
TOMORROW starts here.
cisco
Cisco live!

Advances in BGP

BRKRST-3371

Oliver Boehmer
AS Solutions Architect

Cisco live!

What is BGP?

- What a Google search “bgp abbreviation” finds”
  - Source: http://www.all-acronyms.com/BGP

Border Gateway Protocol
Bacterial Growth Potential
Battlegroup
Becker, Green and Pearson
<sensored entry>
Bermuda grass pollen
Berri Gas Plant
beta-glycerophosphate
biliary glycoprotein
blood group
bone gamma-carboxyglutamic acid protei…
bone gamma-carboxyglutamic acid-contai…
bone gla protein
bone Gla-containing protein
Borders Group, Inc.
brain-type glycogen phosphorylase
Bridge Gateway Protocol
Broader Gateway Protocol
Bureau de Gestion de Projet
Brain Gain Program

- Without BGP the Internet would not exist in its current stable and simple form
- It is the plumbing technology of the Internet

What is BGP? – What it Truly is?

The

Bloody Good Protocol

Agenda

- Motivation to Enhance BGP
- Scale and Performance Enhancements
- What happened in BGP Landscape?
- Some new cool features that may interest you

Agenda

- Motivation to Enhance BGP
- Scale and Performance Enhancements
- What happened in BGP Landscape?
- Some new cool features that may interest you

BGP Started in 1989…

- ***Motivation and Development of BGP:*** When the Internet grew and moved to an autonomous system (AS) mesh architecture it was needed to have stable, non-chatty and low CPU consuming protocol to connect all of these AS's together.
- In June 1989, the first version of this new routing protocol was formalised, with the publishing of RFC 1105, *A Border Gateway Protocol (BGP)*.

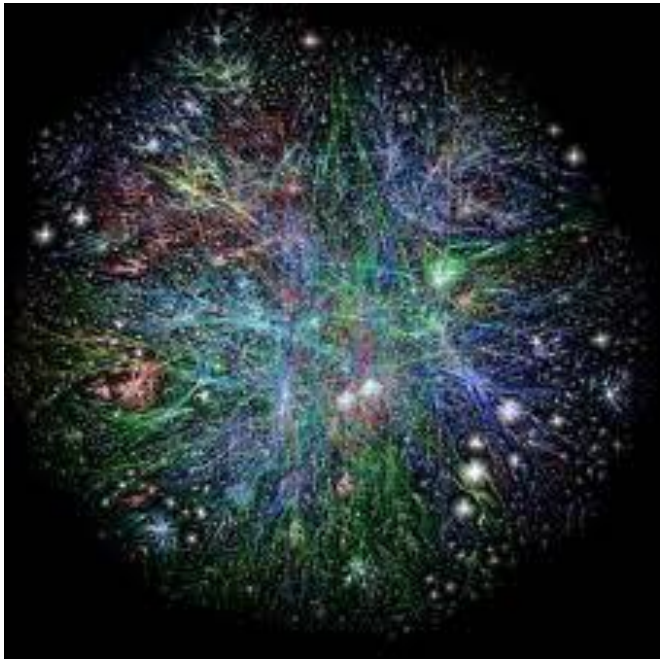

Service Provider Routing and Services Progress

- Multimedia, Mobile Internet and Cloud Services will generate massive bandwidth explosion
- Prefix growth is almost a linear curve
- Evolution of offered BGP services go from basic technologies to very advanced infrastructures

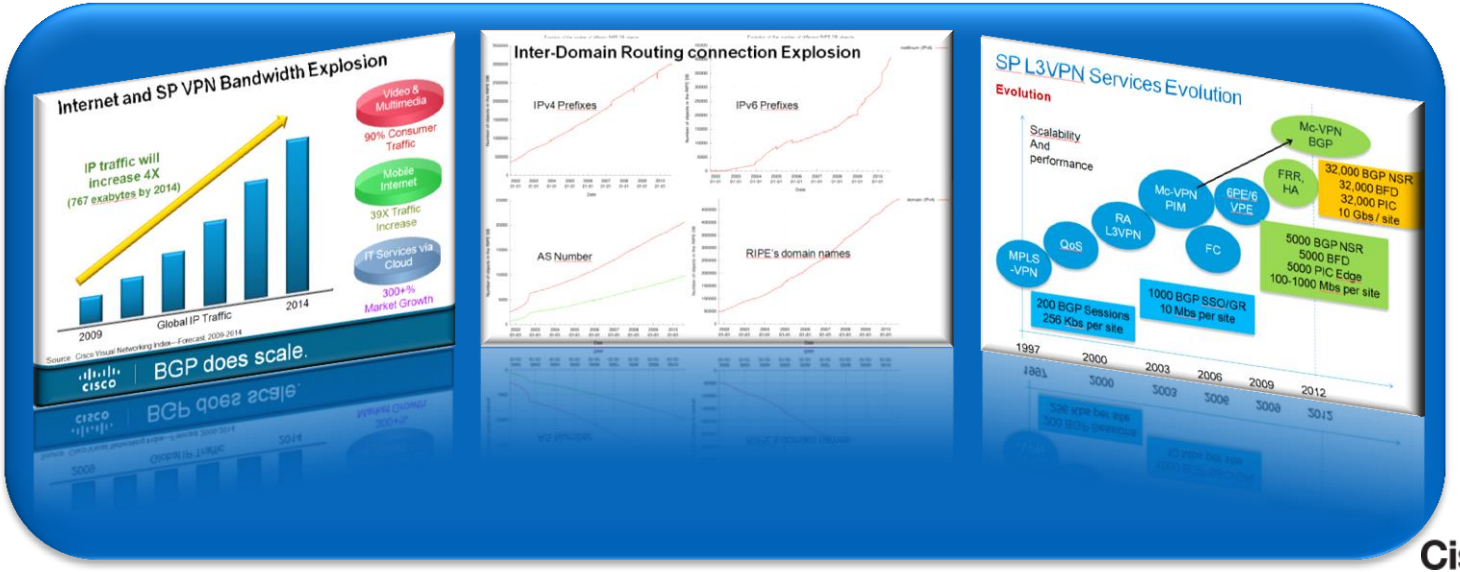

Control-plane Evolution

Most of services are moving towards BGP

Service/transport	200x and before	2013 and future
IDR (Peering)	BGP	BGP (IPv6)
SP L3VPN	BGP	BGP + FRR + Scalability
SP Multicast VPN	PIM	BGP Multicast VPN
DDOS mitigation	CLI	BGP flowspec
Network Monitoring	SNMP	BGP monitoring protocol
Security	Filters	BGP Sec (RPKI), DDoS Mitigation
Proximity		BGP connected app API
SP-L3VPN-DC		BGP Inter-AS, VPN4DC
Business & CE L2VPN	LDP	BGP PW Sign (VPLS)
DC Interconnect L2VPN		BGP MAC Sign (EVPN)
MPLS transport	LDP	BGP+Label (Unified MPLS)
Data Centre	OSPF/ISIS	BGP + Multipath
Massive Scale DMVPN	NHRP / EIGRP	BGP + Path Diversity
Campus/Ent L3VPN	BGP	BGP

Agenda

- Motivation to Enhance BGP
- Scale and Performance Enhancements
- What happened in BGP Landscape?
- Some new cool features that may interest you

Scale & Performance Enhancements

BGP Scaling

- Update Generation Enhancements
  - Update generation is the most important, time-critical task
  - Is now a separate process, to provide more CPU Quantum
- Parallel Route Refresh
  - Significant delay (up to 15-30 minutes) seen in advertising incremental updates while RR is servicing route refresh requests or converging newly established peers
  - Refresh and incremental updates run in parallel
- Keepalive Enhancements
  - Loosing or delayed keep-alive message result in session flaps
  - Hence keep-alive processing is now placed into a separate process using priority queuing mechanism

Scale & Performance Enhancements (contd.)

BGP Scaling

- Adaptive Update Cache Size
  - Instead of using a fixed cache size, the new code dynamically adapts to the address family used, the available router memory and the number of peers in an update group
  - Loosing or delayed keep-alive message result in session flaps
- Route Reflector Scaling - Selective RIB Download
  - A Route-Reflector needs to receive the full RIB, however not all prefixes MUST be in the Forwarding Information Base (FIB)
  - So, we now allow by using user policy to only download selected prefixes in the FIB

Scale & Performance Enhancements (contd.)

Slow Peer Management

Issue: Slow peers in update groups block convergence of other update group members by filling message queues/transmitting slowly

Persistent network issue affecting all BGP routers

- Two components to solution
  - Detection
  - Protection
- Detection
  - BGP update timestamps
  - Peer's TCP connection characteristics

Agenda

- Motivation to Enhance BGP
- Scale and Performance Enhancements
- What happened in BGP Landscape?
- Some new cool features that may interest you

What Happened in XR Landscape?

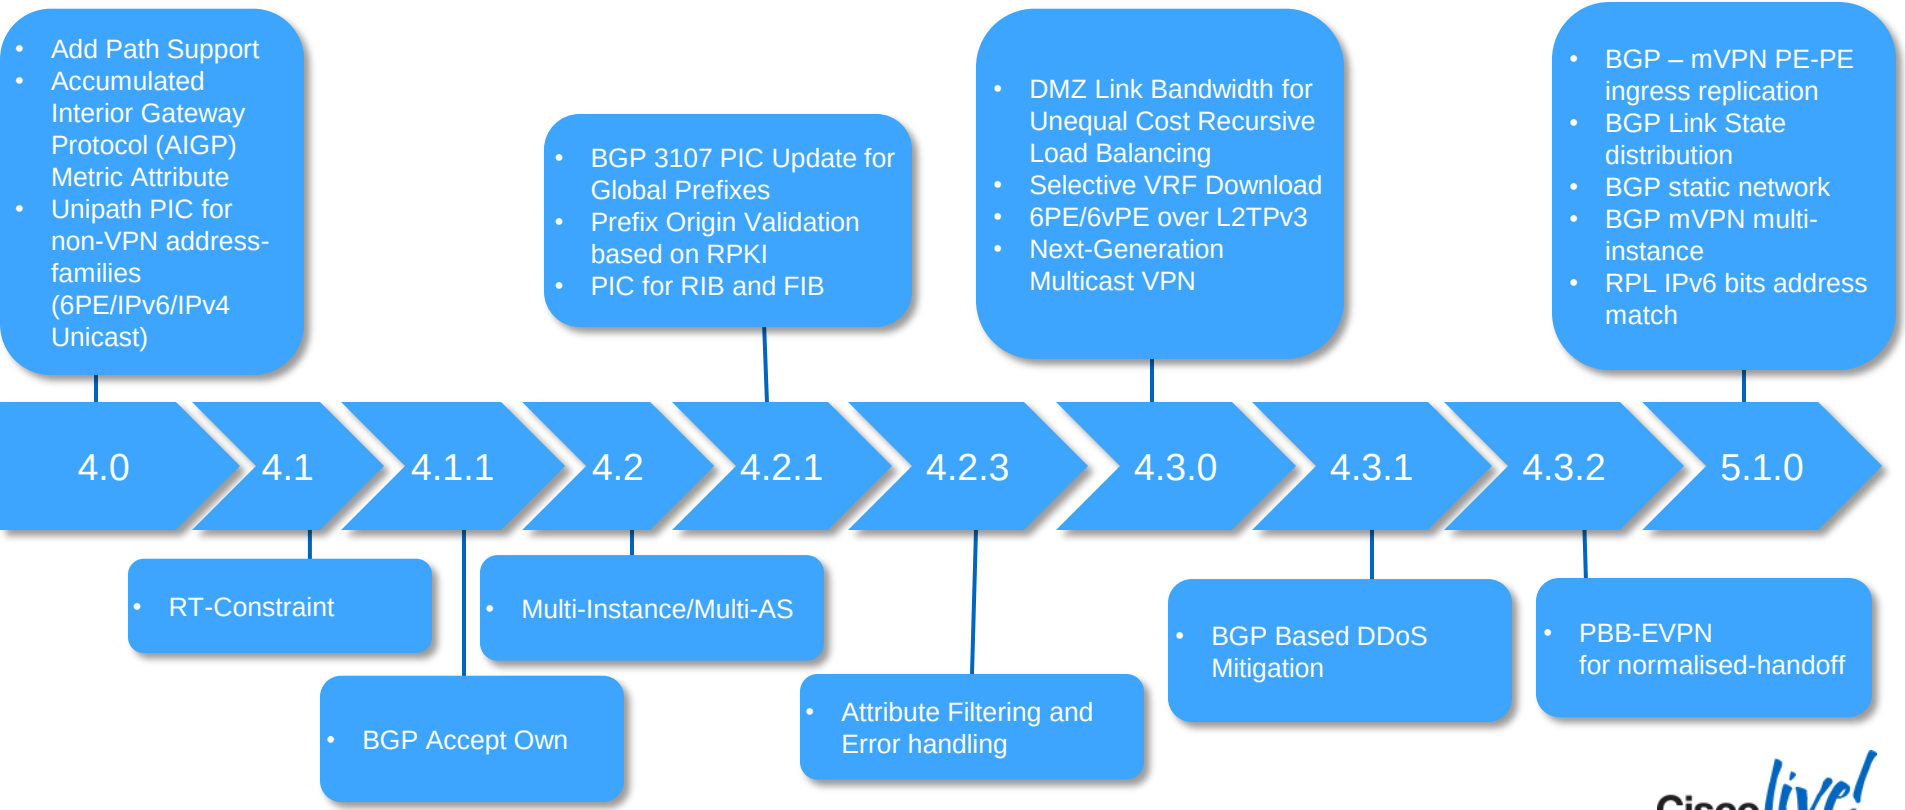

What Happened in IOS Landscape?

