Cisco live!

June 25–29, 2017 • Las Vegas, NV

FlexVPN Remote-Access, IoT & Site-to-Site Advanced Crypto Design

Frederic Detienne – Distinguished Engineer

Piotr Kupisiewicz – Aspiring Technical Leader

BRKSEC-3054

Cisco Spark

Questions?

Use Cisco Spark to chat with the speaker after the session

How

1. Find this session in the Cisco Live Mobile App
2. Click “Join the Discussion”
3. Install Spark or go directly to the space
4. Enter messages/questions in the space

Cisco Spark spaces will be available until July 3, 2017.

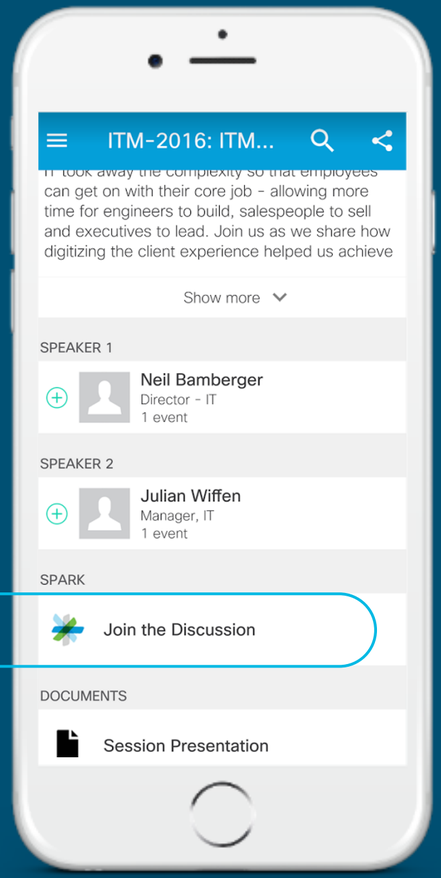

cs.co/ciscolivebot#BRKXXX-xxx

Agenda

- FlexVPN in a nutshell
- Hub & Spoke and Shortcut Switching + MPLS + Multicast
- Routed Backup and High Availability
- AAA and Per Branch Features
- Backup Mechanisms and Load Balancing
- Conclusion

FlexVPN,GET, DMVPN and all that…

- GET VPN does not provide a VPN
  - Overlay needs to be provided by other technologies (e.g. MPLS)
  - cryptography only with header preservation
  - Group key only
- IWAN2 is DMPVN based and IWAN3 is LISP/ESON based
  - Marketing decision to segment the markets
  - ➔ if you need IWAN3, LISP/ESON is the way to go!
- Excluding IWAN, FlexVPN supersedes DMVPN
  - FlexVPN = IKEv2 (EzVPN + DMVPN + Point-to-Point + RA + 3rd Party)
- This presentation shows what FlexVPN can do
  - Use what you see fit

Supported Platforms

- Routers:
  - IOS: ISR-G2
  - IOS-XE: ASR-1000 series, CSR-1000v, ISR-NG (4000 series)
  - ESR Series for IoT
  - IR Series for IoT
- Software clients: AnyConnect
- Interoperates with 3rd party implementing IPsec/IKEv2
- Not supported: ASA, ISR-G1, Switches
  - ASA support high on mind but not committed (yet ?). Contact your account team if you have interest.

Key Platforms

ASR 1000 series

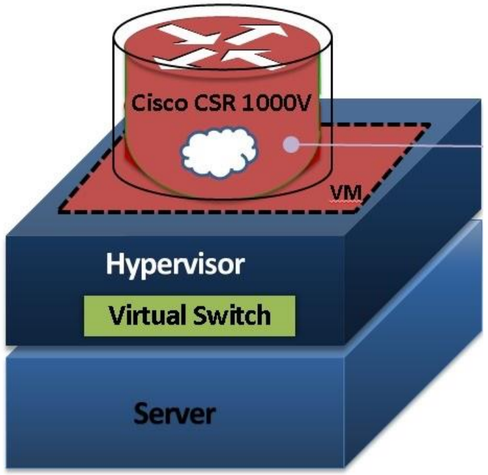

ISR 800 Series

ISR 4000 Series

Cisco Press Book ‘IKEv2 IPsec VPNs’ by Amjad Inamdar & Graham Bartlett

Grab Cisco Press rebate code: ike35

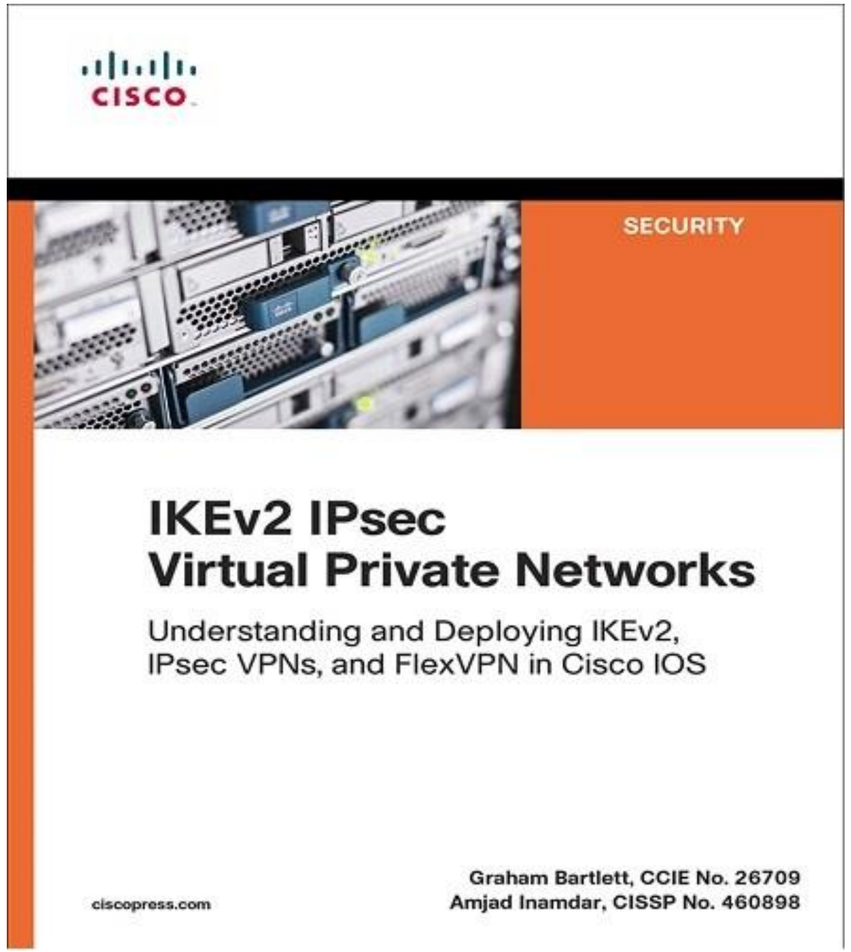

https://www.amazon.com/IKEv2-IPsec-Virtual-PrivateNetworks/dp/1587144603/

Listed in the CCIE Security reading list
https://learningnetwork.cisco.com/community/certifications/ccie_security/written_exam/study-material

Customer Reviews ★★★★★

One of the best technical books I've read
This book is the IKEv2 VPN equivalent of Jeff Doyle's Routing TCP/IP Vol 1 & 2 - a must read for any network security engineer wanting to design and build secure VPN's. One of the best technical books I've read.

Superb book and well worth the money for anyone even thinking about Cisco crypto
This book is the most comprehensive book on IKEv2 for Cisco network engineers that you will find and is all about real-world scenarios.

Definitive guide on modern IPsec VPN theory and practice
Many times I wish I had a book like this to help distill many complex IETF RFCs into “plain English” and provide practical and actionable security best practices.

Brilliant
I bought the Kindle version of this on a bit of an impulse. I'm really glad I did, it's well worth the money. Not only can I establish secure IKEv2 tunnels, I also feel like I know the subject thoroughly now. Even in respect to non-Cisco equipment. The book is a great reference too. I don't usually leave reviews but was motivated to in this instance. Good job, highly recommended.

The best book on IKEv2 IPsec VPNs
The book is awesome! I appreciate authors' work on presenting deeply technical topics in extremely easy to understand manner.

Finally, all you need to know about FLEX in one place!
Well written , concise and accurate. An absolute must for anyone designing, supporting or troubleshooting IKEv2 VPNs. You too can become a FLEX expert!

Very good Book on IPsec VPN for Enterprise networks
Very well Written book, This book touches on most important topic on building Dynamic VPN for enterprise networks.

