Routing Implementation Cisco vs. MikroTik

Lay Minh (Makito)

CCIE # 47682

MikroTik Certified Trainer, MikroTik Consultant

About Me

- **Lay Minh**
  - My nick name is **Makito**
  - CCIE # 47682
  - MikroTik Certified Trainer & Consultant
  - Chief Technology Officer (CTO) at i-BEAM
  - Experiences:
    - 10 years in ISP industry since 2005
    - Billing solutions for service providers
    - ISP core network design and operations
  - Certifications:
    - CCIE Service Provider, JNCIA-Junos, JNCIS-SP
    - MTCNA, MTCRE, MTCWE, MTCTCE, MTCUME, MTCINE
  - Areas of interest: BGP, MPLS, IPv6

About The Presentation

i-BEAM steering ahead

- This presentation **IS** talking about differences between **Cisco IOS**' implementation and **MikroTik RouterOS**' implementation on some standardized technologies
- This presentation **IS NOT** a lecture talks about the technology itself

- There are some prerequisites:
  - General knowledge about routing concepts
  - Basic understanding about what they are and what they do:
    - Open Shortest Path First (OSPF)
    - Border Gateway Protocol (BGP)
    - Multiprotocol Label Switching (MPLS)

Routing Components

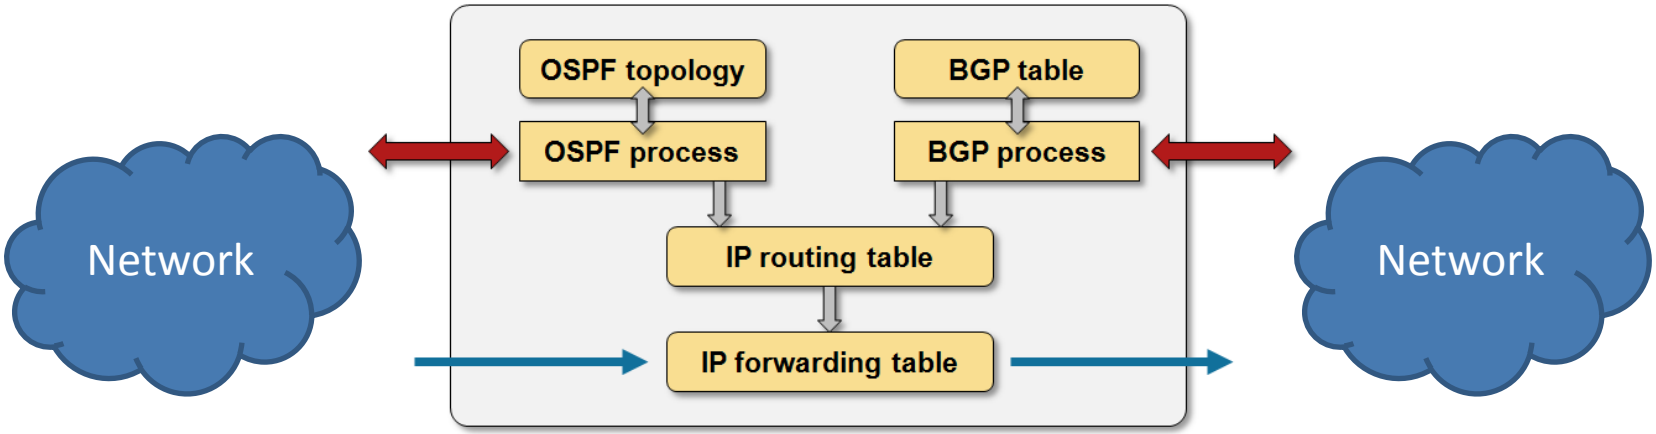

- **RIB (Routing Information Base)**
  - Well-known as “IP Routing Table”
  - For network engineers to read routing information
- **FIB (Forwarding Information Base)**
  - Well-known as “IP Forwarding Table”
  - More hardware-friendly, for the router hardware
  - CEF (Cisco Express Forwarding) in Cisco IOS

General Routing

- Difference between vendor's implementations:

	Cisco IOS	MikroTik RouterOS
FIB Load Balancing	Per Src. and Dst. Address Pair or Per-Packet	Per Src. and Dst. Address Pair (Flush every 10 minutes)
Equal Cost Multi Path (ECMP)	Add multiple routes to same destination with same distance, but different gateway	Add only one route by mentioning multiple gateways in the same route
Recursive Nexthop Lookup	Enabled	Disabled by default, can manually enable by route's Target Scope parameter
Route Filtering Behavior	Implicit deny at the end of each filtering component (access-list, prefix-list, filter-list, route-map…etc.)	Implicit permit at the end of filtering component (Routing Filters)

