Putting IPv6 to work

North American IPv6 Summit

Plaza Tower One Conference Facilities
Greenwood Village, CO
April 22-23, 2015

Rocky Mountain IPv6 Task Force

RMv6TF

6PE: The IPv6 Bolt-On for MPLS

Chris Lenart

Senior Network Architect, Verizon

April 22, 2015

Agenda

- 6PE Definition
- How 6PE Works
- 6PE Use Cases
  - Rapid Deployment
  - BGP Free Core
  - IPv6 Free Core
- 6VPE
- Further Network Evolution
- 6PE Summary and Q&A

What Is 6PE?

6PE

- 6PE is a technology that leverages MPLS to tunnel IPv6 traffic across an IPv4-only core
- Similar to Layer 3 VPN aka BGP/MPLS VPN
  - Same basic mechanisms
- Easily added on to a BGP/MPLS network
  - Only needs minor tweaks at the edge of a network
    - No IPv6 needed toward core
    - Customers can use whatever as long as IPv6 is on edge
  - No virtual routing tables required

How 6PE Works

6PE Building Blocks - BGP

- Multi Protocol Border Gateway Protocol (MP-BGP)
  - Next-Hop Attribute
  - MP_REACH_NLRI
    - Special BGP path attribute
    - Allows BGP to carry more than just IPv4
    - Subsequent/Address Family Identifier (AFI/SAFI)
- Label Distribution
  - MP-BGP can carry label information in NLRI
- Route Reflection
  - Can be done per-address family
  - Can be separate dedicated router
- BGP is “control plane”

6PE Building Blocks - MPLS

- Two Major Protocols
  - Resource Reservation Protocol-Traffic Engineering (RSVP-TE)
  - Label Distribution Protocol (LDP)
    - Can be used in conjunction with RSVP-TE
      - Tunneled over RSVP-TE or parallel
- Label bindings for all prefixes
    - Policies can allow for bindings to only a subset of addresses
- If BGP is control plane, MPLS is forwarding plane

6PE Building Blocks – BGP+MPLS

- Creates control plane/forwarding plane concept
- Control Plane
  - BGP carries full routing table
    - BGP Next-hop is loopback of router
    - Can use external route-reflector or hop-by-hop
  - Carries the IPv6 Information
- Forwarding Plane
  - Interior Gateway Protocol (IGP)
    - Pick your flavor
  - Loopbacks/Router ID must be in IGP
  - Label binding for each loopback address
  - Carries the IPv4 information

6PE Building Blocks

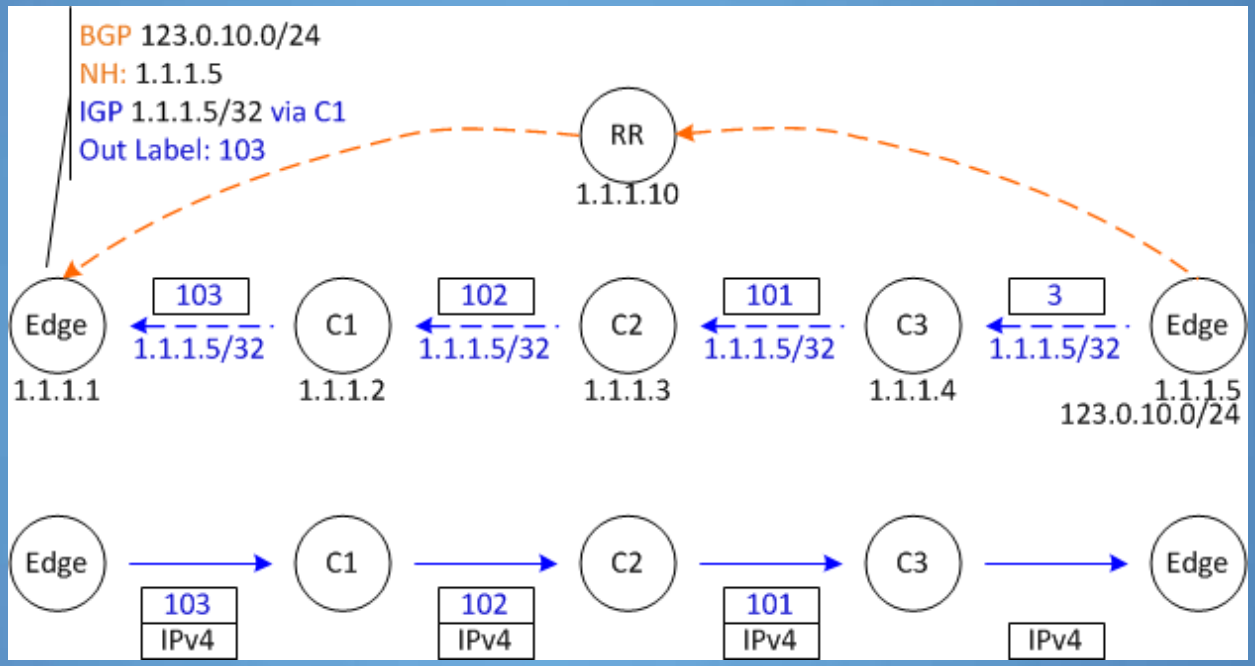

6PE Core Components

- MP-BGP sends label in route advertisement
  - AFI of 2 for IPv6, SAFI of 4 for labeled BGP
    - Configure BGP to send label for IPv6
    - Label is mapped to BGP next-hop
      - Distinct from IGP label
      - May be several IPv6 prefixes within same MP_REACH_NLRI
  - BGP Next-Hop uses special syntax
    - IPv6 BGP table uses ::FFFF: <IPv4> address
    - Router knows to encapsulate in LDP or RSVP

