Segment Routing
(and why you shouldn’t be afraid of it)

RIPE 77, Routing Working Group / October 18, 2018

Florian Hibler <florian@arista.com>

ARISTA

What is Changing?

- More and more Cloud and Service providers are focusing on the need to:
  - Simplify
  - Scale
  - Increased need to control and traffic-engineering (fine grained and macro)
- Increased ECMP in network designs and need to leverage those paths
- More native IPv6 deployments

What about LDP (Label Distribution Protocol)?

- MPLS + LDP + IGP -> Separate labeling protocol needed
- Only follows IGP shortest path
- Slow responsiveness (Sync IGP <-> LDP table/process)
- Lack of IPv6 support

FEC
IP B = Label 100
IP C = Label 200
IP D = Label 300
…

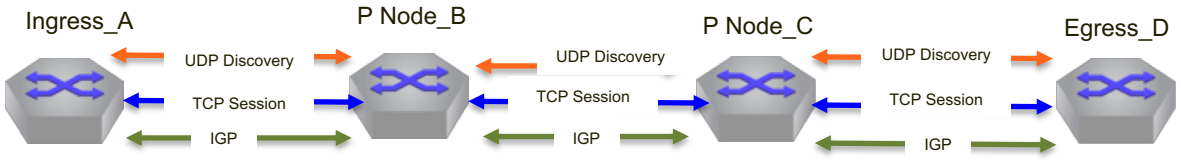

What about RSVP-TE?

- Complex...
- Control plane scale has always been an issue:
  - Per tunnel persistent state at every network hop
  - Soft state refresh/maintenance increases message overhead
- Lack of ECMP, multiple explicit paths needed
- Difference with constraints and path computational options – vendor specific
- Most implementations have not scaled to level required by large operators

ERO
IP A>B>C>D

RECORD_ROUTE
IP B = Label 100
IP C = Label 200
IP D = Label 300

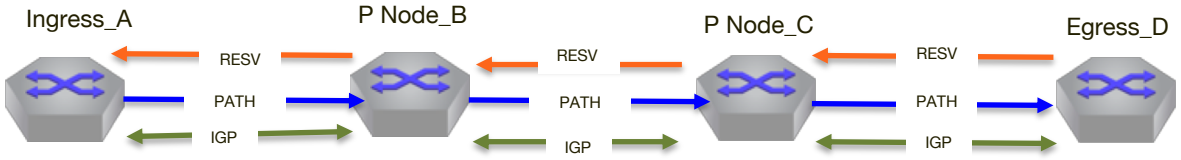

Segment Routing (SR) – Operation Overview

- SR divides the network into “segments” identified by a Segment ID (SID)
  - Global SIDs identify nodes (loopback IP), prefix or Anycast SID (shared IP)
    - All nodes in the SR domain use same SID to identify the prefixes, node or Anycast SID – reducing data plane state
  - Local significant SIDs identify the Adjacency links in the network
    - Only the originating Node understands the advertised Adjacency SID
  - Both local and global SIDs are advertised as TLV extensions to the IGP (ISIS/OSPF)
  - The SID is encoded as an MPLS Label in the forwarding plane

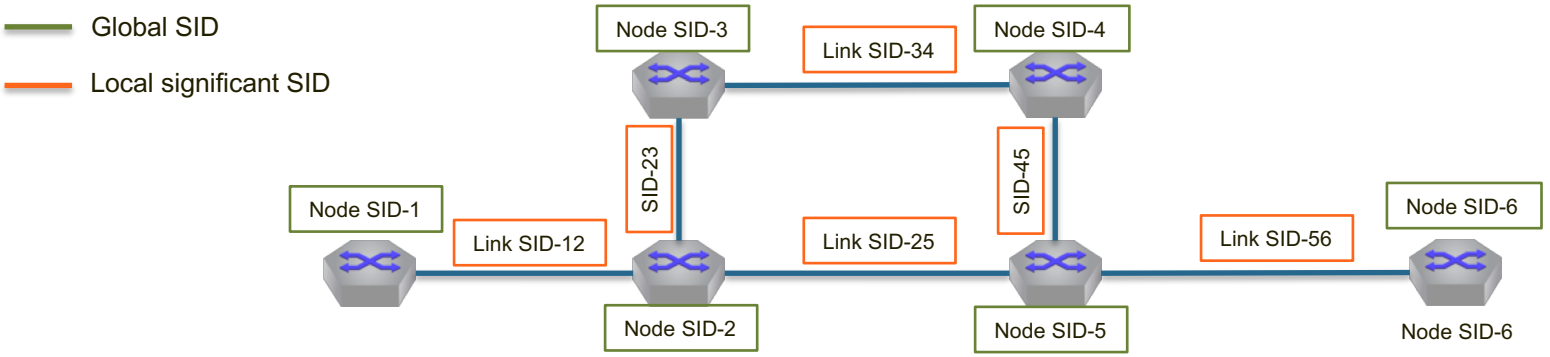

The MPLS Basics

- **Segment Routing is still MPLS…**
- Same concept regards actions
  - PUSH = Add Label
  - SWAP = Change label
  - POP = Remove label

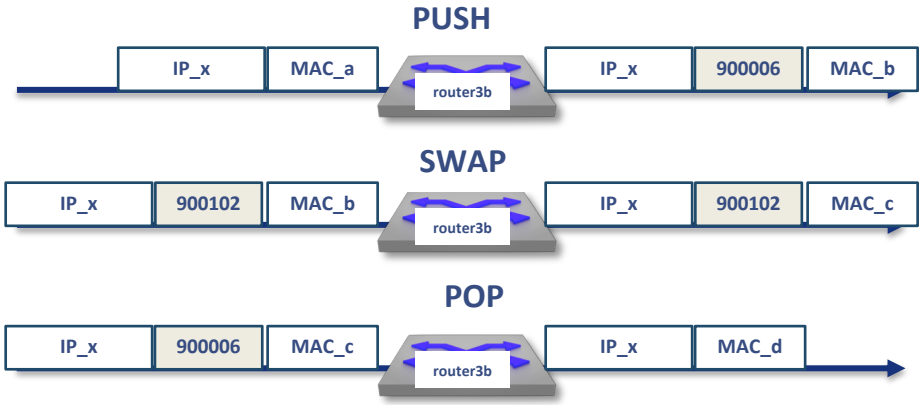


```
router3b#sh mpls ro
(…)
 900006  A[1]
                via M, 1.1.1.30, pop
                    EgressAcl: Apply
                    directly connected, Ethernet3
                    08:00:27:5a:60:e1,
 900007  A[1]
                via M, 1.1.1.41, pop
                    EgressAcl: Apply
                    directly connected, Ethernet4
                    08:00:27:9e:bb:97,
 900102  A[1]
                via M, 1.1.1.30, swap 900102
                    EgressAcl: Apply
                    directly connected, Ethernet3
                    08:00:27:5a:60:e1,
(…)
```


ARISTA

Segment Routing – Global SIDs

- **Segment Routing Operation – Global SIDs**
  - **Rtr 5** advertises it's global node SID, which representing it's loopback IP (5.5.5.5/32)
  - Advertisement via ISIS, simple sub-TLV containing prefix and corresponding SID index
  - All remote nodes install the same node segment ID (90010) to 5.5.5.5 in the MPLS data plane
  - At ingress, top label 900010 and to packet and reaches 5.5.5.5 via the IGP shortest-path

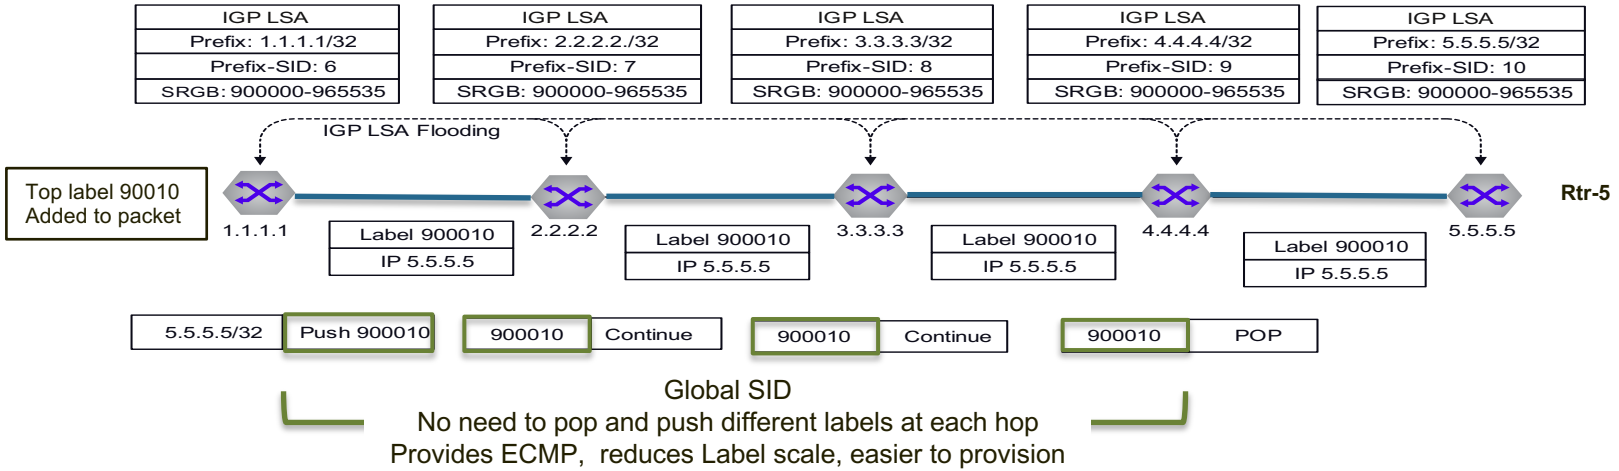

ARISTA

Segment Routing – ECMP aware

- **The forwarding route for the global Node SID represents IGP calculated shortest path to the prefix.**
  - If the PE is dual-homed and supports ECMP, traffic can be load-balanced across the each of the equal cost paths
  - Traffic load-balancing on 5-tuple (Prefix SID mapped to IGP ECMP SPF paths).