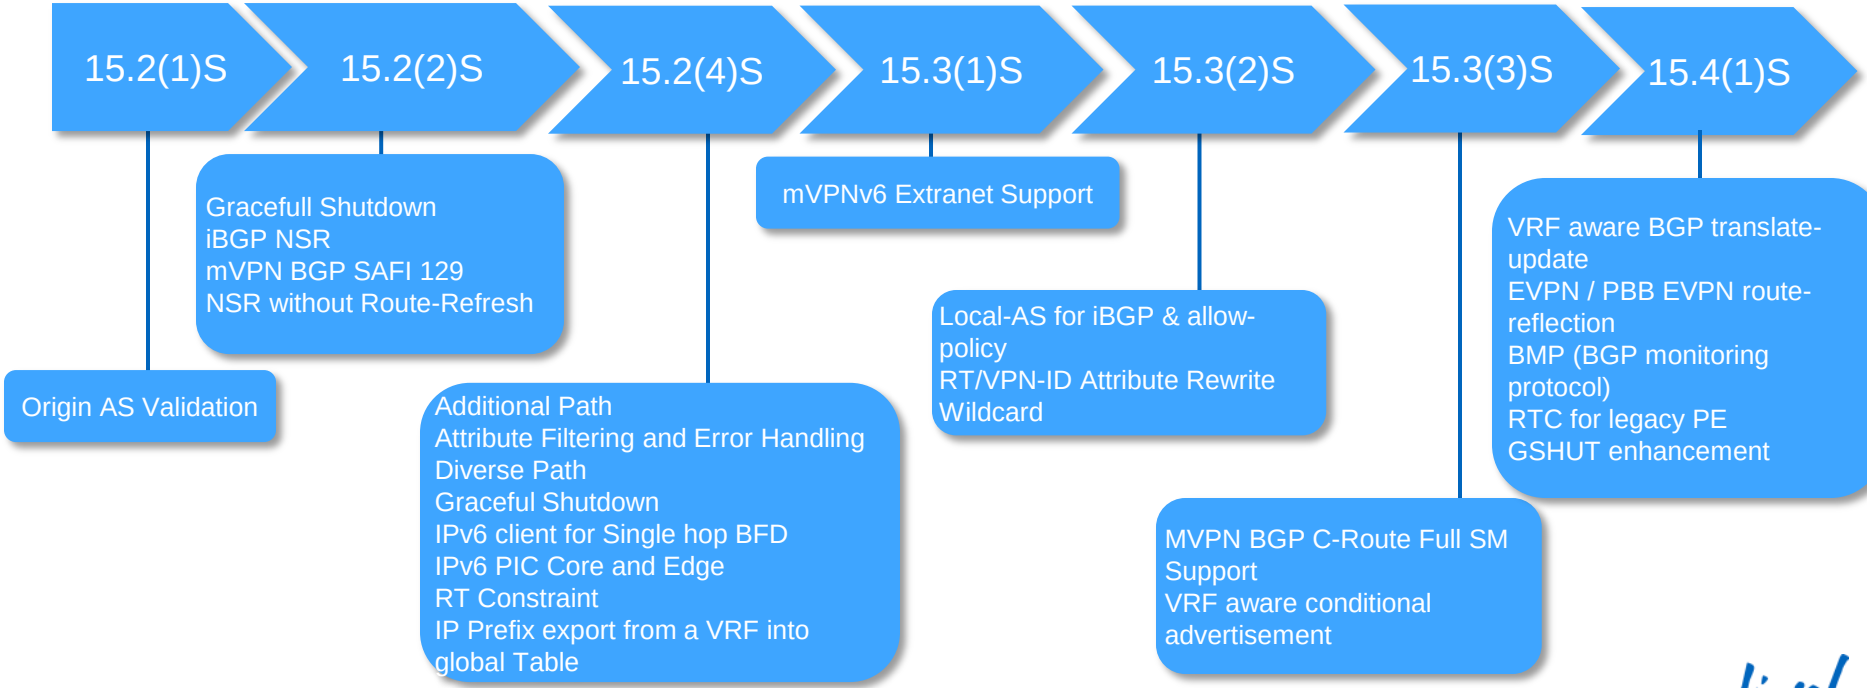

What Happened in XE Landscape?

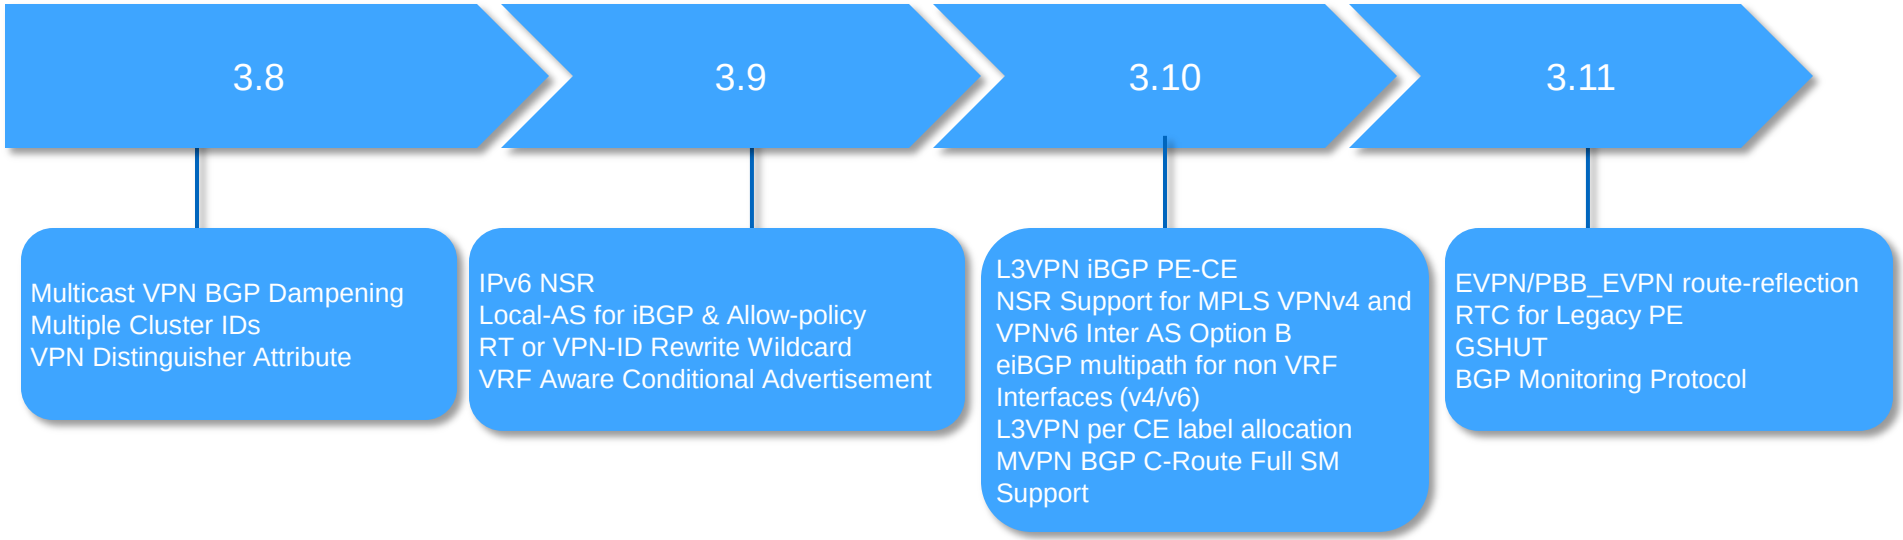

http://www.cisco.com/en/US/docs/routers/asr1000/release/notes/asr1k_rn_rel_notes.pdf

http://www.cisco.com/en/US/docs/routers/asr1000/release/notes/asr1k_rn_rel_notes.html

What Happened NXOS Landscape?

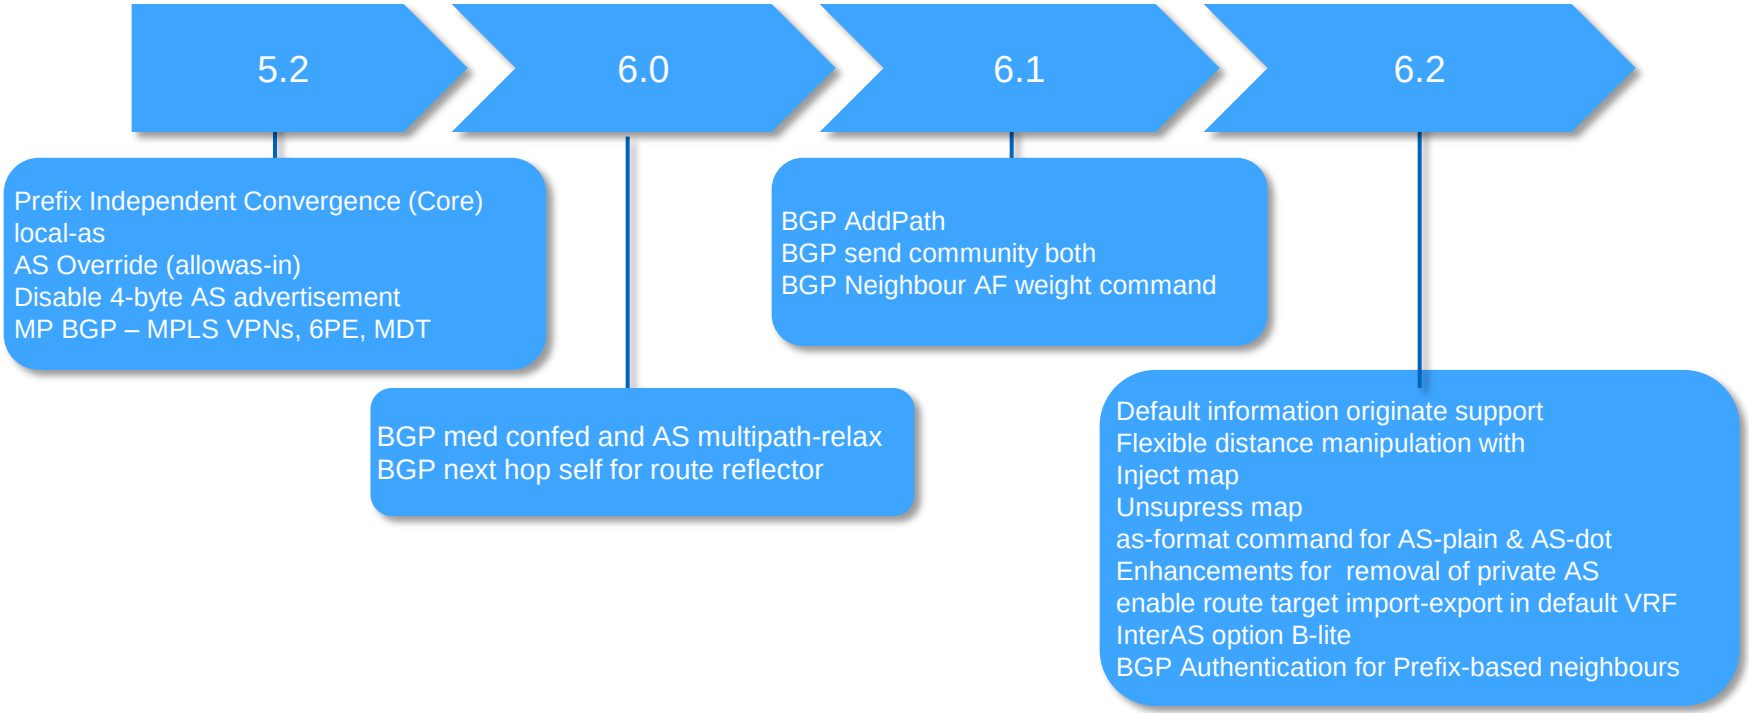

Agenda

- Motivation to Enhance BGP
- Scale and Performance Enhancements
- What happened in BGP Landscape?
- Some new cool features that may interest you

