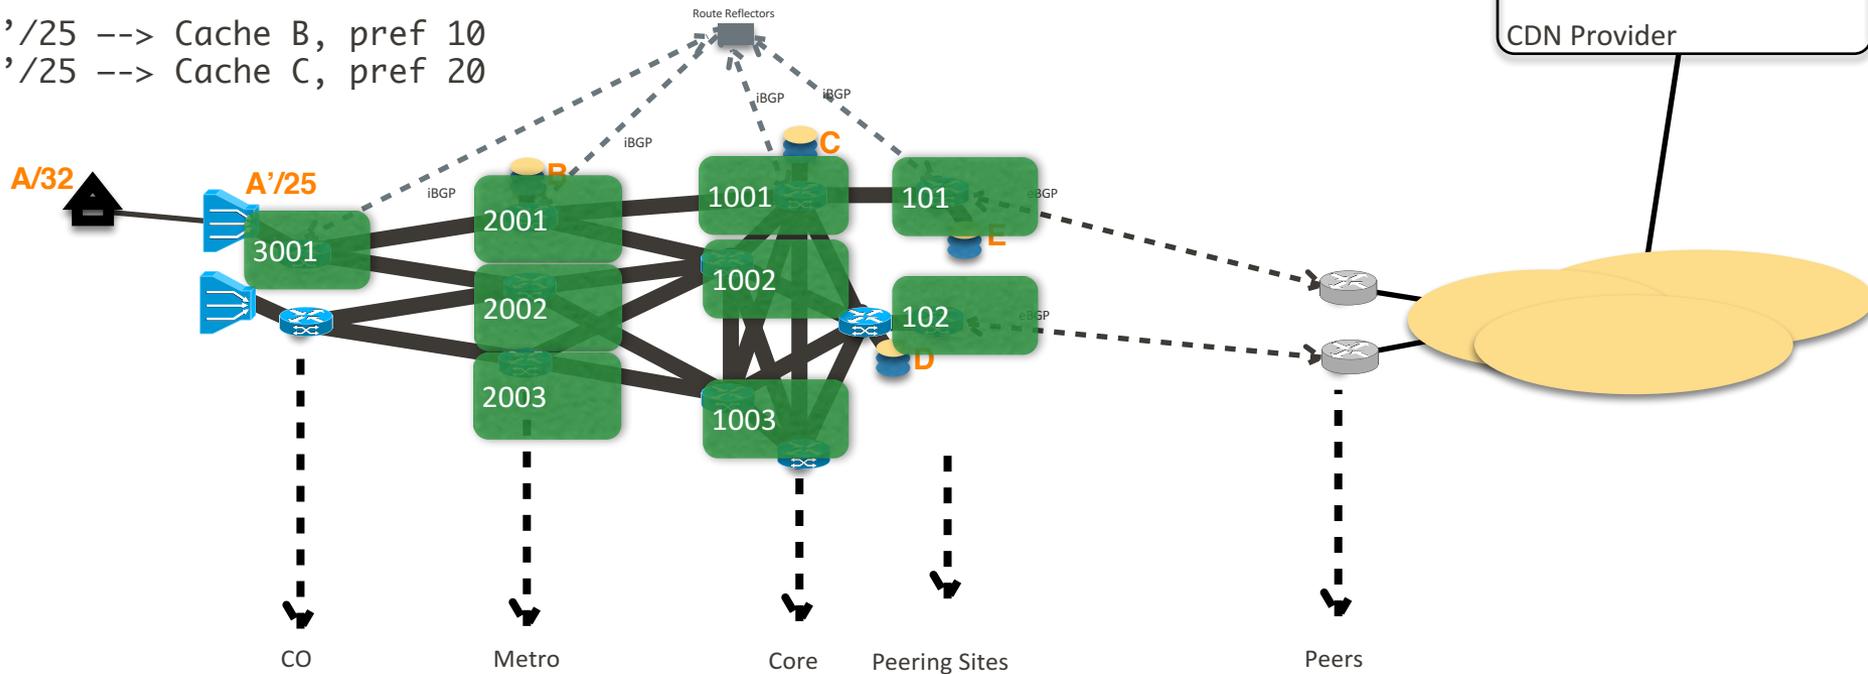
A'/25 --> Cache B, pref 10

A'/25 --> Cache C, pref 20



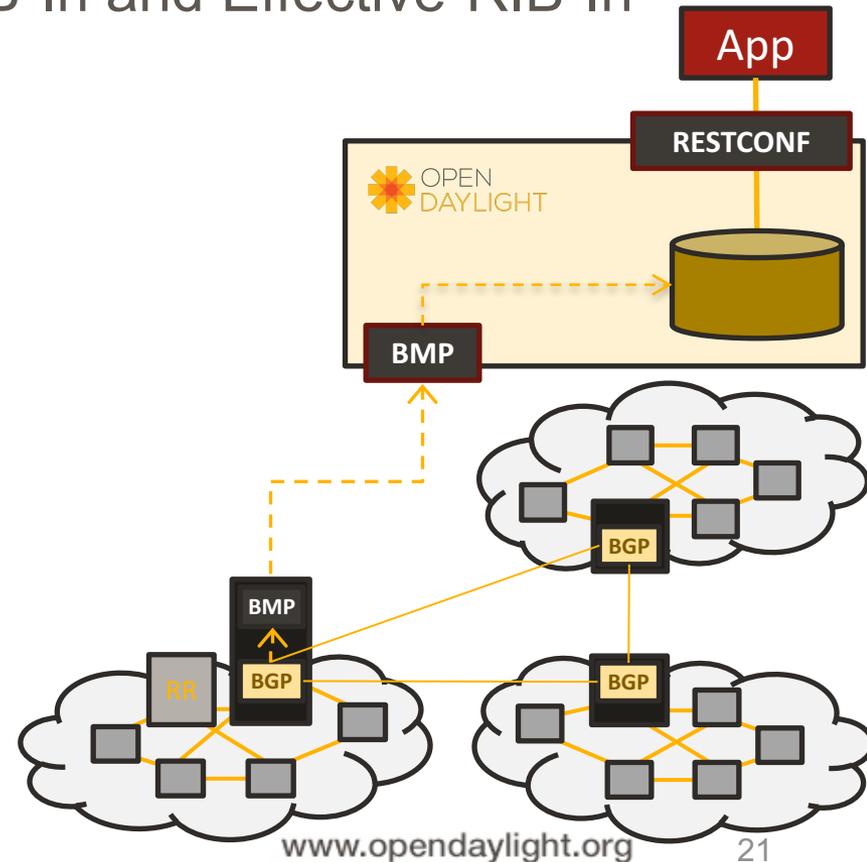
How to build these Community Lists??

One option would be to use link-state topology