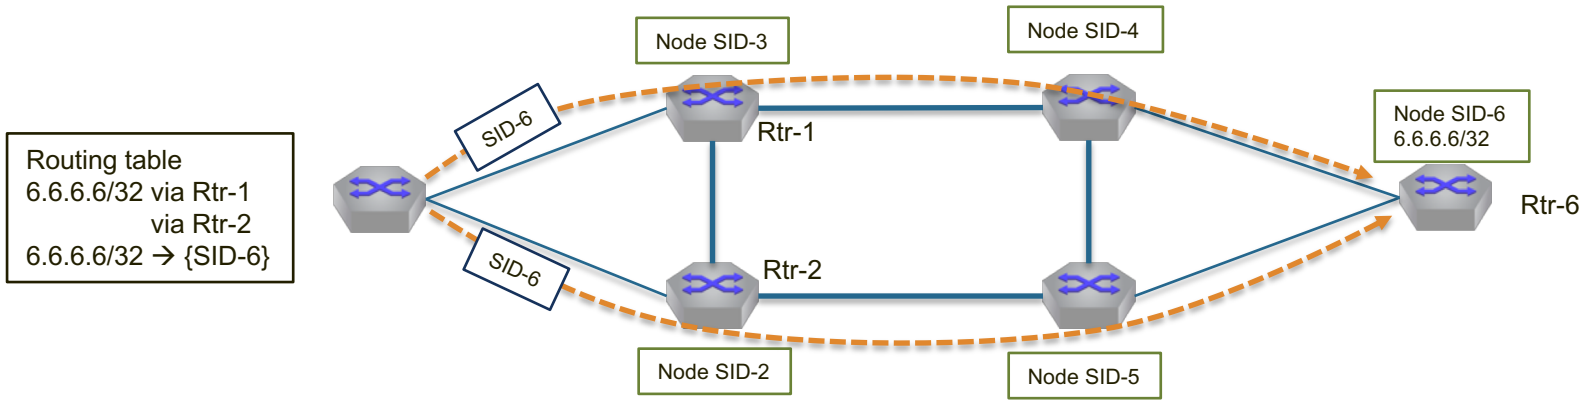

What about mix LDP and SR

- Labels have ranges
  - Example 1-15 reserved range (example Label 3 => IP Tunnel/MPLS Lookup then POP Label)
  - SR can coexist with LDP/RSVP in the implementation phase

```
router3b#sh mpls label ranges
Start     End       Size      Usage
-----------------------------------------------
0         15        16        reserved
16        99999     99984     static mpls
100000    116383    16384     pseudowire (dynamic)
116384    132767    16384     isis (dynamic)
132768    149151    16384     ldp (dynamic)
149152    362143    212992    free (dynamic)
362144    899999    537856    unassigned
900000    965535    65536     isis-sr

router3b#sh mpls tunnel fib
   Tunnel Type          Index       Endpoint           Nexthop        Interface       Labels
------------------- ----------- ------------------ -------------- ----------------- ----------
   LDP                  2           1.1.1.102/32       1.1.1.40       'Ethernet1'     [ 132769 ]
   IS-IS SR IPv4        5           1.1.1.102/32       1.1.1.40       'Ethernet1'     [ 900102 ]
   LDP                  1           1.1.1.4/32         1.1.1.40       'Ethernet1'     [ 3 ]
   IS-IS SR IPv4        4           1.1.1.4/32         1.1.1.40       'Ethernet1'     [ 3 ]
   IS-IS SR IPv4        6           1.1.1.6/32         1.1.1.40       'Ethernet1'     [ 900006 ]
   LDP                  3           1.1.1.6/32         1.1.1.40       'Ethernet1'     [ 132770 ]
```

IGP + extra layers to achieve the simple task map IP=Label

- Let's get it straight…
  - Even if LDP is the controlplane for map labels to routes…
  - It can't do much without example Linkstate IGP, actually only thing LDP does is tell its peer buddy about IP=Label

```
router3b#sh isis topology
IS-IS Instance: backbone VRF: default
  IS-IS paths to level-2 routers
    System Id        Metric   Next-Hop             Interface       SNPA
    veos4            10       veos4            Ethernet1       P2P
    bgp-edge         20       veos4            Ethernet1       P2P
    bgp-edge2        30       veos4            Ethernet1       P2P

router3b#sh isis database bgp-edge2.00-00 detail
IS-IS Instance: backbone VRF: default
  IS-IS Level 2 Link State Database
    LSPID                 Seq Num   Cksum  Life  IS Flags
    bgp-edge2.00-00       7         45013  65524 L2 <>
      NLPID: 0xCC(IPv4)
      Area address: 49.0000
      Hostname: bgp-edge2
      Interface address: 1.1.100.0
      Interface address: 1.1.1.34
      Interface address: 1.1.1.102
      IS Neighbor           : bgp-edge.00         Metric: 10
      Reachability          : 1.1.1.102/32 Metric: 10 Type: 1 Up
      Reachability          : 1.1.100.0/31 Metric: 10 Type: 1 Up
      Reachability          : 1.1.1.32/30 Metric: 10 Type: 1 Up
```

```
router3b#sh mpls ldp bindings
1.1.1.4/32
   Local binding:  Label: 132768
   Remote binding: Peer ID: 1.1.1.4:0, Label: imp-null
1.1.1.6/32
   Local binding:  Label: 132770
   Remote binding: Peer ID: 1.1.1.4:0, Label: 132770
1.1.1.7/32
   Local binding:  Label: imp-null
   Remote binding: Peer ID: 1.1.1.4:0, Label: 132768
1.1.1.102/32
   Local binding:  Label: 132769
   Remote binding: Peer ID: 1.1.1.4:0, Label: 132769
```

IGP do the heavy lifting… without any band-aid

- Why not delegate that also to the Linkstate IGP…

```
router3b#sh isis database detail bgp-edge2.00-00
IS-IS Instance: backbone VRF: default
  IS-IS Level 2 Link State Database
    LSPID                 Seq Num   Cksum  Life  IS Flags
    bgp-edge2.00-00       6         38593  64357 L2 <>
      NLPID: 0xCC(IPv4)
      Area address: 49.0000
      Hostname: bgp-edge2
      Interface address: 1.1.100.0
      Interface address: 1.1.1.34
      Interface address: 1.1.1.102
      IS Neighbor           : bgp-edge.00         Metric: 10
        IPv4 Neighbor Address: 1.1.100.1
        IPv4 Interface Address: 1.1.100.0
        Adj-sid: 116384 flags: [ L V ] weight: 0x0
      Reachability          : 1.1.100.0/31 Metric: 10 Type: 1 Up
      Reachability          : 1.1.1.32/30 Metric: 10 Type: 1 Up
      Reachability          : 1.1.1.102/32 Metric: 10 Type: 1 Up
        SR Prefix-SID: 102 Flags: [ N ] Algorithm: 0
      Router Capabilities: 1.1.1.102 Flags: [ ]
        SR Capability: Flags: [ I ]
          SRGB Base: 900000 Range: 65536
```

```
router3b#sh mpls segment-routing bindings
1.1.1.4/32
   Local binding:  Label: 900004
   Remote binding: Peer ID: 0010.0100.1004, Label: imp-null
1.1.1.6/32
   Local binding:  Label: 900006
   Remote binding: Peer ID: 0010.0100.1004, Label: 900006
1.1.1.7/32
   Local binding:  Label: imp-null
   Remote binding: Peer ID: 0010.0100.1004, Label: 900007
1.1.1.102/32
   Local binding:  Label: 900102
   Remote binding: Peer ID: 0010.0100.1004, Label: 900102
```

Use Case / UCMP + SR

ARISTA

BGP UCMP & Segment Routing

- Handle ECMP peering with mixed peering speed and propagate LBW EXTENDED_COMMUNITIES down to the Leafs
- The path control however is with Segment Routing

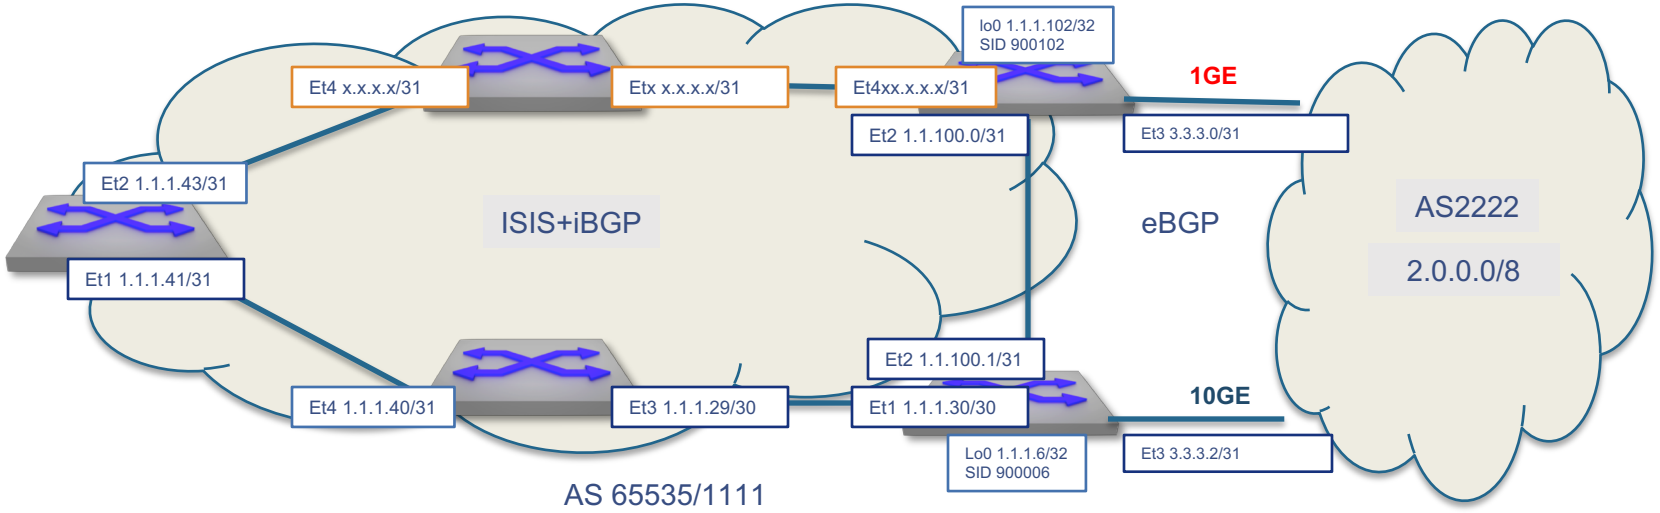

ARISTA

BGP UCMP without Segment Routing…

- Peering are bandwidth and price critical.
- Even if topology (due to link/node down) not longer provide ECMP, expensive peers needs to be used. Bw inside the network not problematic
- However if IP only forwarding, all traffic will egress only one path