The Bloody Good Protocol

New Features Covered Here

- BGP Fast Convergence/BGP PIC
- Add-Path
- Automated Route-Target Filtering
- AIGP
- Multi-AS BGP
- Attribute Filtering & Error Handling
- BGP Origin Validation
- BGP Graceful Shutdown
- VRF-aware Conditional Advertising
- Local-AS for iBGP
- eiBGP for Global Routing Context
- per-CE Label
- L3VPN iBGP PE-CE
- BGP Management Protocol

BGP PIC – Control vs. Data Plane Convergence

- Control Plane Convergence
  - For the topology after the failure, the optimal path is known and installed in the data plane
  - May be extremely long (depends on number of prefixes carried)
- Data Plane Convergence
  - Once IGP convergence has detected the failure, the packets are rerouted onto a valid path to the BGP destination
  - While valid, this path may not be the most optimum one from a control plane convergence viewpoint
  - We want this behaviour, in a prefix-independent way, no matter if BGP carries 1000 or 1,000,000 prefixes!

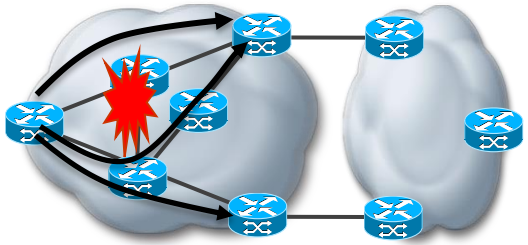

BGP Prefix Independent Convergence (BGP PIC)

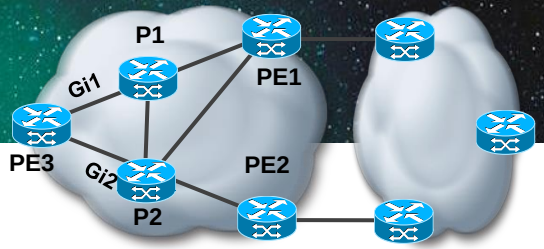

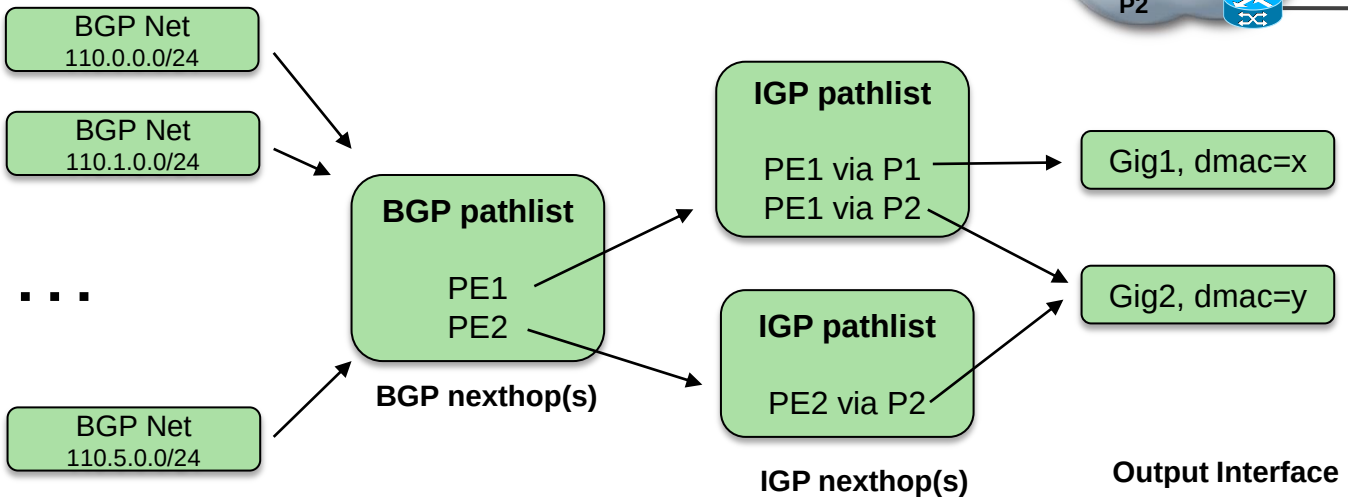

- Pointer Indirection between BGP and IGP entries allow for immediate update of the multipath BGP pathlist at IGP convergence
- Only the parts of FIB actually affected by a change needs to be touched
- Used in newer IOS and IOS-XR (all platforms), enables Prefix Independent Convergence

Characterisation
BGP PIC Core

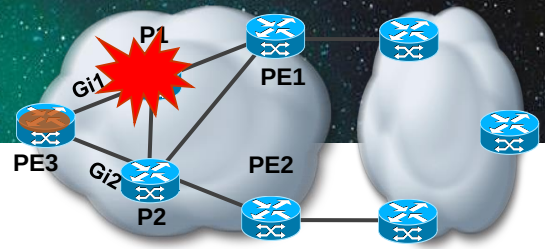

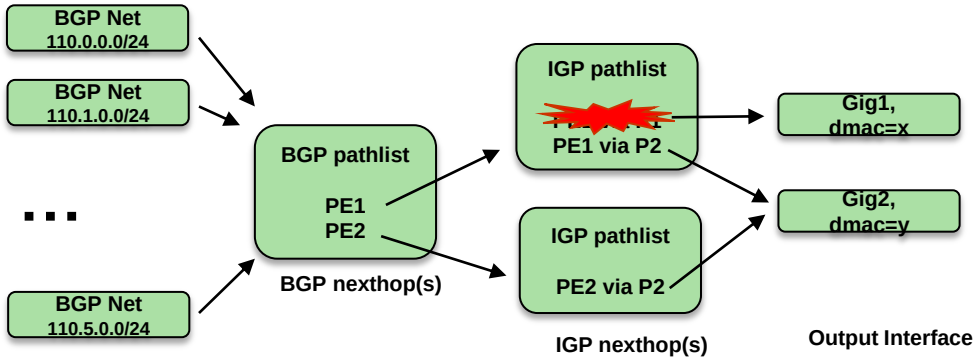

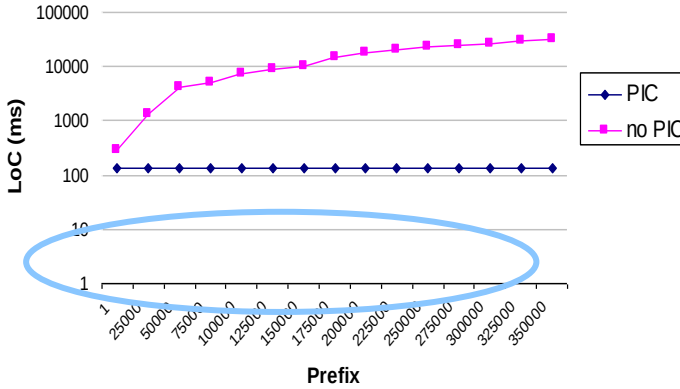

- As soon as IGP converges, the IGP pathlist memory is updated, and hence all children BGP path-lists leverage the new path immediately
- Optimum convergence, Optimum Load-Balancing, Excellent Robustness

Characterisation
BGP PIC Edge

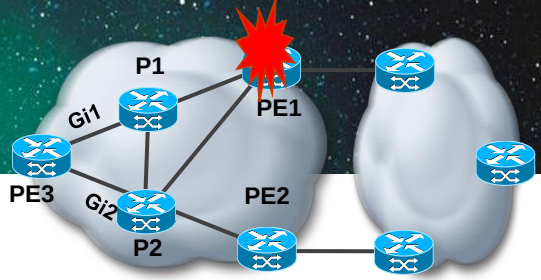

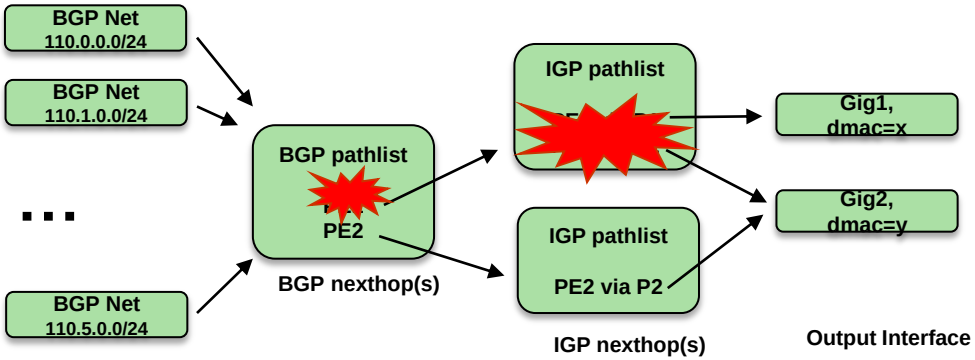

- At IGP Convergence time, the complete IGP pathlist to PE1 is deleted
- FIB updates the affected BGP path lists, traffic converges to the alternate next-hop PE2

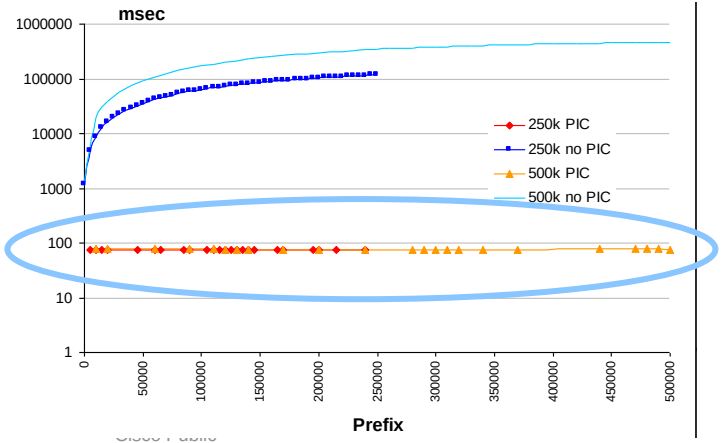

PIC Edge: PE-CE Link Protection

BGP Resiliency/HA Enhancement

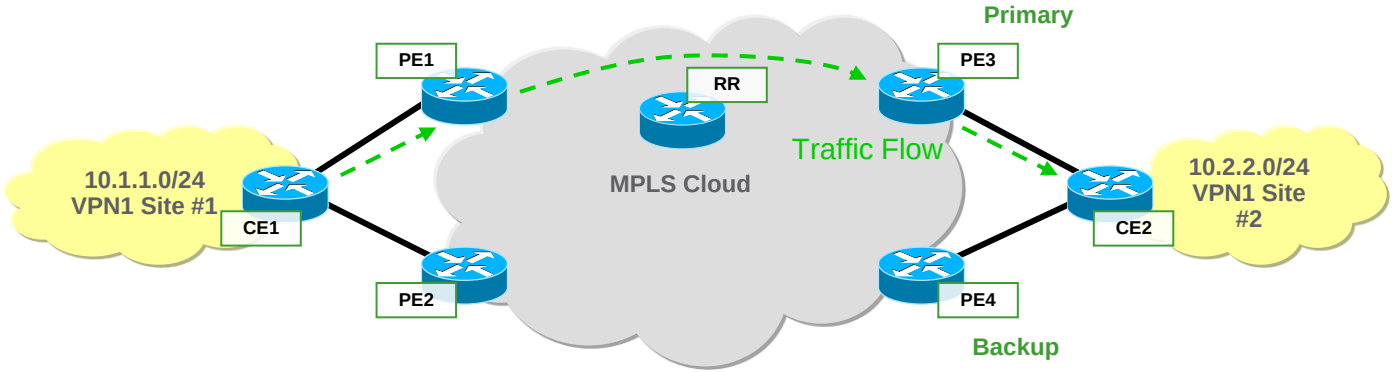

- PE3 configured as primary, PE4 as backup
  - –PE3 preferred over PE4 by local preference
  - –CE2 has different RDs in VRFs on PE3 and PE4
  - –PE4: advertise-best-external, to advertise route received via PE4-CE2 link
  - –PE3: additional-paths install, to install primary and backup path