6PE Core Components

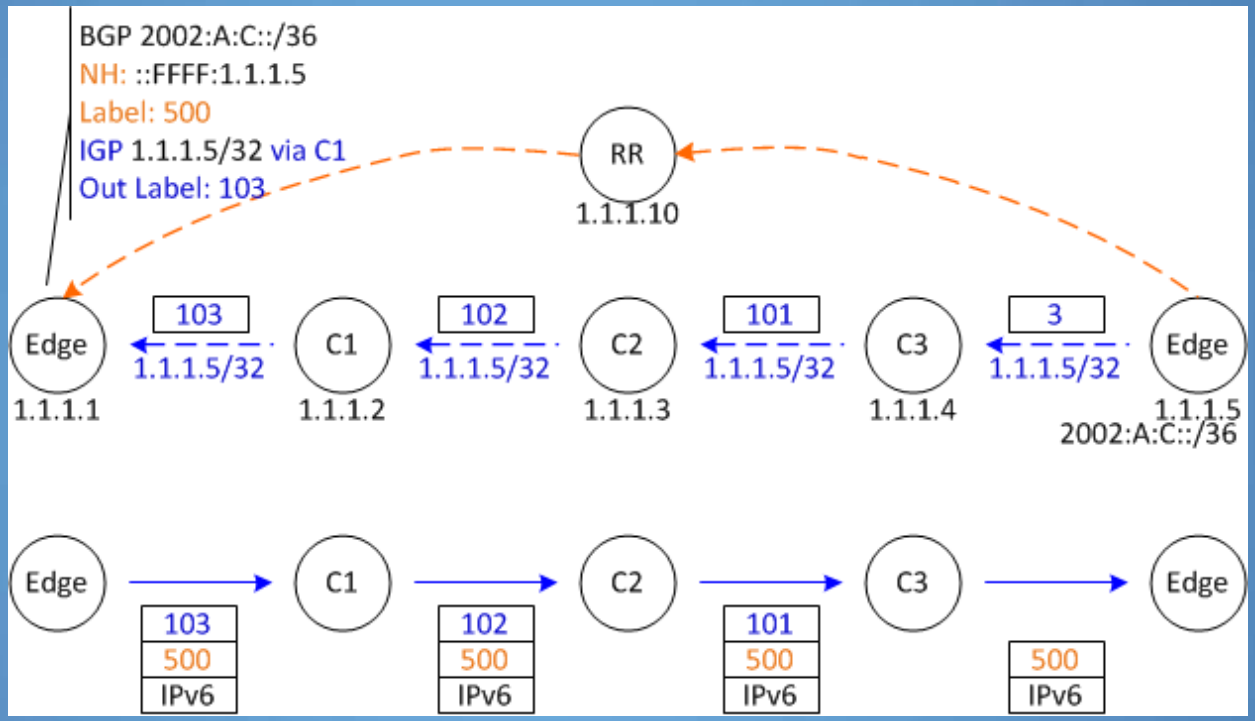

6PE

Cisco IOS-XR (Edge)

```
router bgp 65412
 address-family ipv6 unicast
  network 2002:A:C::/36
  allocate-label all
 !
 neighbor-group 6PE
  use session-group 6PE-SESSION-GROUP
  !
  address-family ipv6 labeled-unicast
   route-policy FOO-IN in
   route-policy BAR-OUT out
 !
 neighbor 1.1.1.10
  use neighbor-group 6PE
  description 6PE-RouteReflector
```

Juniper JUNOS (RR)

```
group 6PE-RR {
    local-address 1.1.1.10;
    family inet6 {
        labeled-unicast {
            explicit-null;
        }
    }
    export FOO-OUT;
    neighbor 1.1.1.5 {
        description 6PE-Client;
    }
}
```

6PE

RP/0/RP0/CPU0:ROUTER#show route ipv6 bgp

B **2002:A:C::/36**
[200/0] via ::, 1w1d, Null0
B **2002:B:D::/36**
[200/0] via **::ffff:1.1.1.1** (nexthop in vrf default), 6d05h

RP/0/RP0/CPU0:ROUTER#show bgp ipv6 labeled-unicast
BGP router identifier 1.1.1.5, local AS number 65412

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 2002:A:C::/36					
	::			32768	i
*>i2002:B:D::/36					
	1.1.1.1	0	100	0	i