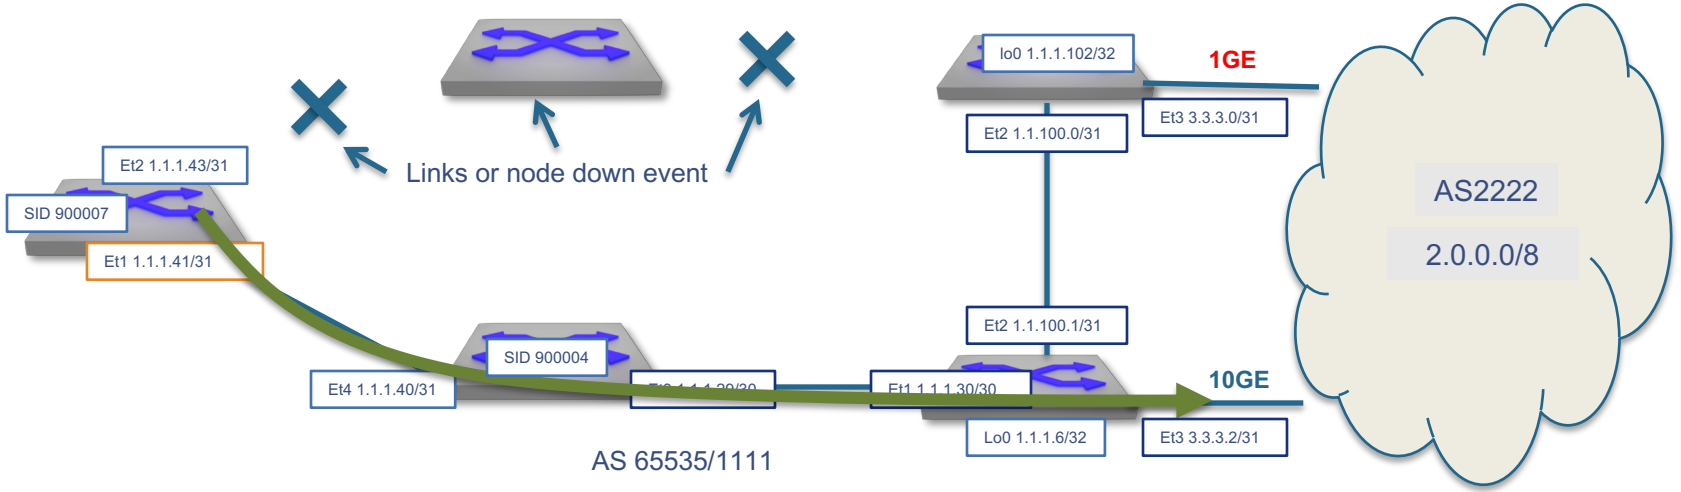

ARISTA

BGP UCMP & Segment Routing

- Hashing towards both peers handle by UCMP
- Segment Routing traffic goes to both BGP Edges and not only to the 10GE Peer, since forwarding is based on labels (900102&900006)

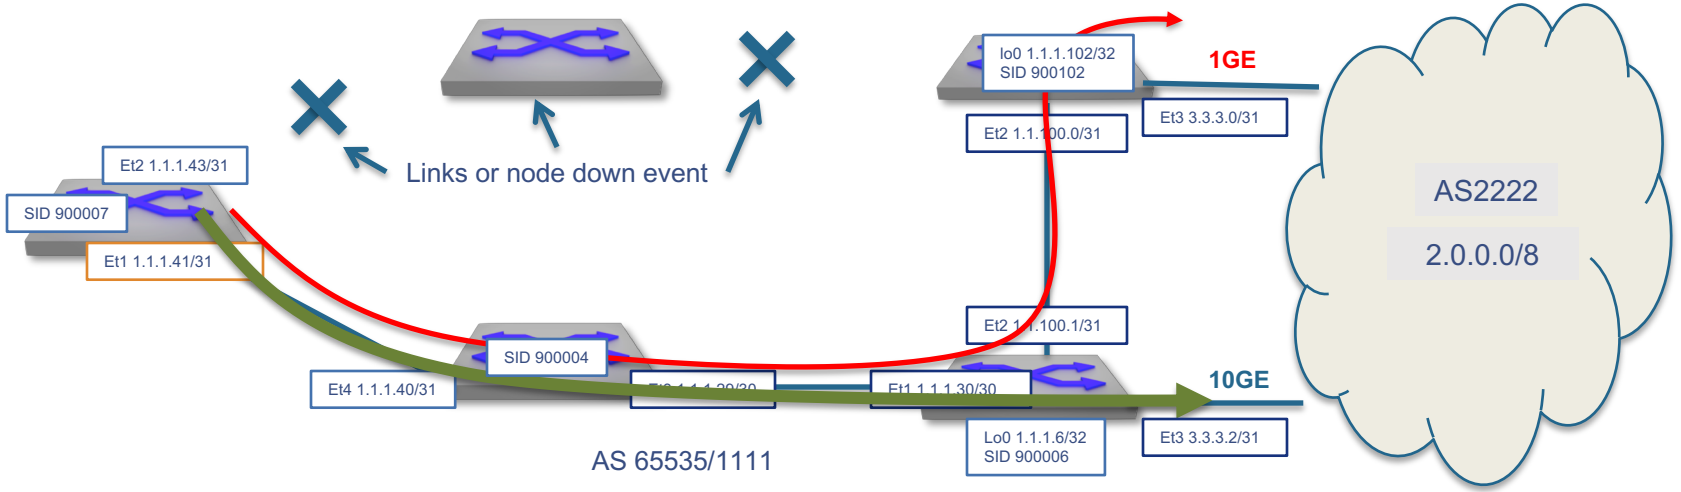

ARISTA

BGP UCMP in a Segment Routing network…

- BGP&ISIS Configuration

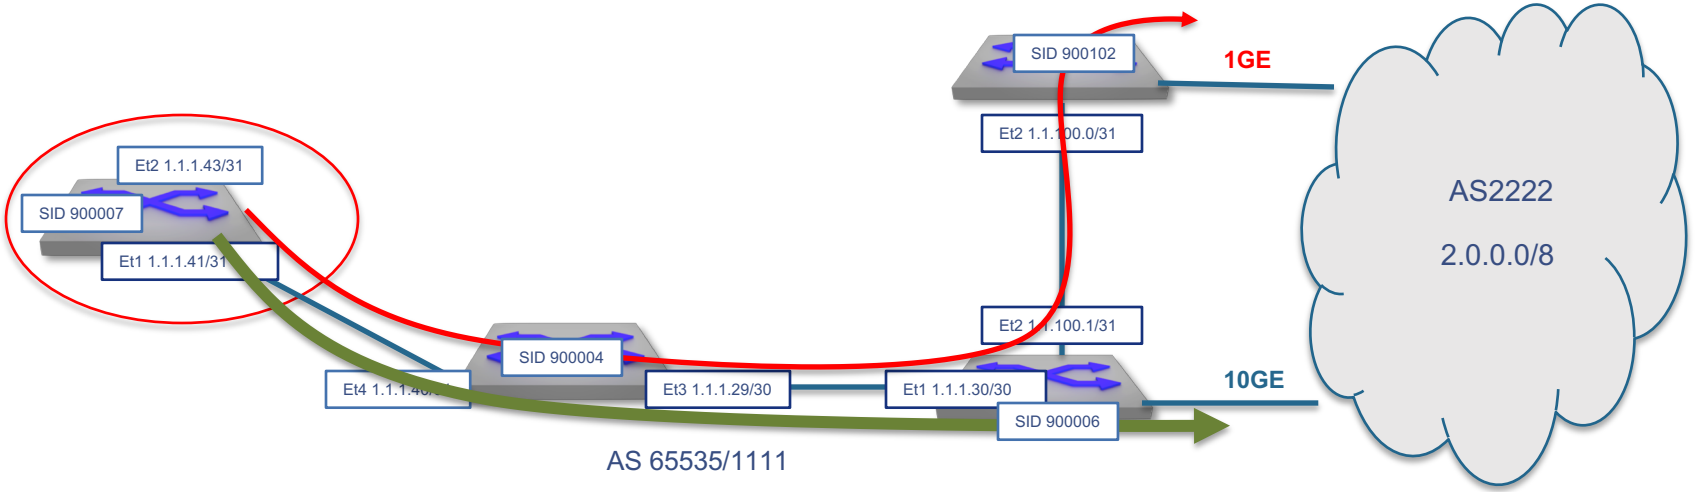

ARISTA

BGP UCMP in a Segment Routing

- BGP&ISIS Configuration

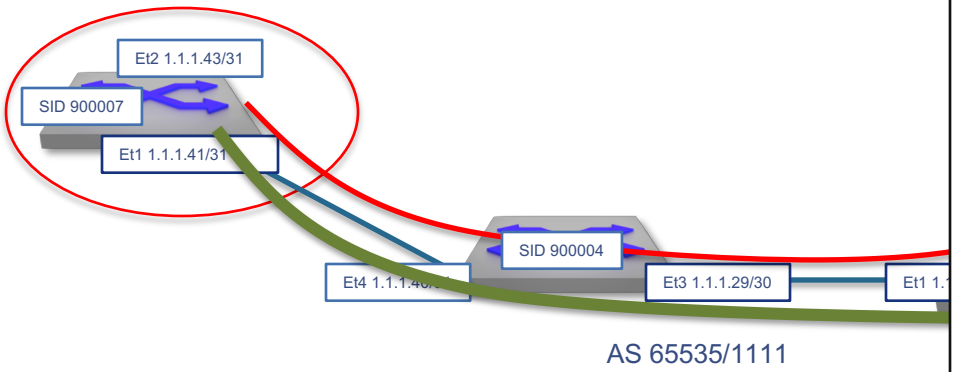


```
(…)
router-id 1.1.1.7
   maximum-paths 32 ecmp 32
   ucmp mode 1 32 0.01
   ucmp link-bandwidth encoding-weighted
   neighbor 65535 peer-group
   neighbor 65535 remote-as 65535
   neighbor 65535 local-as 65535 no-prepend replace-as
   neighbor 65535 update-source Loopback0
   neighbor 65535 soft-reconfiguration inbound all
   neighbor 65535 graceful-restart
   neighbor 65535 send-community
   neighbor 65535 maximum-routes 12000
   neighbor 65535 link-bandwidth
   neighbor 1.1.1.6 peer-group 65535
   neighbor 1.1.1.6 maximum-routes 12000
   neighbor 1.1.1.6 link-bandwidth auto
   neighbor 1.1.1.102 peer-group 65535
   redistribute connected
(…)
interface Loopback0
   ip address 1.1.1.7/32
   isis enable backbone
   node-segment ipv4 index 7
!
interface Ethernet1
   no switchport
   ip address 1.1.1.41/31
   isis enable backbone
   isis network point-to-point
!
mpls ip
!
router isis backbone
   net 49.0000.0010.0100.1101.00
   is-type level-2
   max-lsp-lifetime 65535
   !
   address-family ipv4 unicast
      maximum-paths 32
   !
   segment-routing mpls
      router-id 1.1.1.7
      no shutdown
      adjacency-segment allocation all-interfaces
```