PIC Edge: Link Protection

BGP Resiliency/HA Enhancement

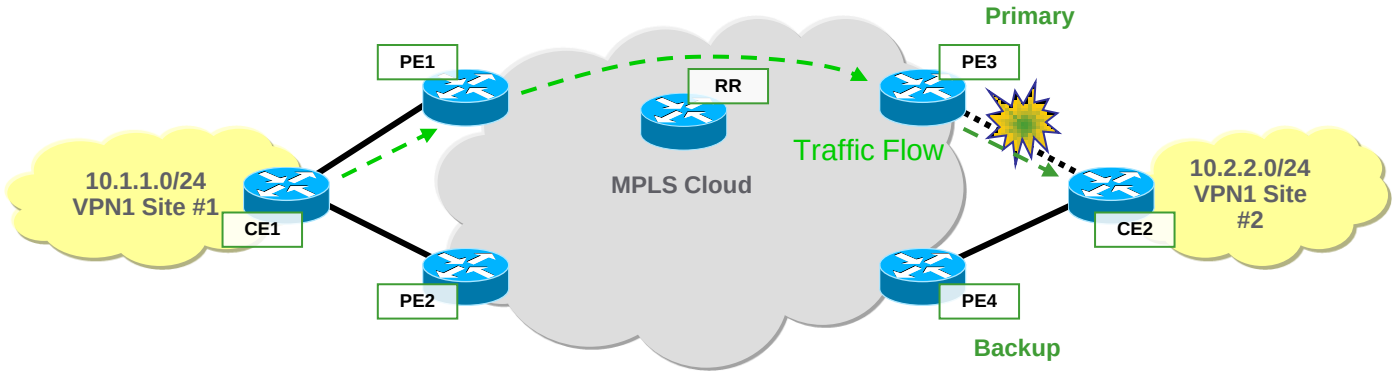

- PE3 has primary and backup path
  - Primary via directly connected PE3-CE2 link
  - Backup via PE4 best external route
- What happens when PE3-CE2 link fails?

PIC Edge: Link Protection

BGP Resiliency/HA Enhancement

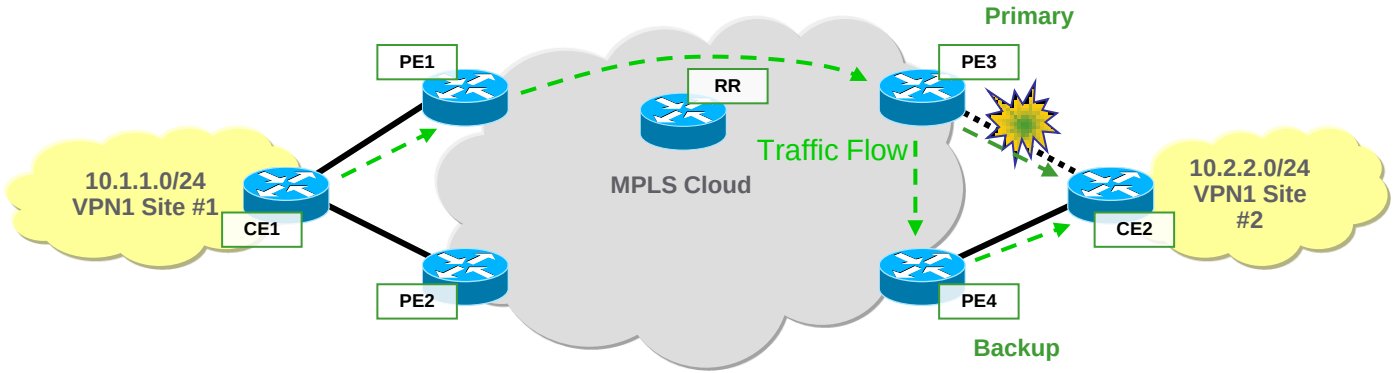

- CEF (via BFD or link layer mechanism) detects PE3-CE2 link failure
  - CEF immediately swaps to repair path label
    Traffic shunted to PE4 and across PE4-CE2 link

PIC Edge: Link Protection

BGP Resiliency/HA Enhancement

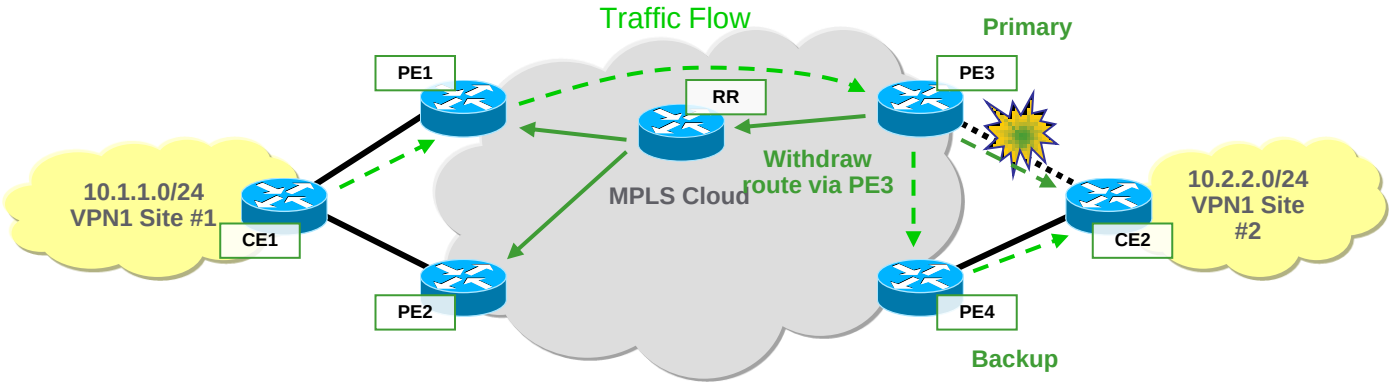

- PE3 withdraws route via PE3-CE2 link
  - Update propagated to remote PE routers

PIC Edge: Link Protection

BGP Resiliency/HA Enhancement

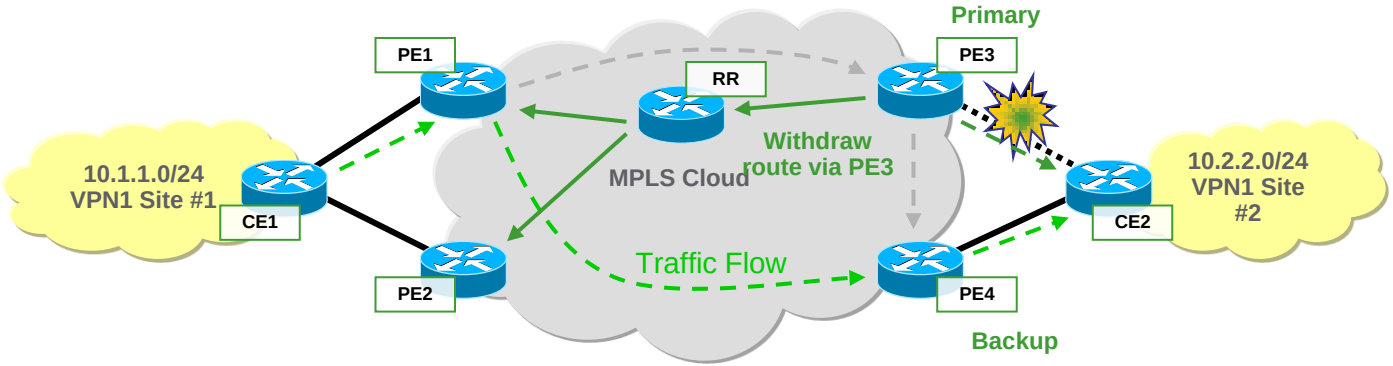

- BGP on remote PEs selects new bestpath
  - New bestpath is via PE4
  - Traffic flows directly to PE4 instead of via PE3

Enabling BGP PIC Edge: IOS-XR

- Two BGP-PIC Edge Flavors: BGP PIC Edge Multipath and Unipath
- **Multipath**: Re-routing router load-balances across multiple next-hops, backup next-hops are actively taking traffic, are active in the routing/forwarding plane, commonly found in active/active redundancy scenarios.
  - No configuration, apart from enabling BGP multipath (maximum-paths ...)
- **Unipath**: Backup path(s) are NOT taking traffic, as found in active/standby scenarios

```
route-policy backup
 ! Currently, only a single backup path is supported
 set path-selection backup 1 install [multipath-protect] [advertise]
end-policy

router bgp ...
 address-family ipv4 unicast
  additional-paths selection route-policy backup
!
 address-family vpnv4 unicast
  additional-paths selection route-policy backup
!
```

Enabling BGP PIC Edge: IOS

- As in IOS-XR, PIC-Edge w/ multipath requires no additional configuration
- PIC-Edge unipath needs to be enabled explicitly ...

```
router bgp ...
 address-family ipv4 [vrf ...]
   or
 address-family vpnv4
  bgp additional-paths install
```

... or implicitly when enabling best external

```
router bgp ...
 address-family ipv4 [vrf ...]
   or
 address-family vpnv4
  bgp advertise-best-external
```

http://www.cisco.com/en/US/docs/ios/iproute_bgp/configuration/guide/irg_bgp_mp_pic.html

http://www.cisco.com/en/US/docs/ios/ios_xe/iproute_bgp/configuration/guide/irg_best_external_xe.html

Question:
How will my PE's Learn about the alternate paths?

