www.iparchitechs.com 1-855-MIKROTIK

BGP as an IGP for Carrier/Enterprise Networks

Presented by:

Kevin Myers, SENIOR NETWORK ENGINEER

IP ArchiTECHS MANAGED SERVICES

Background

- **Kevin Myers**
  - 15 + years in IT/Network Engineering
  - Designed and implemented networks in Service Provider, Enterprise and Government environments
  - Areas of Design Focus:
    - MikroTik/Cisco integration
    - Design of BGP/MPLS/OSPF Service Provider Triple-Play networks
    - Design of Enterprise Data Center networks
  - Certifications
    - Pursuing CCIE Route/Switch
    - Certified – CCNP, CCNA, MCP, MTCRE, MTCTCE, MTCNA

IP ArchiTechs Managed Services

- Exhibitor at 2013 MUM – Please stop by our exhibitor booth and register to win an android tablet
- The first Carrier-Grade 24/7/365 MikroTik TAC (Technical Assistance Center)
  - Three tiers of engineering support
  - Monthly and per incident pricing available
  - 1-855-MIKROTIK or support.iparchitechs.com
- Private Nationwide 4G LTE MPLS backbone
  - Partnership with Verizon Wireless - available anywhere in the Verizon service area
  - Not Internet facing – privately routed over our MPLS infrastructure
  - Point-to-Point or Point-to-MultiPoint
- Proactive Monitoring / Ticketing / Change Control / IPAM (IP Address Management)
- Carrier-Grade Network Engineering / Design in large (10,000+ nodes) environments

24/7/365 MikroTik TAC **Nationwide Private 4G LTE MPLS** **Proactive Network Monitoring** **Design / Engineering / Operations**

Introduction - IGP vs EGP – which is which?

- **IGP** - Interior Gateway Protocol
  - is a routing protocol by which elements comprising an autonomous system (AS) exchange routing information
  - Examples: RIP, OSPF, EIGRP, ISIS
  - Mikrotik: RIP, OSPF (OSPF is recommended)
- **EGP** – Exterior Gateway Protocol
  - is a routing protocol by which elements in different autonomous systems (AS) exchange routing information
  - Example: BGP

IGP vs EGP (Continued)

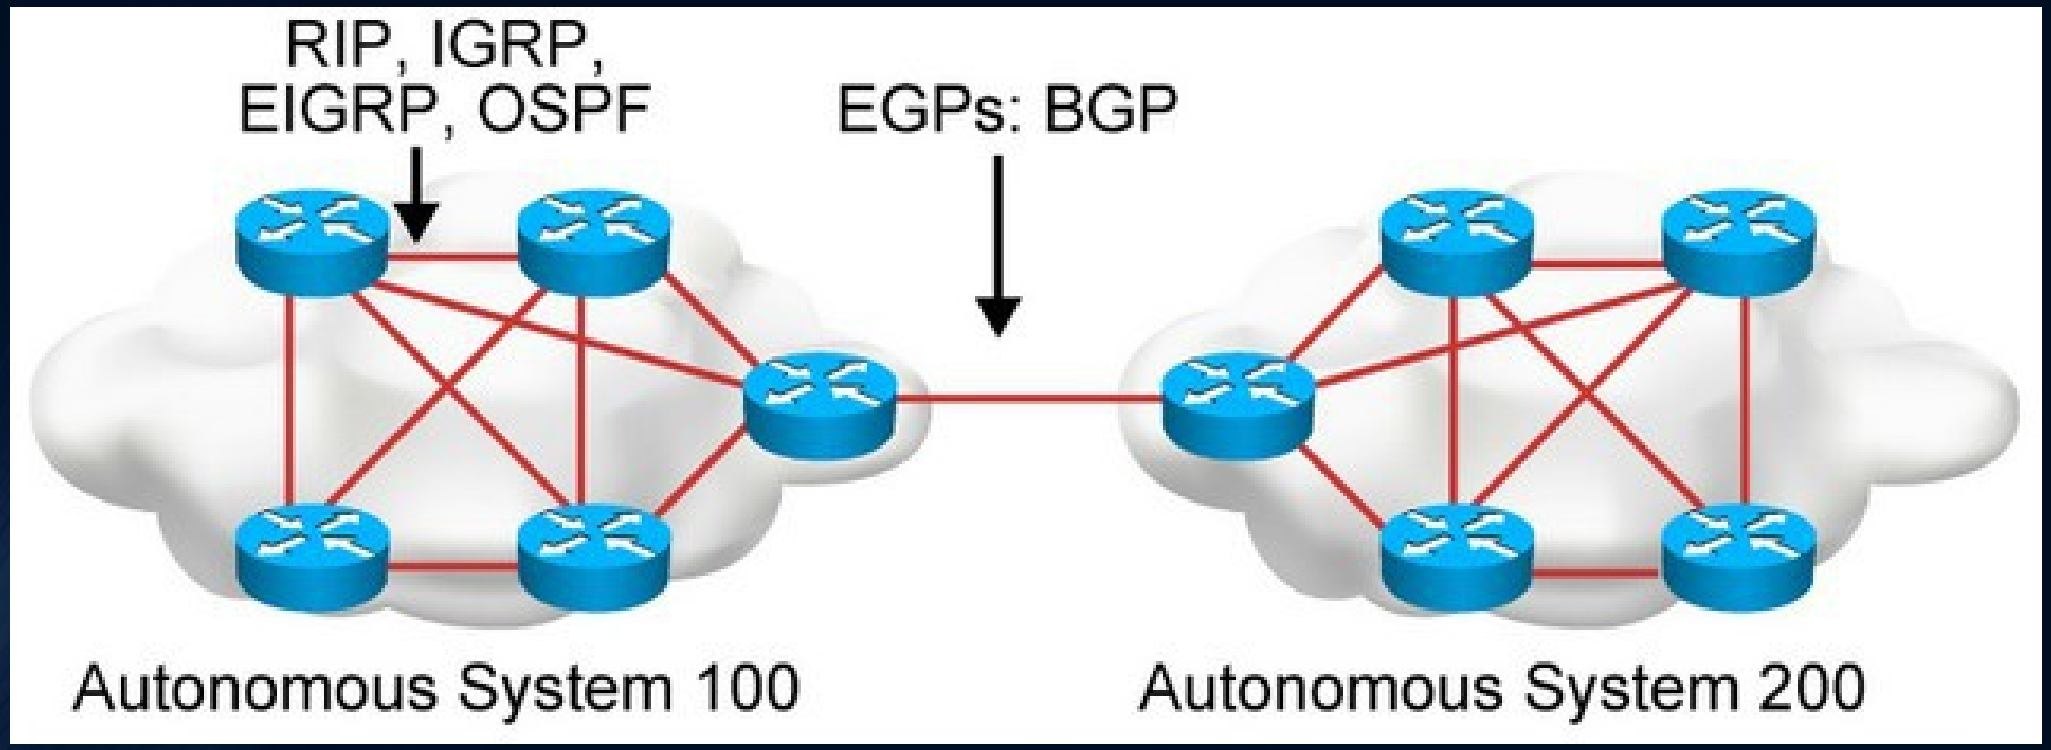

Why use BGP to replace or complement a standalone IGP?

- BGP Provides more control of routing policy than any other routing protocol
- BGP throughout the network simplifies new routing changes – end-to-end architecture
- Limits redistribution which increases network stability/uptime
- Two different approaches
  - BGP only on all routers (less common)
  - BGP hybrid with an IGP (OSPF) underneath to advertise loopbacks and transit subnets
  - The BGP/OSPF hybrid model is preferred and is frequently used as a global standard to build MPLS

Exploring the hybrid - BGP over OSPF

- Understanding Transit vs. Traffic subnets
  - Transit (OSPF) - subnets that connect routers and links (also loopbacks)
  - Traffic (BGP) - subnets that are the origination/destination for traffic
- Configuration/Verification – configuration is NOT verifcation
- Examine the use of OSPF underneath BGP (Transit Layer)
- Examine the use of BGP on top of OSPF (Traffic Layer)
- Variations of BGP Design (iBGP, eBGP, Confederations)

Benefits of hybrid - BGP over OSPF

- Reduction of routing information in the IGP which increases stability
- Fast convergence and bandwidth aware capabilities of OSPF to build multiple network paths and connect transit subnets
- Policy control and "nerd knobs" of BGP can be tuned to provide extremely powerful administrative control of the flow of traffic subnets
- Adds ECMP capability to BGP by utilizing the underlying OSPF ECMP routes
- Peering via Loopbacks that have multiple OSPF paths provides higher availability for BGP sessions