Load Balancing with ECMP

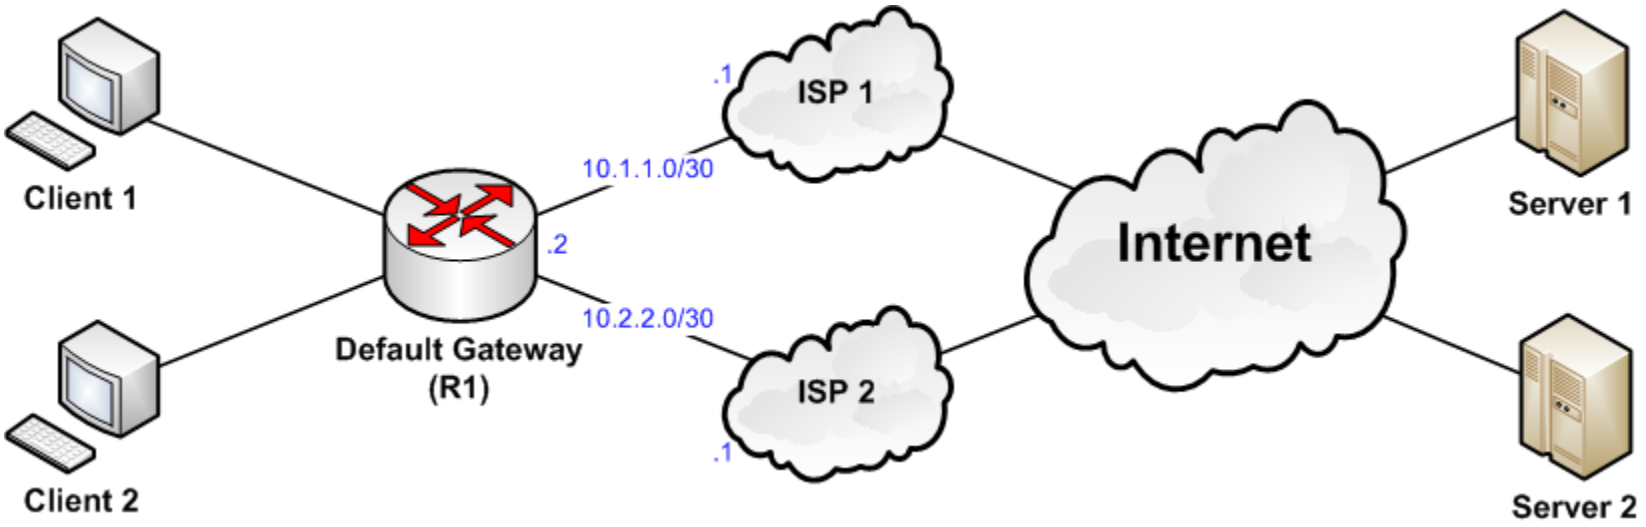

- Cisco IOS configuration:

```
R1(config)#ip route 0.0.0.0 0.0.0.0 10.1.1.1
R1(config)#ip route 0.0.0.0 0.0.0.0 10.2.2.1
```

- MikroTik RouterOS configuration:

```
[admin@R1] > /ip route add dst-address=0.0.0.0/0 \
        gateway=10.1.1.1,10.2.2.1
```

Recursive Nexthop

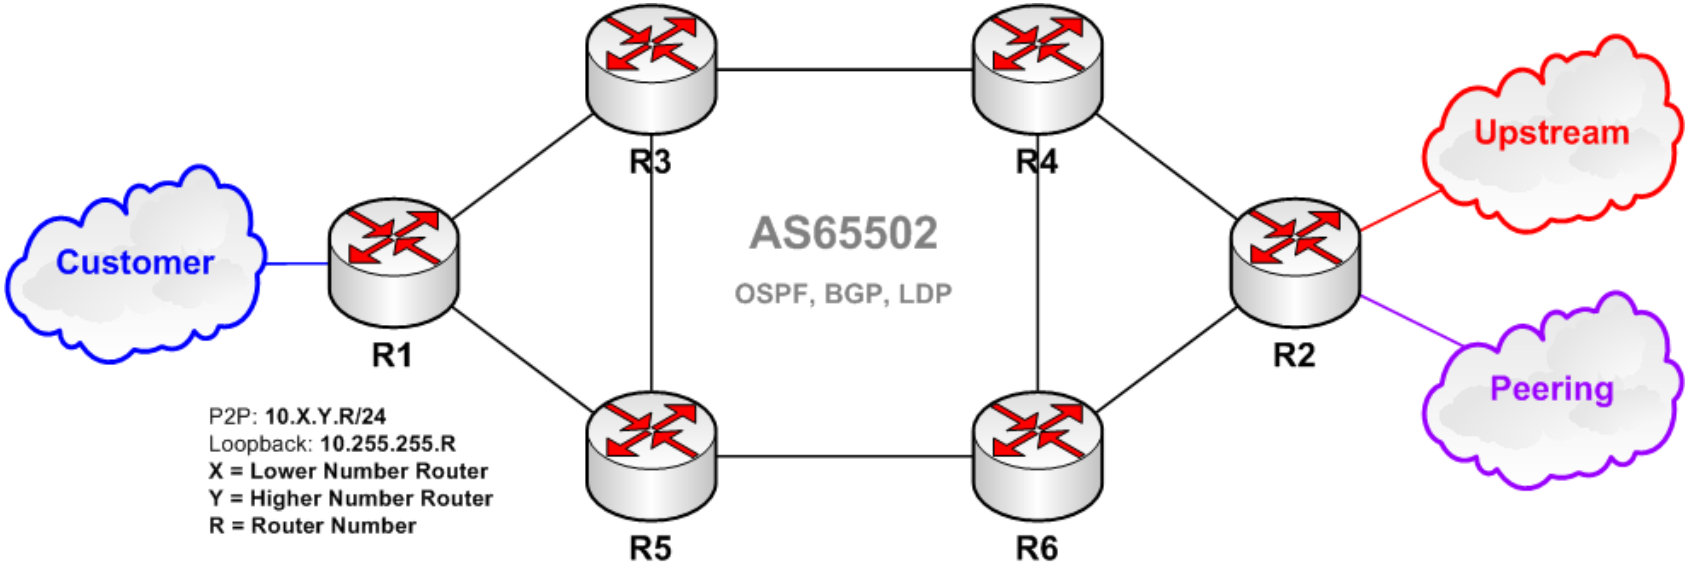

- In common ISP practice, R1 and R2 will have iBGP peering
- R2 advertises routes to R1 with its Loopback as nexthop
  - **R2:** "Wanna go to 8.8.8.0/24? Come to me! Then I'll send you there! "
- R1 performs recursive lookup to find out ways to reach R2
  - **R1:** "Well…how to go to **R2**? let me take a look in my routing table…"
  - **R1:** "Ahha! I got it! It is via **R3** or **R5**, thanks to my lovely OSPF! ☺"

Open Shortest Path First (OSPF)

- RFC 2328 for OSPFv2, RFC 5340 for OSPFv3
- Difference between vendor's implementations:

	Cisco IOS	MikroTik RouterOS
Router ID	Highest active Loopback IP, then Highest active Interface IP	Lowest active interface IP
Link Cost	Vary depends on Link BW Ref. BW (bps) / Link BW (bps) Default Ref. BW is 100Mbps	Fixed Cost 10 for any links, changeable in OSPF interface configuration
OSPF Timer	Vary depends on network type "broadcast" and "point-to-point", Hello interval is 10, Dead interval is 40 For "nbma", Hello interval is 30, Dead interval is 120	Fixed Hello interval 10, and Dead interval 40, changeable in OSPF interface configuration *** CAUTION: When changed network type, intervals are not changed!!