- By default my RR Only-Reflects the Best-Route

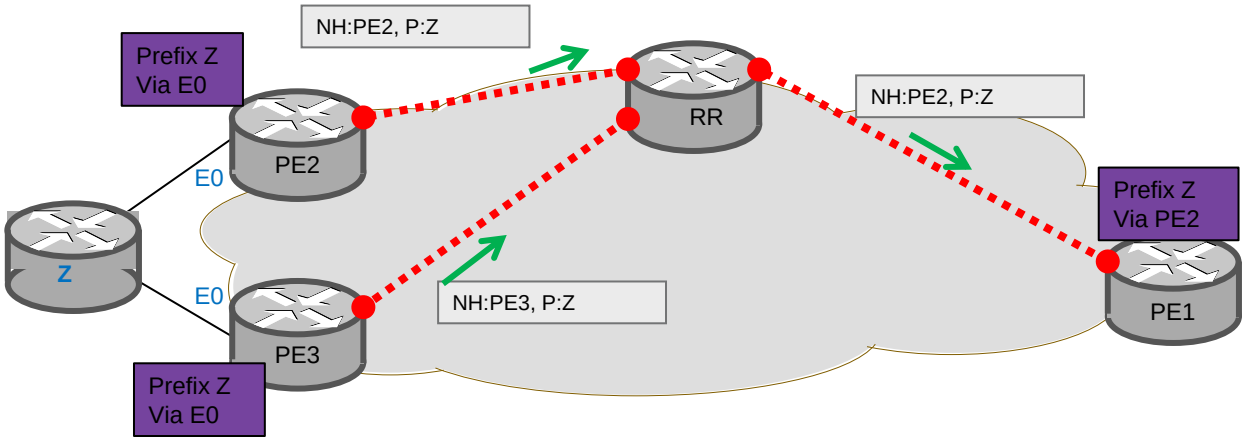

Diverse BGP Path Distribution
Shadow Session

- Easy deployment – no upgrade of any existing router is required, just new iBGP session per each extra path (CLI knob in RR1)
- Diverse iBGP session does announce the 2nd best path

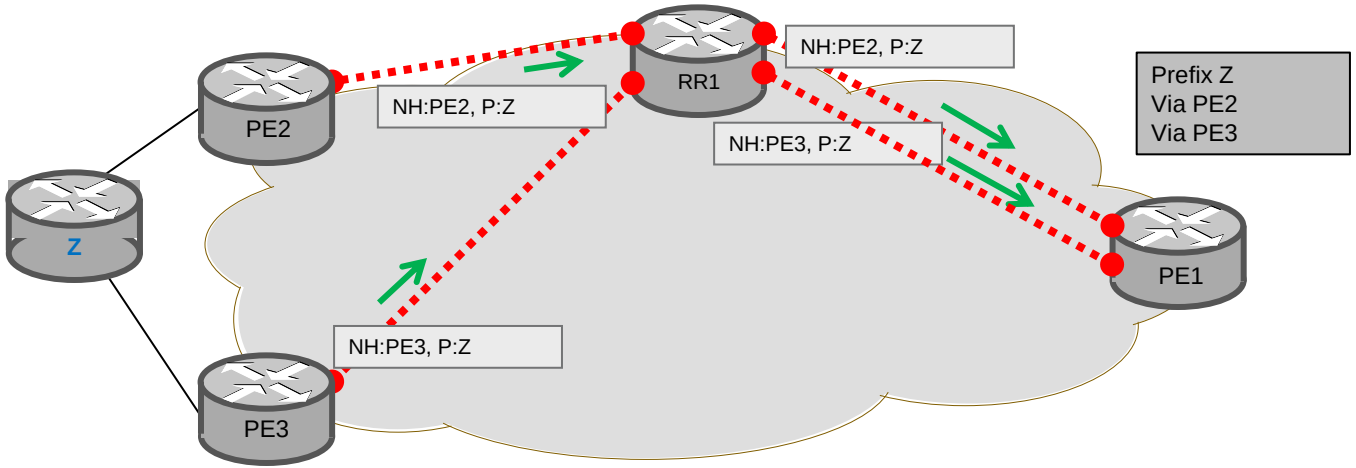

BGP Add-Path

- Add-Path will signal diverse paths from 2 to X paths
- Required all Add-Path receiver BGP router to support Add-Path capability.

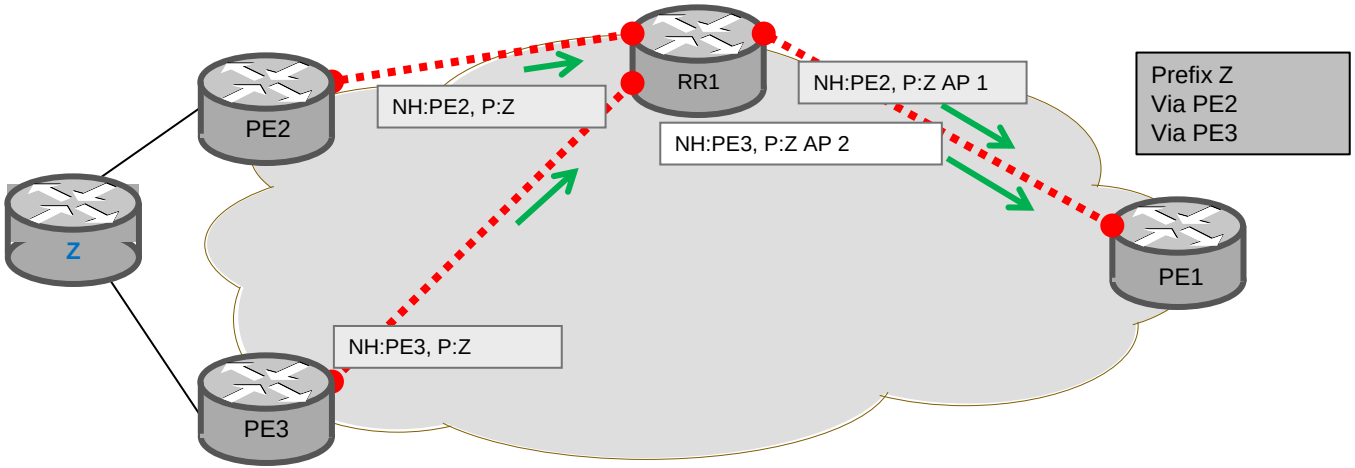

BGP Add-path Flavours

- IETF defines 5 flavors of Add-x-Path. 2 are implemented by Cisco:
- **Add-n-path:** with add-n-path the route reflector will do best path computation for all paths and send n best to BR/PE.
  - Use case: Primary + n Backup scenario
- **Add-all-path:** with add-all-path, the route reflector will do the primary best path computation (only on first path) and then send all path to BR/PE.
  - Use case: Large DC ECMP load balancing, hot potato routing scenario

- Cisco innovation: Add-all-multipath and Add-all-multipath+backup in XR 4.3.1

Automated Route Target Filtering

- Increased VPN service deployment increases load on VPN routers
  - 10% Year-over-Year VPN table growth
  - Highly desirable to filter unwanted VPN routes before sending them to the PEs for them to discard them
- Multiple filtering approaches
  - Extended community ORF
  - New RT filter address family

Automated Route Target Filtering

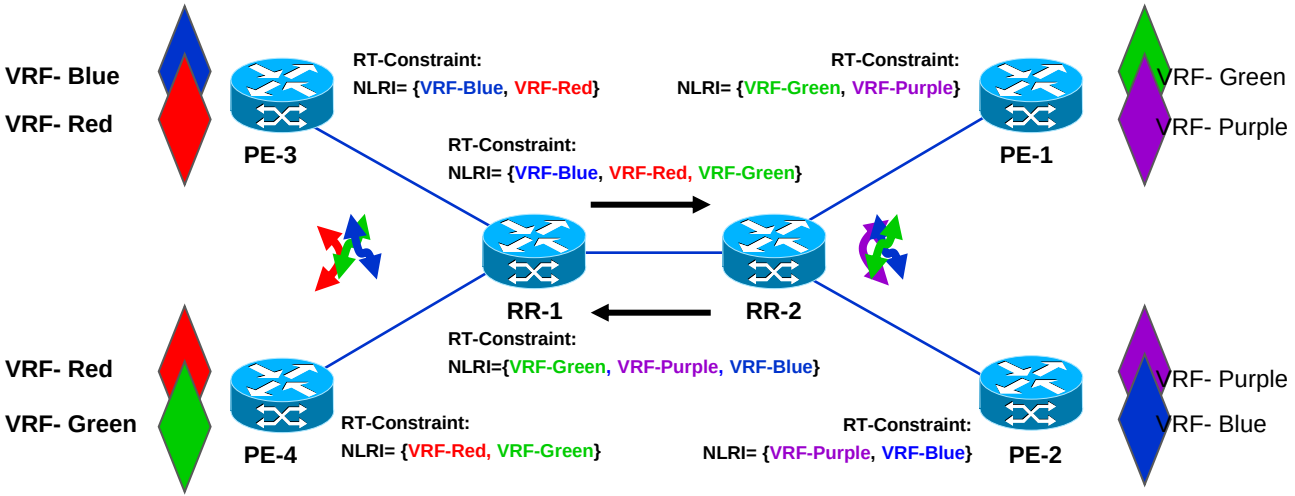

- Improves PE and RR scaling and performance by sending only relevant VPN routes

```
router bgp as-number
 address-family rtfilter unicast
  neighbor {ip-address | peer-group-name} activate
  neighbor {ip-address | peer-group-name} send-community extended
```

RTC for Legacy PE

- BGP Route-Target Constraint expect that routers involved support new RTC AF
- This enhancement will make use of VPN unicast route exchange from the legacy PEs to a new BGP speaker (e.g. an RR) to signal RT membership. The legacies PEs announce a set of "special" routes with mapped RTs to the RR along with a standard community
- The presence of the community triggers the RR to extract the RTs and build RT membership information.

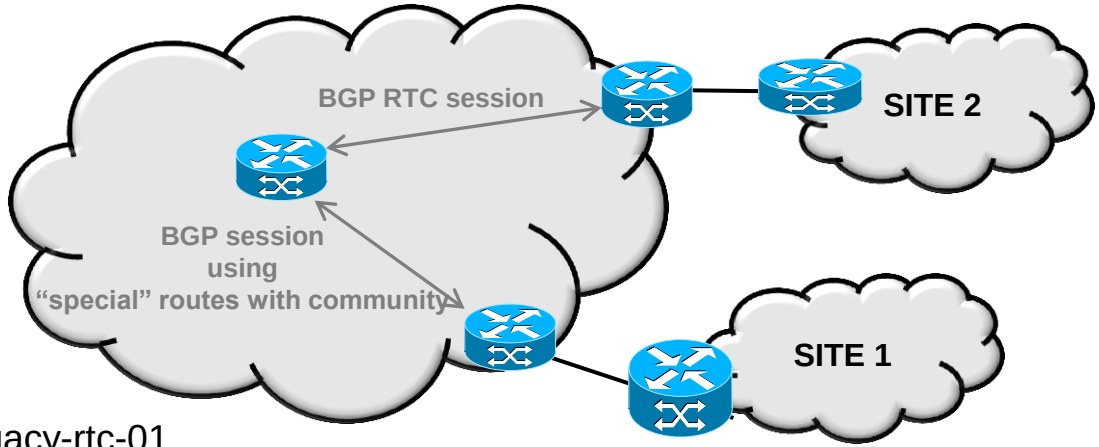

draft-ietf-idr-legacy-rtc-01

RTC for Legacy PE

- How does RTC for legacy PE work?
  - Uses traditional VPN route exchange to signal RT membership
  - Legacy PE announces “special routes” with mapped RT’s on the PE
  - In addition it tags these routes with a community
  - The presence of the community triggers the RR to extract the RTs and build RT membership information

Route-reflector

```
!
 router bgp 1
 address-family vpnv4 unicast
  neighbor 10.1.1.1 accept-route-legacy-rt
!
```

Legacy-PE

```
ip vrf route-filter
 rd 55:1111
 export map SET_RT

route-map SET_RT permit 10
 match ip address prefix-list RT_NET1
 set community 65535:2 (0xFFFF0002)
 set extcommunity rt  255.220.0.0:12241 255.220.0.0:12242 additive
 set extcommunity rt  255.220.0.0:12243 255.220.0.0:12244 additive
 set extcommunity rt  255.220.0.0:12245 255.220.0.0:12246 additive
 set extcommunity rt  255.220.0.0:12247 255.220.0.0:12248 additive
 set extcommunity rt  255.220.0.0:12249 255.220.0.0:12250 additive
!
route-map SET_RT permit 20
 match ip address prefix-list RT_NET2
 set community 65535:2 (0xFFFF0002)
 set extcommunity rt  255.220.0.0:12251 255.220.0.0:12252 additive
 set extcommunity rt  255.220.0.0:12253 255.220.0.0:12254 additive
 set extcommunity rt  255.220.0.0:12255 additive
!

ip route vrf route-filter 1.1.1.1 255.255.255.255 Null0 – (matching prefix-set R
ip route vrf route-filter 1.1.1.2 255.255.255.255 Null0 –(matching prefix-set RT_NET2)

route-map LEG_PE permit 10
 match ip address prefix-list RT_NET1 RT_NET2
 set community no-advertise additive

router bgp 55
 address-family vpnv4 unicast
 neighbor x.x.x.x route-map LEG_PE out
```