Where is FlexVPN used – Some Examples

tail-f

Network Service Orchestrator

Cloud Web Security

ARCANA NETWORKS

MEVO

Network management platform for

Enterprises, Government,

Service Providers / SMB

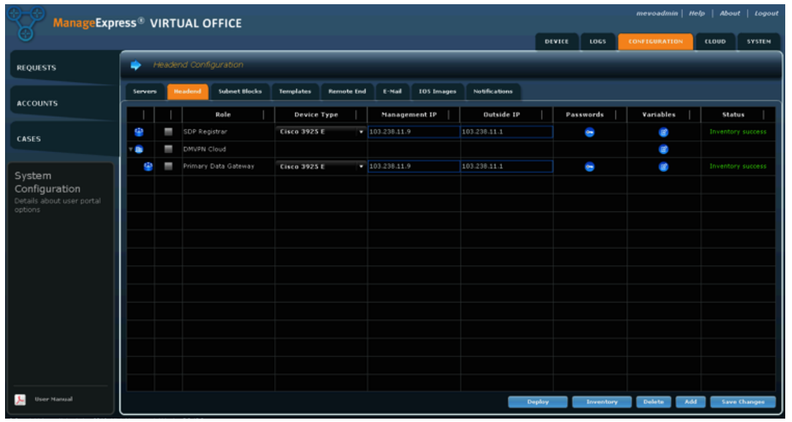

Field Network Director

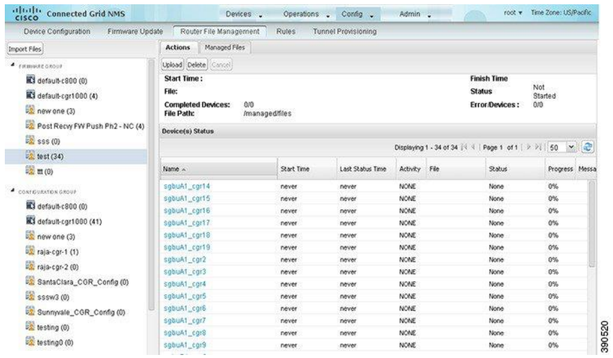

CISCO

CSIRT Intercloud

FlexVPN Quick Recap

IKEv2 CLI Overview

IKEv2 Profile – extensive CLI

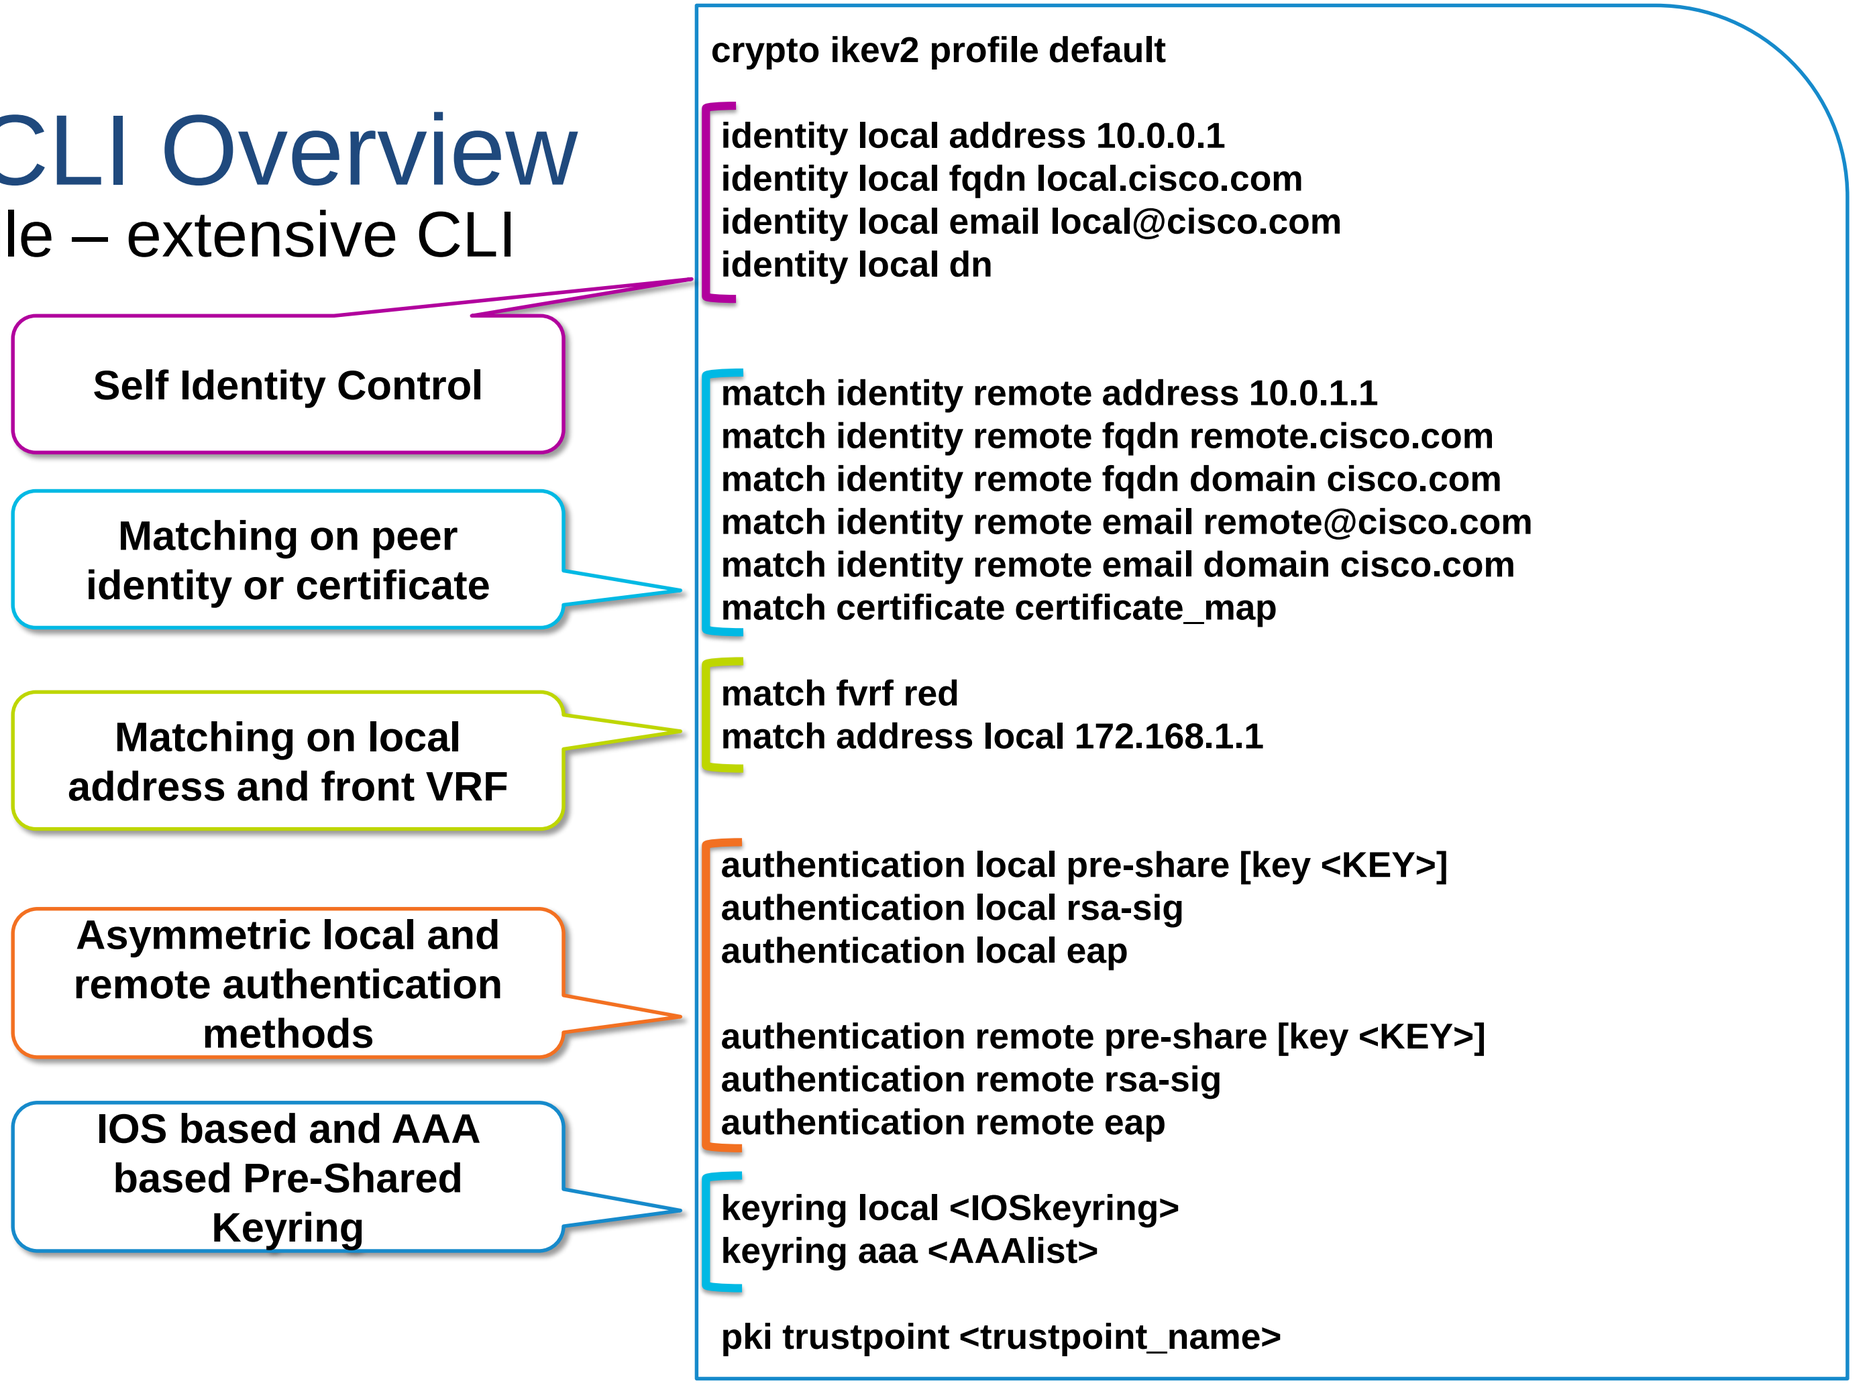

IKEv2 Basic Negotiation

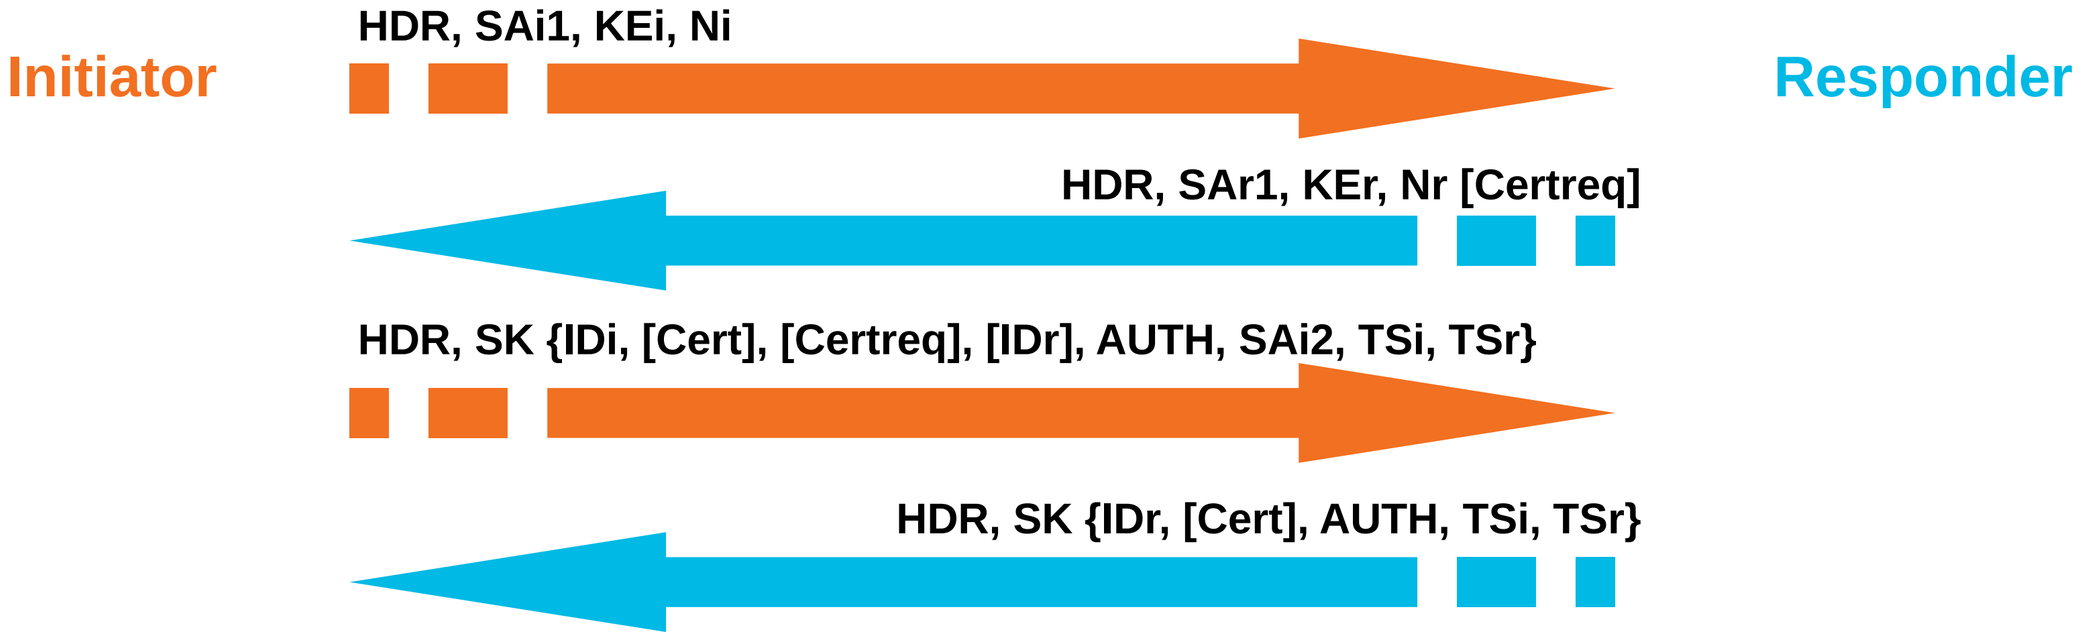

HDR – IKE Header

SA[i/r] – cryptographic algorithms the peer proposes/accepts

KE[i/r] – Initator Key Exchange material

N[i/r] – Initiator/Responder Nonce

SK– payload encrypted and integrity protected

ID[i/r] – Initiator/Responder Identity

Cert(req) – Certificate (request)