BGP UCMP in a Segment Routing network…

- Recursive Lookup
  - Pls Follow the BGP NEXT_HOP double recursion

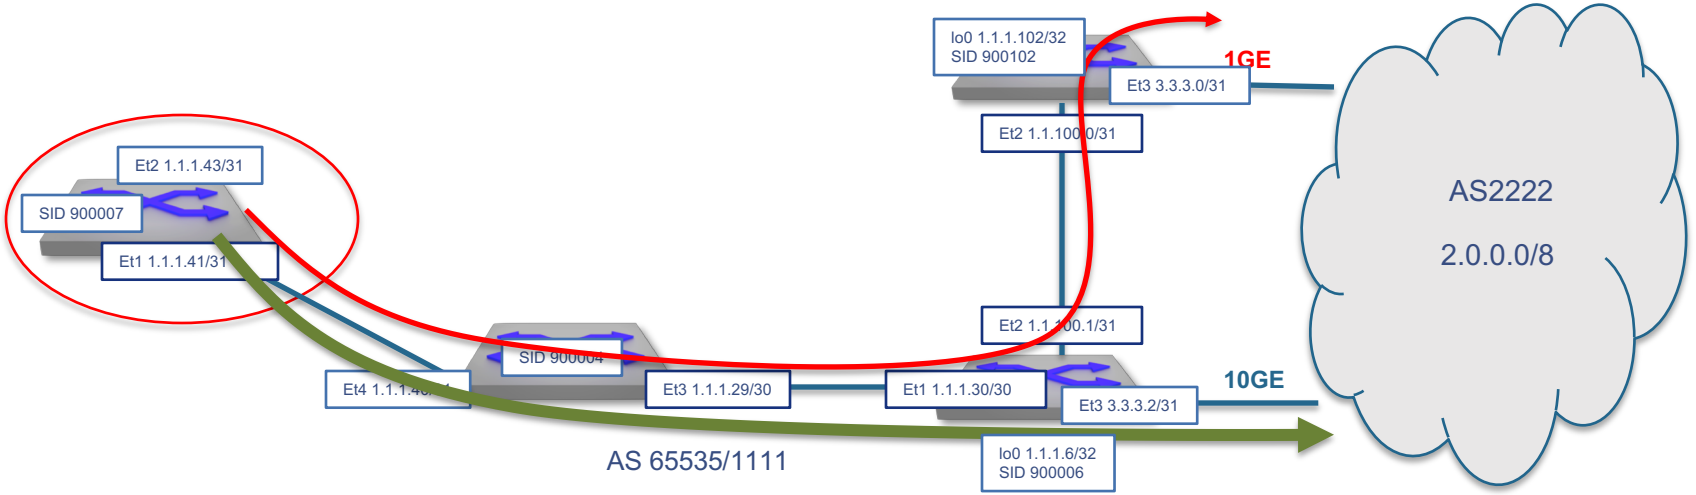

ARISTA

BGP UCMP in a Seg

- Recursive Lookup
  - Pls Follow the BGP NEXT_HOP double recursion

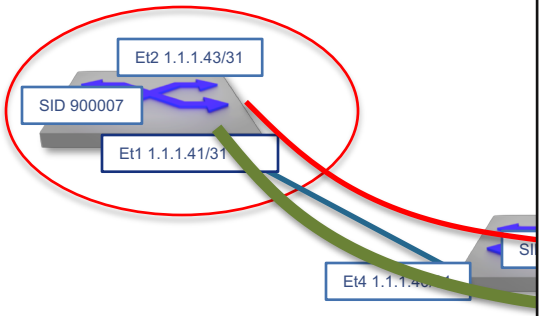


```
router3b#sh ip bgp 2.0.0.0
BGP routing table information for VRF default
Router identifier 1.1.1.7, local AS number 64515
BGP routing table entry for 2.0.0.0/8
 Paths: 2 available
  2222
    3.3.3.3 from 1.1.1.6 (1.1.1.6)
      Origin INCOMPLETE, metric 0, localpref 100, IGP metric 0, weight 0, received 00:05:24 ago, valid,
      internal, ECMP head, ECMP, UCMP, best, ECMP contributor
      Extended Community: Link-Bandwidth-AS:65535:10000000000.000000
      Rx SAFI: Unicast
  2222
    3.3.3.0 from 1.1.1.102 (1.1.1.102)
      Origin INCOMPLETE, metric 0, localpref 100, IGP metric 0, weight 0, received 00:05:24 ago, valid,
      internal, ECMP, UCMP, ECMP contributor
      Extended Community: Link-Bandwidth-AS:65535:1000000000.000000
      Rx SAFI: Unicast

router3b#sh ip bgp 3.3.3.3
(…)
BGP routing table entry for 3.3.3.2/31
 Paths: 1 available
  Local
    1.1.1.6 from 1.1.1.6 (1.1.1.6)
      Origin IGP, metric 0, localpref 100, IGP metric 30, weight 0, received 00:06:38 ago, valid,
      internal, best
      Rx SAFI: Unicast

router3b#sh ip bgp 3.3.3.0
(…)
BGP routing table entry for 3.3.3.0/31
 Paths: 1 available
  Local
    1.1.1.102 from 1.1.1.102 (1.1.1.102)
      Origin IGP, metric 0, localpref 100, IGP metric 40, weight 0, received 00:06:50 ago, valid,
      internal, best
      Rx SAFI: Unicast

router3b#sh ip ro 1.1.1.102
(...)
 I L2    1.1.1.102/32 [115/40] via 1.1.1.40, Ethernet1

router3b#sh ip ro 1.1.1.6
(...)
 I L2    1.1.1.6/32 [115/30] via 1.1.1.40, Ethernet1
```

BGP UCMP in a Segment Routing network…

- Ingress node
  MPLS table

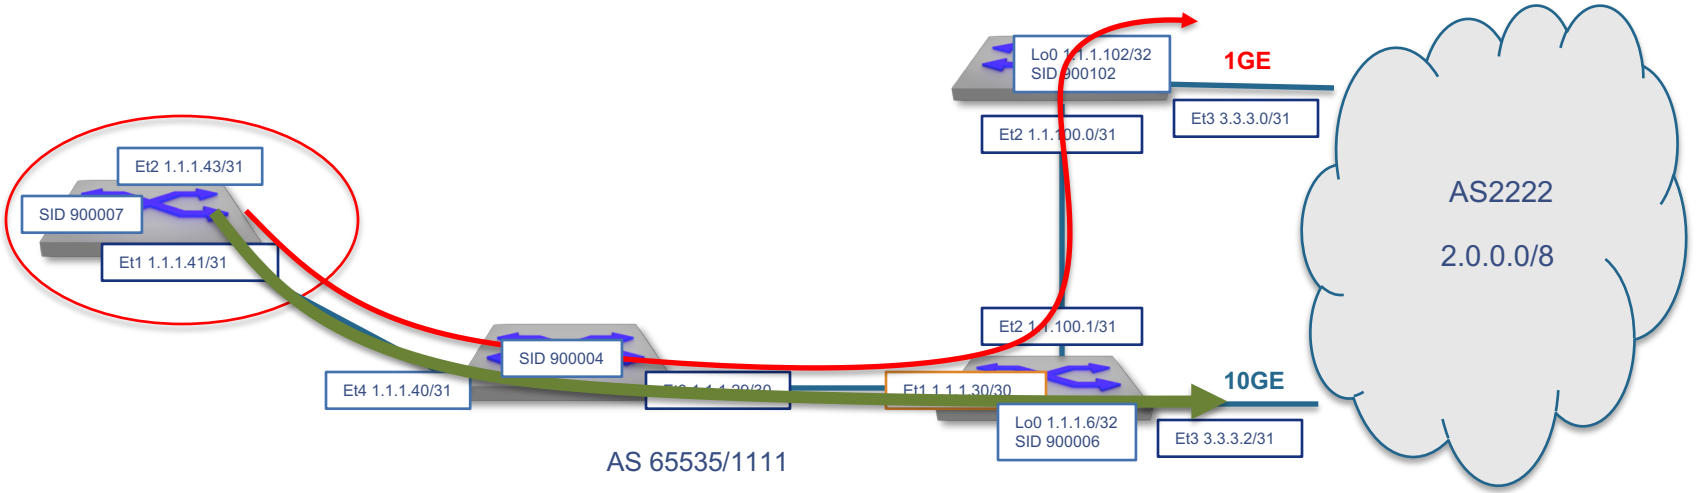

ARISTA

BGP UCMP in a Segment Routing network…

- Ingress node MPLS table

```
router3b#sh mpls segment-routing bi
1.1.1.4/32
   Local binding:  Label: 900004
   Remote binding: Peer ID: 0010.0100.1004, Label: imp-null
1.1.1.6/32
   Local binding:  Label: 900006
   Remote binding: Peer ID: 0010.0100.1004, Label: 900006
1.1.1.7/32
   Local binding:  Label: imp-null
   Remote binding: Peer ID: 0010.0100.1004, Label: 900007
1.1.1.102/32
   Local binding:  Label: 900102
   Remote binding: Peer ID: 0010.0100.1004, Label: 900102

router3b#sh mpls tunnel fib
   Tunnel Type         Index       Endpoint           Nexthop        Interface      Labels
------------------- ----------- ------------------ -------------- ----------------- ----------
   IS-IS SR IPv4       3           1.1.1.102/32       1.1.1.40       'Ethernet1'    [ 900102 ]
   IS-IS SR IPv4       1           1.1.1.4/32         1.1.1.40       'Ethernet1'    [ 3 ]
   IS-IS SR IPv4       2           1.1.1.6/32         1.1.1.40       'Ethernet1'    [ 900006 ]
(…)
```