Centralised Extranet VRF Policy Control
Accept own

- Currently, Extranet VPN import requires individual “route-target import …” on all PEs where customer/client and extranet VPNs are connected to.. Could be cumbersome to maintain/scale
- How about doing this centrally on a route reflector? Attaching route-target communities to vpnv4/vpnv6 routes?
- But we might need to send prefixes with the modified RT communities back to originating PE so it can import them
- Use special/new community “accept-own” to ignore originator check on the receiving end

PE

```
router bgp 1
 neighbor <route-reflector>
  accept-own
```

RR

```
route-policy rr-pe-out
 …
 if  extcommunity rt matches-any CUSTOMERS then
    set extcommunity rt (xxx:yyyy) additive
    set extcommunity rt (‘accept-own’) additive
 endif
end-policy
```

Overview – AIGP

Accumulated IGP Metric Attribute for BGP (draft-ietf-idr-aigp-09)

- Optional, non-transitive BGP path attribute
- BGP attribute to provide BGP a way to make its routing decision based on the IGP metric, to choose the “shortest” path between two nodes across different AS
- The main driving force for this feature is to solve the IGP scale issue seen in some ISP core network (aka “Unified MPLS” architecture)
- Mainly to be deployed to carry nexthop prefixes/labels across different AS within the same administrative domain
- The remote ingress PE select its best path using the modified best path selection process using AIGP metric

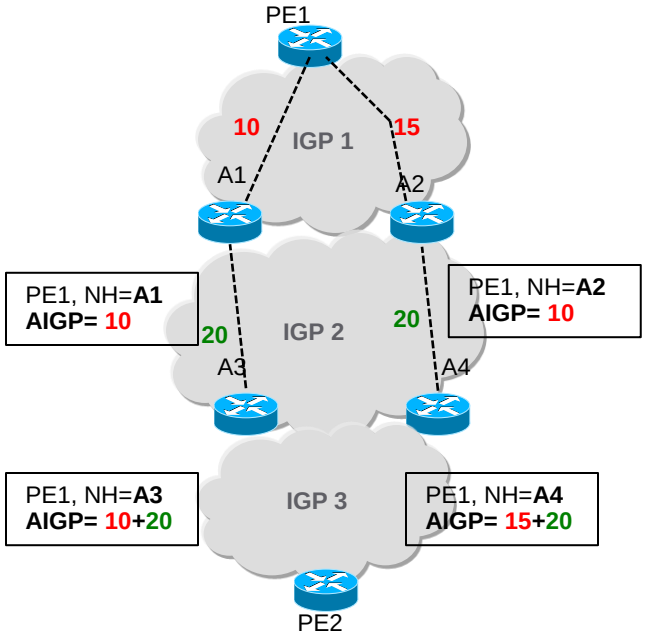

AIGP: Originating AIGP

- AIGP is enabled between iBGP neighbours by default
- AIGP between eBGP neighbuors need to be enabled
- AIGP can be originated by using redistribute ospf, redistribute isis, redistribute static or the BGP network command.
- AIGP can also be originated using neighbour address-family inbound or outbound policy to set AIGP to be the IGP cost or to a fixed value.

```
router bgp 1
   address-family ipv4 unicast
      redistribute ospf 1 route-policy set_aigp_1
```

```
route-policy set_aigp_1
  if destination in (…) then
    set aigp-metric 111
 elseif destination in (..) then
    set aigp-metric igp-cost
 endif
end-policy
```

What is Multi-AS BGP?

- Run multiple instances of BGP on a router (possibly on different RP instances)
- It will be possible to configure each instance of a multi-instances BGP with a different AS number
- Global address families can't be configured under more than one AS except vpnv4 and vpnv6
- VPN address-families may be configured under multiple AS instances that do not share any VRFs

Scale/Fault Isolation

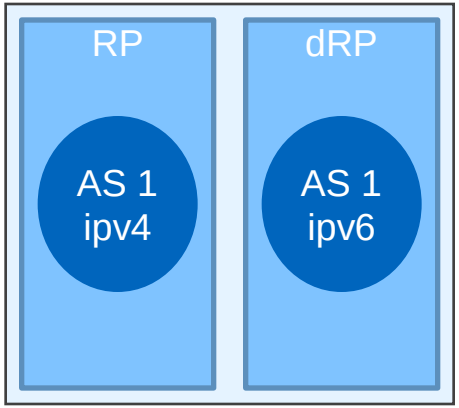

PE-CE/VRF Scale

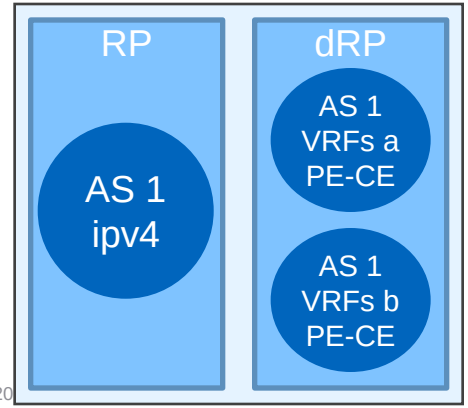

Service Integration

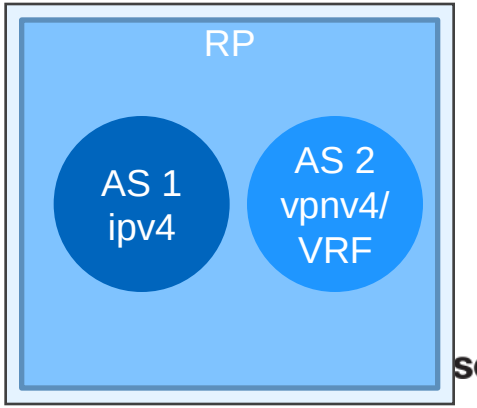

Configuration Example


```
router bgp 1 instance ipv4
    bgp router-id 10.0.0.1
    address-family ipv4 unicast
    neighbor 10.0.101.1
       remote-as 1
       address-family ipv4 unicast
          route-policy inbound in
          route-policy outbound out
       !
    !
  !
!

router bgp 1 instance ipv6
    bgp router-id 10.0.0.2
    address-family ipv6 unicast
    neighbor 10.0.101.2
       remote-as 1
       address-family ipv6 unicast
          route-policy inbound in
          route-policy outbound out
       !
    !
  !
!
```

```
router bgp 3 instance vpn1
    bgp router-id 20.0.0.1
    address-family vpnv4 unicast
    neighbor 20.0.101.1
       remote-as 200
       address-family vpnv4 unicast
          route-policy inbound in
          route-policy outbound out
       !
    !
  !
  vrf foo
  !
!
router bgp 3 instance vpn2
    bgp router-id 20.0.0.2
    address-family vpnv4 unicast
    neighbor 20.0.101.2
       remote-as 200
       address-family vpnv4 unicast
          route-policy inbound in
          route-policy outbound out
       !
    !
  !
  vrf bar
  !
!
```

Attribute Filtering and Error-handling

- Attribute filtering
  - Unwanted optional transitive attribute such as ATTR_SET, CONFED segment in AS4_PATH causing outage in some equipments.
  - Prevent unwanted/unknown BGP attributes from hitting legacy equipment
    - Block specific attributes
    - Block a range of non-mandatory attributes
- Error-handling
  - draft-ietf-idr-optional-transitive-04.txt
  - *Punishment should not exceed the crime*
  - Gracefully fix or ignore non-severe errors
  - Avoid session resets for most cases
  - Never discard update error, as that can lead to inconsistencies

Architecture

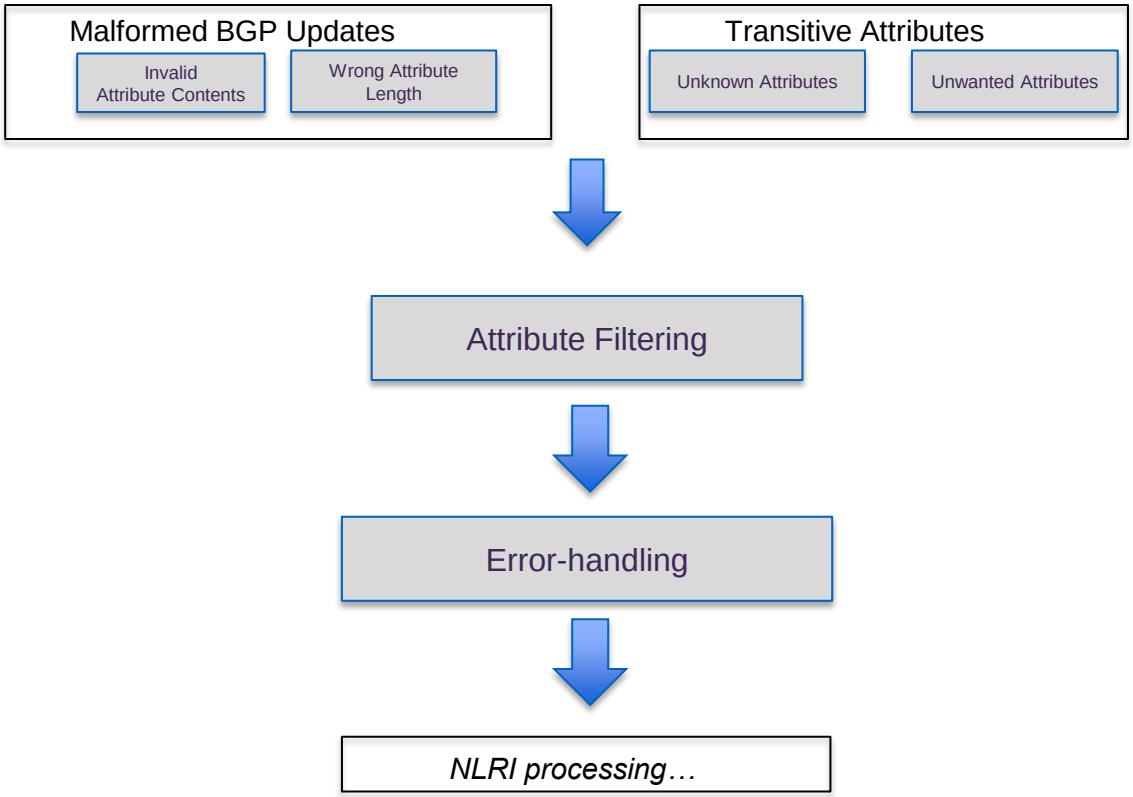

BGP Origin Validation

- Support client functionality of RPKI RTR protocol
  Separate database to store record entries from the cache
- Support to announce path validation state to IBGP neighbours using a well known path validation state extended community
- Modified route policies to incorporate path validation states

Prefix Hijacking

- Announce someone else's prefix
- Announce a more specific of someone else's prefix
- Either way, you are trying to "steal" someone else's traffic by getting it routed to you
  - Capture, sniff, redirect, manipulate traffic as you wish

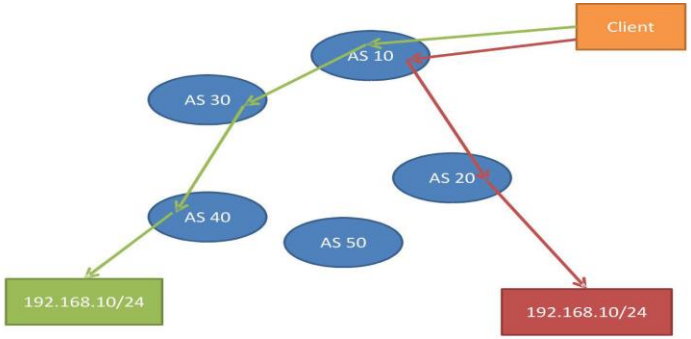

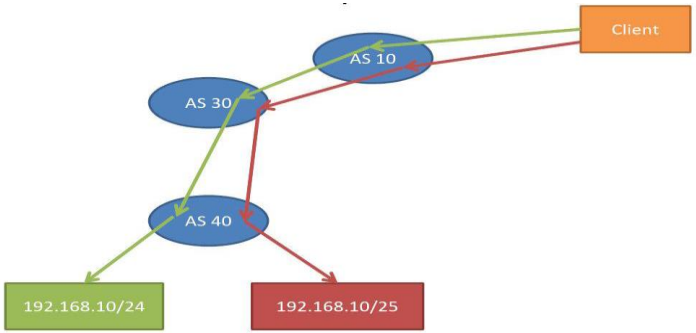

Source: nanog 46 preso

How does the Solution Look Like?