Use of an underlying IGP to connect Transit Subnets

- OSPF (Preferred IGP in RouterOS)
  - used for rapid convergence
  - advertises transit subnets only, not traffic subnets
  - also used to advertise "loopback" addresses
  - Config Recommendations
    - varies by network design and physical connectivity – RF / Fiber / Copper etc
    - Hello Timer - 1 second
    - Dead Timer - 4 seconds
    - OPSF Network Type - PtP if possible to eliminate DR/BDR election

```
/routing ospf interface
add dead-interval=4s hello-interval=1s interface=bonding1 network-type=point-to-point
```

  - MikroTik Loopback - bridge interface with no other bridge port members

```
/interface bridge
add name=Loopback0
```

Use of an underlying IGP to connect Transit Subnets

- Transit (OSPF) - subnets that connect routers and links

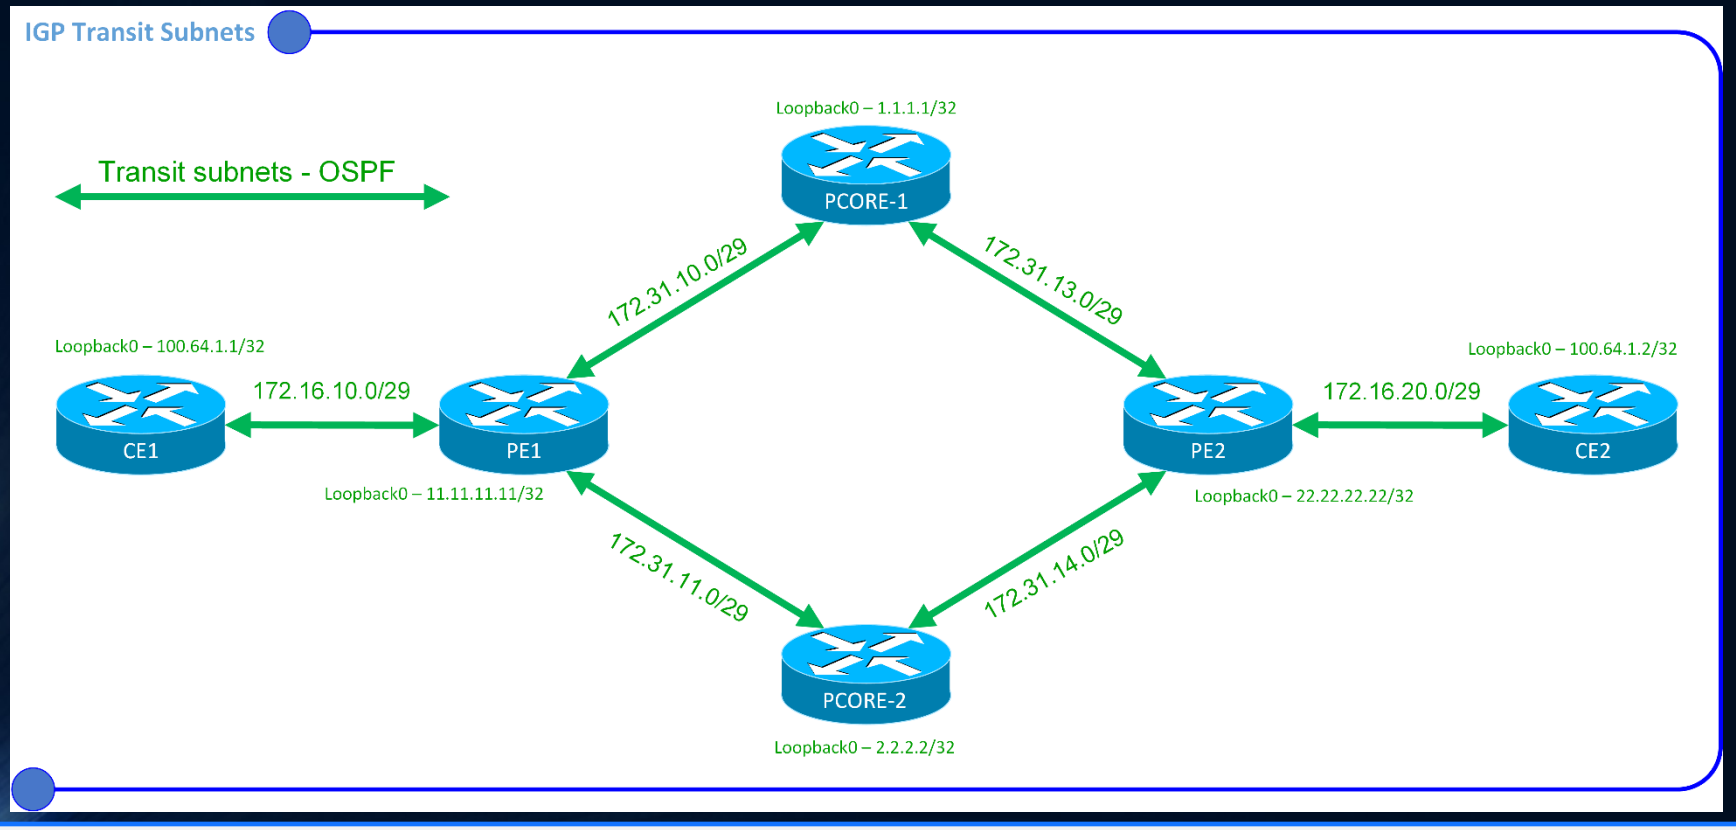

Use of an underlying IGP - Configuration

- Configure (example is PE1 – full configs/presentation will be posted on mum.iparchitechs.com)

```
[admin@PE1] > routing ospf export
# sep/15/2013 22:58:21 by RouterOS 5.19
# software id = HDT2-V13L
#
/routing ospf instance
set [ find default=yes ] disabled=no distribute-default=never in-filter=ospf-in metric-bgp=auto metric-connected=20 \
    metric-default=1 metric-other-ospf=auto metric-rip=20 metric-static=20 name=default out-filter=ospf-out \
    redistribute-bgp=no redistribute-connected=no redistribute-other-ospf=no redistribute-rip=no redistribute-static=no \
    router-id=0.0.0.0
/routing ospf area
set [ find default=yes ] area-id=0.0.0.0 disabled=no instance=default name=backbone type=default
/routing ospf interface
add authentication=none authentication-key="" authentication-key-id=1 cost=10 dead-interval=4s disabled=no \
    hello-interval=1s instance-id=0 interface=bonding1 network-type=point-to-point passive=no priority=1 \
    retransmit-interval=5s transmit-delay=1s use-bfd=no
add authentication=none authentication-key="" authentication-key-id=1 cost=10 dead-interval=4s disabled=no \
    hello-interval=1s instance-id=0 interface=bonding2 network-type=point-to-point passive=no priority=1 \
    retransmit-interval=5s transmit-delay=1s use-bfd=no
add authentication=none authentication-key="" authentication-key-id=1 cost=10 dead-interval=4s disabled=no \
    hello-interval=1s instance-id=0 interface=ether2 network-type=point-to-point passive=no priority=1 \
    retransmit-interval=5s transmit-delay=1s use-bfd=no
/routing ospf network
add area=backbone disabled=no network=11.11.11.11/32
add area=backbone disabled=no network=172.31.10.0/29
add area=backbone disabled=no network=172.31.11.0/29
add area=backbone disabled=no network=172.16.10.0/29
```