Open Shortest Path First (OSPF) (Cont.)

	Cisco IOS	MikroTik RouterOS
Stub Area	Type 3 LSAs are advertised into Stub Area by default, unless configured as Totally Stubby Area	Behaves like Totally Stubby Area in IOS by default. Enable "Inject Summary LSAs" option if want to advertise Type 3 LSAs into Stub Area
Route Filtering	Use "distribute-list" command to permit/deny routes to be installed into RIB	Use "Routing Filters" to permit/deny routes to be installed into RIB, but can filter only Type 5 LSAs
Advertising Loopback interface into OSPF	By default, subnet mask of the interface route is forced to be /32 Set the interface network type as "point-to-point" to advertise exact subnet mask	By default, exact subnet mask is advertised, no action required

OSPF Cost (IOS)

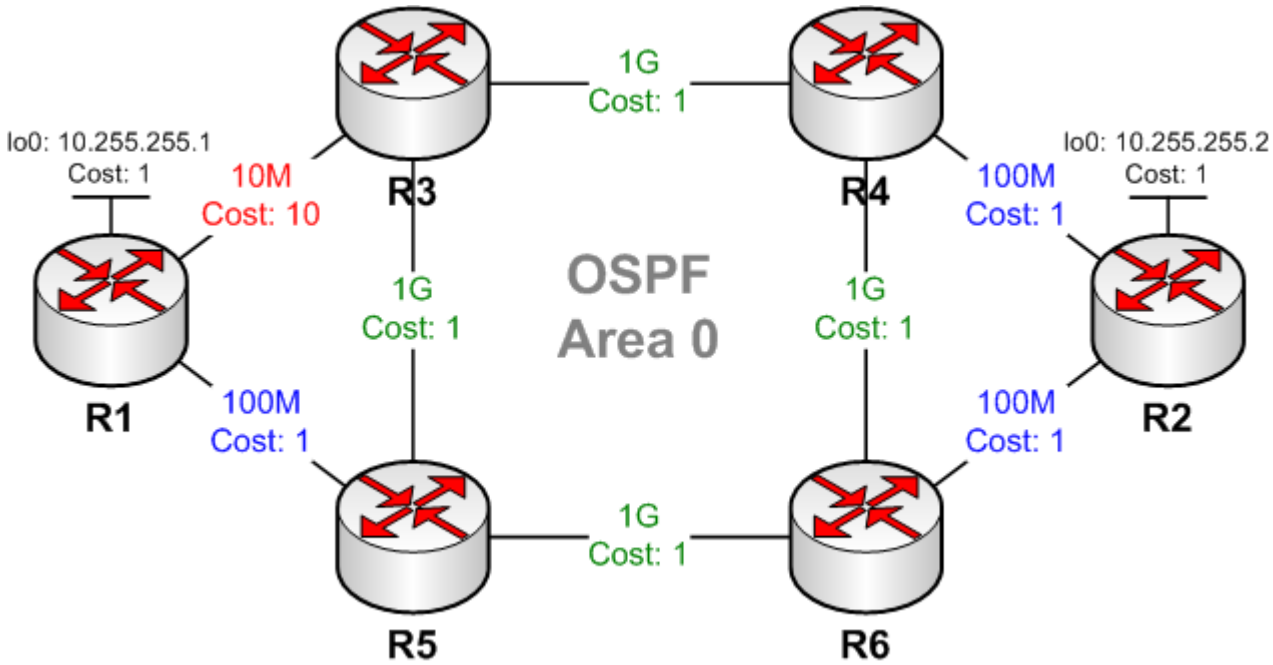

- Link Cost = Reference Bandwidth in bps / Link Bandwidth in bps
  - By default, Reference Bandwidth is 100Mbps
- For R1 to reach R2's Loopback (10.255.255.2)
  - Shortest path is **R5-R6-R2**, total cost is **4** (**1**+**1**+**1**+**1**)
  - Path via **R3** has higher cost **13** (**10**+**1**+**1**+**1**)

OSPF Cost (RouterOS)