AUTH – Authentication data

SA - Includes SA, Proposal and Transform Info to Create the 1st CHILD_SA

Ts[i/r] – Traffic Selector as src/dst proxies

IKEv2 Profile Match Statements

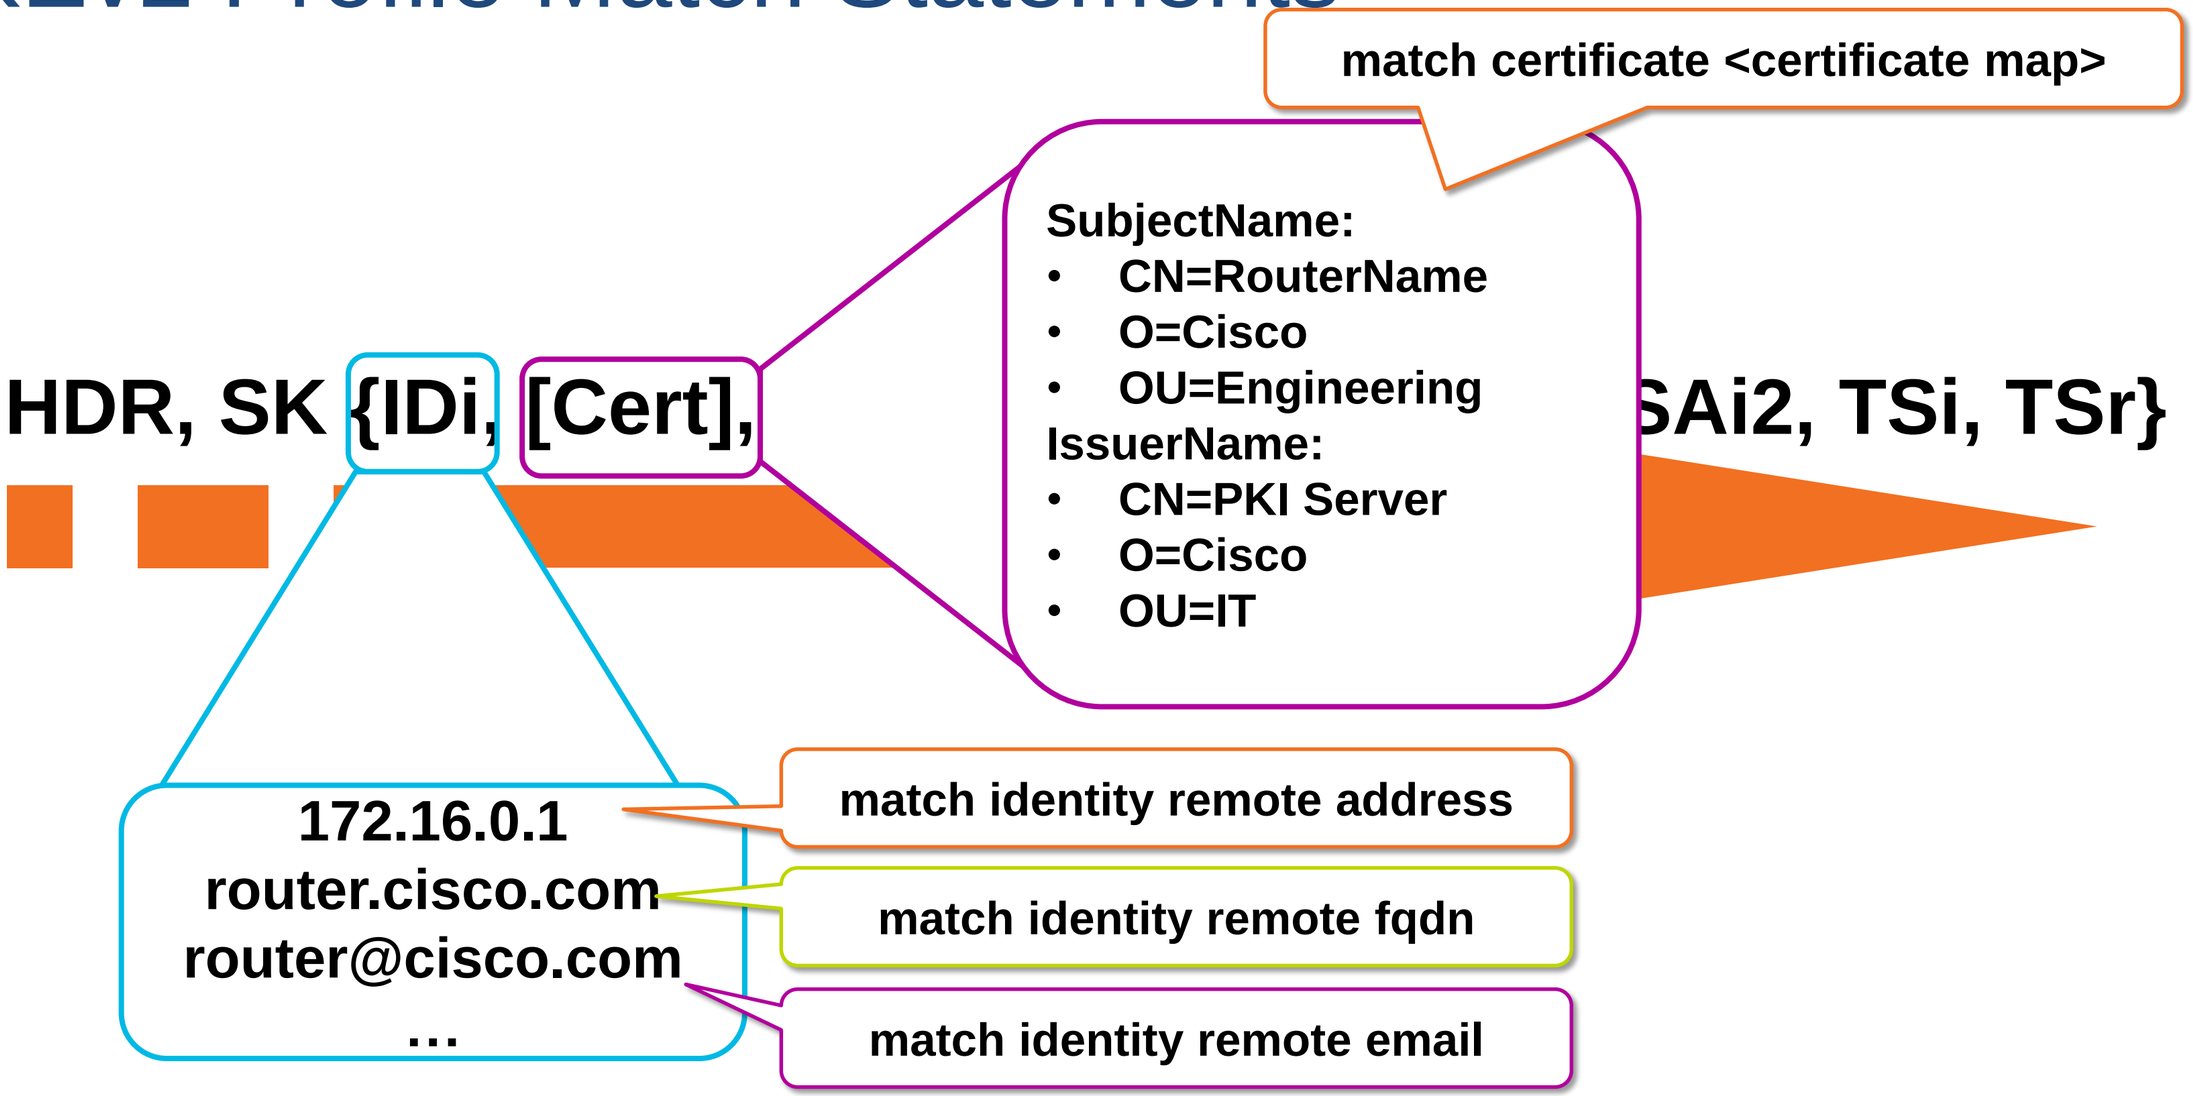

IPsec CLI Overview

Tunnel Protection

IPsec transform

IPsec profile defines SA parameters and points to IKEv2 profile

Dynamic and Static point-to-point interfaces

Static point-to-point interfaces

Tunnel protection links to IPsec profile

```
crypto ipsec transform-set default esp-aes 128 esp-sha-hmac

crypto ipsec profile default
 set transform-set default
 set crypto ikev2 profile default

interface Virtual-Template1 type tunnel
 ip unnumbered Loopback0
 tunnel protection ipsec profile default

interface Tunnel0
 ip address 10.0.0.1 255.255.255.252
 tunnel source Ethernet0/0
 tunnel destination 172.16.2.1
 tunnel protection ipsec profile default
```

Introducing Smart Defaults

Intelligent, reconfigurable defaults

```
crypto ipsec transform-set default
            esp-aes 128 esp-sha-hmac

crypto ipsec profile default
 set transform-set default
 set crypto ikev2 profile default

crypto ikev2 proposal default
  encryption aes-cbc-256 aes-cbc-128 3des
  integrity sha512 sha 256 sha1 md5
  group 5 2

crypto ikev2 policy default
  match fvrf any
  proposal default

crypto ikev2 authorization policy default
 route set interface
 route accept any
```

These constructs are the Smart Defaults

```
crypto ikev2 profile default
 match identity remote address 10.0.1.1
 authentication local rsa-sig
 authentication remote rsa-sig
 aaa authorization user cert list default default
 pki trustpoint TP
!
interface Tunnel0
 ip address 192.168.0.1 255.255.255.252
 tunnel protection ipsec profile default
```