Use of an underlying IGP - Verification

- Verify LSA Database (abbrev.) and Routing Table (on PCORE-1 and PE2)

```
[admin@PCORE-1] > routing ospf lsa print detail
 instance=default area=backbone type=router id=172.16.20.2
   originator=172.16.20.2 sequence-number=0x80000006 age=831 checksum=0x3009
   options="E" body=
     flags=
         link-type=Stub id=100.64.1.2 data=255.255.255.255 metric=10
         link-type=Transit id=172.16.20.1 data=172.16.20.2 metric=10

 instance=default area=backbone type=router id=172.31.10.1
   originator=172.31.10.1 sequence-number=0x80000005 age=1290 checksum=0x7DE0
   options="E" body=
     flags=
         link-type=Stub id=100.64.1.1 data=255.255.255.255 metric=10
         link-type=Point-To-Point id=172.31.11.1 data=172.16.10.2 metric=10
         link-type=Stub id=172.16.10.0 data=255.255.255.248 metric=10

 instance=default area=backbone type=router id=172.31.11.1
   originator=172.31.11.1 sequence-number=0x80000010 age=1396 checksum=0xBFA5
   options="E" body=
     flags=
         link-type=Point-To-Point id=172.31.10.1 data=172.16.10.1 metric=10
         link-type=Stub id=172.16.10.0 data=255.255.255.248 metric=10
         link-type=Stub id=11.11.11.11 data=255.255.255.255 metric=10
         link-type=Point-To-Point id=172.31.13.2 data=172.31.10.1 metric=10
         link-type=Stub id=172.31.10.0 data=255.255.255.248 metric=10
         link-type=Point-To-Point id=172.31.14.2 data=172.31.11.1 metric=10
         link-type=Stub id=172.31.11.0 data=255.255.255.248 metric=10
```

```
[admin@PE2] > ip route print where ospf
Flags: X - disabled, A - active, D - dynamic,
C - connect, S - static, r - rip, b - bgp, o - ospf, m - mme,
B - blackhole, U - unreachable, P - prohibit
 #      DST-ADDRESS        PREF-SRC        GATEWAY            DISTANCE
 9 ADo  1.1.1.1/32                         172.31.13.2             110
10 ADo  2.2.2.2/32                         172.31.14.2             110
11 ADo  11.11.11.11/32                     172.31.13.2             110
                                           172.31.14.2
12 ADo  100.64.1.1/32                      172.31.13.2             110
                                           172.31.14.2
13 ADo  100.64.1.2/32                      172.16.20.2             110
14 ADo  172.16.10.0/29                     172.31.13.2             110
                                           172.31.14.2
15 ADo  172.31.10.0/29                     172.31.13.2             110
16 ADo  172.31.11.0/29                     172.31.14.2             110
```

Use BGP to advertise Traffic Subnets

- BGP Designs
  - Can be iBGP, eBGP or iBGP/eBGP hybrid
    - Differences between these BGP types will be covered later in the presentation
    - iBGP/eBGP hybrid is the model we will explore
  - advertises traffic subnets only, not transit subnets (That's OSPF's Job!)
  - Peerings should be sourced from "loopback" addresses
  - Use BGP Multihop when peering eBGP with loopbacks
  - Route Filters can be used to manage traffic, advertisements and prefixes

Use iBGP and/or eBGP advertise Traffic Subnets

- Traffic (iBGP or eBGP) - advertised at the edge and carry end user traffic

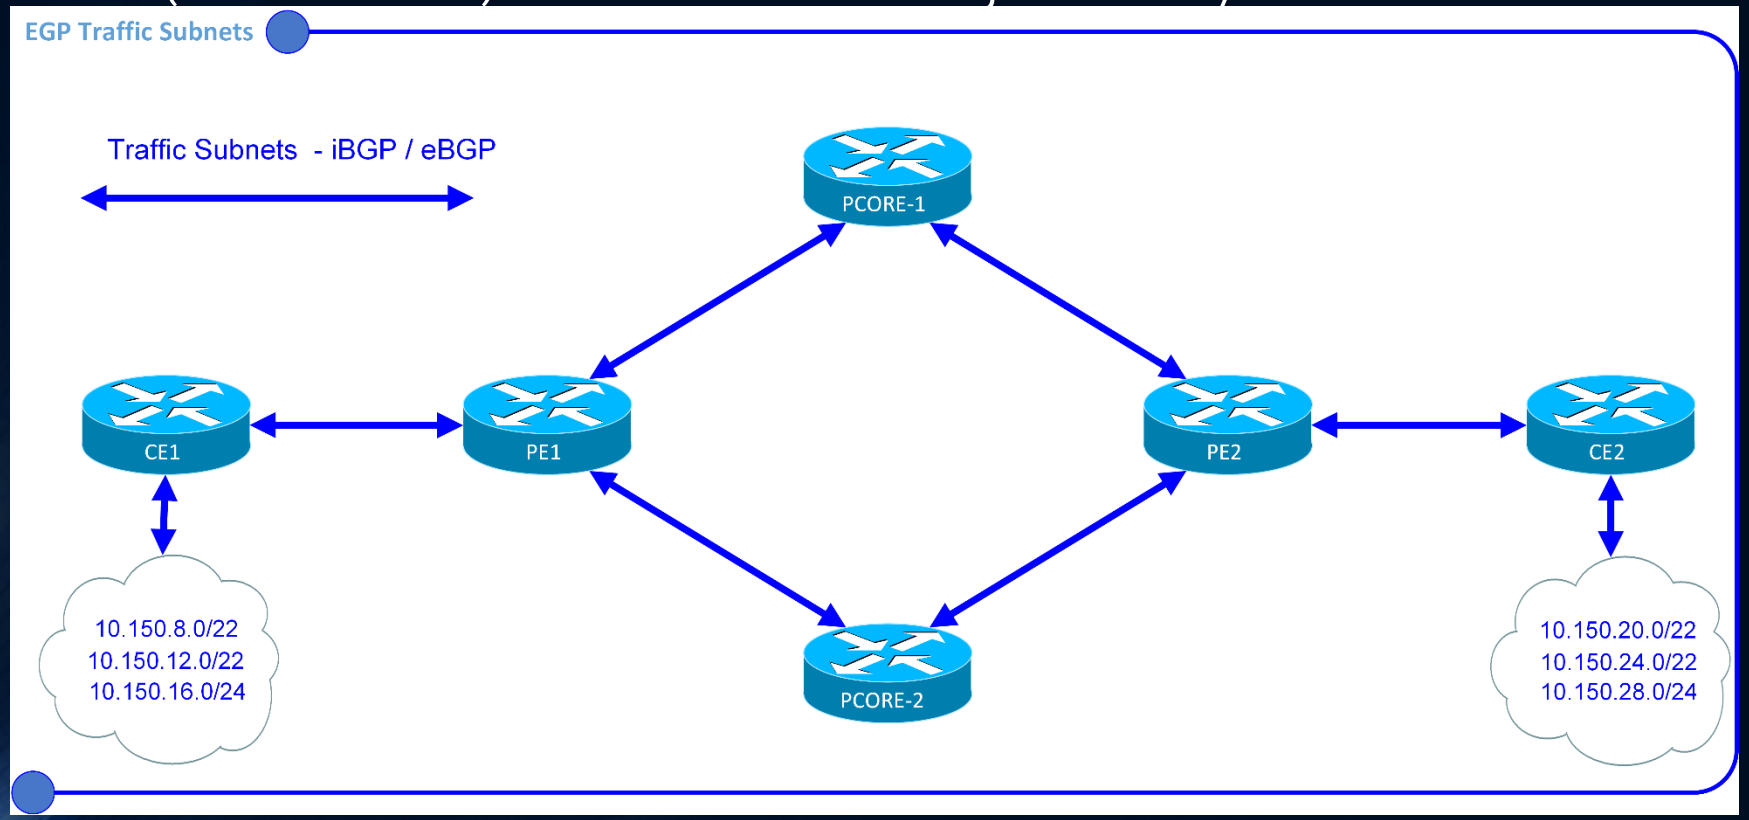

Use iBGP and/or eBGP - Configuration

- Configure (example is CE1 – full configs will be posted on mum.iparchitechs.com)

```
[admin@CE1] > routing bgp export
# sep/15/2013 23:15:31 by RouterOS 5.19
# software id = HDT2-V13L
#
/routing bgp instance
set default as=10 client-to-client-reflection=yes disabled=no ignore-as-path-len=no name=default \
    out-filter="" redistribute-connected=no redistribute-ospf=no redistribute-other-bgp=no \
    redistribute-rip=no redistribute-static=no router-id=0.0.0.0 routing-table=""
/routing bgp network
add disabled=no network=10.150.8.0/22 synchronize=no
add disabled=no network=10.150.12.0/22 synchronize=no
add disabled=no network=10.150.16.0/24 synchronize=no
/routing bgp peer
add address-families=ip as-override=no default-originate=never disabled=no hold-time=3m in-filter="" \
    instance=default multihop=yes name=peer1 nexthop-choice=default out-filter="" passive=no \
    remote-address=11.11.11.11 remote-as=100 remove-private-as=no route-reflect=no tcp-md5-key="" ttl=\
    default update-source=Loopback0 use-bfd=no
```