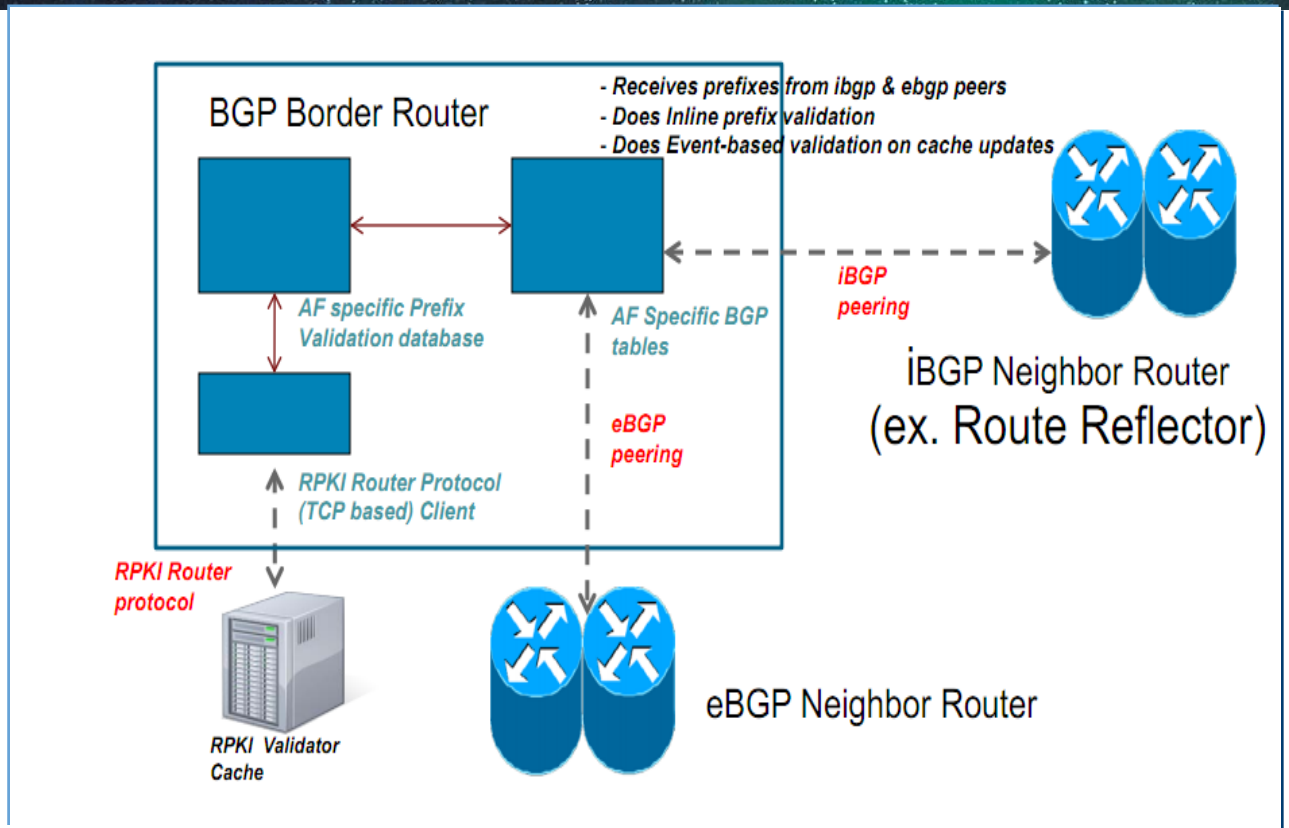

RPKI Service information from APNIC at
https://www.apnic.net/services/services-apnic-provides/resource-certification/RPKI

Sample Configuration


```
router bgp 64726
 bgp rpki server tcp 10.1.2.3 port 30000 refresh 60
 bgp bestpath prefix-validate allow-invalid
 neighbor 10.9.9.9 remote-as 64209
 neighbor 10.9.9.9 route-map FOO in
!
route-map FOO permit 10
 match rpki invalid
 set local-preference 50
route-map FOO permit 20
 match rpki not-found
 set local-preference 100
route-map FOO permit 30
 match rpki valid
 set local-preference 200
route-map FOO permit 40
!
```

BGP Graceful Shutdown

- RFC 6198 – April 2011
- Old Behaviour
  - If session drops then BGP will withdraw all prefixes learned over that session
  - BGP has no mechanism to signal prefix will soon be unreachable (for maintenance for example)
- Historically RR's have worsened the issue as they tend to hide the alternate path as they only forward the best path

BGP Graceful Shutdown allows to do maintenance on router without service disruption.

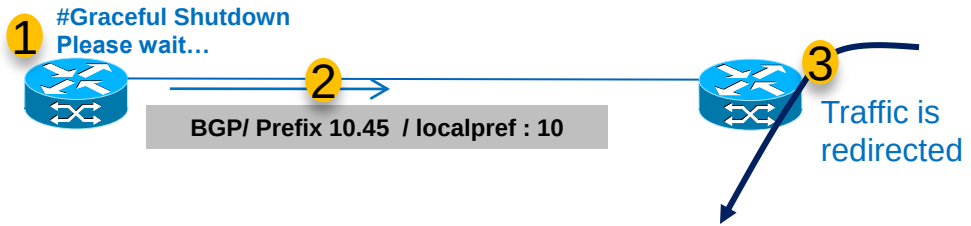

This new knob allows a router to notify neighbour to redirect traffic to other paths and after some time will drop BGP sessions.

The notification could be done using Local Preference attribute or user community attribute

BGP Graceful Shutdown

- GSHUT well-known community
- The GSHUT community attribute is applied to a neighbour specified by the **neighbour shutdown graceful** command, thereby gracefully shutting down the link in an expected number of seconds
- The GSHUT community is specified in a community list, which is referenced by a route map and then used to make policy routing decisions.

```
neighbor {ipv4-address | ipv6-address | peer-group-name} shutdown graceful seconds {community value [local-preference value] | local-preference value}
```

http://www.cisco.com/en/US/docs/ios-xml/ios/iproute_bgp/configuration/15-s/irg-15-s-book.pdf

VRF Aware Conditional Advertisement

- Before: Conditional advertisement supported in IPv4 Unicast/Multicast address-family
- New: Support for IPv4 VRF, IPv6 Unicast and IPv6 VRF

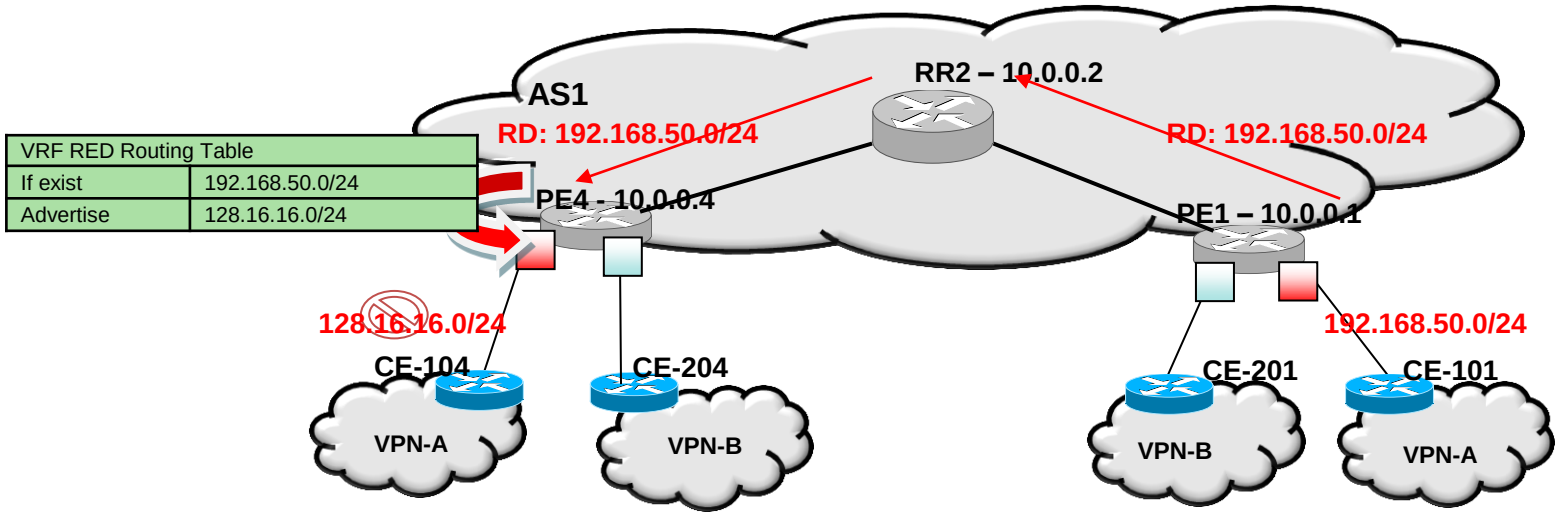

RT / VPN Distinguisher Attribute Rewrite

- A single VPN could span both ISP1 and ISP2
- RTs are preferred to be kept private within an AS
- VPN Distinguisher (VD) Attribute exchanged via eBGP allows to keep private per AS
- Enhancement also allows RANGE statement for VD <-> RT mapping

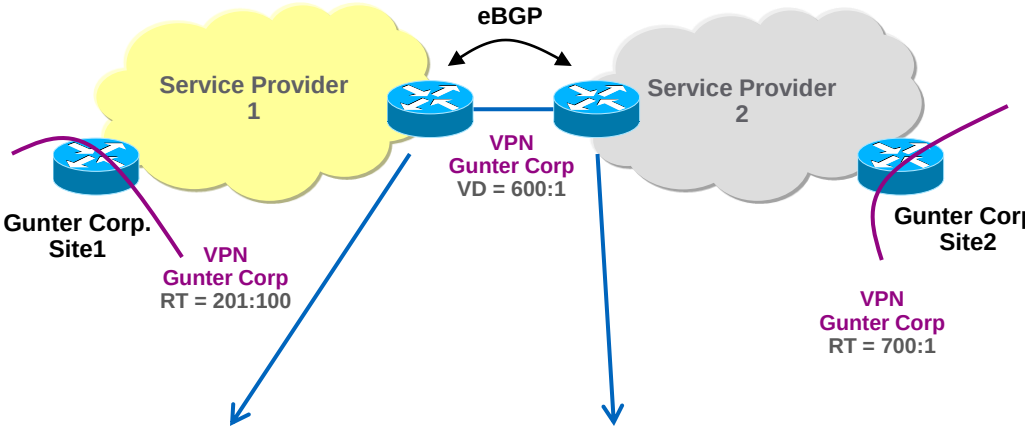


```
ip extcommunity-list 22 permit rt 201:100
!
route-map rt-mapping permit 10
 match extcommunity 22
 set extcomm-list 22 delete
 set extcommunity vpn-distinguisher 600:1
!
route-map rt-mapping permit 20
!
router bgp 3000
 neighbor 192.168.103.1 remote-as 3000
 address-family vpnv4
  neighbor 192.168.103.1 activate
  neighbor 192.168.103.1 route-map rt-mapping out
  exit-address-family
!
```

```
ip extcommunity-list 101 permit VD:600:1
!
route-map rtmap2 permit 10
 match extcommunity 101
 set extcomm-list 101 delete
 set extcommunity rt 700:1 additive
!
route-map rtmap2 permit 20
!
router bgp 4000
 neighbor 192.168.0.50 remote-as 4000
 address-family vpnv4
  neighbor 192.168.0.50 activate
  neighbor 192.168.0.50 route-map rtmap2 in
  exit-address-family
!
```