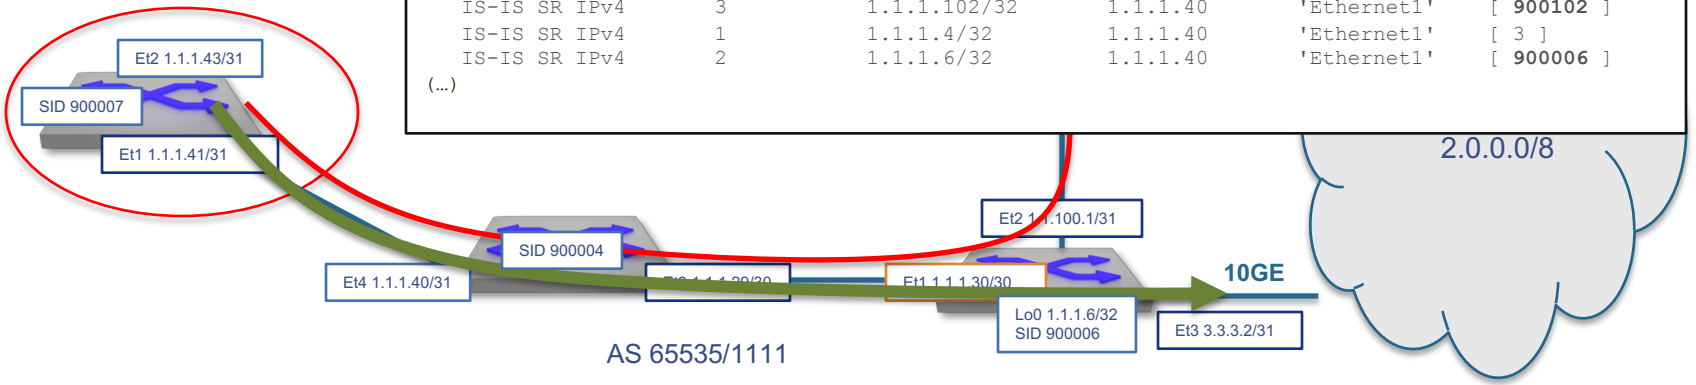

ARISTA

BGP UCMP in a Segment Routing network…

- Transit node MPLS table
  - Since traffic towards BGP_Edge with NEXT_HOP 3.3.3.0 and SID 9000102 been SWAP on Transit node, it will be label forward all the way thereby traffic will not just enter the first BGP_edge box with SID 900006
  - Result both peerings in use

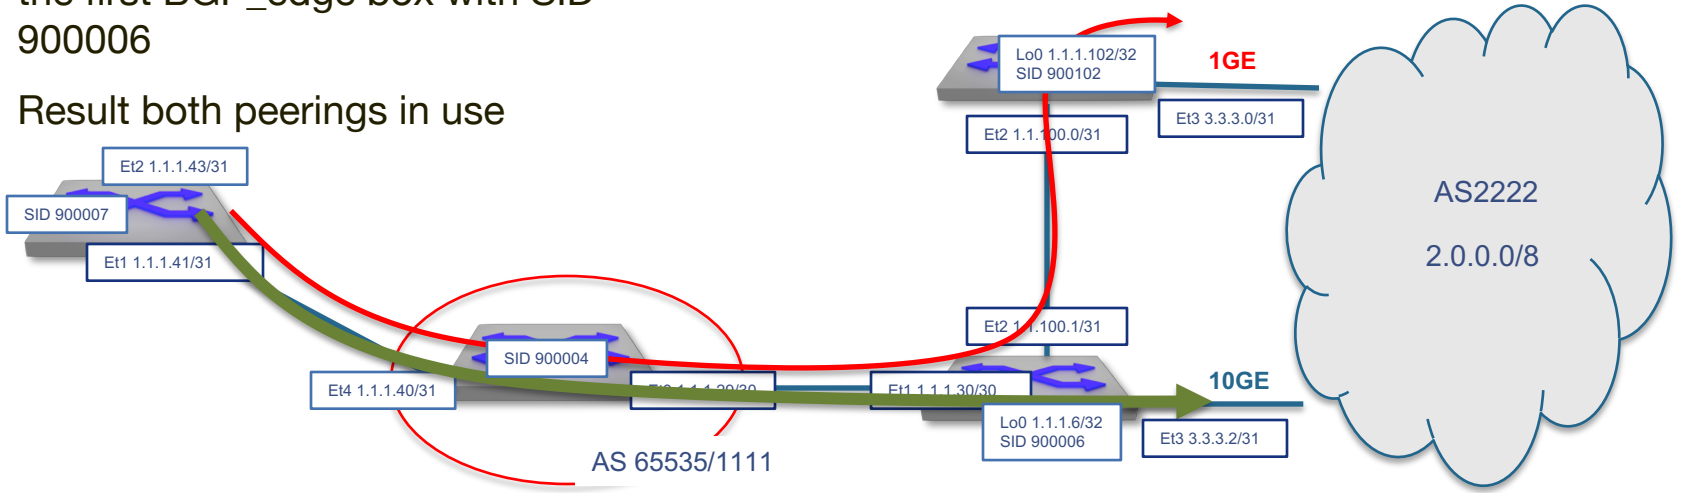

ARISTA

BGP UCMP in a Segment Routing network…

- Transit node MPLS table
  - Since traffic towards BGP_Edge with NEXT_HOP 3.3.3.0 and SID 9000102 been SWAP on Transit node, it will be label forward all the way thereby traffic will not just enter the first BGP_edge box with SID 900006
  - Result both peerings in use

```
router4#sh mpls lfib route
MPLS forwarding table (Label [metric] Vias) - 5 routes
MPLS next-hop resolution allow default route: False
Via Type Codes:
          M - Mpls Via, P - Pseudowire Via,
          I - IP Lookup Via, V - Vlan Via,
          VA - EVPN Vlan Aware Via, ES - EVPN Ethernet Segment Via,
          VF - EVPN Vlan Flood Via, AF - EVPN Vlan Aware Flood Via
Source Codes:
          S - Static MPLS Route, B2 - BGP L2 EVPN,
          B3 - BGP L3 VPN, P - Pseudowire,
          L - LDP, IP - IS-IS SR Prefix Segment,
          IA - IS-IS SR Adjacency Segment, IL - IS-IS SR Segment to LDP,
          LI - LDP to IS-IS SR Segment, BL - BGP LU,
          DE - Debug LFIB

(…)
IP  900006    [1], 1.1.1.6/32
                 via M, 1.1.1.30, pop
                     payload autoDecide, ttlMode autoDecide, apply egress-acl
                     interface Ethernet3
IP  900102    [1], 1.1.1.102/32
                 via M, 1.1.1.30, swap 900102
                     payload autoDecide, ttlMode autoDecide, apply egress-acl
                     interface Ethernet3

(…)
```