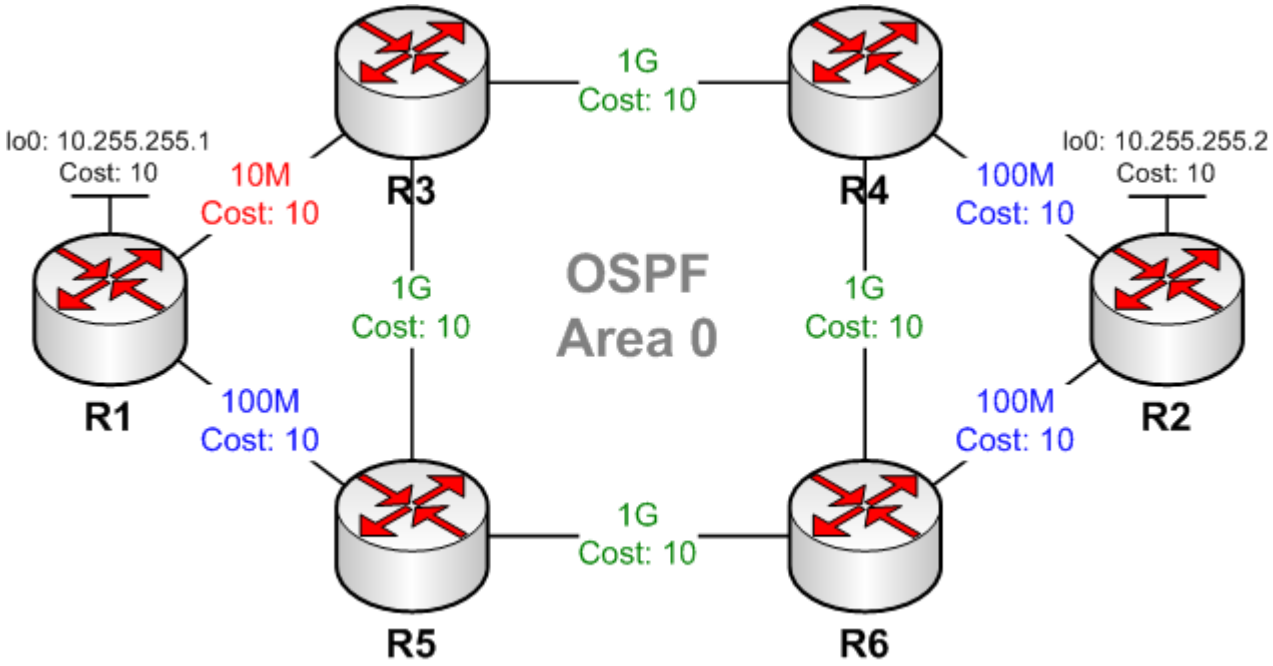

- Link Cost = 10
- For R1 to reach R2's Loopback (10.255.255.2)
  - There are two shortest paths, **R3-R4-R2** and **R5-R6-R2**
  - Each path's total cost is **40** (**10**+**10**+**10**+**10**)
  - **R1** will install ECMP route and load balance over both links

Border Gateway Protocol (BGP)

i-BEAM steering ahead

- Described in RFC 4271
- Difference between vendor's implementations:

	Cisco IOS	MikroTik RouterOS
Router ID	Highest active Loopback IP, Highest active Interface IP	Lowest active interface IP
Routes Received	Stores in BGP table, best path will be installed to RIB	Stores in RIB, best path will be active
Best Path Selection	IGP Metric to Nexthop is considered when multiple equal-cost paths exist	IGP Metric to Nexthop is ignored
Soft Reconfiguration	Unused routes are stored in BGP table and marked with "(receive-only)" flag	Unused routes are installed in RIB, but inactive

Border Gateway Protocol (BGP) (Cont.)

	Cisco IOS	MikroTik RouterOS
BGP Multi Path	Up to 8 ~ 32 paths depends on IOS version	Not supported For dual-homed scenarios, can tweak BGP load sharing by peering eBGP multi-hop with Loopbacks if there are multiple links to neighbor AS
Route Advertisement	Based on BGP table, can advertise best path installed in BGP table but not in RIB	Based on RIB, can advertisement only best path installed and active in RIB
BGP Community	Receive only, community sending is disabled by default	Send and receive are enabled by default
IPv6 Prefix over IPv4 BGP Session	Nexthop is IPv4 address, needs to correct with “route-map”	Nexthop is IPv6 address, no action required