Use iBGP and/or eBGP - Verification

- Verify BGP Advertisements and Routing Table (on PCORE-1 and PE2)

```
[admin@PCORE-1] > routing bgp advertisements print
PEER     PREFIX              NEXTHOP          AS-PATH     ORIGIN     LOCAL-PREF
peer1    10.150.8.0/22       100.64.1.1       10          igp               100
peer1    10.150.12.0/22      100.64.1.1       10          igp               100
peer1    10.150.16.0/24      100.64.1.1       10          igp               100
peer2    10.150.28.0/24      100.64.1.2       20          igp               100
peer2    10.150.24.0/22      100.64.1.2       20          igp               100
peer2    10.150.20.0/22      100.64.1.2       20          igp               100
```

```
[admin@PE2] > ip route print where bgp
Flags: X - disabled, A - active, D - dynamic,
C - connect, S - static, r - rip, b - bgp, o - ospf, m - mme,
B - blackhole, U - unreachable, P - prohibit
 #      DST-ADDRESS        PREF-SRC        GATEWAY            DISTANCE
 0 ADb  10.150.8.0/22                      100.64.1.1              200
 1 ADb  10.150.12.0/22                     100.64.1.1              200
 2 ADb  10.150.16.0/24                     100.64.1.1              200
 3 ADb  10.150.20.0/22                     100.64.1.2               20
 4  Db  10.150.20.0/22                     100.64.1.2              200
 5 ADb  10.150.24.0/22                     100.64.1.2               20
 6  Db  10.150.24.0/22                     100.64.1.2              200
 7 ADb  10.150.28.0/24                     100.64.1.2               20
 8  Db  10.150.28.0/24                     100.64.1.2              200
```

Use BGP/OSPF together to complete the design

- Transit and Traffic - Traffic (OPSF/BGP) – Complete routing topology

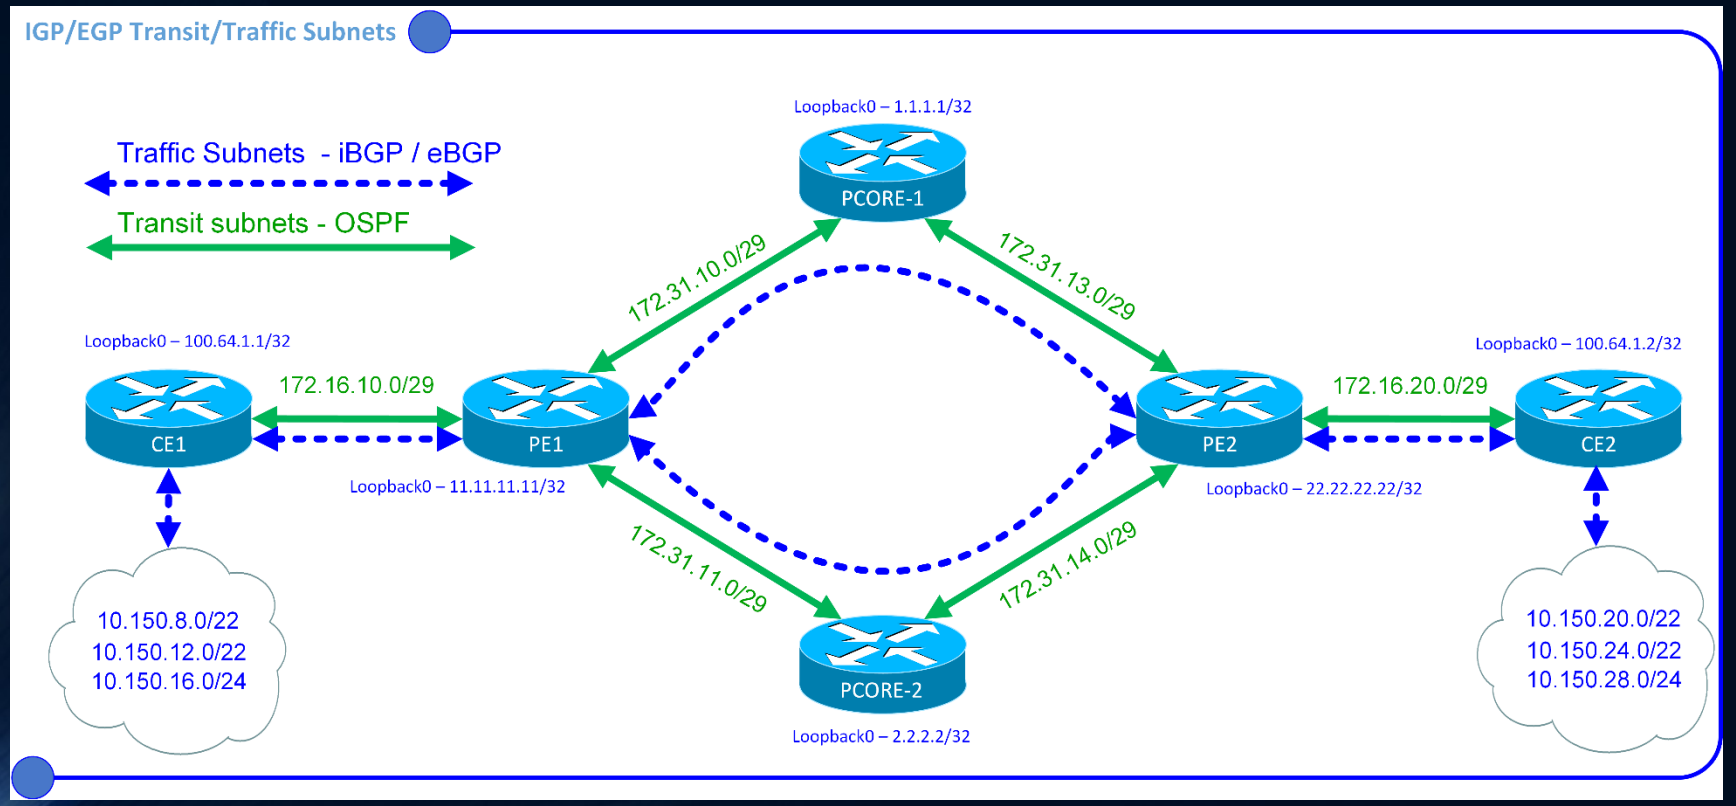

Use BGP/OSPF together - Verification

- Verify full BGP/OSPF routing tables (CE1 and CE2)

```
[admin@CE1] > ip route print
Flags: X - disabled, A - active, D - dynamic,
C - connect, S - static, r - rip, b - bgp, o - ospf, m - mme,
B - blackhole, U - unreachable, P - prohibit
 #      DST-ADDRESS        PREF-SRC        GATEWAY            DISTANCE
 0 ADo  1.1.1.1/32                         172.16.10.1             110
 1 ADo  2.2.2.2/32                         172.16.10.1             110
 2 ADC  10.150.8.0/22      10.150.8.1      TrafficLAN-1              0
 3 ADC  10.150.12.0/22     10.150.12.1     TrafficLAN-2              0
 4 ADC  10.150.16.0/24     10.150.16.1     TrafficLAN-3              0
 5 ADb  10.150.20.0/22                     100.64.1.2               20
 6 ADb  10.150.24.0/22                     100.64.1.2               20
 7 ADb  10.150.28.0/24                     100.64.1.2               20
 8 ADo  11.11.11.11/32                     172.16.10.1             110
 9 ADo  22.22.22.22/32                     172.16.10.1             110
10 ADC  100.64.1.1/32      100.64.1.1      Loopback0                 0
11 ADo  100.64.1.2/32                     172.16.10.1             110
12 ADC  172.16.10.0/29     172.16.10.2     ether2                    0
13 ADo  172.16.20.0/29                    172.16.10.1             110
14 ADo  172.31.10.0/29                    172.16.10.1             110
15 ADo  172.31.11.0/29                    172.16.10.1             110
16 ADo  172.31.13.0/29                    172.16.10.1             110
17 ADo  172.31.14.0/29                    172.16.10.1             110
```

```
[admin@CE2] > ip route print
Flags: X - disabled, A - active, D - dynamic,
C - connect, S - static, r - rip, b - bgp, o - ospf, m - mme,
B - blackhole, U - unreachable, P - prohibit
 #      DST-ADDRESS        PREF-SRC        GATEWAY            DISTANCE
 0 ADo  1.1.1.1/32                         172.16.20.1             110
 1 ADo  2.2.2.2/32                         172.16.20.1             110
 2 ADb  10.150.8.0/22                      100.64.1.1               20
 3 ADb  10.150.12.0/22                     100.64.1.1               20
 4 ADb  10.150.16.0/24                     100.64.1.1               20
 5 ADC  10.150.20.0/22     10.150.20.1     TrafficLAN-1              0
 6 ADC  10.150.24.0/22     10.150.24.1     TrafficLAN-2              0
 7 ADC  10.150.28.0/24     10.150.28.1     TrafficLAN-3              0
 8 ADo  11.11.11.11/32                     172.16.20.1             110
 9 ADo  22.22.22.22/32                     172.16.20.1             110
10 ADo  100.64.1.1/32                     172.16.20.1             110
11 ADC  100.64.1.2/32     100.64.1.2      Loopback0                 0
12 ADo  172.16.10.0/29                    172.16.20.1             110
13 ADC  172.16.20.0/29    172.16.20.2     ether2                    0
14 ADo  172.31.10.0/29                    172.16.20.1             110
15 ADo  172.31.11.0/29                    172.16.20.1             110
16 ADo  172.31.13.0/29                    172.16.20.1             110
17 ADo  172.31.14.0/29                    172.16.20.1             110
```