BMP (BGP Monitoring Protocol)

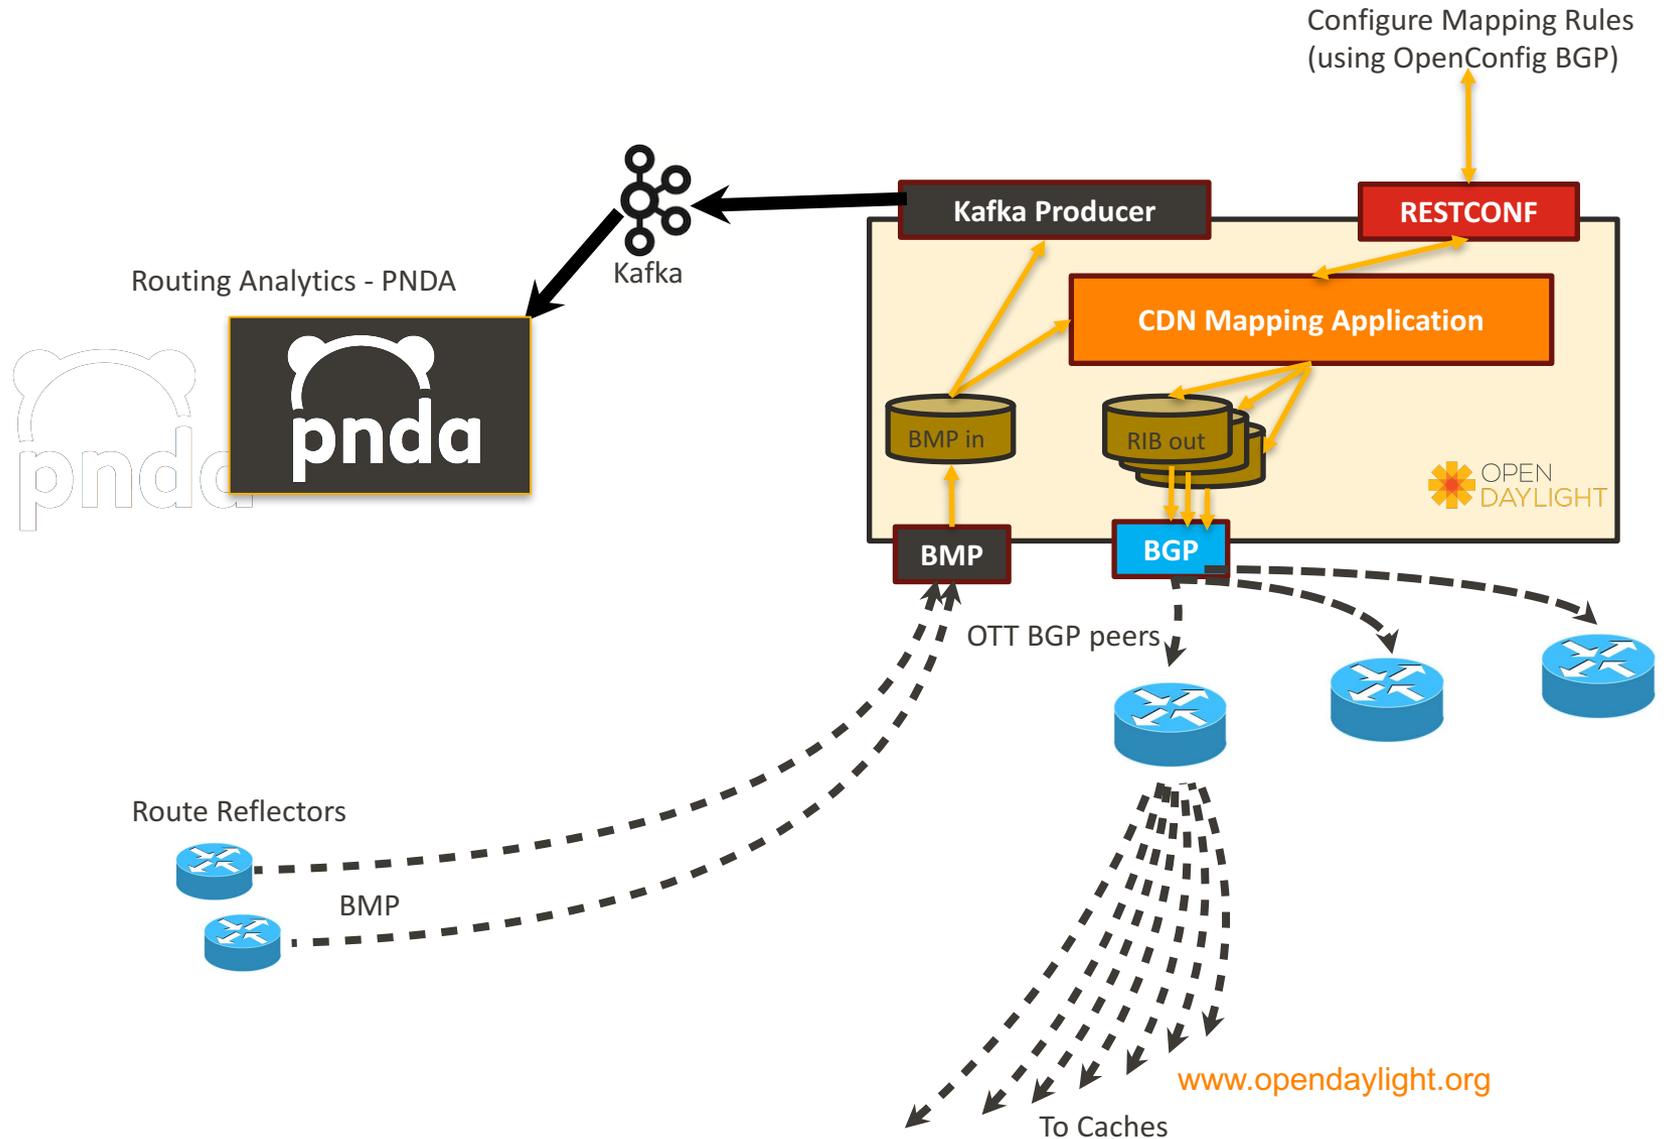
- RFC 7854
- Enables access to peer's Adj-RIB-In and Effective-RIB-In
- You can get a LOT of prefixes:
 1. all peer/upstream routes at IXP
 2. all routes from Route Reflector