BGP Best Path Selection

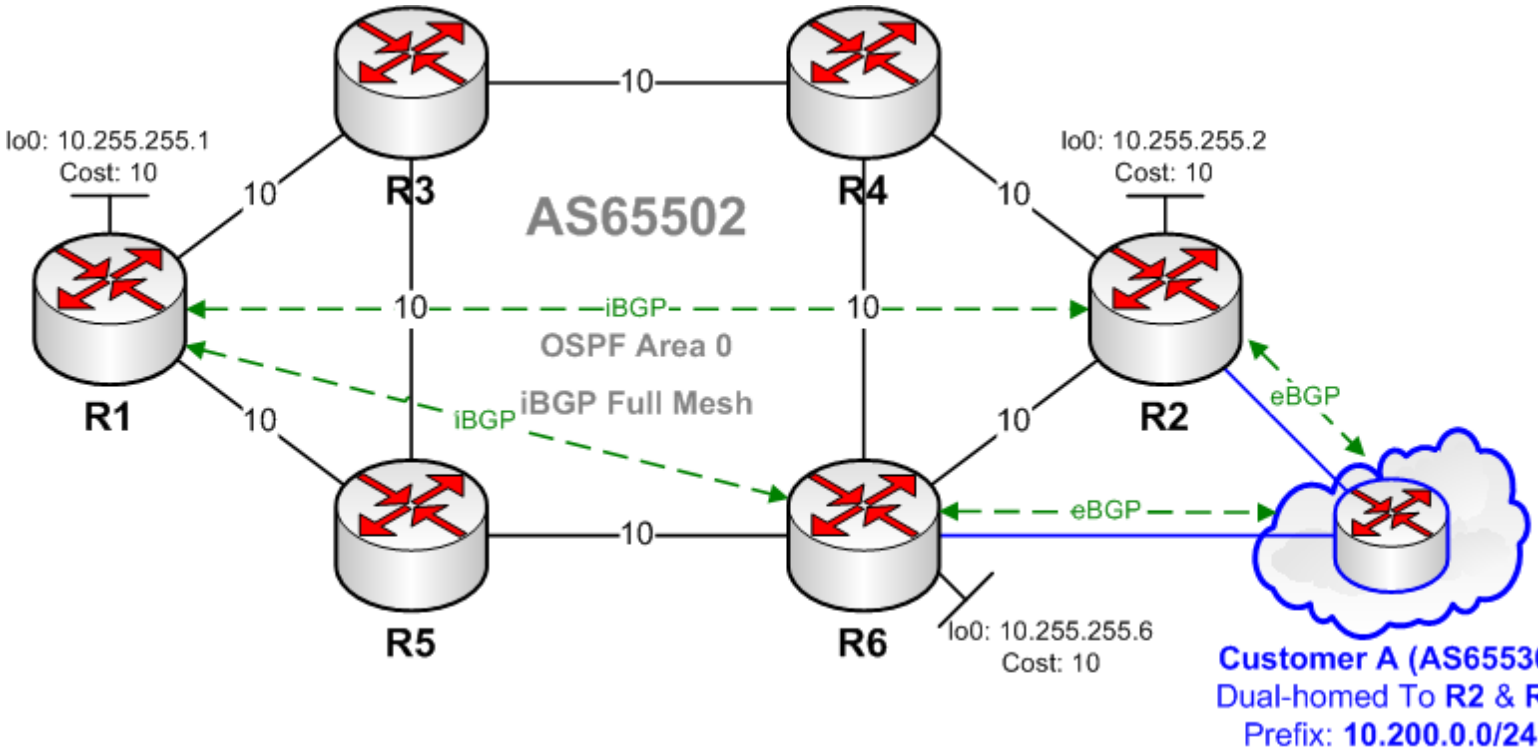

- OSPF link costs in AS65502 are set to 10 for ease of understanding
- Customer A connects to 2 provider routers (dual-homed)
  - **eBGP** peering with **R2** and **R6**
  - Advertises prefix **10.200.0.0/24** with exact same BGP attributes

BGP Best Path Selection (Cont.)

1. Nexthop must be reachable.
2. Highest Weight (default 0).
3. Highest Local Pref. (default 100).
4. Shortest AS Path.
5. Locally originated path (aggregated route or BGP network).
6. Lowest origin type (IGP < EGP < Incomplete).
7. Lowest MED (default 0).
8. Prefer eBGP over iBGP.
9. Lowest Router ID.
10. Lowest Originator ID.
11. Shortest route reflection cluster (default 0).
12. Lowest neighbor address.

BGP Multi Path (IOS)

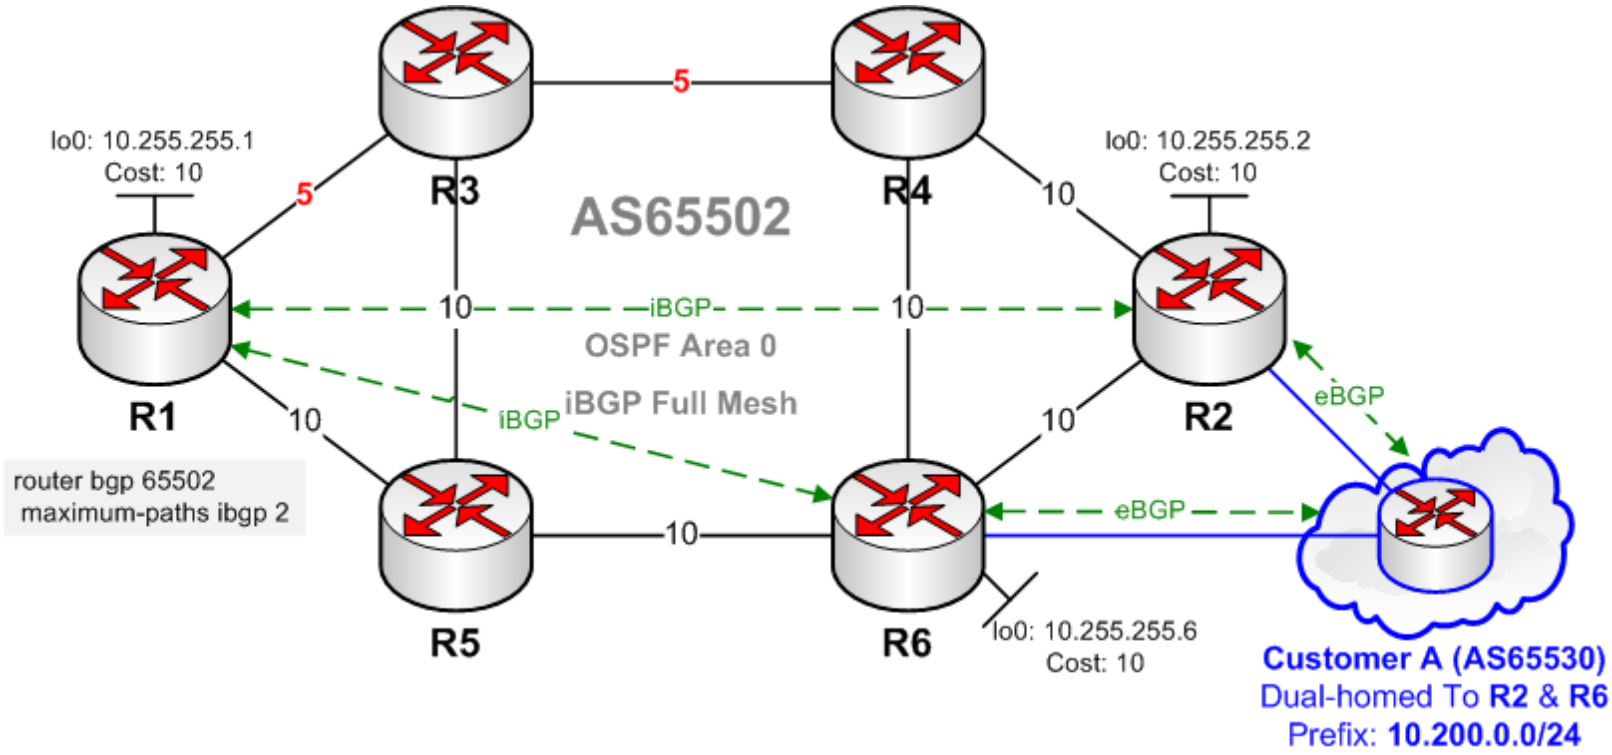

- If R1-R3 link and R3-R4 link both have OSPF link cost 5
- From R1's point of view, there are 2 equal cost paths (via R2 & R6) in BGP to reach 10.200.0.0/24
  - **R1** can install both of them into RIB and perform load balancing