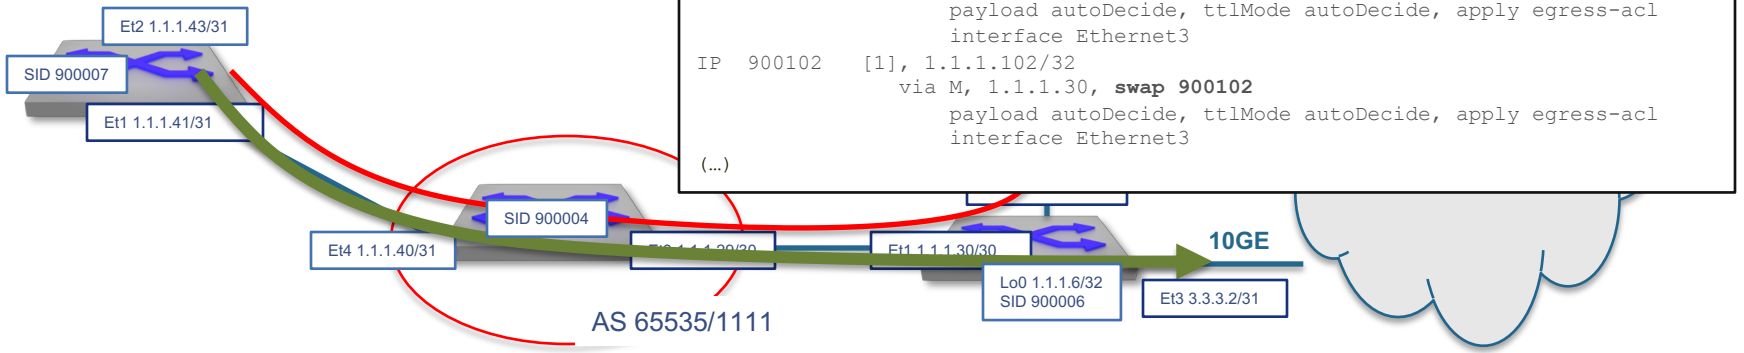

ARISTA

BGP UCMP in a Segment Routing network…

- The result,
  BGP Weights(*)
  applied

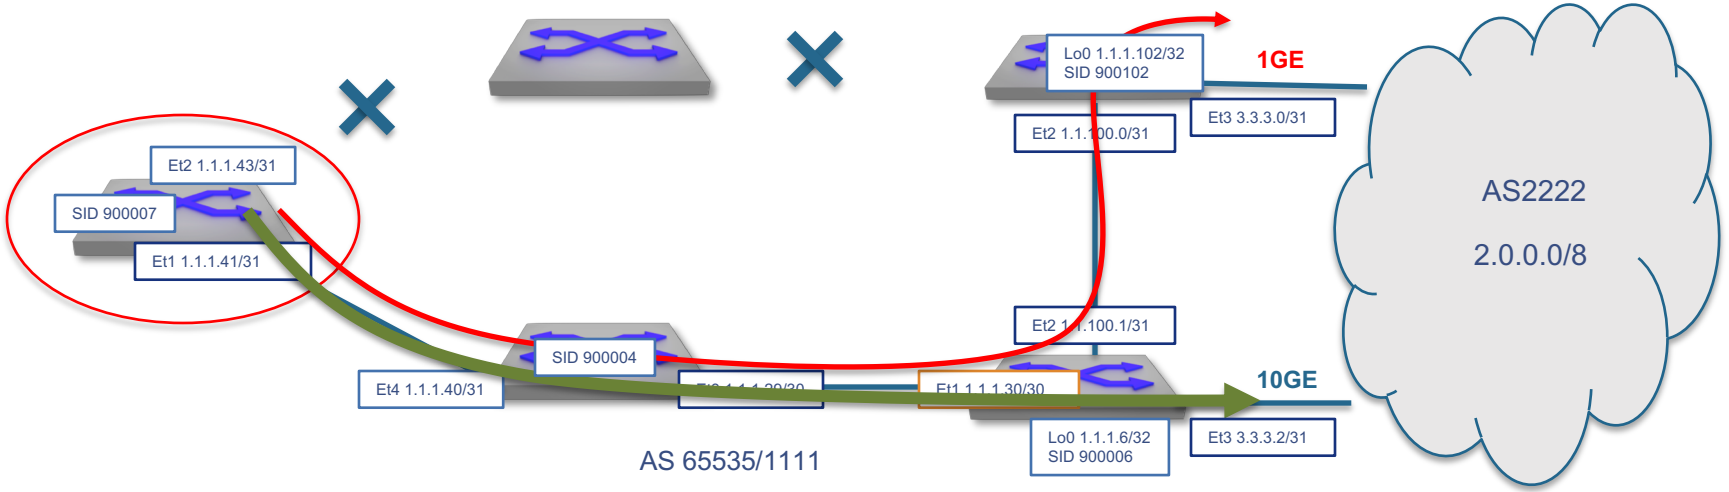

BGP UCMP in a Segment Routing network…

- The result, BGP Weights(*) applied

```
router3b#sh ip ro 2.0.0.0

VRF: default
Codes: C - connected, S - static, K - kernel,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type2, B I - iBGP, B E - eBGP,
       R - RIP, I L1 - IS-IS level 1, I L2 - IS-IS level 2,
       O3 - OSPFv3, A B - BGP Aggregate, A O - OSPF Summary,
       NG - Nexthop Group Static Route, V - VXLAN Control Service,
       DH - Dhcp client installed default route

 B I    2.0.0.0/8 [100/0] via 1.1.1.40, Ethernet1 label 900006, weight 10/11
                          via 1.1.1.40, Ethernet1 label 900102, weight 1/11
```

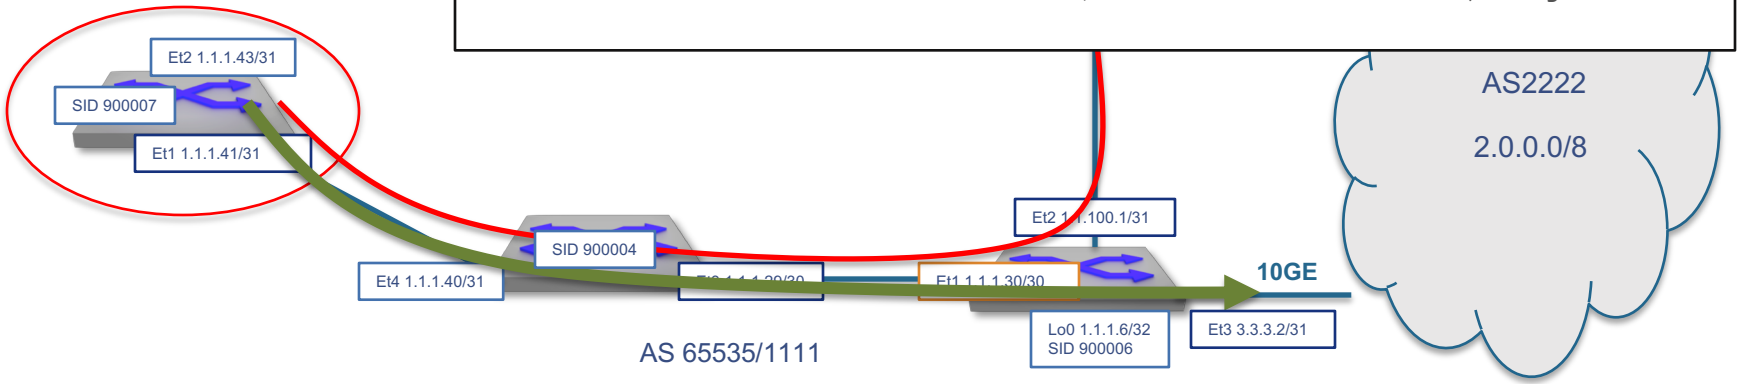

ARISTA

BGP UCMP in a Segment Routing network…

- Of course UCMP principle can be applied to generated routes like 0.0.0.0/0
- NEXT_HOP set to BGP Peer address

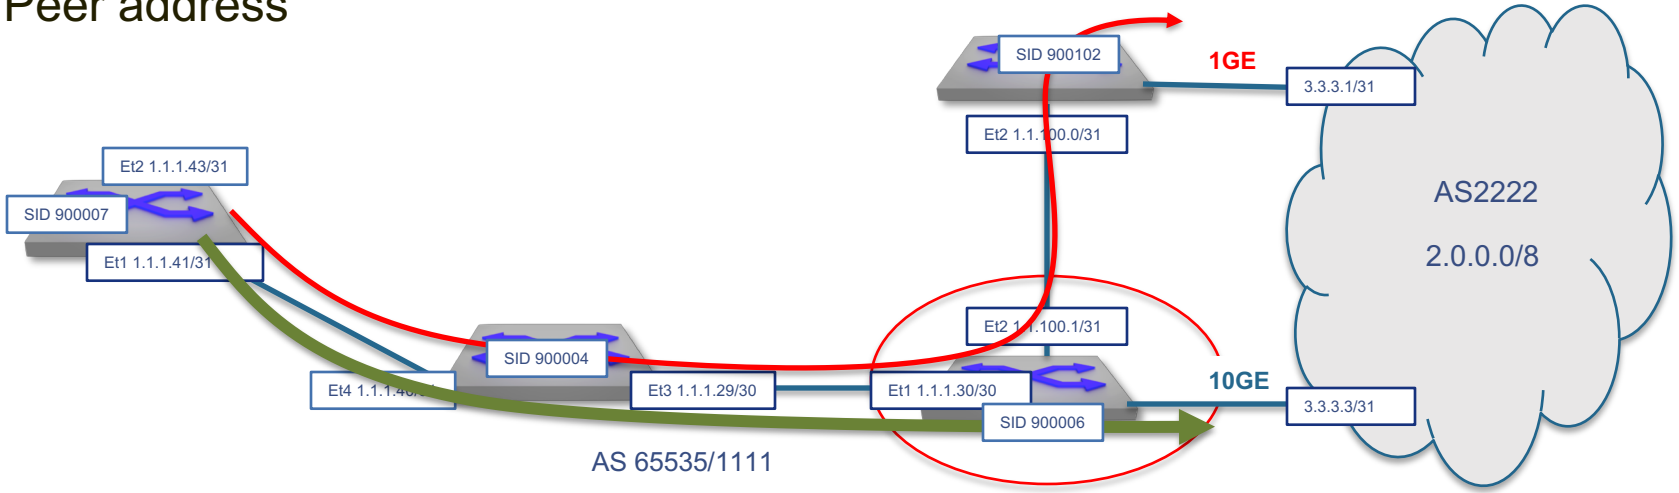

BGP UCMP in a Segment Routing network

- Of course UCMP principle can be applied to generated routes like 0.0.0.0/0
- NEXT_HOP set to BGP Peer address

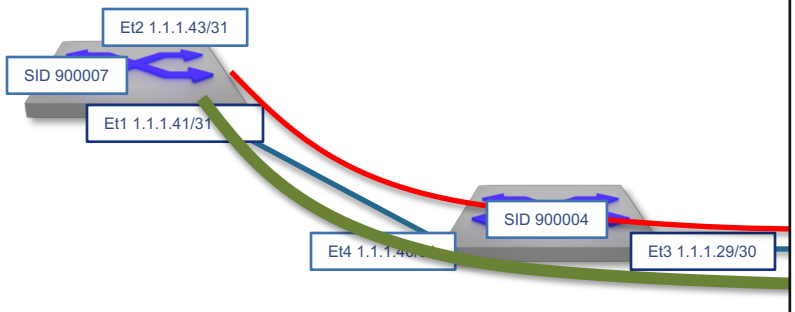


```
(…)
!
route-map default permit 10
   match ip address prefix-list as2222
   set origin incomplete
   set extcommunity lbw 10G
   set ip next-hop 3.3.3.3
!
router bgp 1111
   router-id 1.1.1.6
   maximum-paths 64 ecmp 64
   ucmp mode 1 64 0.01
   ucmp link-bandwidth encoding-weighted
   neighbor 65535 peer-group
   neighbor 65535 remote-as 65535
   neighbor 65535 link-bandwidth update-delay 1
   neighbor 65535 local-as 65535 no-prepend replace-as
   neighbor 65535 update-source Loopback0
   neighbor 65535 soft-reconfiguration inbound all
   neighbor 65535 graceful-restart
   neighbor 65535 route-map isis-sr out
   neighbor 65535 default-originate route-map default
   neighbor 65535 send-community
   neighbor 65535 maximum-routes 12000
   neighbor 65535 link-bandwidth
(…)
   neighbor 1.1.1.7 peer-group 65535
   neighbor 1.1.1.7 send-community link-bandwidth aggregate
   neighbor 1.1.1.7 link-bandwidth auto
(…)
```