What you need to specify

Packet Forwarding Simple Example

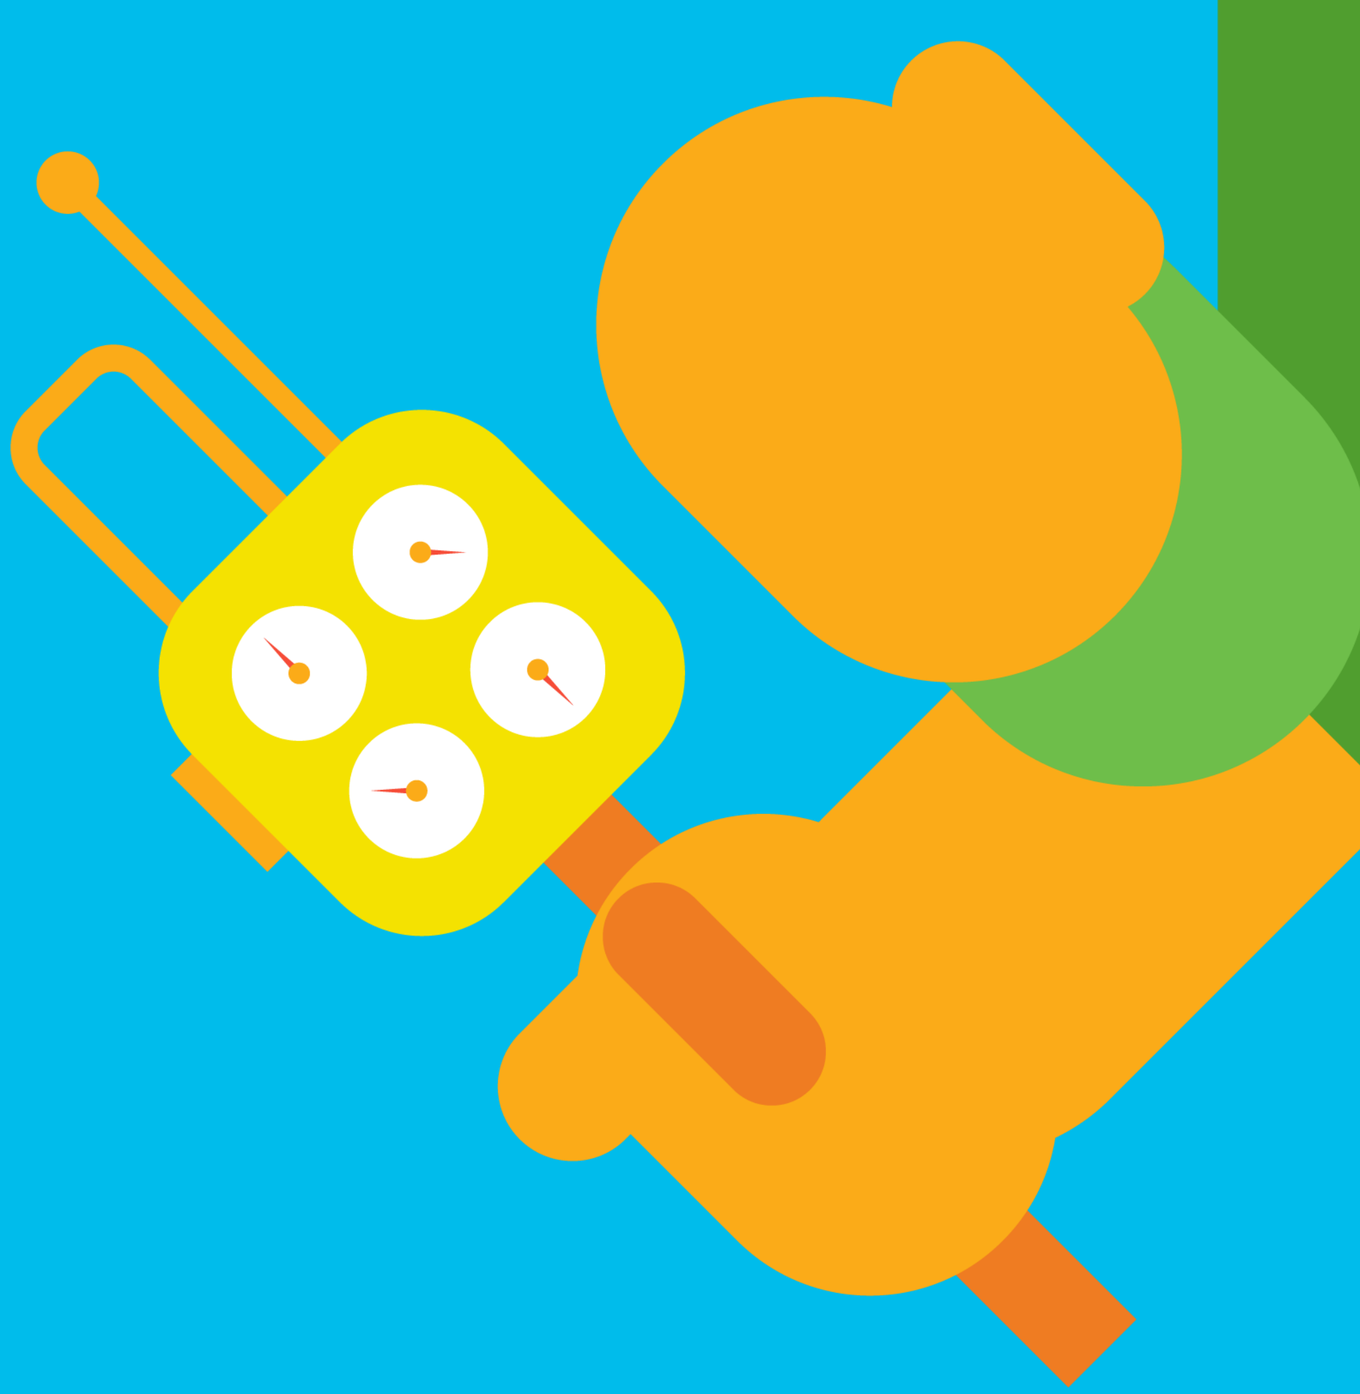

Basic Packet Forwarding

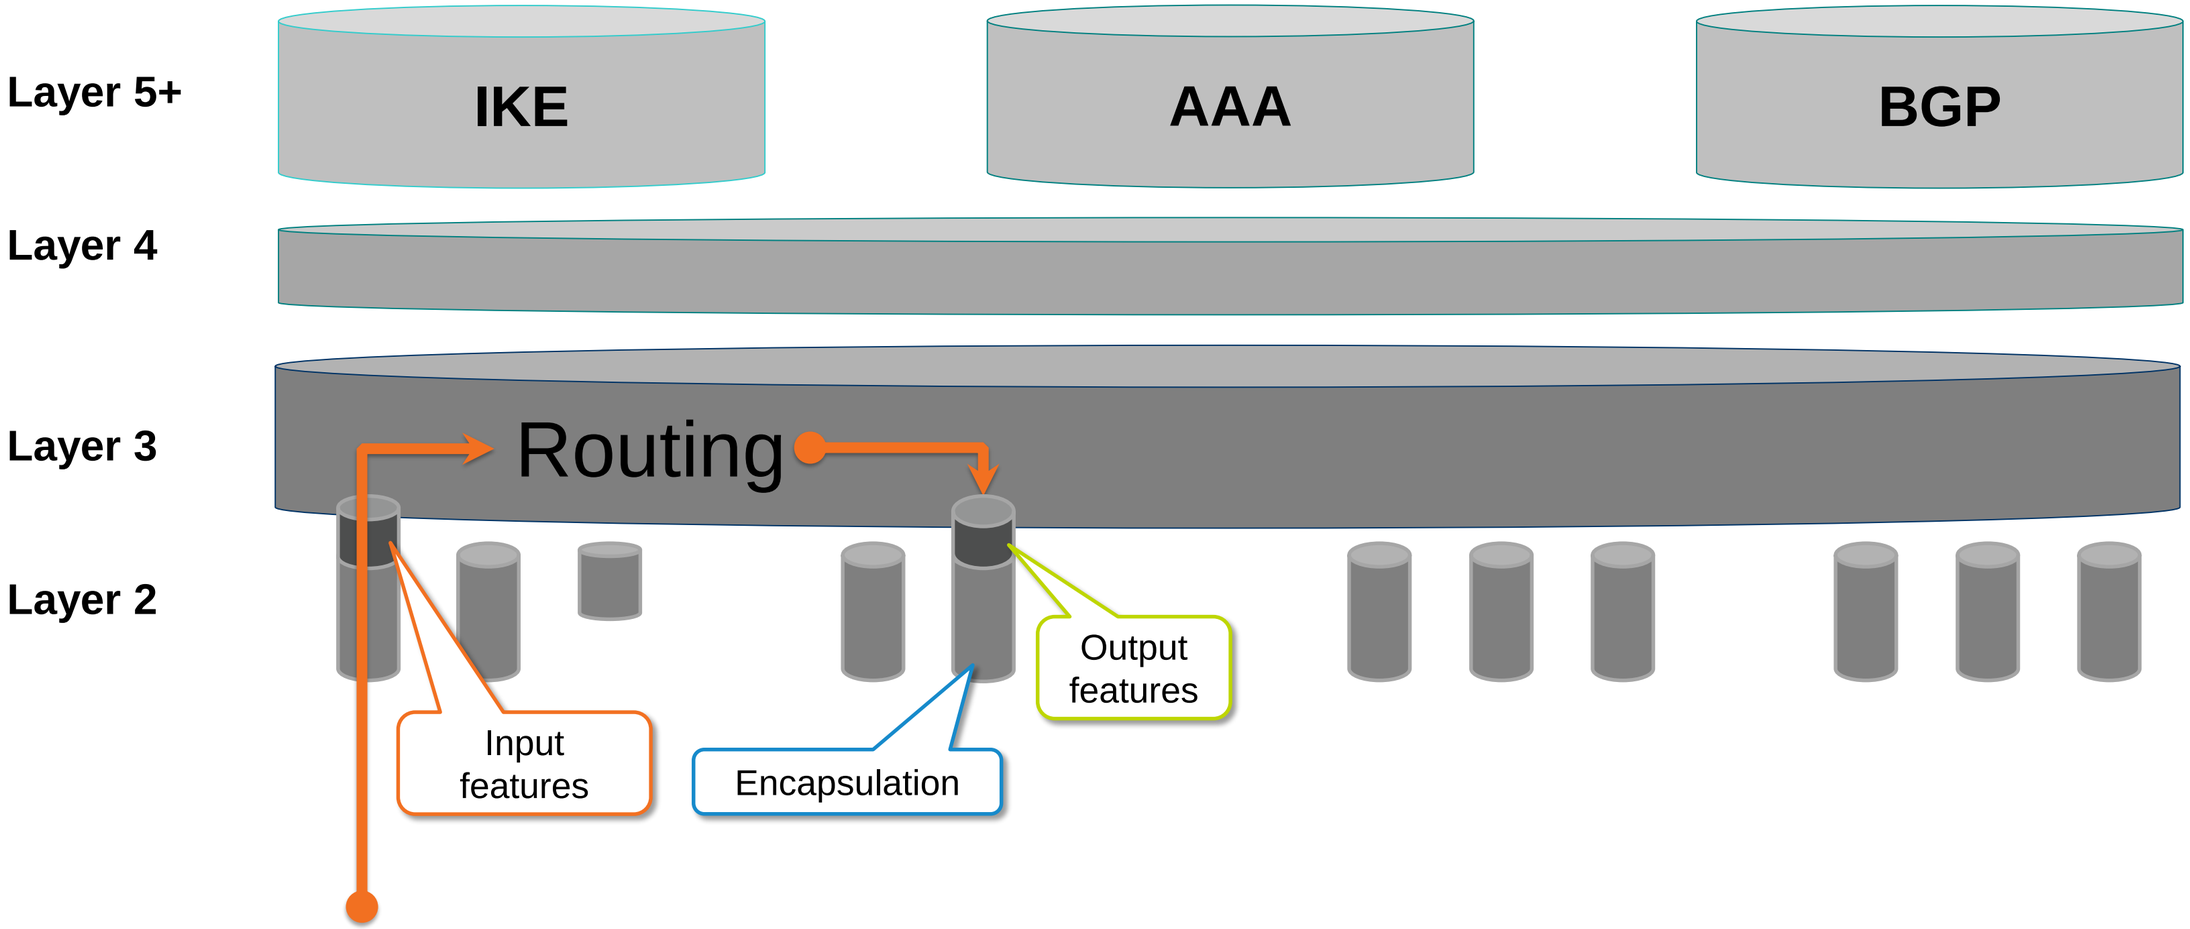

Packet Forwarding – Tunnels & Features

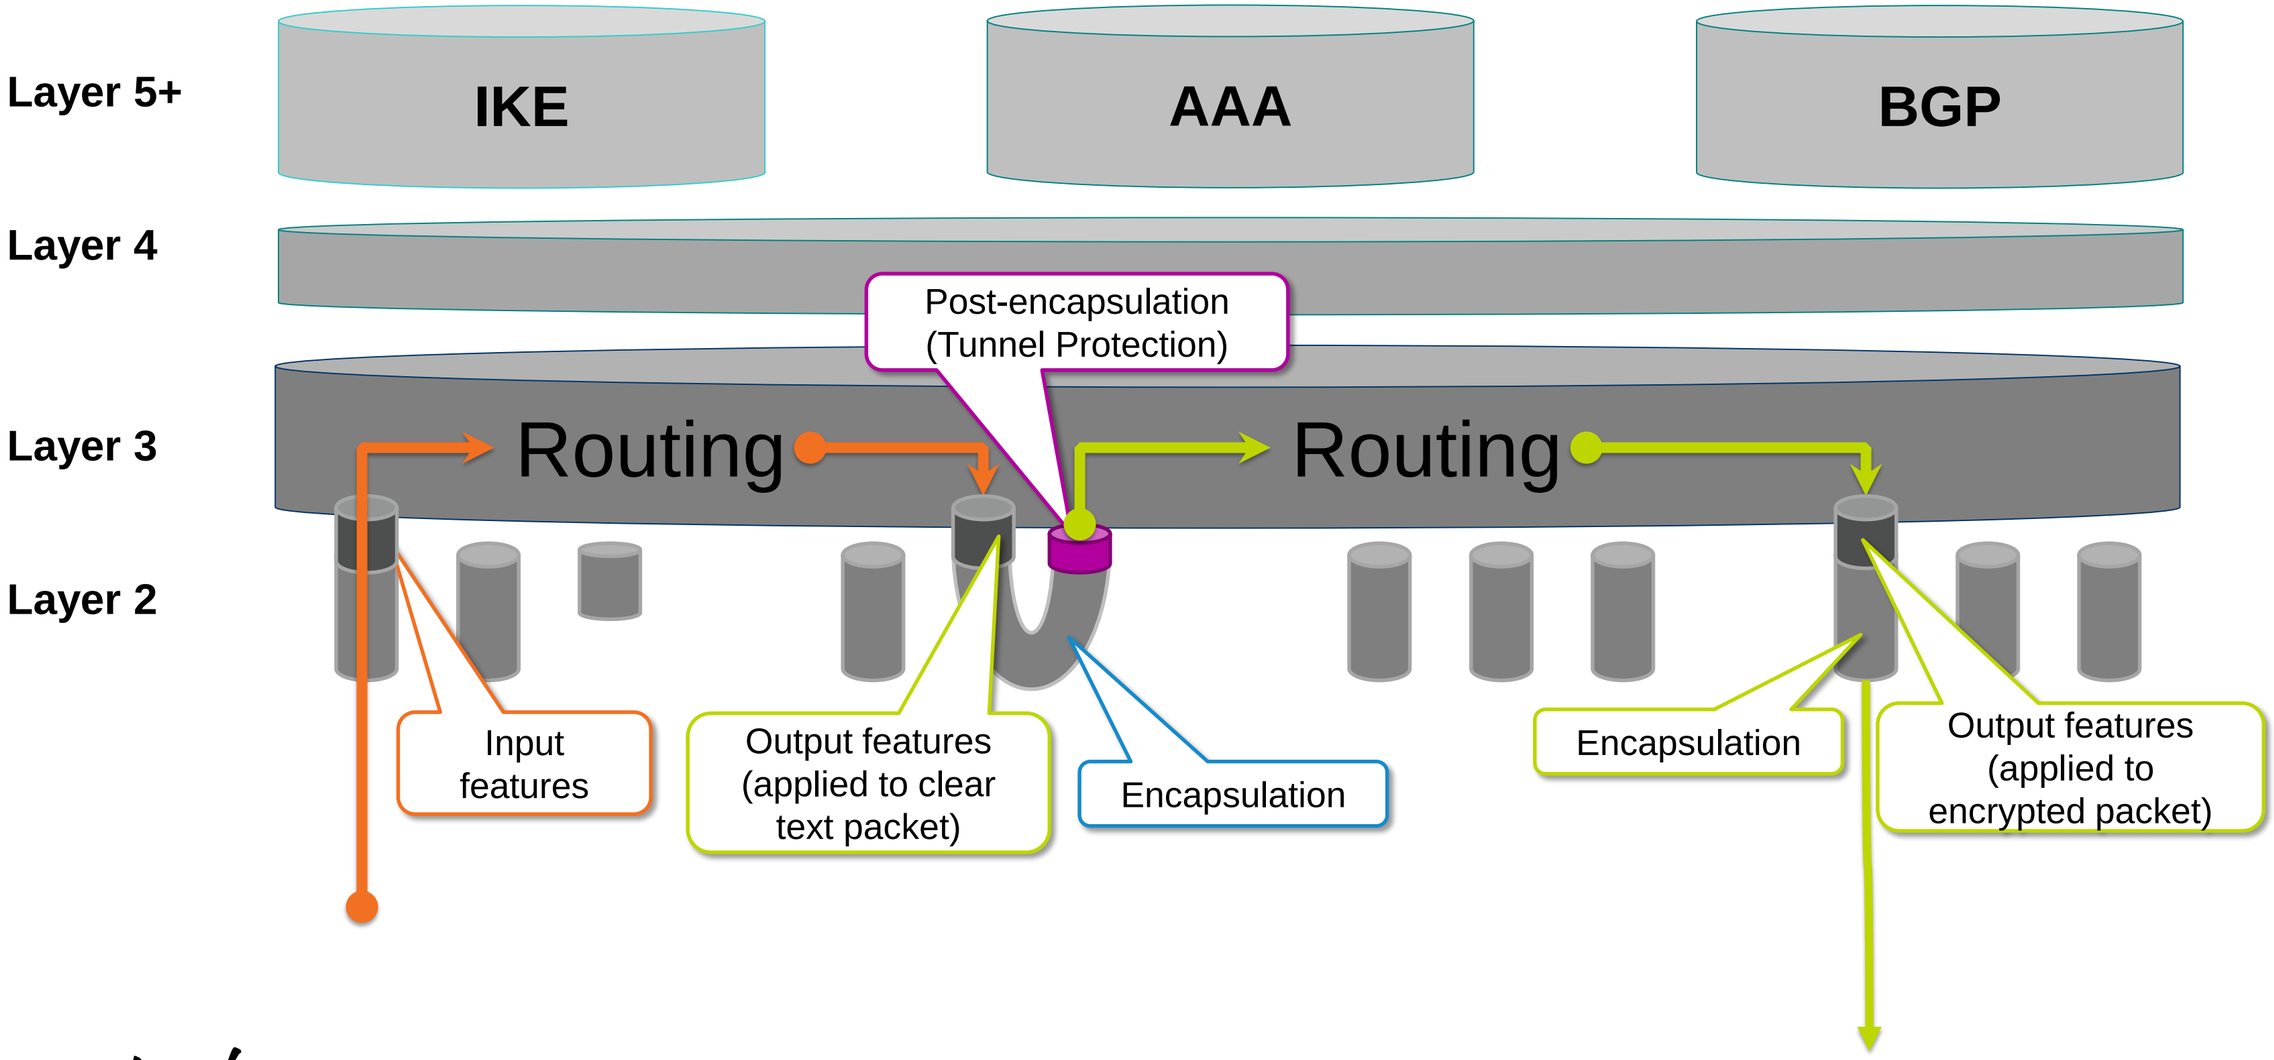

A Simple Site-to-Site Configuration

Example with IPv6 over IPv4 tunneling

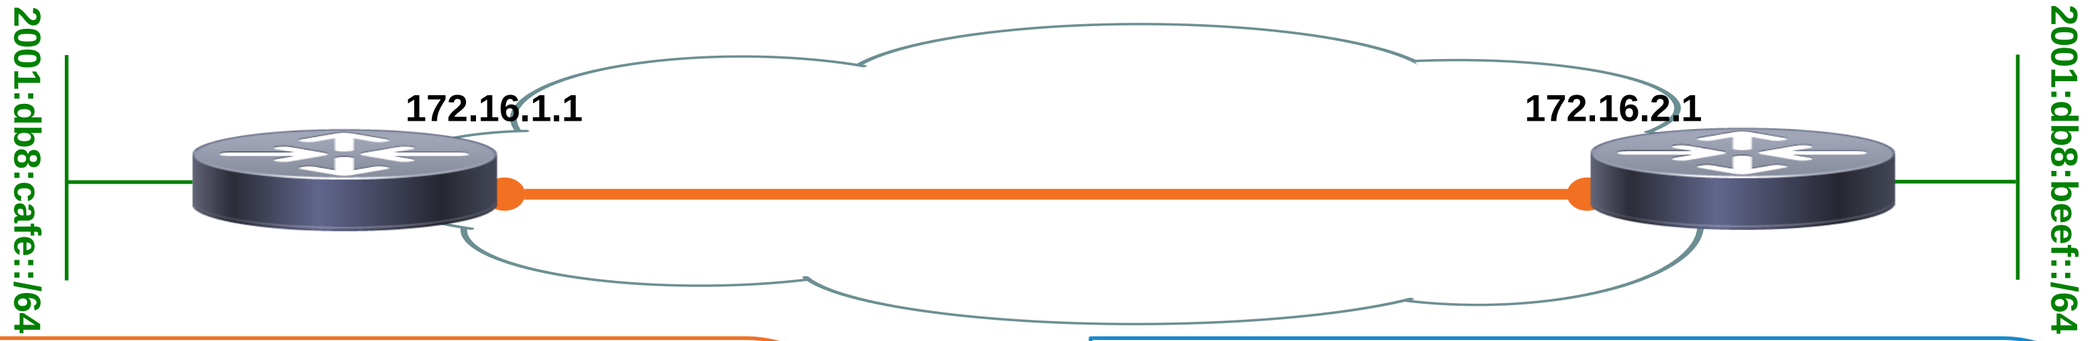


```
crypto ikev2 profile default
  match identity any
  authentication local pre-shared key cisco123
  authentication remote pre-shared key cisco123

ipv6 unicast-routing

interface Tunnel0
 ipv6 ospf 1 area 0
 tunnel source FastEthernet0/0
 tunnel destination 172.16.2.1