BGP Multi Path (RouterOS)

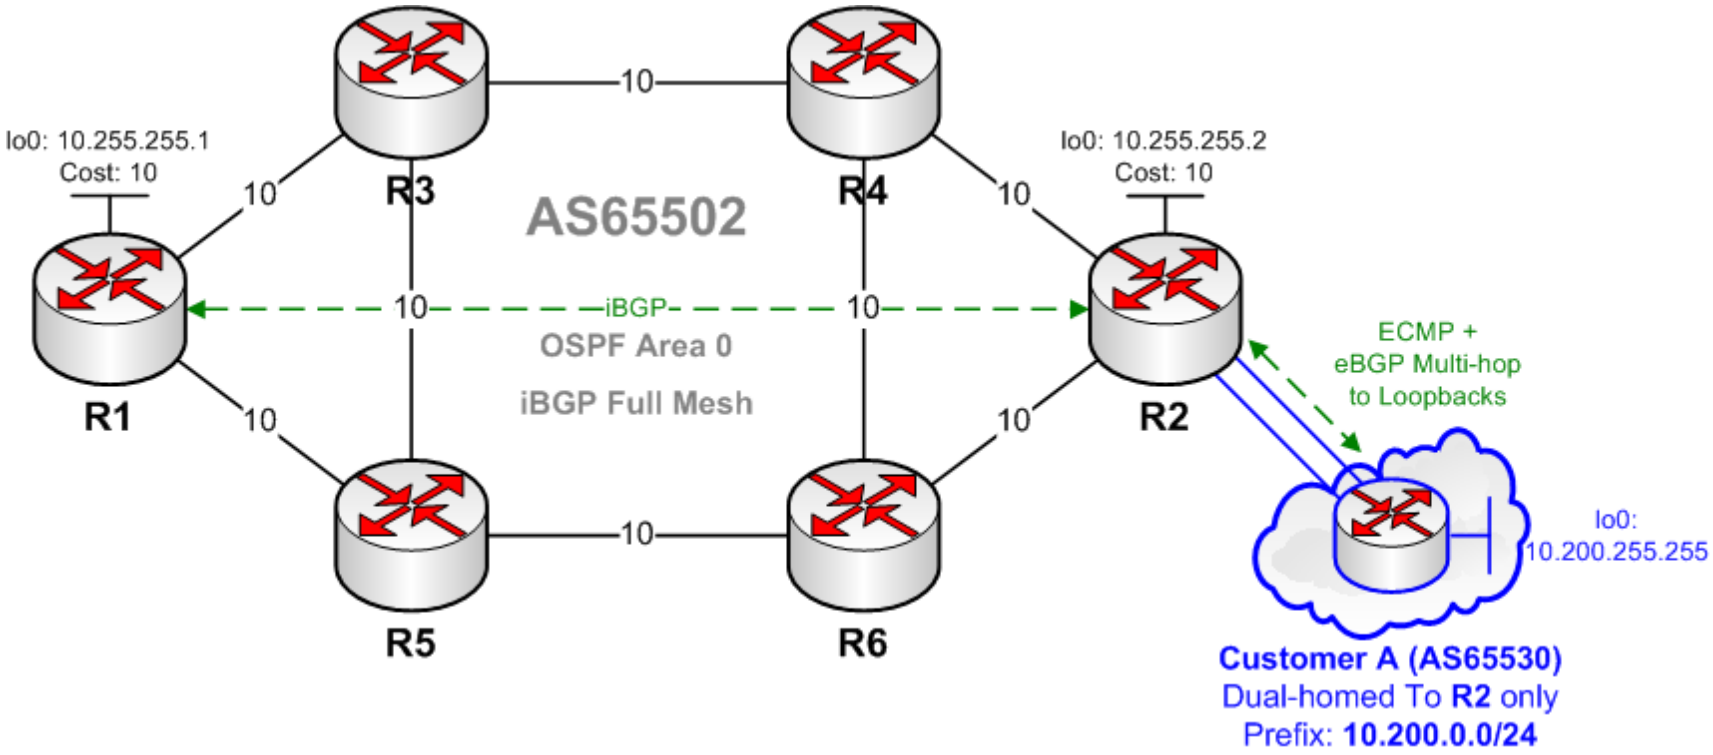

- RouterOS does not support BGP Multi Path feature as IOS
- However, it is possible to do load balancing when customer has multiple links to the same Provider Edge (PE) router
  - Point ECMP routes to PE and CE's Loopbacks, then peer multi-hop eBGP

Route Advertisement (IOS)