ARISTA

BGP UCMP in a Segment Routing network…

- The result, BGP Weights applied

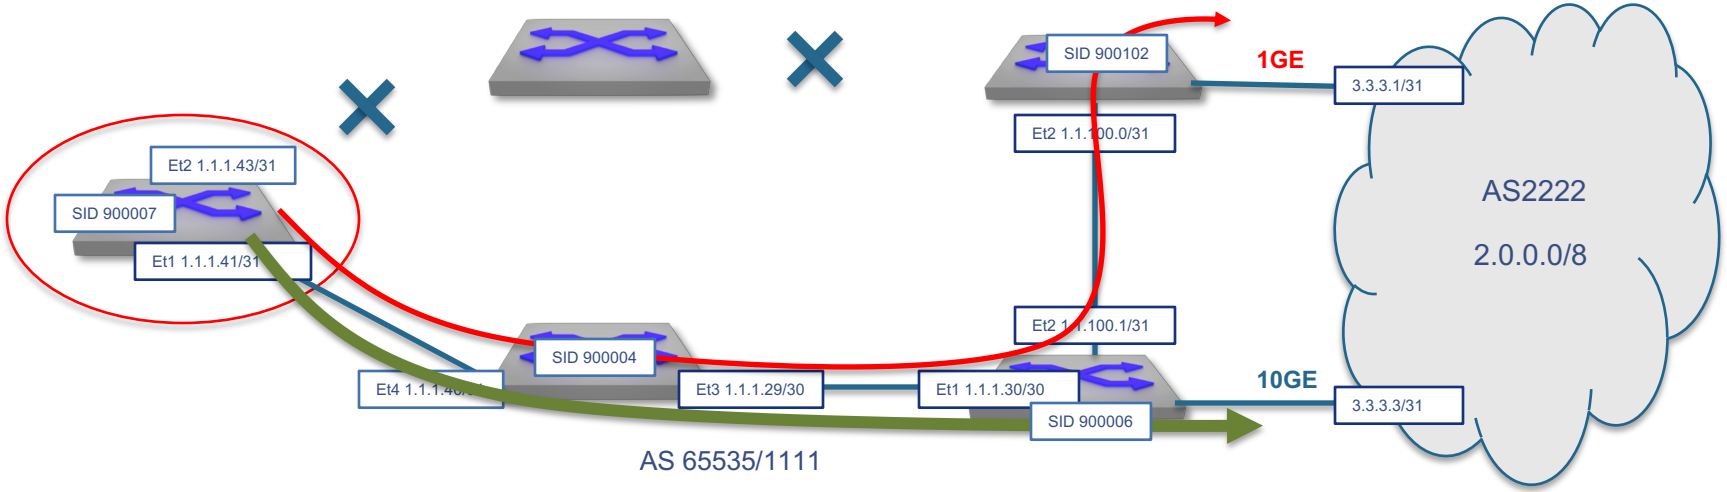

ARISTA

BGP UCMP in a Segment Routing network…

- The result, BGP Weights applied

```
router3b#sh ip bgp 0.0.0.0
BGP routing table information for VRF default
Router identifier 1.1.1.7, local AS number 64515
BGP routing table entry for 0.0.0.0/0
 Paths: 2 available
  Local
    3.3.3.3 from 1.1.1.6 (1.1.1.6)
      Origin INCOMPLETE, metric 0, localpref 100, IGP metric 0, weight 0, received 00:02:16 ago,
      valid, internal, ECMP head, ECMP, UCMP, best, ECMP contributor
      Extended Community: Link-Bandwidth-AS:65535:10000000000.000000
      Rx SAFI: Unicast
  Local
    3.3.3.0 from 1.1.1.102 (1.1.1.102)
      Origin INCOMPLETE, metric 0, localpref 100, IGP metric 0, weight 0, received 00:02:16 ago,
      valid, internal, ECMP, UCMP, ECMP contributor
      Extended Community: Link-Bandwidth-AS:65535:1000000000.000000
      Rx SAFI: Unicast

router3b#sh ip ro 0.0.0.0/0
(...)
Gateway of last resort:
 B I    0.0.0.0/0 [100/0] via 1.1.1.40, Ethernet1 label 900006, weight 10/11
                          via 1.1.1.40, Ethernet1 label 900102, weight 1/11
```

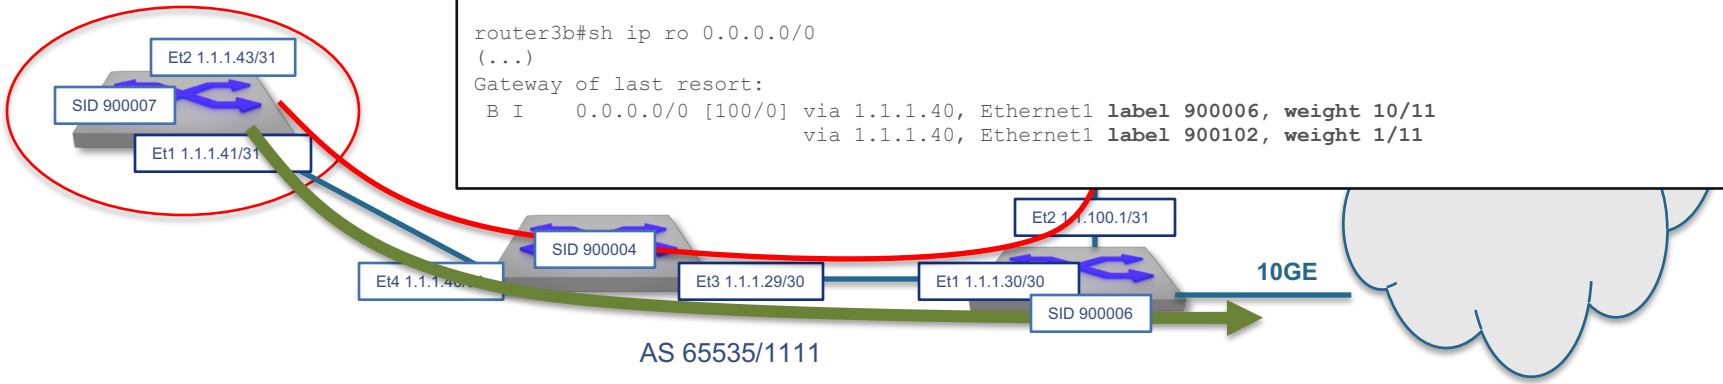

ARISTA

BGP UCMP in a Segment Routing network…

- For more specific path engineering, BGP-LU can be used as controller
- Then its careless what topology or metric

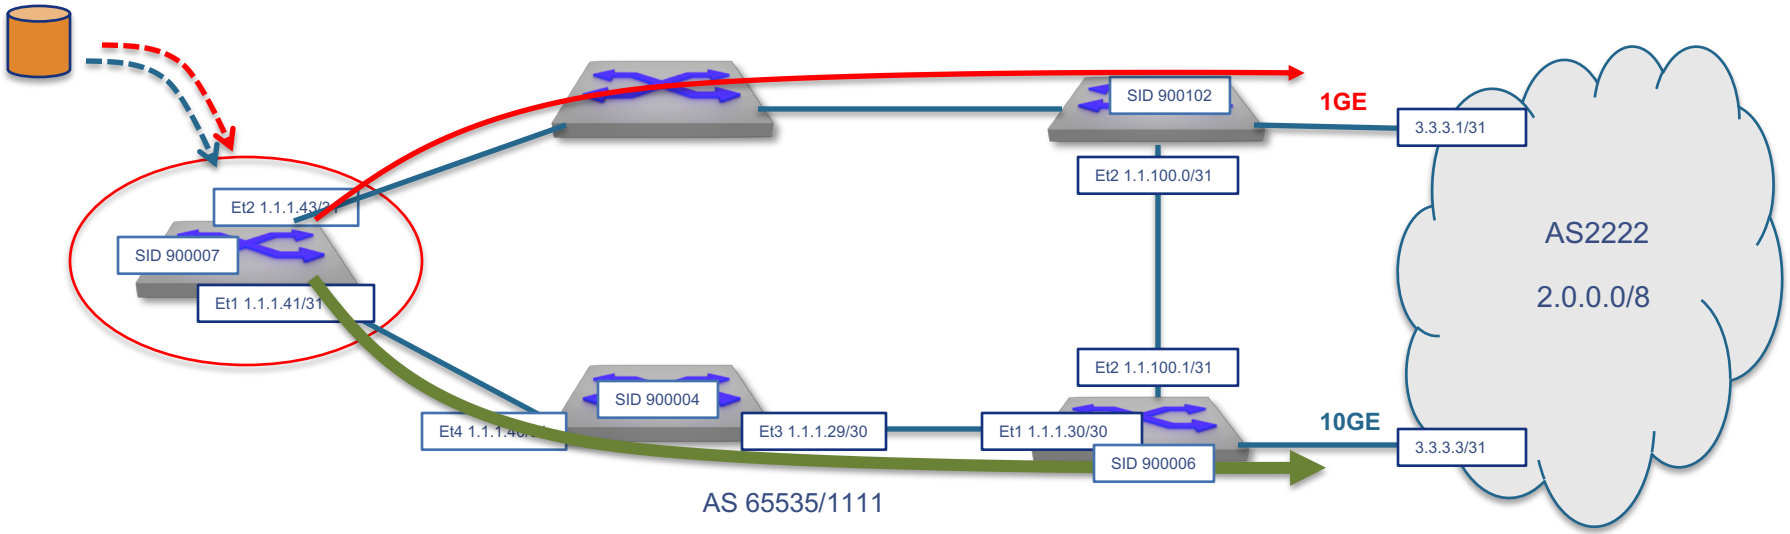

ARISTA

BGP UCMP in a Segment Routing network…

- For more specific
  controller
- Then its careless

```
poc@testhost:~$ /usr/sbin/exabgp /etc/exabgp/exabgp.conf
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | generate it using "exabgp --fi > /usr/etc/exabgp/exabgp.env"
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | reactor       | Performing reload of exabgp 3.4.12
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading | #
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading | # IPv4&Label-unicast template
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading | neighbor 192.168.10.101 {
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading |    router-id 192.168.10.70;
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading |    group-updates true;
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading |    local-address 192.168.10.70;
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading |    local-as 64515;
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading |    peer-as 64515;
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading |    hold-time 180;
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading |    capability {
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading |       graceful-restart 120;
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading |       route-refresh;
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading |       add-path send;
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading |       aigp enable;
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading |    }
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading |    family {
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading | #     ipv4 flow;
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading | #     ipv4 unicast;
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading |       ipv4 nlri-mpls;
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading |    }
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading |    static {
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading |       route 3.3.3.3/32 {
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading |          next-hop 1.1.1.40;
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading |          label [ 900004 900006 ];
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading |       }
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading |       route 3.3.3.1/32 {
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading |          next-hop 1.1.1.40;
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading |          label [ 900004 900006 900102];
Fri, 29 Sep 2017 02:15:30 | INFO     | 2082   | configuration | loading |       }
(...)
```