BGP – iBGP vs eBGP

- What is iBGP?
  - iBGP – Internal Border Gateway Protocol
    - BGP Peering between routers in the same Autonomous System (AS)
    - Does NOT refer to public vs private IP Addresses
    - iBGP can carry public and/or private IP routing information
    - iBGP Full Mesh – requires peerings between all routers or route reflection
- What is eBGP?
  - eBGP – External Border Gateway Protocol
    - BGP Peering between routers in different Autonomous Systems (AS)
    - Does NOT refer to public vs private IP Addresses
    - eBGP can carry public and/or private IP routing information

iBGP in the Network Core and Provider Edge

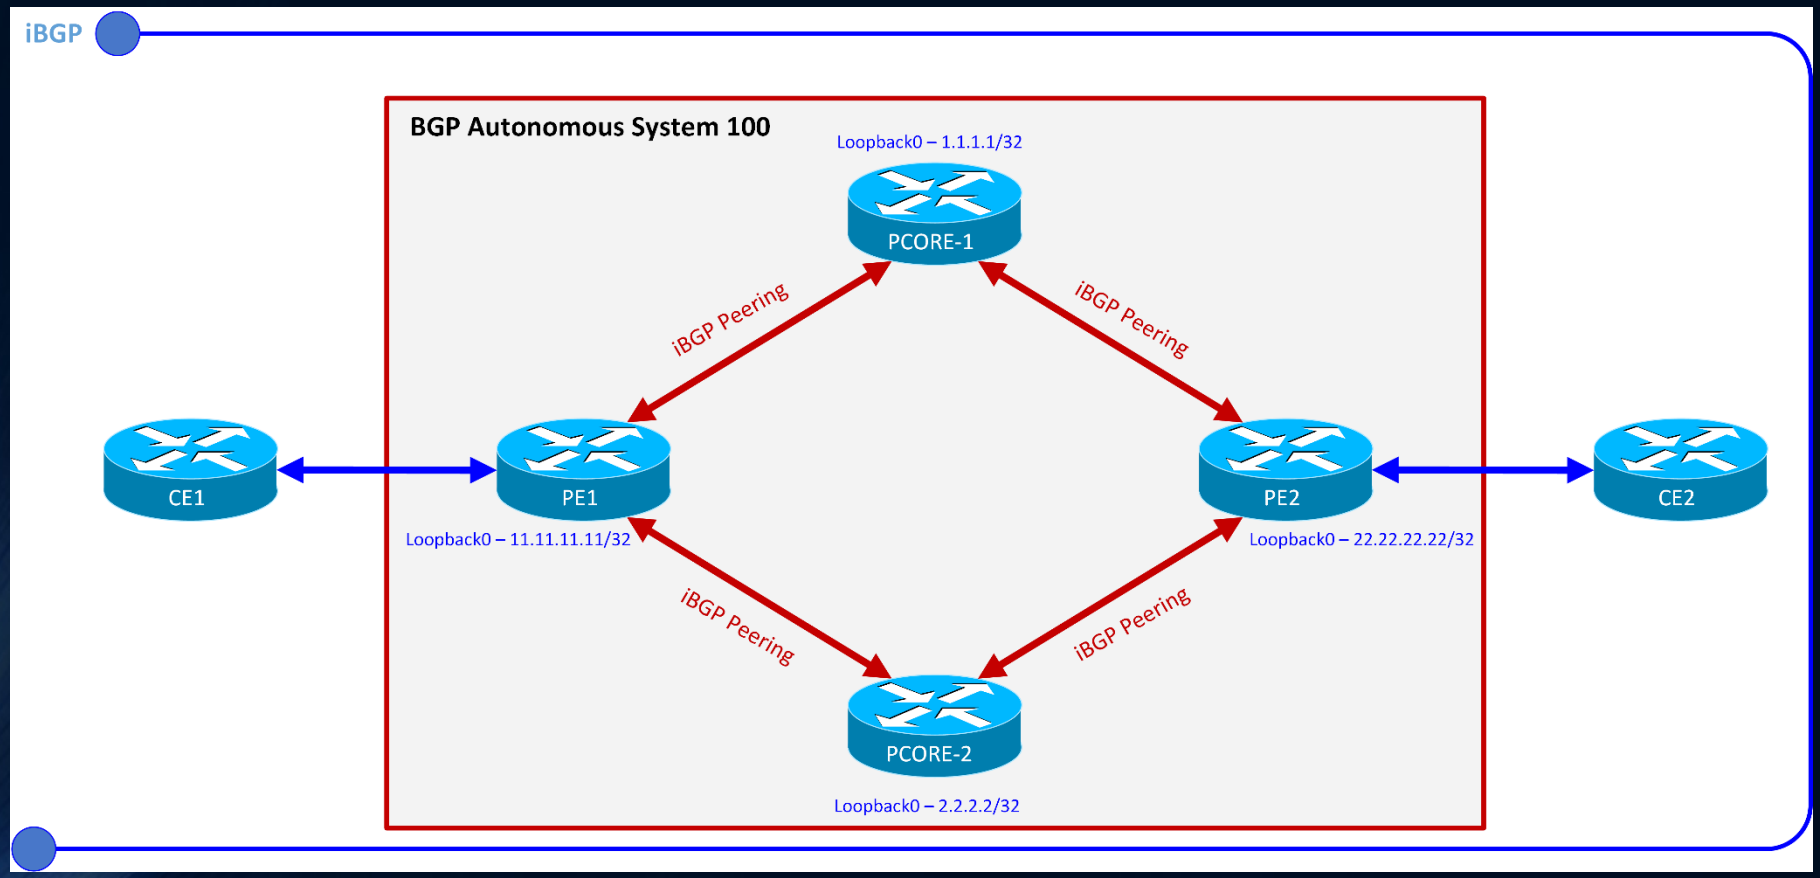

iBGP in the Network Core and Provider Edge

- Configuration - PE1

```
[admin@PE1] > routing bgp peer export
# sep/16/2013 04:06:50 by RouterOS 5.19
# software id = HDT2-V13L
#
/routing bgp peer
add address-families=ip as-override=no default-originate=\
    never disabled=no hold-time=3m in-filter="" instance=\
    default multihop=no name=peer1 nexthop-choice=default \
    out-filter="" passive=no remote-address=1.1.1.1 \
    remote-as=100 remove-private-as=no route-reflect=yes \
    tcp-md5-key="" ttl=default update-source=Loopback0 \
    use-bfd=no
add address-families=ip as-override=no default-originate=\
    never disabled=no hold-time=3m in-filter="" instance=\
    default multihop=no name=peer2 nexthop-choice=default \
    out-filter="" passive=no remote-address=2.2.2.2 \
    remote-as=100 remove-private-as=no route-reflect=yes \
    tcp-md5-key="" ttl=default update-source=Loopback0 \
    use-bfd=no
add address-families=ip as-override=no default-originate=\
    never disabled=no hold-time=3m in-filter="" instance=\
    default multihop=yes name=peer3 nexthop-choice=default \
    out-filter="" passive=no remote-address=100.64.1.1 \
    remote-as=10 remove-private-as=no route-reflect=no \
    tcp-md5-key="" ttl=default update-source=Loopback0 \
    use-bfd=no
```