Processed 2 prefixes, 2 paths

6PE

clenart@RR# run show route table inet6.0 detail

2002:A:C::/36 (1 entry, 1 announced)
*BGP Preference: 170/-101
Next hop type: Indirect
Next-hop reference count: 6
Source: 1.1.1.5
Next hop type: Discard
Protocol next hop: **::ffff:1.1.1.5**
State: <Active Int Ext>
Local AS: 65412 Peer AS: 65412
Age: 1w1d 9:21:35 Metric2: 0
Task: BGP_65412.1.1.1.1+179
Announcement bits (3): 0-KRT 4-BGP RT Background 5-Resolve tree 3
AS path: I (Atomic) Aggregator: 65412 1.1.1.5
Localpref: 100
Router ID: 1.1.1.5
[edit]

6PE Core Components


```
Path attributes
  Path Attribut - MP_REACH_NLRI
    Flags: 0x90: Optional, Non-transitive, Complete, Extended Length
    Type Code: MP_REACH_NLRI (14)
    Length: 2041
    Address family: IPv6 (2)
    Subsequent address family identifier: Labeled Unicast (4)
    Next hop network address (16 bytes)
      Next hop: ::ffff:172.16.64.10 (16)
    Subnetwork points of attachment: 0
    Network layer reachability information (2020 bytes)
      Label Stack=16001 (bottom), IPv6=2002:202::1/128
      Label Stack=16001 (bottom), IPv6=2002:202::2/128
      Label Stack=16001 (bottom), IPv6=2002:202::3/128
      Label Stack=16001 (bottom), IPv6=2002:202::4/128
      Label Stack=16001 (bottom), IPv6=2002:202::5/128
      Label Stack=16001 (bottom), IPv6=2002:202::6/128
      Label Stack=16001 (bottom), IPv6=2002:202::7/128
```

6PE Use Cases

Rapid Deployment

- When a customer needs IPv6 immediately
- Takes advantage of ease of deployment
  - Simple edge configuration for end-to-end IPv6
  - Only edge routers are aware of 6PE
  - Only minor changes to BGP are necessary
    - Could be as simple as adding new address-family
- Customers still use native IPv6 or dual stack
- Only “rapid” if BGP/MPLS network exists

IPv6 Free Core

- Absolutely no IPv6 in core network
- Possibly because of roadblocks
  - Older hardware
  - Dated software
- Reduce scope of IPv6 work
- 6PE provides tunneling for IPv6

BGP Free Core

- IPv6 can be in core, but no BGP
- Primarily MPLS label switch routers
  - Run LDP and/or RSVP-TE with limited IGP
  - Leverages cost optimized hardware, less memory
  - Route-reflectors required to exchange BGP between edges
- BGP Free Core Challenge
  - No BGP, no native IPv4 or IPv6 routing
  - MPLS is absolutely necessary to forward IPv6
  - 6PE is used to tunnel IPv6 using IPv4 and MPLS

6VPE

- Customer routing table on PE is in VRF
- Uses VPNv6 MP-iBGP
  - AFI/SAFI of 2/128
  - Route distinguisher and route targets
  - BGP Label
- Uses mapped ::FFFF:<IPv4> address
- Very similar to IPv4 BGP/MPLS L3VPNs

6PE MPLS Alternatives

- Psuedowires (AToM)
  - Backbone does not need any IPv6
  - Not scalable
- Tunnel IPv6 via IPv4 BGP/MPLS VPN
  - CE is dual stack router
  - CE creates tunnel IPv6 via GRE, 6to4, Manual, etc.
  - No IPv6 in backbone
- Tunnel IPv6 over RSVP-TE
  - Via forwarding-adjacency or static routes

Steps Forward

- LDPv6
  - Back to Dual Stack Native w/ MPLS
  - Creates forwarding separation for IPv4 and IPv6
  - Uses IPv4 or IPv6 address for single transport connection
  - “Somewhat undeployable in existing production networks”
- RSVPv6
  - Could be just as useful to use RSVPv4
  - Could add extra state to network
- Keep 6PE for as long as possible
  - Already there, why not?
  - May hide IPv6 traffic statistics
  - Only have control over IPv6 on edges

6PE Summary

- Control and Forwarding Plane
- Leverages MP-BGP
  - AFI/SAFI of 2 (IPv6)/4 (Labeled BGP)
  - ::FFFF:<IPv4 Address>
    - Used as Next-Hop for BGP
    - Best to use loopback
- MPLS and IGP tunnel IPv6 through backbone
- Only minimal configuration on edges
- Can “Bolt On” to BGP/MPLS network

6PE Summary

- Advantages
  - Mature Technology
  - Easily added to existing BGP/MPLS network
  - Allows for flexibility at the edge
  - No NAT necessary
    - Any mapping (e.g. FFFF) done automatically
  - Less work in core
  - Scalable
- Disadvantages
  - May be stuck with it forever
  - Need BGP and MPLS already or not worth it
- Can coexist with native IPv6

Q&A

Q&A