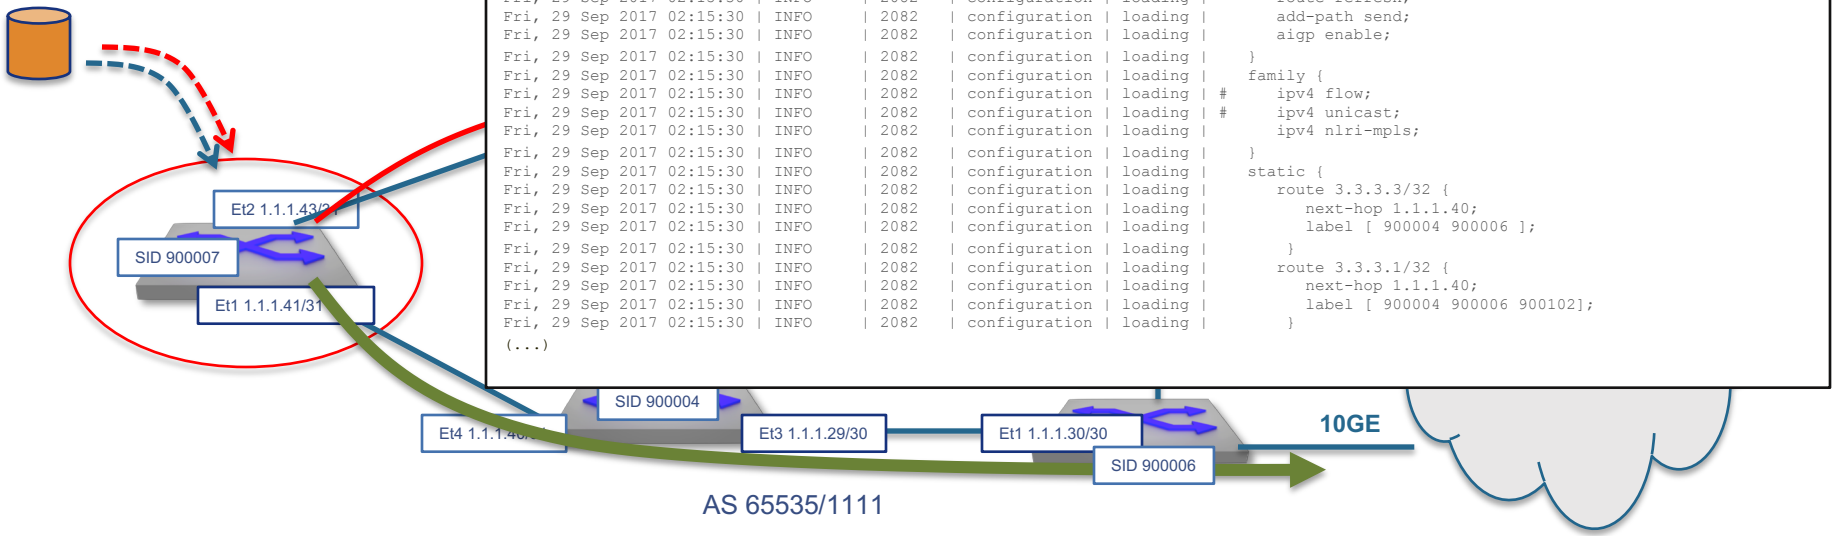

ARISTA

BGP UCMP in a Segment Routing network…

- The result, BGP-LU programs path to the NEXT_HOP and weights applied to routes

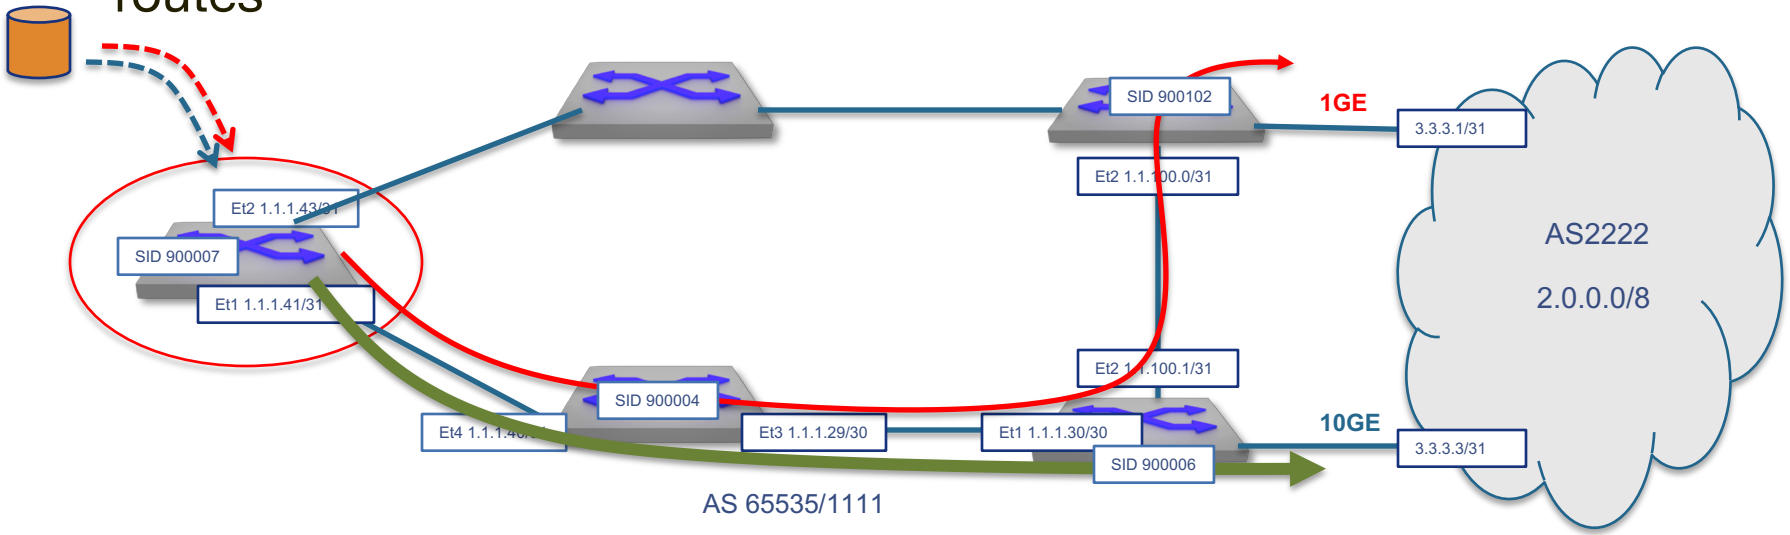

ARISTA

BGP UCMP in a Segment Routing network…

- The result: BGP LU

```
router3b#sh mpls tunnel fib
   Tunnel Type          Index       Endpoint           Nexthop        Interface      Labels
------------------- ----------- ------------------ -------------- ----------------- -------------------------
   IS-IS SR IPv4        3           1.1.1.102/32       1.1.1.40       'Ethernet1'    [ 900102 ]
   IS-IS SR IPv4        1           1.1.1.4/32         1.1.1.40       'Ethernet1'    [ 3 ]
   IS-IS SR IPv4        2           1.1.1.6/32         1.1.1.40       'Ethernet1'    [ 900006 ]
   BGP LU               2           3.3.3.1/32         1.1.1.40       'Ethernet1'    [ 900004 900006 900102 ]
   BGP LU               1           3.3.3.3/32         1.1.1.40       'Ethernet1'    [ 900004 900006 ]

veos3b# sh ip ro | grep "900004 900006"
 B I    0.0.0.0/0 [100/0] via 1.1.1.40, Ethernet1 label 900004 900006, weight 10/11
                          via 1.1.1.40, Ethernet1 label 900004 900006 900102, weight 1/11
 B I    2.0.0.0/8 [100/0] via 1.1.1.40, Ethernet1 label 900004 900006, weight 10/11
                          via 1.1.1.40, Ethernet1 label 900004 900006 900102, weight 1/11
```

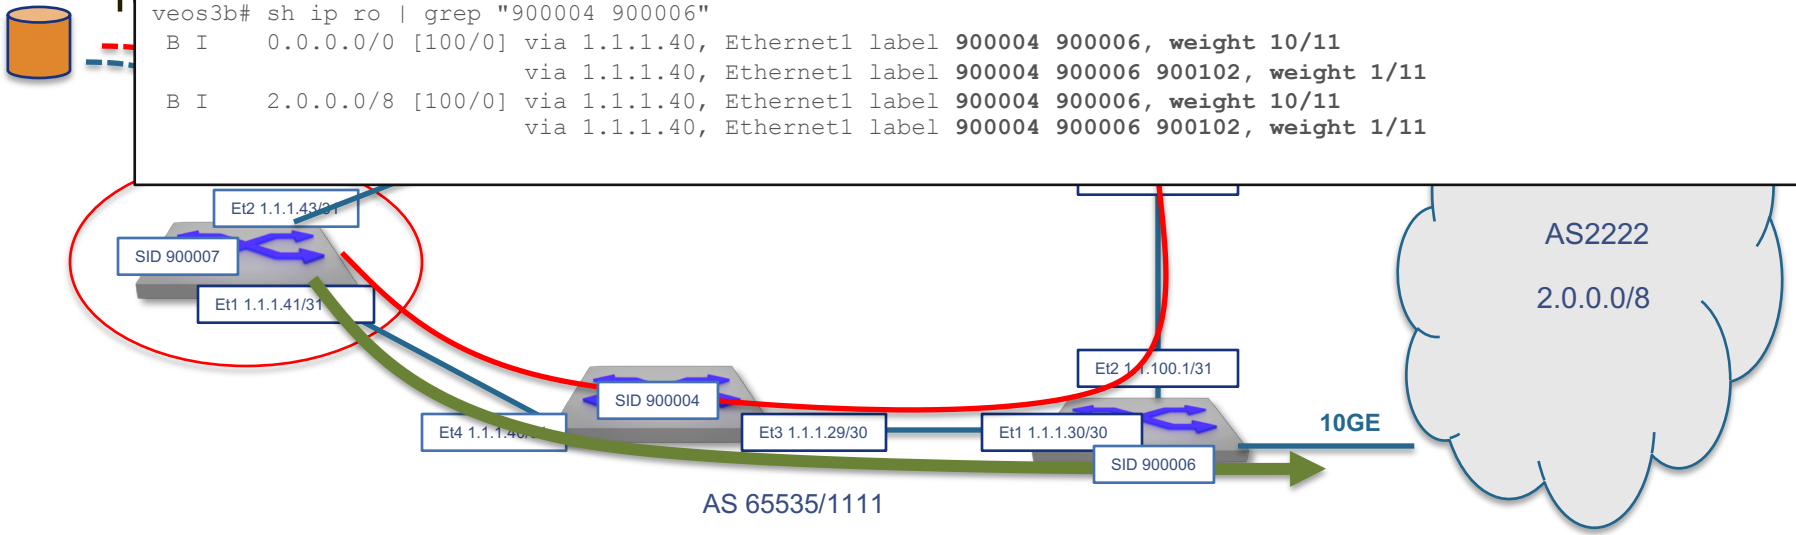

ARISTA

Thank You
For Your Attention

Florian Hibler
Systems Engineer
Arista Networks, Inc.

(e) florian@arista.com
(m) +49 171 7576089
(w) http://www.arista.com

ARISTA