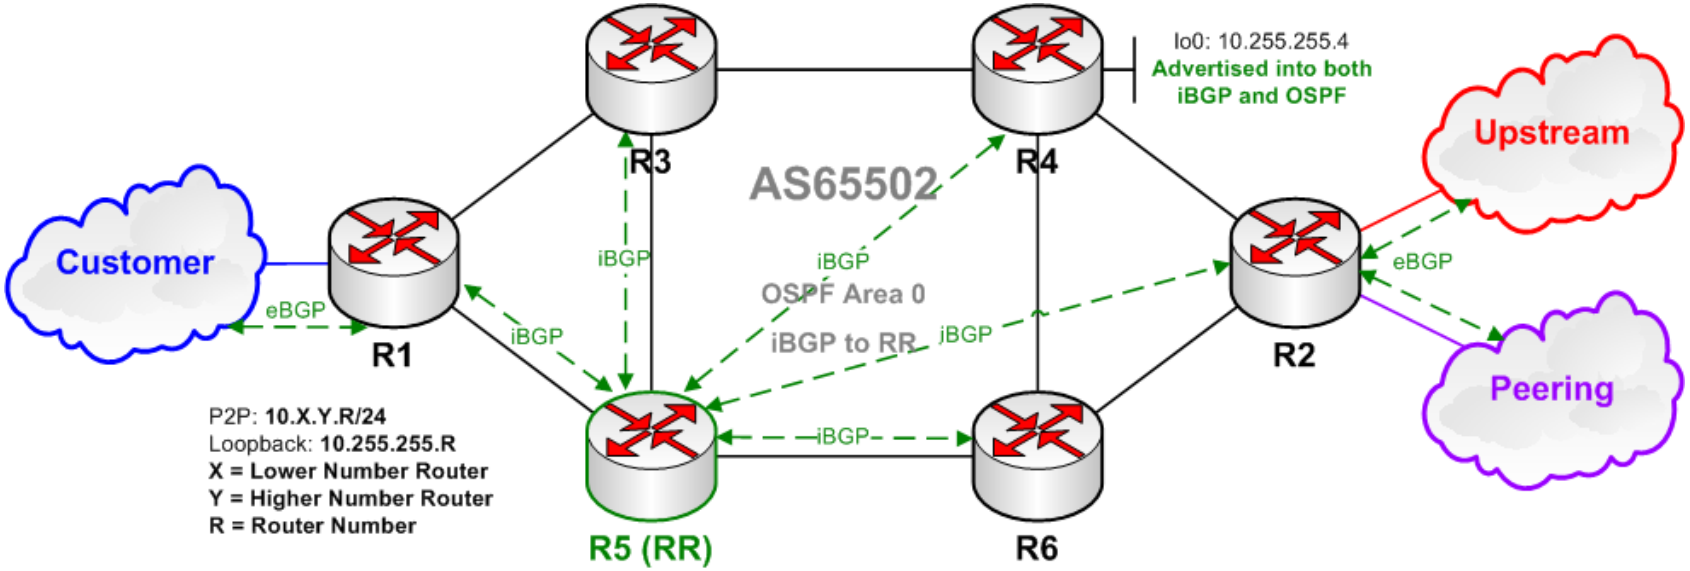

- R4 advertises 10.255.255.4/32 into both OSPF and BGP
- R5 receives 10.255.255.4/32 via both protocols
  - According to AD, R5 installs only OSPF route into RIB
  - However, R5 will reflect BGP prefix 10.255.255.4/32 to other RR clients
- Other routers receive 10.255.255.4/32 via both protocols
  - According to AD, they install only OSPF route into RIB
  - But they will also re-advertise it to eBGP peers (if not filtered)

Route Advertisement (RouterOS)

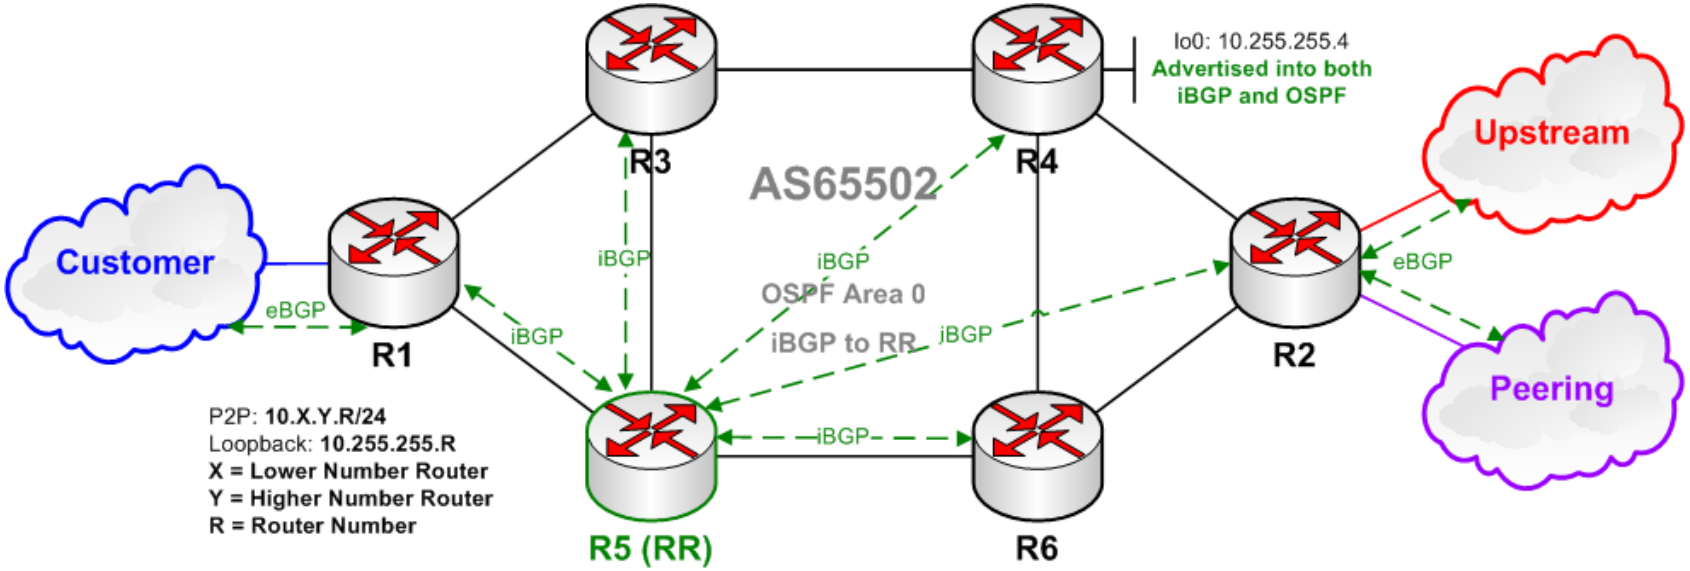

- R4 advertises 10.255.255.4/32 into both OSPF and BGP
- R5 receives 10.255.255.4/32 via both protocols
  - According to AD, R5 installs only OSPF route into RIB
  - R5 will not advertise BGP prefix 10.255.255.4/32 further, because the received BGP prefix is inactive in RIB
- Other routers receive 10.255.255.4/32 via OSPF only
  - OSPF route is installed