 tunnel protection ipsec profile default

interface E0/0
 ipv6 address 2001:db8:cafe::1/64
 ipv6 ospf 1 area 0
```

```
crypto ikev2 profile default
  match identity any
  authentication local pre-shared key cisco123
  authentication remote pre-shared key cisco123

ipv6 unicast-routing

interface Tunnel0
 ipv6 ospf 1 area 0
 tunnel source FastEthernet0/0
 tunnel destination 172.16.1.1
 tunnel protection ipsec profile default

interface E0/0
 ipv6 address 2001:db8:beef::1/64
 ipv6 ospf 1 area 0
```

Hub & Spoke and Shortcut Switching with IKEv2 Routing

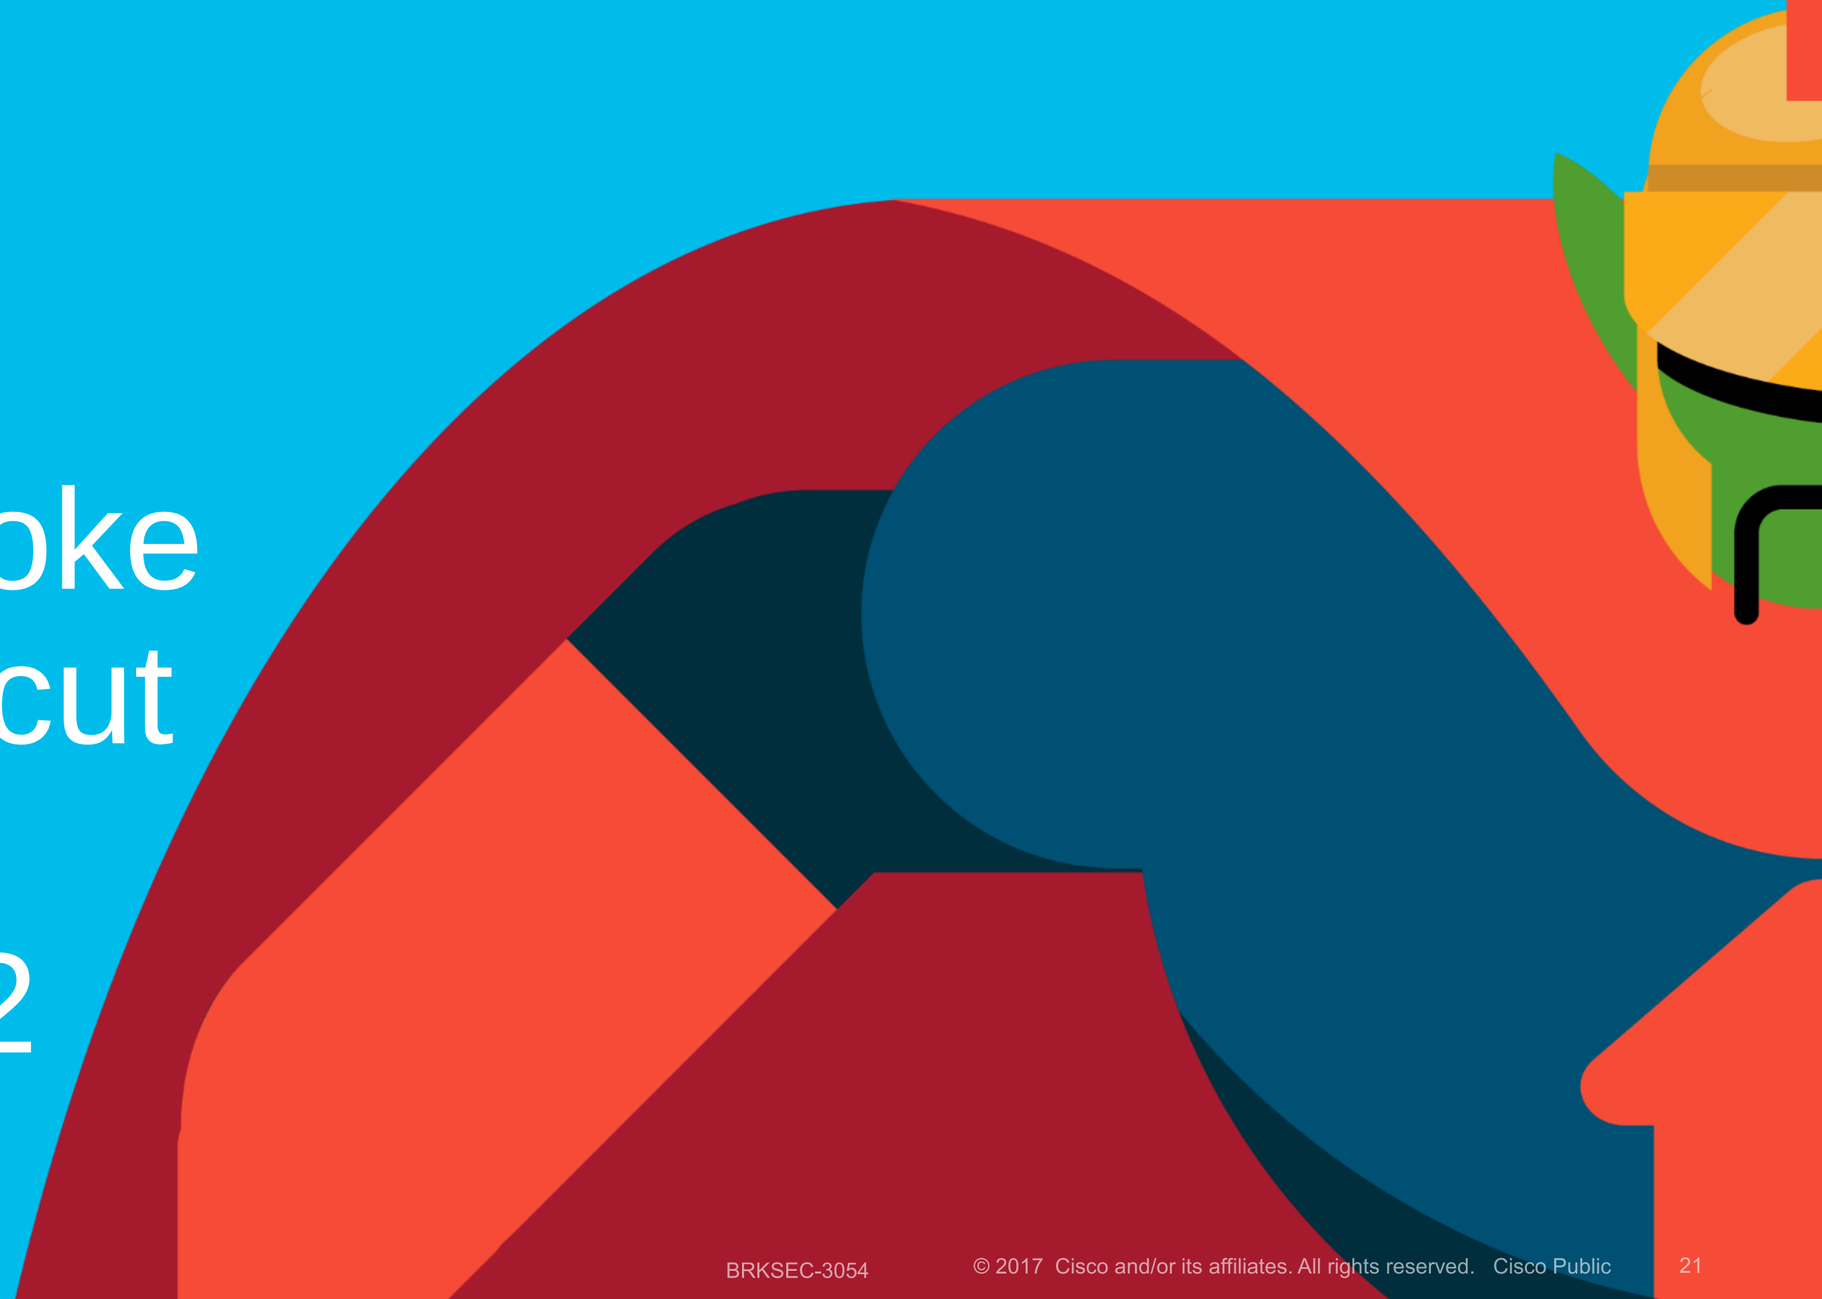

Cisco live!