- Configuration - PCORE-1

```
[admin@PCORE-1] > routing bgp peer export
# sep/16/2013 04:07:29 by RouterOS 5.19
# software id = HDT2-V13L
#
/routing bgp peer
add address-families=ip as-override=no default-originate=\
    never disabled=no hold-time=3m in-filter="" instance=\
    default multihop=no name=peer1 nexthop-choice=default \
    out-filter="" passive=no remote-address=22.22.22.22 \
    remote-as=100 remove-private-as=no route-reflect=yes \
    tcp-md5-key="" ttl=default update-source=Loopback0 \
    use-bfd=no
add address-families=ip as-override=no default-originate=\
    never disabled=no hold-time=3m in-filter="" instance=\
    default multihop=no name=peer2 nexthop-choice=default \
    out-filter="" passive=no remote-address=11.11.11.11 \
    remote-as=100 remove-private-as=no route-reflect=yes \
    tcp-md5-key="" ttl=default update-source=Loopback0 \
    use-bfd=no
[admin@PCORE-1] >
```

iBGP in the Network Core and Provider Edge

- Verification - PE1

```
[admin@PE1] > routing bgp peer print status
Flags: X - disabled, E - established
 0 E name="peer1" instance=default remote-address=1.1.1.1
     remote-as=100 tcp-md5-key="" nexthop-choice=default
     multihop=no route-reflect=yes hold-time=3m ttl=default
     in-filter="" out-filter="" address-families=ip
     update-source=Loopback0 default-originate=never
     remove-private-as=no as-override=no passive=no
     use-bfd=no remote-id=172.31.13.2
     local-address=11.11.11.11 uptime=11h28m52s
     prefix-count=3 updates-sent=11 updates-received=4
     withdrawn-sent=4 withdrawn-received=1
     remote-hold-time=3m used-hold-time=3m
     used-keepalive-time=1m refresh-capability=yes
     as4-capability=yes state=established

 1 E name="peer2" instance=default remote-address=2.2.2.2
     remote-as=100 tcp-md5-key="" nexthop-choice=default
     multihop=no route-reflect=yes hold-time=3m ttl=default
     in-filter="" out-filter="" address-families=ip
     update-source=Loopback0 default-originate=never
     remove-private-as=no as-override=no passive=no
     use-bfd=no remote-id=172.31.14.2
     local-address=11.11.11.11 uptime=11h28m56s
     prefix-count=3 updates-sent=13 updates-received=9
     withdrawn-sent=5 withdrawn-received=10
     remote-hold-time=3m used-hold-time=3m
     used-keepalive-time=1m refresh-capability=yes
     as4-capability=yes state=established
```

- Configuration - PCORE-1

```
[admin@PCORE-1] > routing bgp peer print status
Flags: X - disabled, E - established
 0 E name="peer1" instance=default remote-address=22.22.22.22
     remote-as=100 tcp-md5-key="" nexthop-choice=default
     multihop=no route-reflect=yes hold-time=3m ttl=default
     in-filter="" out-filter="" address-families=ip
     update-source=Loopback0 default-originate=never
     remove-private-as=no as-override=no passive=no
     use-bfd=no remote-id=172.31.13.1 local-address=1.1.1.1
     uptime=11h32m6s prefix-count=3 updates-sent=7
     updates-received=5 withdrawn-sent=4 withdrawn-received=1
     remote-hold-time=3m used-hold-time=3m
     used-keepalive-time=1m refresh-capability=yes
     as4-capability=yes state=established

 1 E name="peer2" instance=default remote-address=11.11.11.11
     remote-as=100 tcp-md5-key="" nexthop-choice=default
     multihop=no route-reflect=yes hold-time=3m ttl=default
     in-filter="" out-filter="" address-families=ip
     update-source=Loopback0 default-originate=never
     remove-private-as=no as-override=no passive=no
     use-bfd=no remote-id=172.31.11.1 local-address=1.1.1.1
     uptime=11h32m15s prefix-count=3 updates-sent=4
     updates-received=9 withdrawn-sent=1 withdrawn-received=6
     remote-hold-time=3m used-hold-time=3m
     used-keepalive-time=1m refresh-capability=yes
     as4-capability=yes state=established
```

iBGP – Route Reflection vs Full Mesh Peering

- Design Problem: iBGP cannot advertise any prefix it learns from an iBGP peer to any other iBGP peer (Also known as the BGP Split Horizon rule)
- Solution - Full Mesh
  - iBGP peering to every other iBGP router in the AS
  - 4 routers in an AS = 3 iBGP Peerings
  - 8 routers in an AS = 7 iBGP Peerings
  - 32 routers in an AS = 31 iBGP peerings
  - As the number of routers increases, iBGP peering to all routers becomes impractical and cumbersome.

iBGP – Route Reflection vs Full Mesh Peering

- Design Problem: iBGP cannot advertise any prefix it learns from an iBGP peer to any other iBGP peer (Also known as the BGP Split Horizon rule)
- Solution – Route Reflector
  - Can be deployed as a single router or multiple routers (known as a cluster)
  - Learns routes from one iBGP peer and reflects them to another iBGP peer
  - Achieves a Full Mesh topology without the extra peerings
  - Route Reflectors can be peered to each other to form a highly available cluster
  - Simple config in RouterOS