Route Redistribution

- Difference between vendor's implementations:

	Cisco IOS	MikroTik RouterOS
Redistribute from BGP into OSPF	By default, prefixes are redistributed as classful subnet Use "subnets" keyword to redistribute as CIDR	By default, prefixes are redistributed as CIDR, no action required
Redistribute from OSPF into BGP	Origin code of the prefix is "Incomplete" External routes are not redistributed by default unless "external" or "nssa-external" keyword is specified	Origin code of the prefix is "IGP" External routes are redistributed by default

Multiprotocol Label Switching (MPLS)

- MPLS Architecture described in RFC 3031
- MPLS is covering more than 130 RFCs
- Difference between vendor's implementations:

	Cisco IOS	MikroTik RouterOS
Multi Path with Label Distribution Protocol (LDP)	Yes	No, only first gateway will be used in MPLS forwarding table (MFIB)
MPLS Fast Reroute	Link Protection (~50ms) Node Protection	Not supported
MPLS Applications	6PE, 6VPE, L3VPN (Unicast and Multicast), AToM, VPLS	L3VPN (Unicast), VPLS
MPLS QoS with EXP bit	Possible on P routers and PE routers by utilizing Modular QoS CLI (MQC)	Only possible on PE routers, P routers will not apply any policy to MPLS packets

MPLS QoS

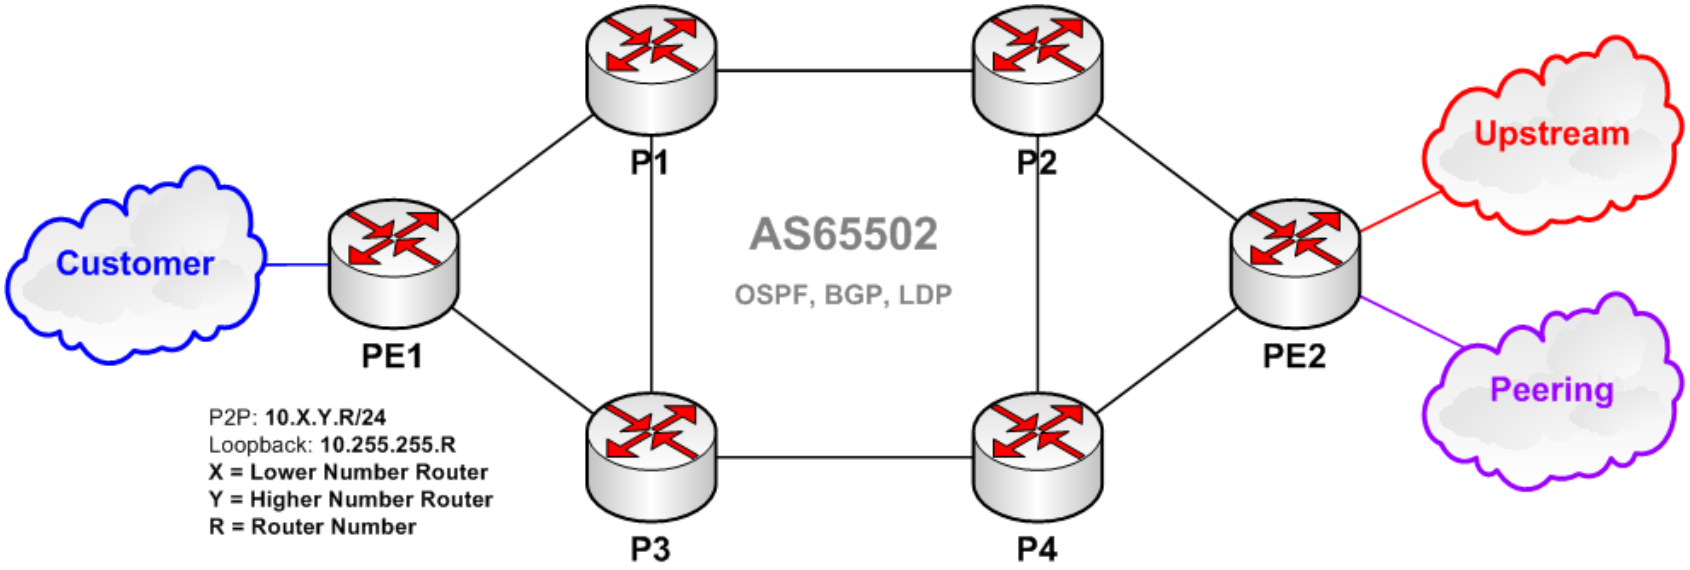

- Provider Edge (PE) = Access Router or Border Router
- Provider (P) = Core Router
- In IOS, QoS is possible everywhere along the path
- In RouterOS, P router ignores all QoS treatments and firewall filters, policies have to be implemented on PEs
  - **Newer versions of RouterOS do not have this limitation anymore**

Questions & Answers

If you have any questions, please feel free to ask!

The End

Thanks for your attention!

Contact Me

makito@informationbeam.net

Skype: akn_makito

Phone: (+95) 09 799 799 282