FlexVPN Mesh – Scalable Network Diagram

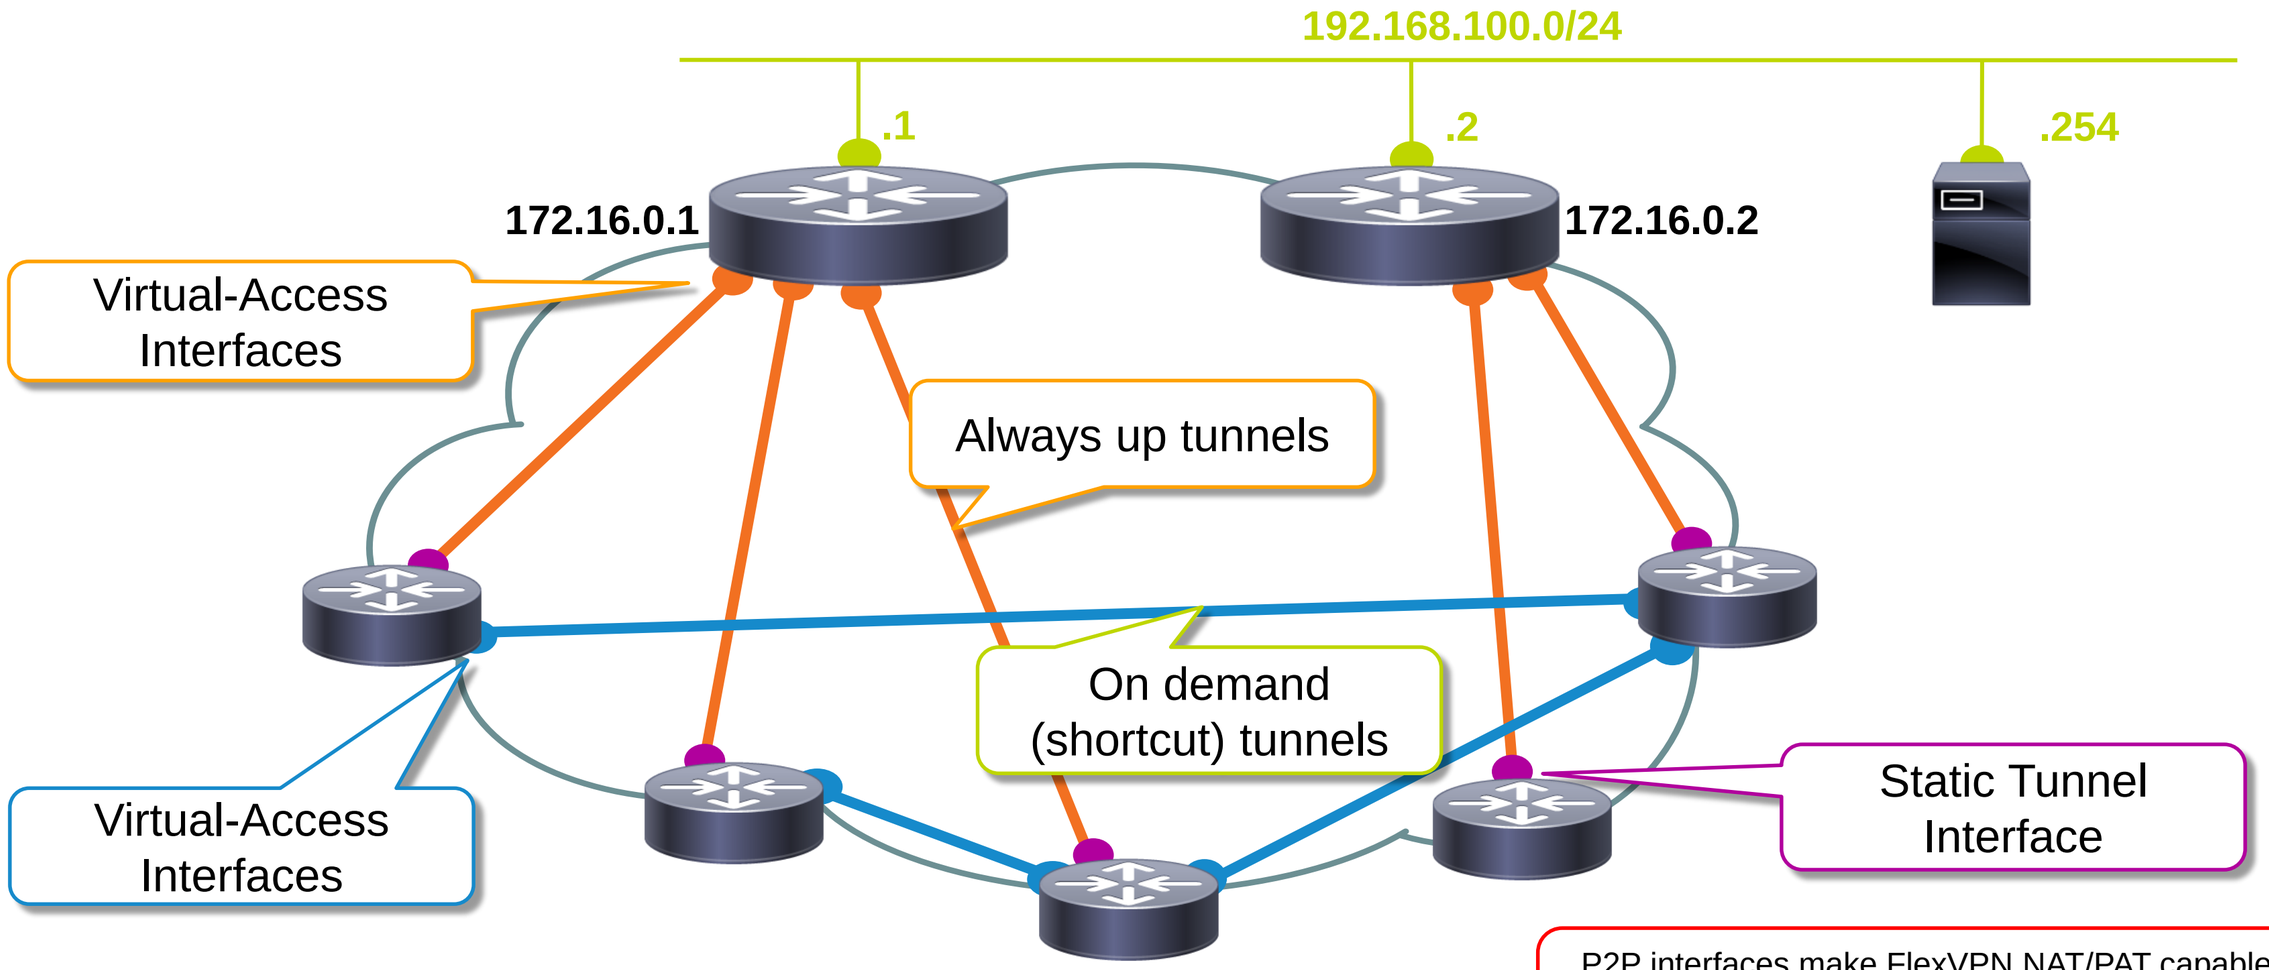

P2P interfaces make FlexVPN NAT/PAT capable
Hub&Spoke fully supported
Shortcuts may be blocked by PAT

Hub & Spoke Bootstrap – Config Exchange

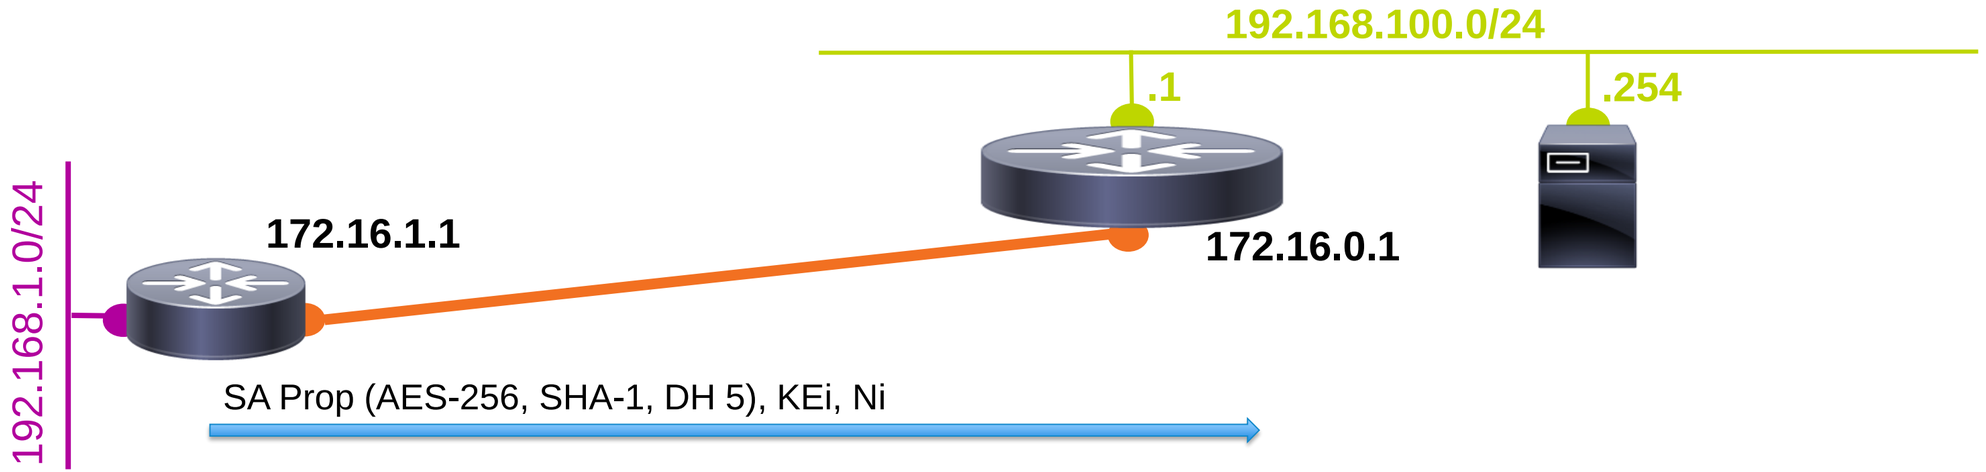

SA Prop (AES-256, SHA-1, DH 5), KEi, Ni

SA Prop (AES-256, SHA-1, DH 5), KEr, Nr

IDi=Spoke1.cisco.com, Auth, TSi, TSr,

CFG_Req(IP4_SUBNET…)

IDr, cert, Auth, TSi, TSr,

CFG_Reply(IP4_SUBNET=10.0.0.254/32, 192.168.0.0/16;
IP4_ADDRESS=10.0.0.1)

CFG_set(IP4_SUBNET=10.0.0.1/32, 192.168.1.0/24)

CFG_ack()

Interfaces

Ethernet0/0: 172.16.1.1
Ethernet0/1: 192.168.1.1
Tunnel0: **10.0.0.1**

Spoke Assigned Address (optional)

Routing Table

172.16.0.1/32 → 172.16.1.254 (E0/0)
192.168.1.0/24 → Ethernet 0/1
10.0.0.254/32 → Tunnel 0
192.168.0.0/16 → Tunnel 0

Supernet covering all spokes LAN prefixes

Interfaces

Ethernet0/0: 172.16.0.1
Ethernet0/1: 192.168.100.1
Loopback0: **10.0.0.254/32**
VirtualAccess1: 10.0.0.254/32

Routing Table

0.0.0.0/0 → 172.16.0.254 (E0/0)
192.168.100.0/24 → Ethernet 0/1
10.0.0.1/32 → VirtualAccess1
192.168.1.0/24 → VirtualAccess1