iBGP – Route Reflection

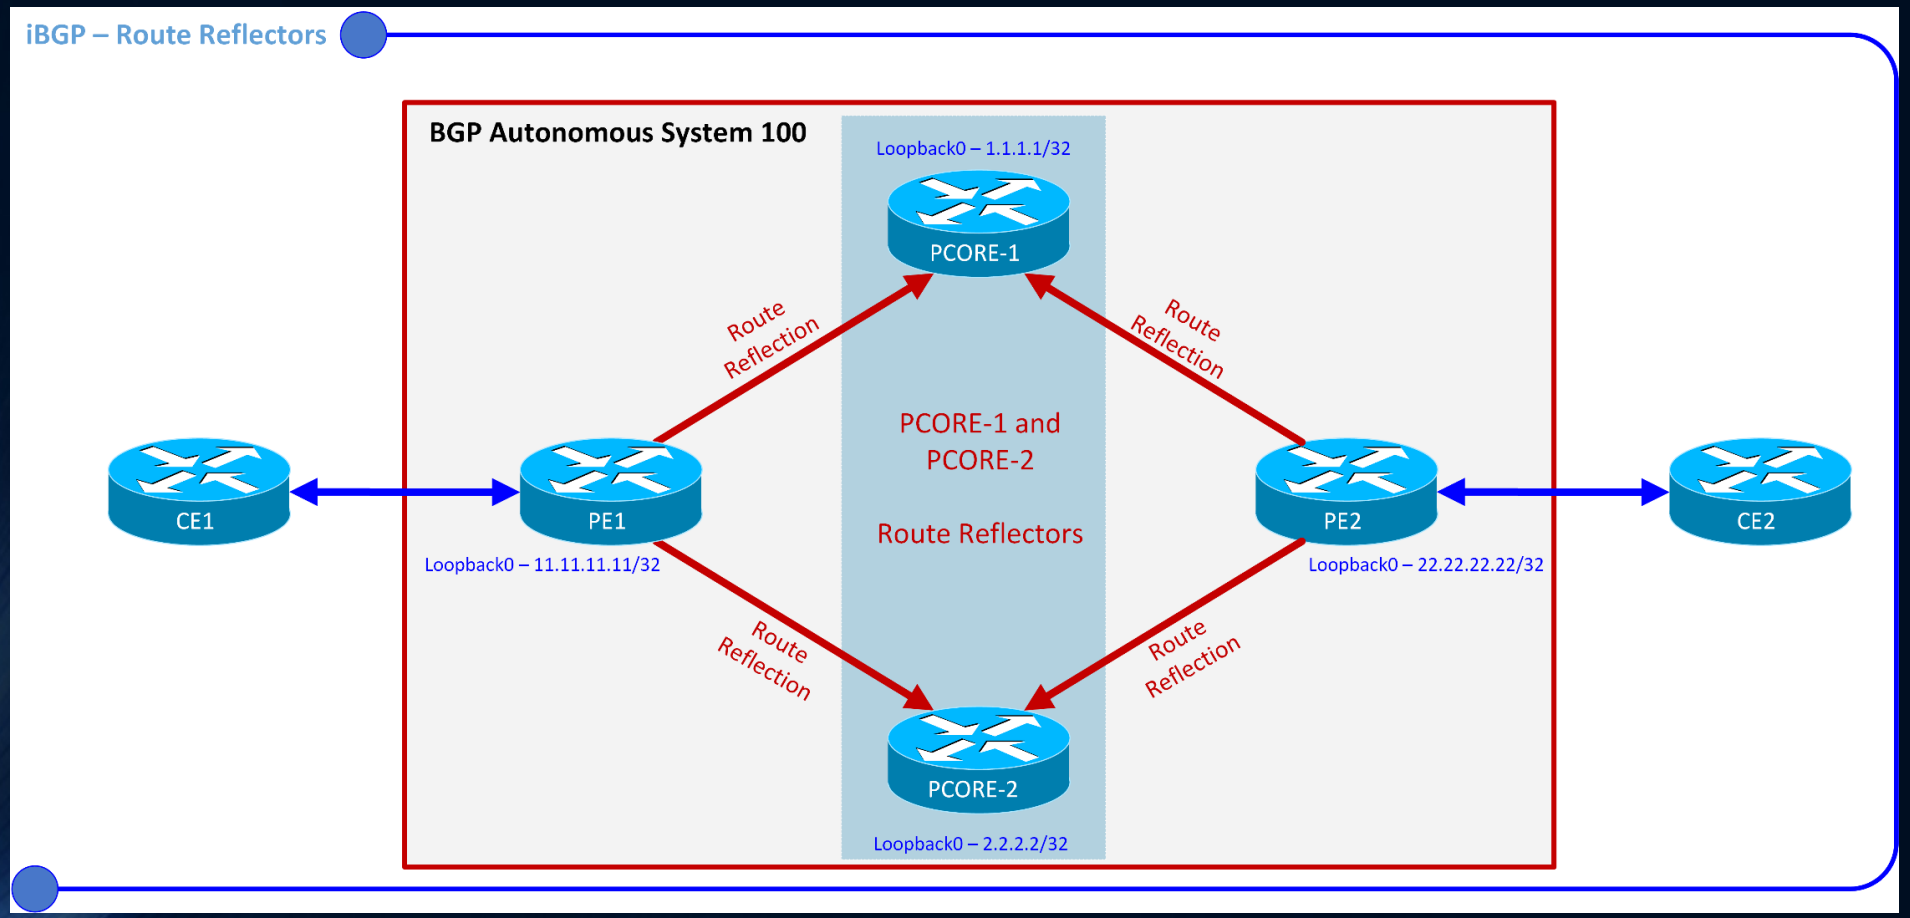

iBGP – Route Reflection - Configuration

- Route Reflector

```
[admin@PCORE-1] > routing bgp instance export
# sep/16/2013 10:46:44 by RouterOS 5.19
# software id = HDT2-V13L
#
/routing bgp instance
set default as=100 client-to-client-reflection=yes disabled=no ignore-as-path-len=no name=\
    default out-filter="" redistribute-connected=no redistribute-ospf=no redistribute-other-bgp=\
    no redistribute-rip=no redistribute-static=no router-id=0.0.0.0 routing-table=""
[admin@PCORE-1] >
```

- Route Reflector Client (Peer)

```
[admin@PCORE-1] > routing bgp peer export
# sep/16/2013 10:49:44 by RouterOS 5.19
# software id = HDT2-V13L
#
/routing bgp peer
add address-families=ip as-override=no default-originate=never disabled=no hold-time=3m \
    in-filter="" instance=default multihop=no name=peer1 nexthop-choice=default out-filter="" \
    passive=no remote-address=22.22.22.22 remote-as=100 remove-private-as=no route-reflect=yes \
    tcp-md5-key="" ttl=default update-source=Loopback0 use-bfd=no
add address-families=ip as-override=no default-originate=never disabled=no hold-time=3m \
    in-filter="" instance=default multihop=no name=peer2 nexthop-choice=default out-filter="" \
    passive=no remote-address=11.11.11.11 remote-as=100 remove-private-as=no route-reflect=yes \
    tcp-md5-key="" ttl=default update-source=Loopback0 use-bfd=no
[admin@PCORE-1] >
```

iBGP – Route Reflection - Verification

- Route Reflector

```
[admin@PCORE-1] > routing bgp peer print detail
Flags: X - disabled, E - established
 0 E name="peer1" instance=default remote-address=22.22.22.22 remote-as=100 tcp-md5-key="" nexthop-choice=default multihop=no
     route-reflect=yes hold-time=3m ttl=default in-filter="" out-filter="" address-families=ip update-source=Loopback0
     default-originate=never remove-private-as=no as-override=no passive=no use-bfd=no

 1 E name="peer2" instance=default remote-address=11.11.11.11 remote-as=100 tcp-md5-key="" nexthop-choice=default multihop=no
     route-reflect=yes hold-time=3m ttl=default in-filter="" out-filter="" address-families=ip update-source=Loopback0
     default-originate=never remove-private-as=no as-override=no passive=no use-bfd=no
[admin@PCORE-1] >
```

- Route Reflector Client (Peer)

```
[admin@PCORE-1] > routing bgp advertisements print peer1 detail
 peer="peer1" prefix=10.150.8.0/22 nexthop=100.64.1.1 as-path="10" origin=igp local-pref=100 originator-id=172.31.13.2
   cluster-list="{ 34414508 }"

 peer="peer1" prefix=10.150.12.0/22 nexthop=100.64.1.1 as-path="10" origin=igp local-pref=100 originator-id=172.31.13.2
   cluster-list="{ 34414508 }"
```

iBGP Traffic Engineering

- Route Filters can be created and applied to influence the flow of traffic.
- Weight
  - Weight – can be set to prefer a specific next hop out of the router.
  - Not passed as a value to other BGP Peers. Local setting only
- Local Preference
  - Local Preference – set to prefer a specific next-hop out of the AS
  - Passed as a value to other BGP Peers – does not leave the current AS

eBGP in the Provider Edge / Customer Edge

- Used to connect to the Internet Edge
- Used to connect remote networks to an iBGP core
- Can use eBGP for Traffic Engineering via prepending / communities
- eBGP requires multihop if peering with loopbacks
- Changes the next hop by default (iBGP does not)
- Does not require Full Mesh