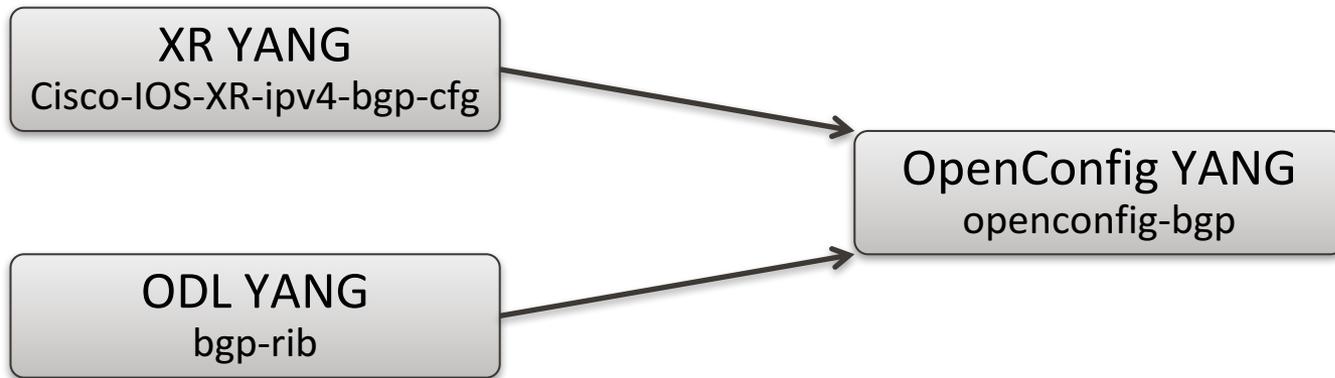
CDN Localisation



Potential Architecture – ODL App



How to configure the mapping? OpenConfig BGP YANG model!



Links

- <https://github.com/CiscoDevNet/opendaylight-setup>
- <https://github.com/CiscoDevNet/pathman-sr>
- <https://www.youtube.com/watch?v=-20bklzKcZ4>

Thanks