Virtual-Access (Tunnel) Instantiation

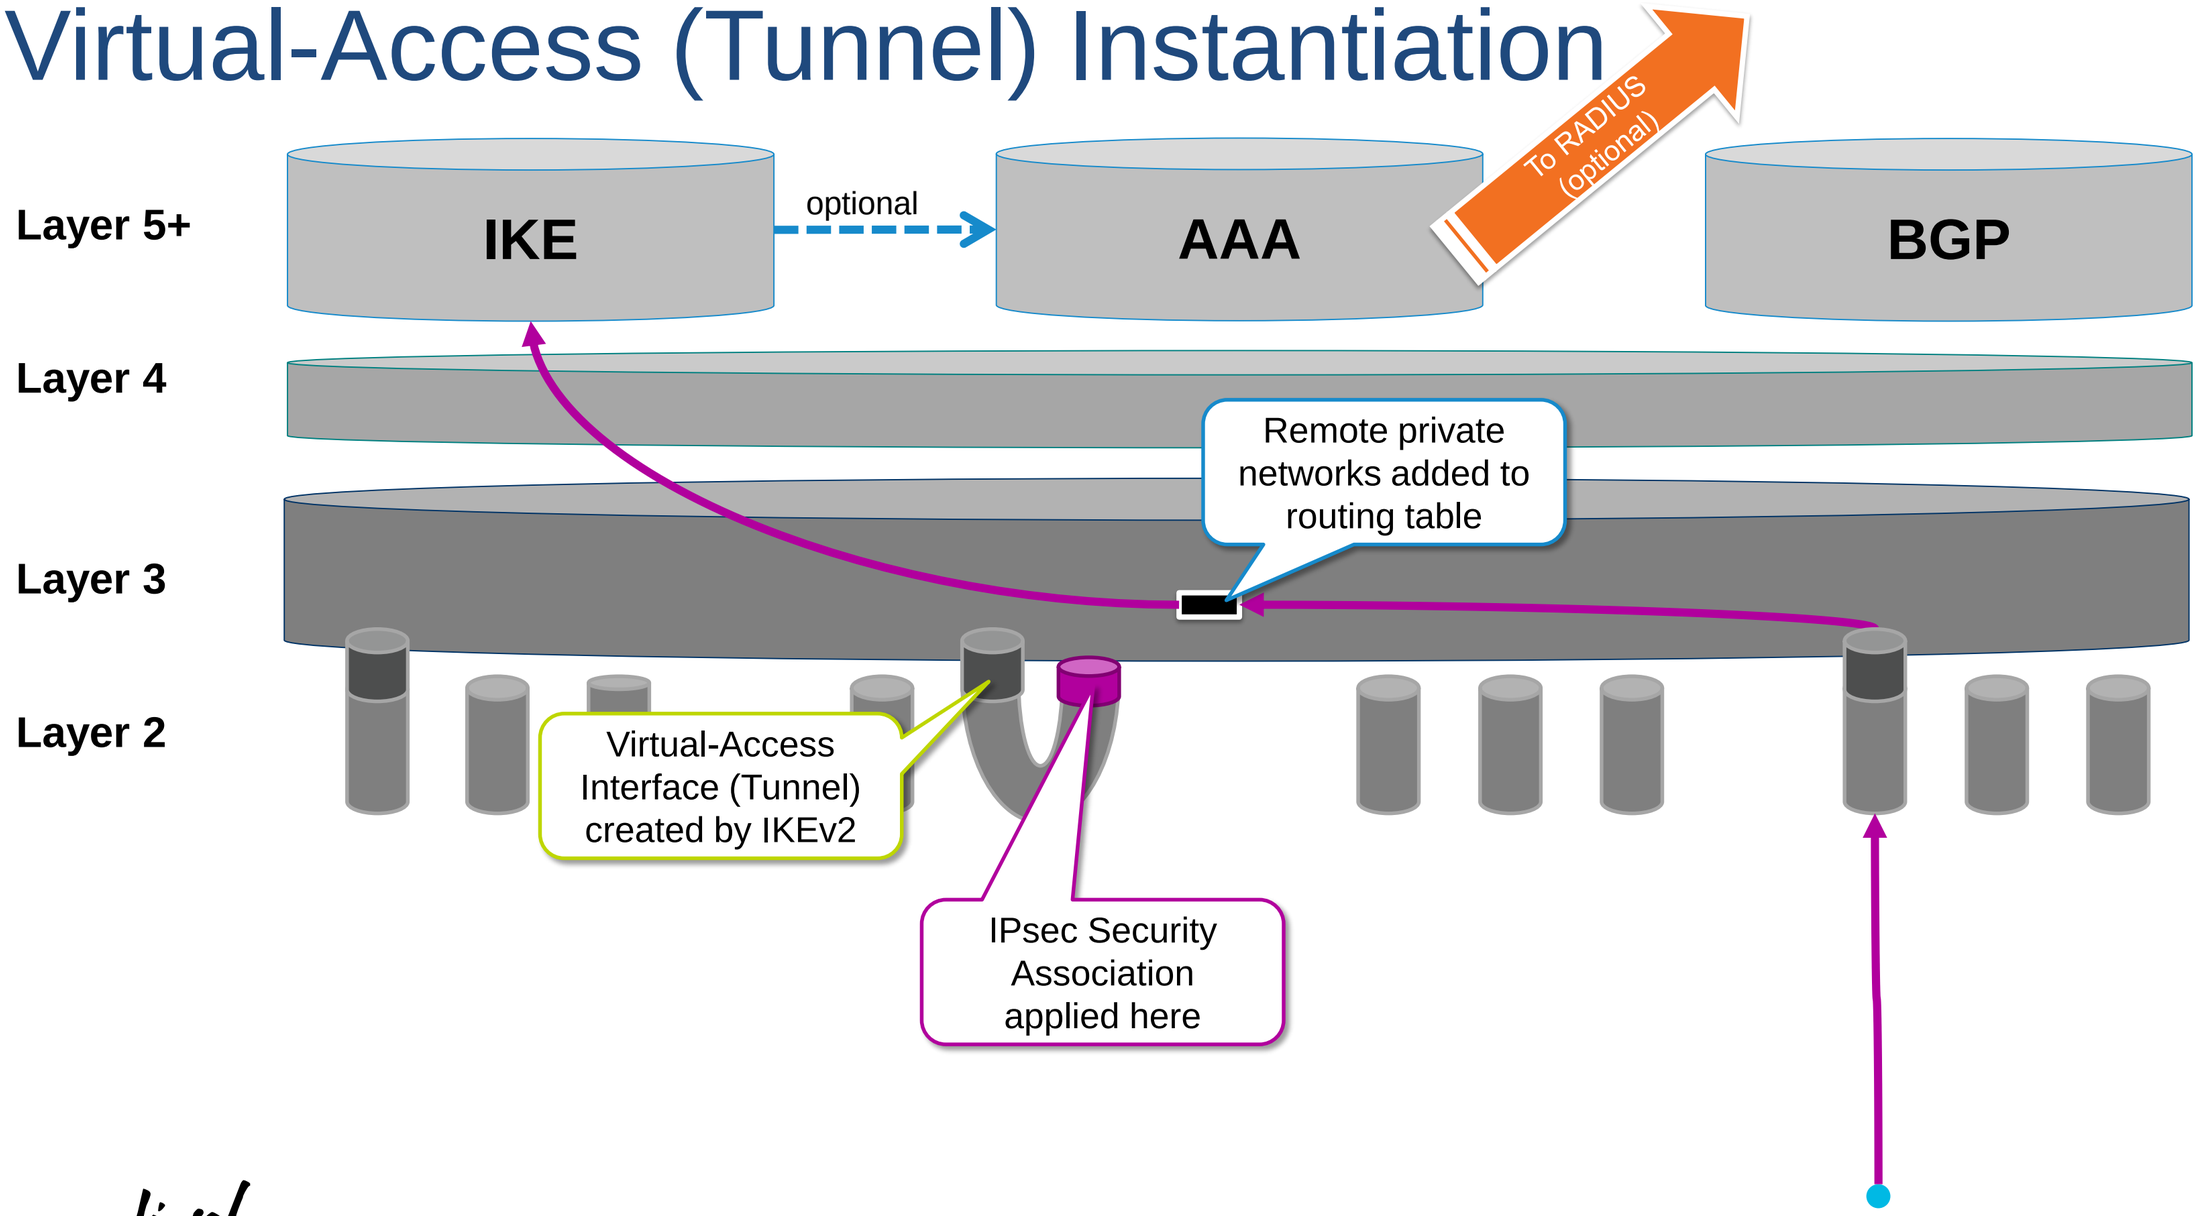

FlexVPN Hub and Spoke – IKE Route Exchange

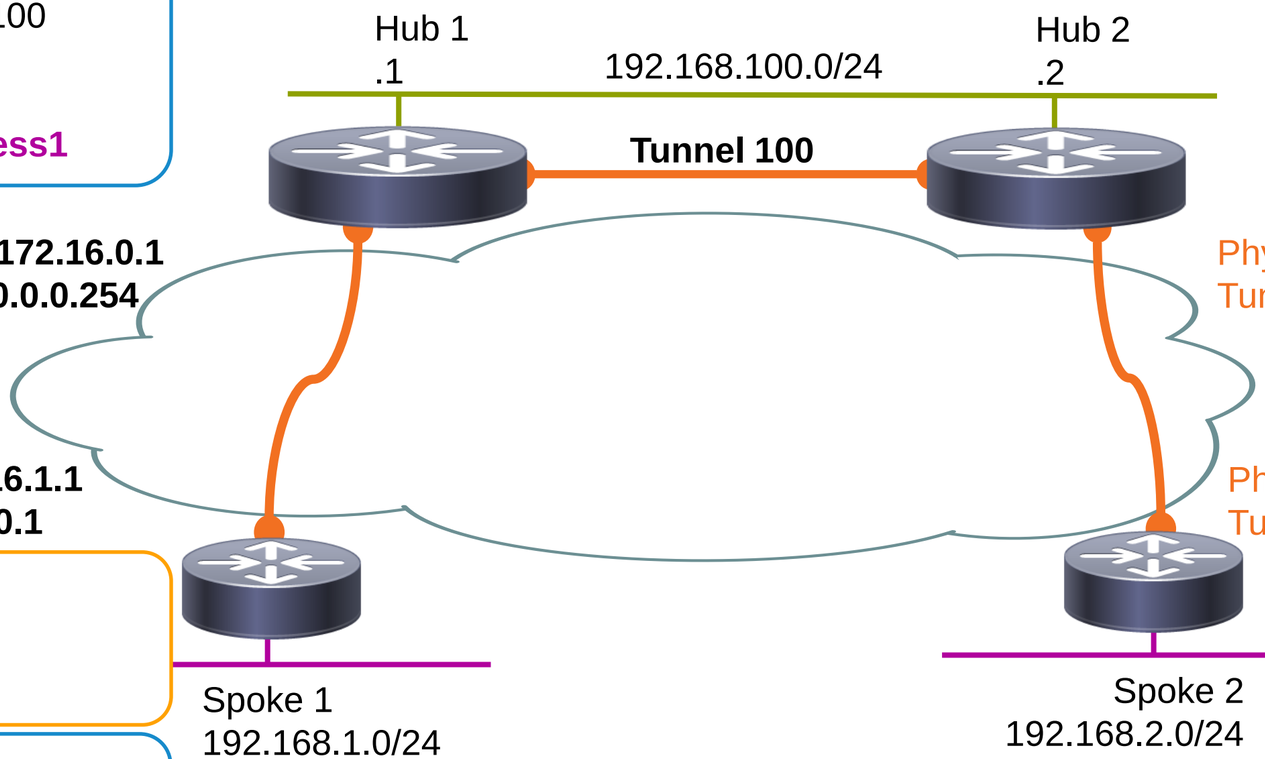

Routing Table

C 10.0.0.254 → Loopback0
C 192.168.100.0/24 → Eth0
S 192.168.0.0/16 → Tunnel100
S 10.0.0.0/8 → Tunnel100
S 10.0.0.1 → V-Access1
S 192.168.1.0/24 → V-Access1

Routing Table

C 10.0.0.253 → Loopback0
C 192.168.100.0/24 → Eth0
S 192.168.0.0/16 → Tunnel100
S 10.0.0.0/8 → Tunnel100
S 10.0.0.2 → V-Access1
S 192.168.2.0/24 → V-Access1

NHRP Table

-

NHRP Table

-

Routing Table

C 192.168.1.0/24 → Eth0
C 10.0.0.1 → Tunnel0
S 0.0.0.0/0 → Dialer0
S 10.0.0.254/32 → Tunnel0
S 192.168.0.0/16 → Tunnel0

Routing Table

C 192.168.2.0/24 → Eth0
C 10.0.0.2 → Tunnel1
S 0.0.0.0/0 → Dialer0
S 10.0.0.253/32 → Tunnel1
S 192.168.0.0/16 → Tunnel1

FlexVPN Mesh – Indirection

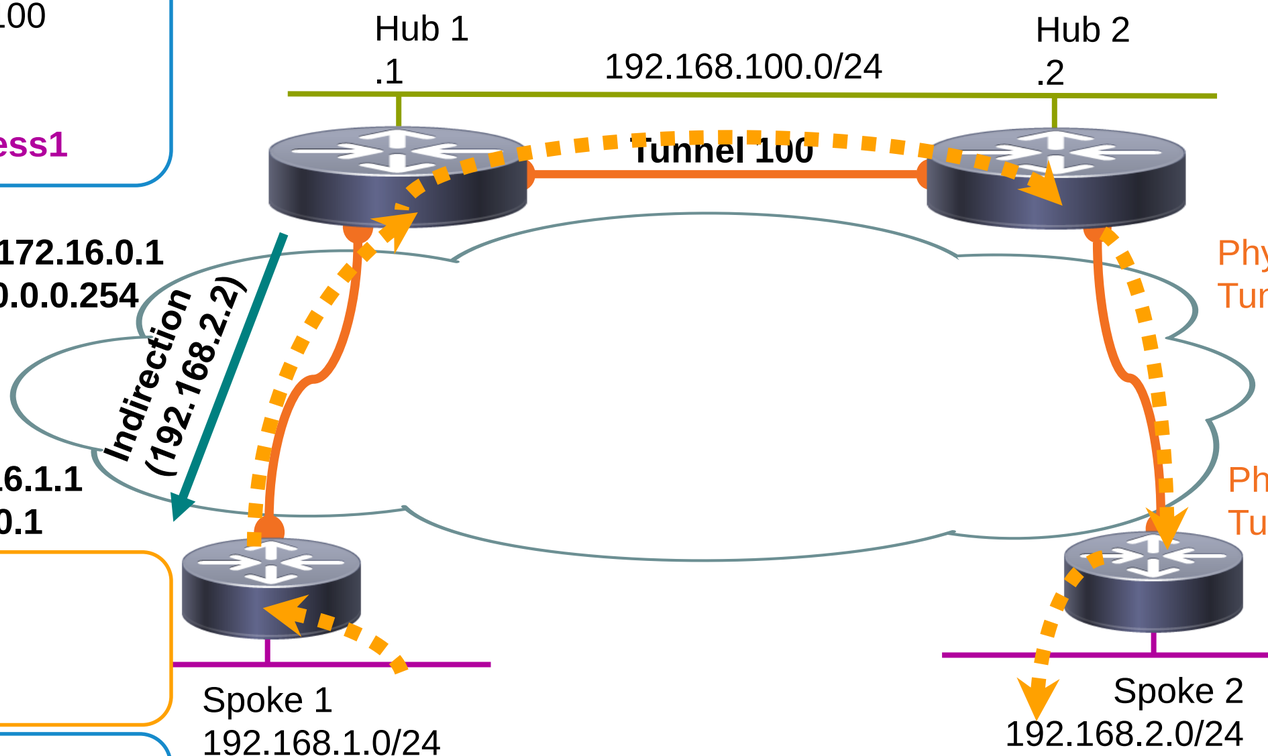

Routing Table

C 10.0.0.254 → Loopback0
C 192.168.100.0/24 → Eth0
S 192.168.0.0/16 → Tunnel100
S 10.0.0.0/8 → Tunnel100
S 10.0.0.1 → V-Access1
S 192.168.1.0/24 → V-Access1