eBGP in the Provider Edge / Customer Edge

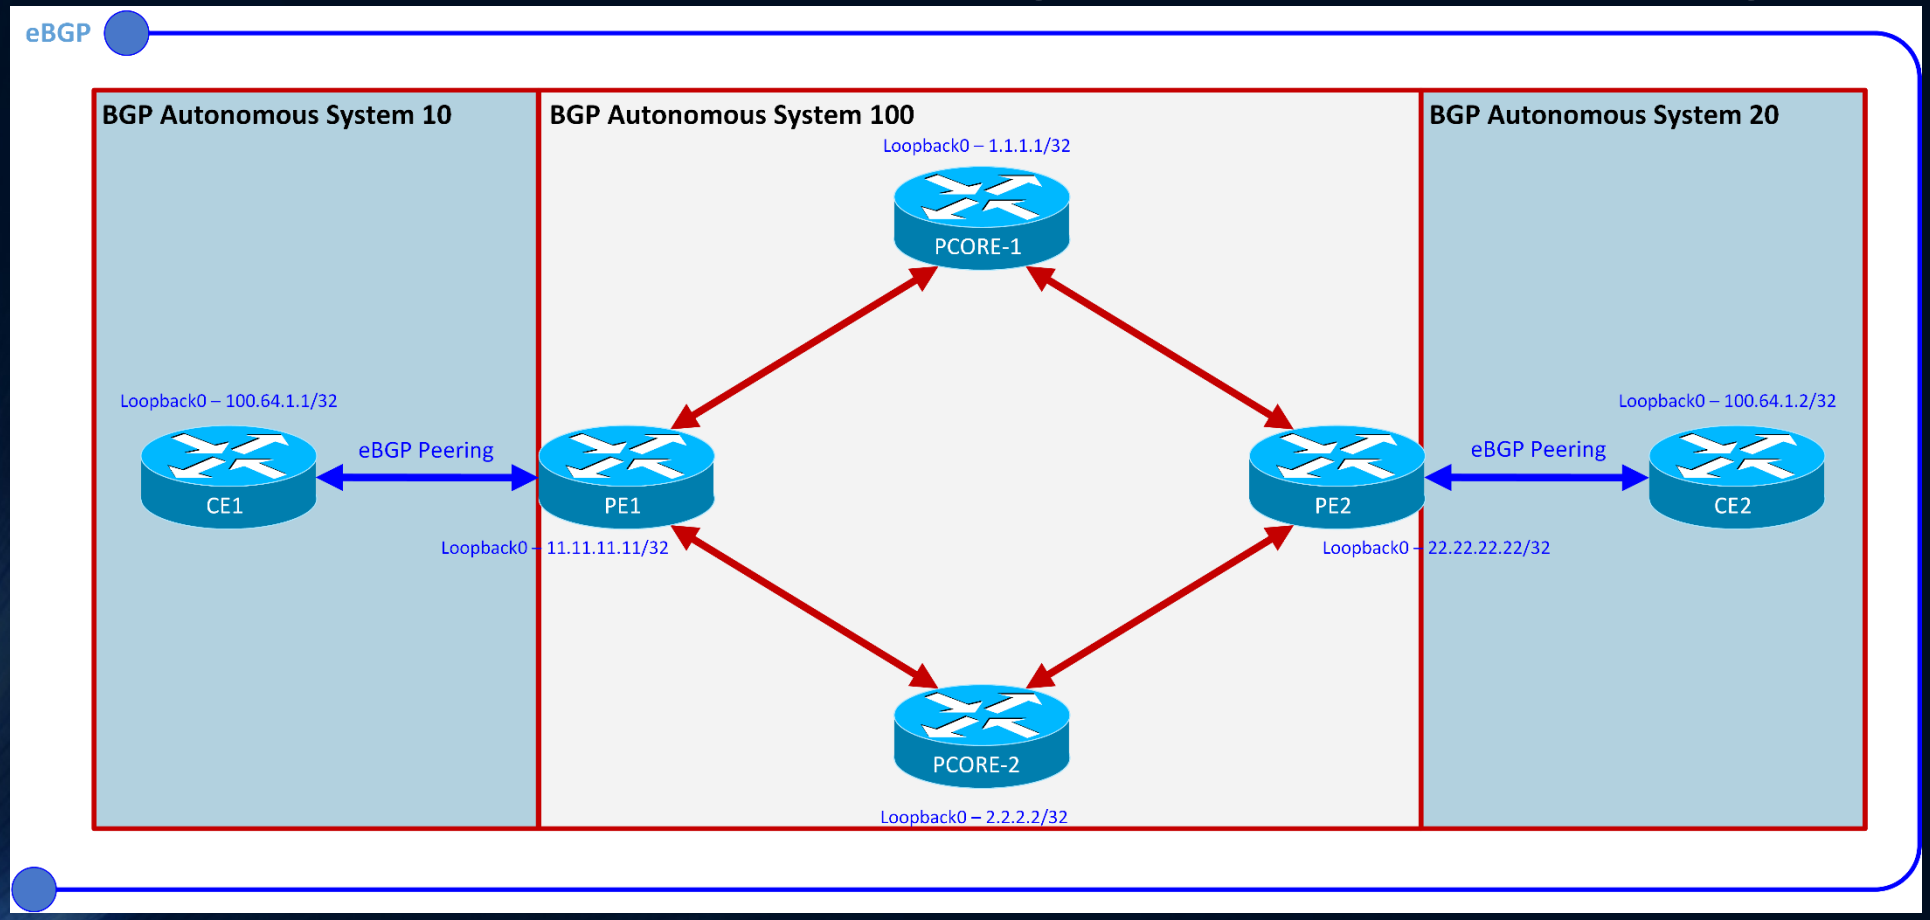

eBGP in the PE/CE - Configuration

- Configuration - PE1

```
[admin@PE1] > routing bgp peer print detail
Flags: X - disabled, E - established
 0 E name="peer1" instance=default remote-address=1.1.1.1
     remote-as=100 tcp-md5-key="" nexthop-choice=default
     multihop=no route-reflect=yes hold-time=3m ttl=default
     in-filter="" out-filter="" address-families=ip
     update-source=Loopback0 default-originate=never
     remove-private-as=no as-override=no passive=no use-bfd=no

 1 E name="peer2" instance=default remote-address=2.2.2.2
     remote-as=100 tcp-md5-key="" nexthop-choice=default
     multihop=no route-reflect=yes hold-time=3m ttl=default
     in-filter="" out-filter="" address-families=ip
     update-source=Loopback0 default-originate=never
     remove-private-as=no as-override=no passive=no use-bfd=no

 2 E name="peer3" instance=default remote-address=100.64.1.1
     remote-as=10 tcp-md5-key="" nexthop-choice=default
     multihop=yes route-reflect=no hold-time=3m ttl=default
     in-filter="" out-filter="" address-families=ip
     update-source=Loopback0 default-originate=never
     remove-private-as=no as-override=no passive=no use-bfd=no
[admin@PE1] >
```

- Configuration - PCORE-1

```
[admin@PCORE-1] > routing bgp peer print detail
Flags: X - disabled, E - established
 0 E name="peer1" instance=default remote-address=22.22.22.22 remote-as=100
     tcp-md5-key="" nexthop-choice=default multihop=no route-reflect=yes
     hold-time=3m ttl=default in-filter="" out-filter="" address-families=ip
     update-source=Loopback0 default-originate=never remove-private-as=no
     as-override=no passive=no use-bfd=no

 1 E name="peer2" instance=default remote-address=11.11.11.11 remote-as=100
     tcp-md5-key="" nexthop-choice=default multihop=no route-reflect=yes
     hold-time=3m ttl=default in-filter="" out-filter="" address-families=ip
     update-source=Loopback0 default-originate=never remove-private-as=no
     as-override=no passive=no use-bfd=no
[admin@PCORE-1] >
```

eBGP in the PE/CE - Verification

```
[admin@PE1] > ip route print detail
Flags: X - disabled, A - active, D - dynamic,
C - connect, S - static, r - rip, b - bgp, o - ospf, m - mme,
B - blackhole, U - unreachable, P - prohibit
 0 ADo  dst-address=1.1.1.1/32 gateway=172.31.10.2
        gateway-status=172.31.10.2 reachable via  bonding1
        distance=110 scope=20 target-scope=10 ospf-metric=20
        ospf-type=intra-area

 1 ADo  dst-address=2.2.2.2/32 gateway=172.31.11.2
        gateway-status=172.31.11.2 reachable via  bonding2
        distance=110 scope=20 target-scope=10 ospf-metric=20
        ospf-type=intra-area

 2 ADb  dst-address=10.150.8.0/22 gateway=100.64.1.1
        gateway-status=100.64.1.1 recursive via 172.16.10.2
              ether2
        distance=20 scope=40 target-scope=30 bgp-as-path="10"
        bgp-origin=igp received-from=peer3

 3  Db  dst-address=10.150.8.0/22 gateway=100.64.1.1
        gateway-status=100.64.1.1 recursive via 172.16.10.2
              ether2
        distance=200 scope=40 target-scope=30 bgp-as-path="10"
        bgp-local-pref=100 bgp-origin=igp received-from=peer2

 4 ADb  dst-address=10.150.12.0/22 gateway=100.64.1.1
        gateway-status=100.64.1.1 recursive via 172.16.10.2
              ether2
        distance=20 scope=40 target-scope=30 bgp-as-path="10"
        bgp-origin=igp received-from=peer3

 5  Db  dst-address=10.150.12.0/22 gateway=100.64.1.1
        gateway-status=100.64.1.1 recursive via 172.16.10.2
```

```
[admin@PE1] > routing bgp advertisements print
PEER     PREFIX              NEXTHOP          AS-PATH
peer1    10.150.8.0/22       100.64.1.1       10
peer1    10.150.12.0/22      100.64.1.1       10
peer1    10.150.16.0/24      100.64.1.1       10
peer2    10.150.8.0/22       100.64.1.1       10
peer2    10.150.20.0/22      100.64.1.2       20
peer2    10.150.28.0/24      100.64.1.2       20
peer2    10.150.12.0/22      100.64.1.1       10
peer2    10.150.24.0/22      100.64.1.2       20
peer2    10.150.16.0/24      100.64.1.1       10
peer3    10.150.20.0/22      100.64.1.2       20
peer3    10.150.28.0/24      100.64.1.2       20
peer3    10.150.24.0/22      100.64.1.2       20
[admin@PE1] >
```

Enterprise considerations

- Topology will easily translate into the Enterprise with a few modifications
- Terminology shift

Provider terminology	Enterprise terminology
CE - Customer Edge	Access Layer
PE – Provider Edge	Distribution Layer
P-CORE – Provider Core	Enterprise Core Layer

- Many Enterprise firewalls have started supporting BGP which allows for end-to-end BGP in the Enterprise.

2013 St Louis MUM –Tablet Giveaway !!

- One 7" Android .TAB Nero will be given away on Sep 19th and one on Sep 20th
- Stop by the IP ArchiTechs exhibition booth, guess the right number and WIN!

Questions?

- The content of this presentation will be available at mum.iparchitechs.com
- Please come see us at the IP ArchiTechs booth in the Exhibitor Hall
- Email: kevin.myers@iparchitechs.com
- Office: (303) 590-9943
- Web: www.iparchitechs.com
- Thank you for your time and enjoy the MUM!!