Local-AS for iBGP & Allow-Policy

- Existing
  - **neighbor local-as** was used on a route-reflector to customise AS_PATH attributes from an eBGP neighbor
- New
  - **neighbor local-as** can now be used to enable sending of iBGP attributes (LOCAL_PREF, ORIGINATOR_ID, CLUSTER_ID and CLUSTER_LIST) over an iBGP local-AS session
  - It seems reasonable that by using route-map statements these attributes can be modified
  - This flexibility is achieved by configuring the **neighbor allow-policy** command

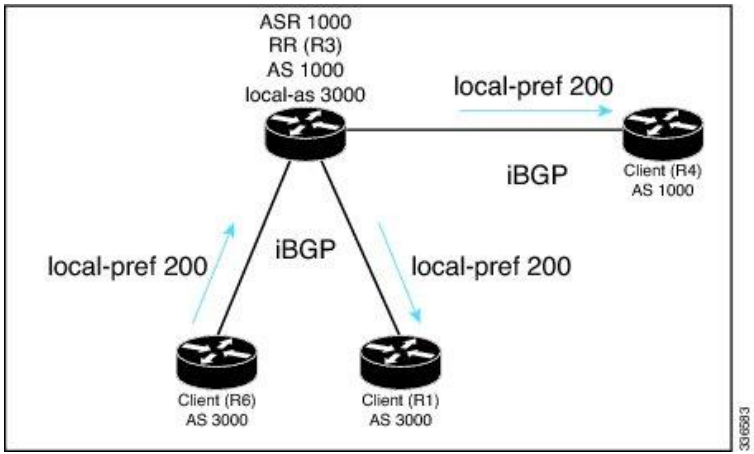

eiBGP Multi-path for non VRF Interface

- Existing
  - BGP best path selection prefers eBGP over iBGP
- New
  - eiBGP multipath removes this criteria in BGP path selection mechanism
  - ECMP hash load balancing mechanism will forward 50% of traffic over iBGP and 50% over eBGP

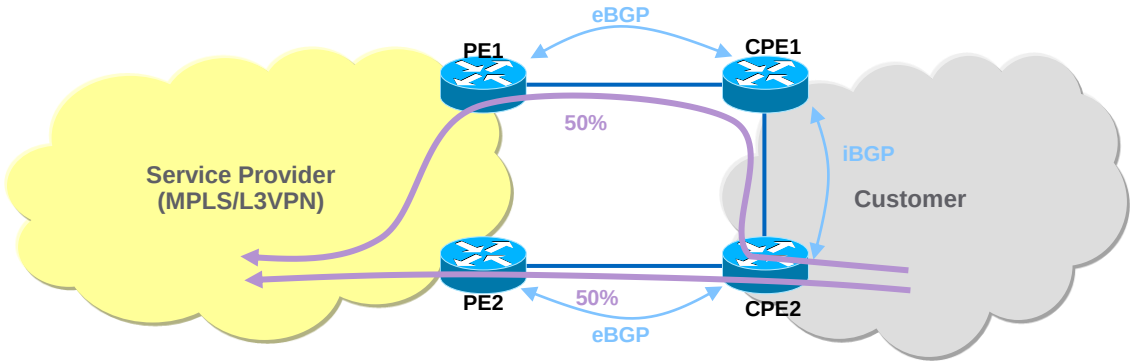


```
!
router bgp 64496
 address-family ipv4 unicast
  maximum-paths eibgp 4
!
 address-family ipv6 unicast
  maximum-paths eibgp 4
!
```

Per CE Label

- Allows you to configure a single VPN label at the provider edge (PE) instead of per route
- Why?
  - Reduce the number of MPLS labels because label space is limited in large deployments
- Caveats
  - Doesn't work together BGP Multipath and CsC because they make usage of the label diversity
  - BGP PIC support requires resilient per-CE label feature, also addresses eBGP multipath
  - No 6PE support
  - Only single hop eBGP supported, no multihop
  - More restrictions to find in the release-notes

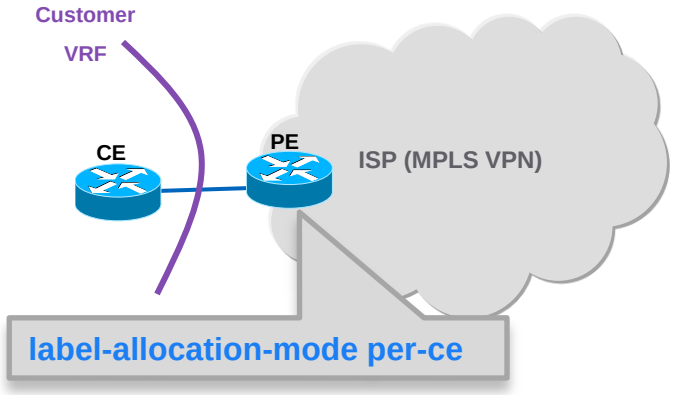

L3VPN iBGP PE-CE (RFC6368)

- Support for iBGP between Customer and Service Provider
- Why?
  - Customer iBGP attributes are retained across the VPN
  - No insertion of ISP BGP AS number in the AS_PATH
- How?
  - Support for RFC6368
  - The PE will place the received iBGP attributes in a new attribute ATTR_SET and transport them over the ISP backbone
  - Identify within the VPN the iBGP L3VPN characteristic
  - Note: requires unique RDs per PE

```
router bgp 100
 address-family ipv4 vrf blue
  neighbor 10.0.0.1 remote-as 200
  neighbor 10.0.0.1 local-as 200
  neighbor 10.0.0.1 internal-vpn-client
  neighbor 10.0.0.1 route-reflector-client
```

EVPN/PBB-EVPN Route-Reflection

- Support for L2 VPNs using BGP as control-plane technology
- EVPN Address Family is allowed on both iBGP as well as eBGP neighbours under default VRF for both IPv4 and IPv6 neighbours

Configuration on RR:

```
router bgp 1
  address-family l2vpn evpn
    neighbor 192.168.1.1 remote-as 1
    neighbor 192.168.1.1 route-reflector-client
    neighbor 192.168.1.1 send-community extended
    .......
```

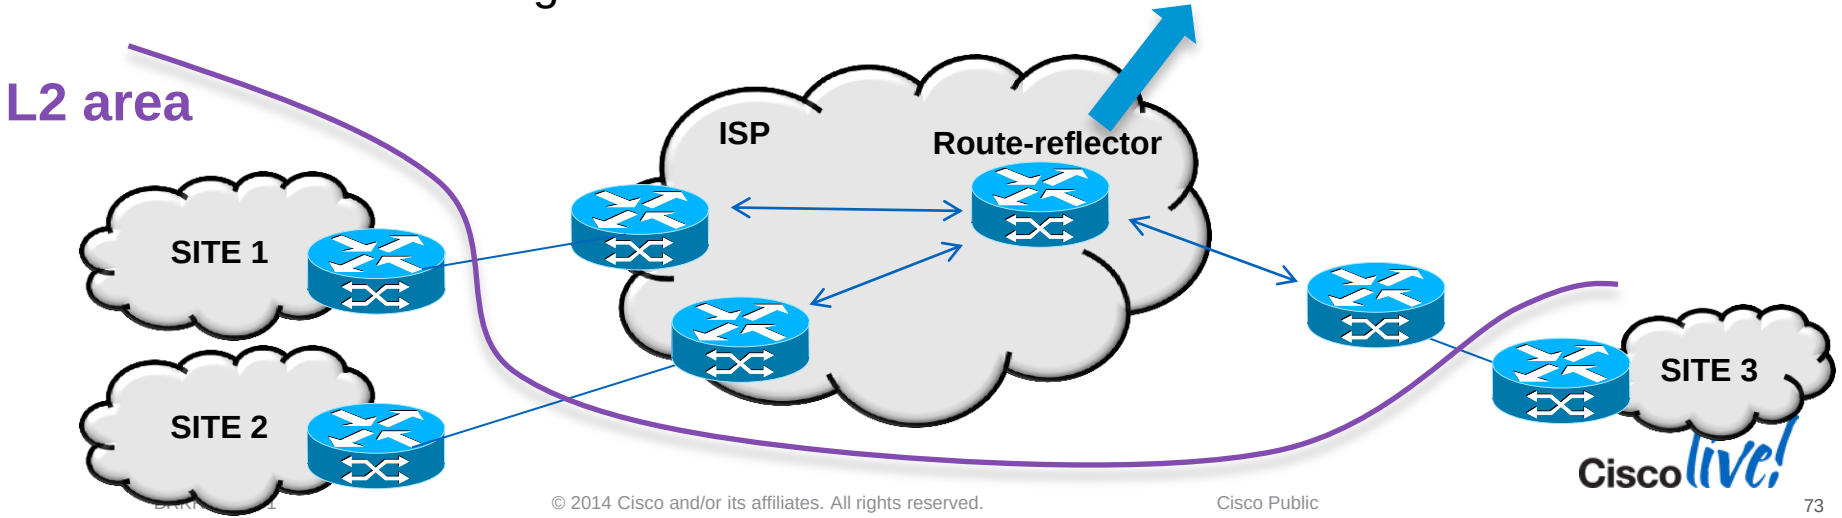

BGP Management Protocol (BMP)

- BMP is intended to be used for monitoring BGP sessions
- BMP is not impacting the routing decision process and is only used to provide monitoring information
- BMP provides access to the Adj-RIB-In of a BGP peer on an ongoing basis and provides a periodic dump of statistical information. A monitoring station can use this for further analysis

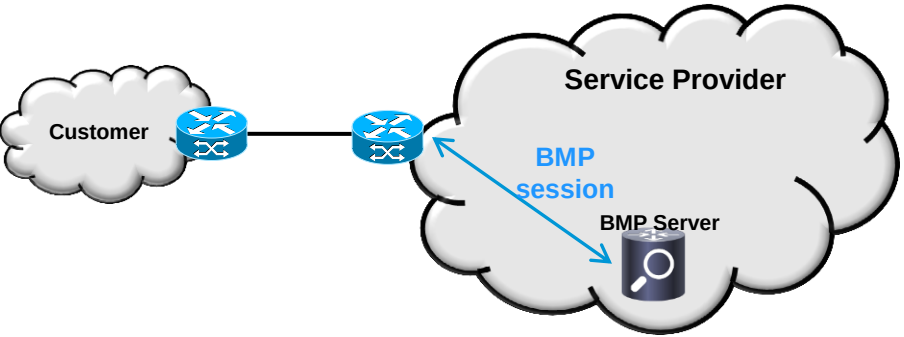

- Configuration
  - Enable monitoring per neighbour
  - Configure the BMP servers

```
!
router bgp 65000
 neighbor 30.1.1.1 bmp-activate server 1
!
bmp server 1
 activate
 address 10.1.1.1 port-number 8000
 description LINE SERVER1
 failure-retry-delay 40
 flapping-delay 120
 initial-delay 20
 set ip dscp 5
 stats-reporting-period 30
 update-source ethernet 0/0
 exit-bmp-server-mode
!
```

http://tools.ietf.org/html/draft-ietf-grow-bmp-07

Summary

- Motivation to Enhance BGP
- Scale and Performance Enhancements
- What happened in BGP Landscape?
- Some new Cool features that may interest you

Bloody Good Protocol




Q & A

Complete Your Online Session Evaluation

Give us your feedback and receive a Cisco Live 2014 Polo Shirt!

Complete your Overall Event Survey and 5 Session Evaluations.

- Directly from your mobile device on the Cisco Live Mobile App
- By visiting the Cisco Live Mobile Site www.ciscoliveaustralia.com/mobile
- Visit any Cisco Live Internet Station located throughout the venue

Polo Shirts can be collected in the World of Solutions on Friday 21 March 12:00pm - 2:00pm

Learn online with Cisco Live!

Visit us online after the conference for full access to session videos and presentations.
www.CiscoLiveAPAC.com


CISCO
TM