Routing Table

C 10.0.0.253 → Loopback0
C 192.168.100.0/24 → Eth0
S 192.168.0.0/16 → Tunnel100
S 10.0.0.0/8 → Tunnel100
S 10.0.0.2 → V-Access1
S 192.168.2.0/24 → V-Access1

NHRP Table

-

NHRP Table

-

Routing Table

C 192.168.1.0/24 → Eth0
C 10.0.0.1 → Tunnel0
S 0.0.0.0/0 → Dialer0
S 10.0.0.254/32 → Tunnel0
S 192.168.0.0/16 → Tunnel0

Routing Table

C 192.168.2.0/24 → Eth0
C 10.0.0.2 → Tunnel1
S 0.0.0.0/0 → Dialer0
S 10.0.0.253/32 → Tunnel1
S 192.168.0.0/16 → Tunnel1

FlexVPN Mesh – Resolution

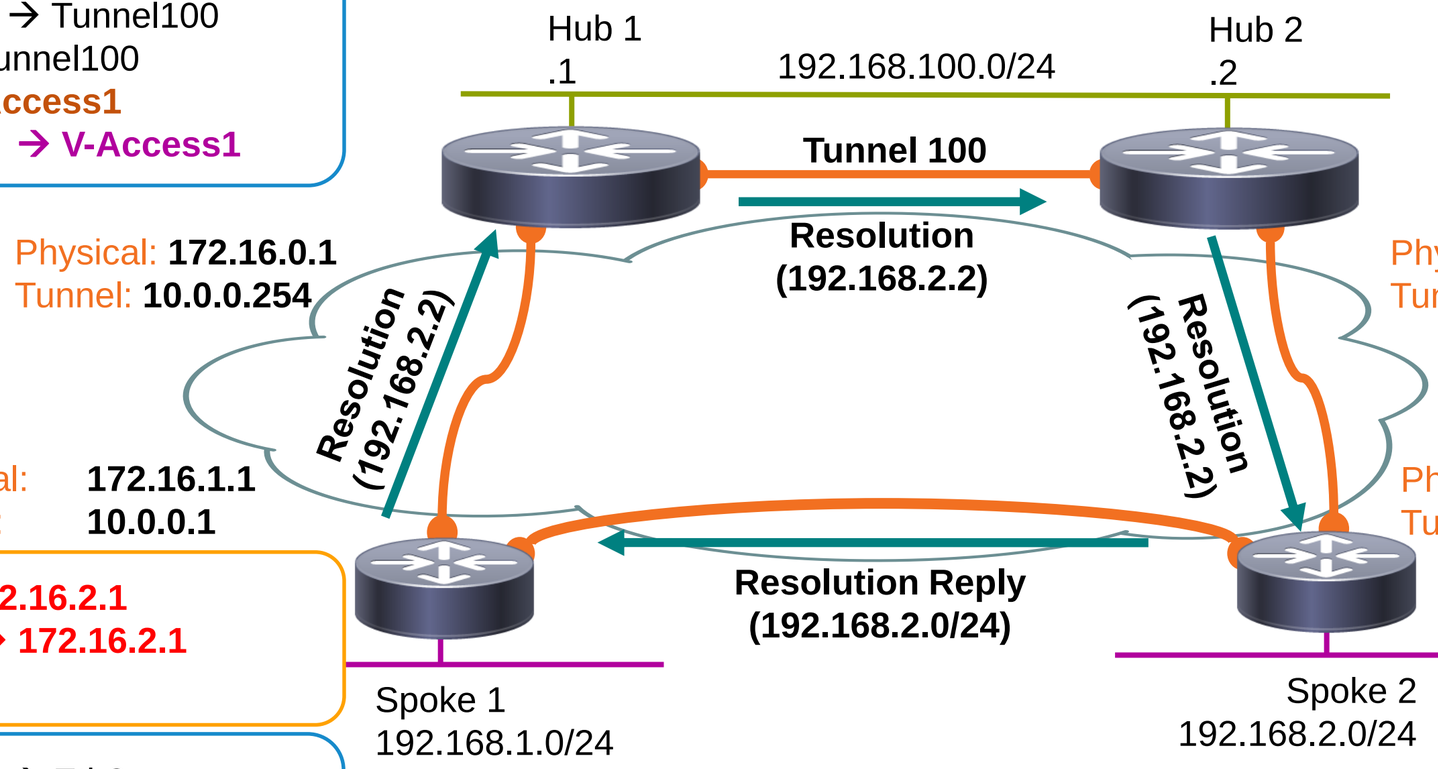

Routing Table

C 10.0.0.254 → Loopback0
C 192.168.100.0/24 → Eth0
S 192.168.0.0/16 → Tunnel100
S 10.0.0.0/8 → Tunnel100
S 10.0.0.1 → V-Access1
S 192.168.1.0/24 → V-Access1

Routing Table

C 10.0.0.253 → Loopback0
C 192.168.100.0/24 → Eth0
S 192.168.0.0/16 → Tunnel100
S 10.0.0.0/8 → Tunnel100
S 10.0.0.2 → V-Access1
S 192.168.2.0/24 → V-Access1

NHRP Table

10.0.0.2/32 → 172.16.2.1
192.168.2.0/24 → 172.16.2.1

NHRP Table

10.0.0.1 → 172.16.1.1

Routing Table

C 192.168.1.0/24 → Eth0
C 10.0.0.1 → Tunnel0
S 0.0.0.0/0 → Dialer0
S 10.0.0.254/32 → Tunnel0
S 192.168.0.0/16 → Tunnel0
H/S 10.0.0.2/32 → V-Access1
H/S 192.168.2.0/24 → V-Access1

Routing Table

C 192.168.2.0/24 → Eth0
C 10.0.0.2 → Tunnel1
S 0.0.0.0/0 → Dialer0
S 10.0.0.253/32 → Tunnel1
S 192.168.0.0/16 → Tunnel1
H/S 10.0.0.1/32 → V-Access1

FlexVPN Mesh – Shortcut Forwarding

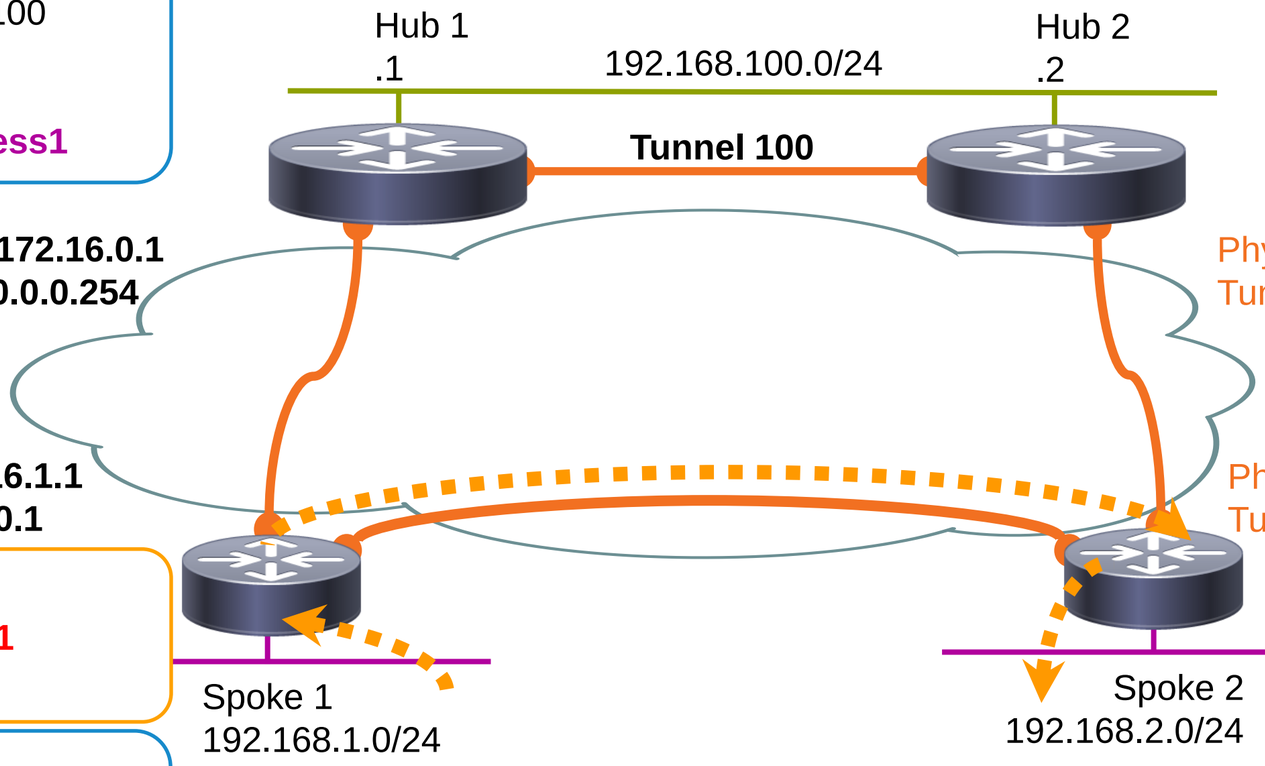

Routing Table

C 10.0.0.254 → Loopback0
C 192.168.100.0/24 → Eth0
S 192.168.0.0/16 → Tunnel100
S 10.0.0.0/8 → Tunnel100
S 10.0.0.1 → V-Access1
S 192.168.1.0/24 → V-Access1

Routing Table

C 10.0.0.253 → Loopback0
C 192.168.100.0/24 → Eth0
S 192.168.0.0/16 → Tunnel100
S 10.0.0.0/8 → Tunnel100
S 10.0.0.2 → V-Access1
S 192.168.2.0/24 → V-Access1

NHRP Table

10.0.0.2/32 → 172.16.2.1
192.168.2.0/24 → 172.16.2.1

NHRP Table

10.0.0.1 → 172.16.1.1

Routing Table

C 192.168.1.0/24 → Eth0
C 10.0.0.1 → Tunnel0
S 0.0.0.0/0 → Dialer0
S 10.0.0.254/32 → Tunnel0
S 192.168.0.0/16 → Tunnel0
H/S 10.0.0.2/32 → V-Access1
H/S 192.168.2.0/24 → V-Access1

Routing Table

C 192.168.2.0/24 → Eth0
C 10.0.0.2 → Tunnel1
S 0.0.0.0/0 → Dialer0
S 10.0.0.253/32 → Tunnel1
S 192.168.0.0/16 → Tunnel1
H/S 10.0.0.1/32 → V-Access1

FlexVPN Mesh (IKEv2 Routing) Hub 1 Configuration

```
crypto ikev2 profile default
 match identity remote fqdn domain cisco.com
 identity local fqdn Hub1.cisco.com
 authentication remote rsa-sig
 authentication local rsa-sig
 pki trustpoint TP
 dpd 10 2 on-demand
 aaa authorization group cert list default default
 virtual-template 1
 !
crypto ikev2 authorization policy default
 route set remote 10.0.0.0 255.0.0.0
 route set remote 192.168.0.0 255.255.0.0
```

Accept connections from Spokes

Local spoke profile

These prefixes can also be set by RADIUS

Defines which prefixes should be protected

```
interface Virtual-Template1 type tunnel
 ip unnumbered Loopback0
 ip nhrp network-id 1
 ip nhrp redirect
 ip access-group AllowMyBGP in
 tunnel protection ipsec profile default
!
interface Loopback0
 ip address 10.0.0.254 255.255.255.255
!
interface Tunnel100
 ip unnumbered Loopback0
 ip nhrp network-id 1
 ip nhrp redirect
 tunnel source Ethernet0/1
 tunnel destination 192.168.100.2
```

Static per-spoke features applied here

All V-Access will be in the same network-id

Hub 1 dedicated overlay address

Inter-Hub link (not encrypted)

Same network-id on V-Access and inter-hub link

FlexVPN Mesh (IKEv2 Routing) Spoke Configuration

```
crypto ikev2 profile default
 match identity remote fqdn domain cisco.com
 identity local fqdn Spoke2.cisco.com
 authentication remote rsa-sig
 authentication local rsa-sig
 pki trustpoint TP
 dpd 10 2 on-demand
 aaa authorization group cert list default default
 virtual-template 1
```

Needed for address and prefix exchange

```
crypto ikev2 authorization policy default
 route set interface
 route set interface e0/0
```

Send tunnel address and private lan address. “route set remote” can also be used.

V-Template to clone for spoke-spoke tunnels

```
interface Loopback0
 ip address 10.0.0.2 255.255.255.255

interface Tunnel0
 ip unnumbered Loopback0
 ip nhrp network-id 1
 ip nhrp shortcut virtual-template 1
 tunnel source Ethernet0/0
 tunnel destination 172.16.0.1
 tunnel protection ipsec profile default
!
interface Tunnel1
 ip unnumbered Loopback0
 ip nhrp network-id 1
 ip nhrp shortcut virtual-template 1
 tunnel source Ethernet0/0
 tunnel destination 172.16.0.2
 tunnel protection ipsec profile default

interface Virtual-Template1 type tunnel
 ip unnumbered Loopback0
 ip nhrp network-id 1
 ip nhrp shortcut virtual-template 1
 tunnel protection ipsec profile default
```

Tunnel to Hub 1

Tunnel1 to Hub 2

QoS can be applied here

Shortcut Switching With a routing protocol (BGP)

FlexVPN Mesh with BGP Routing

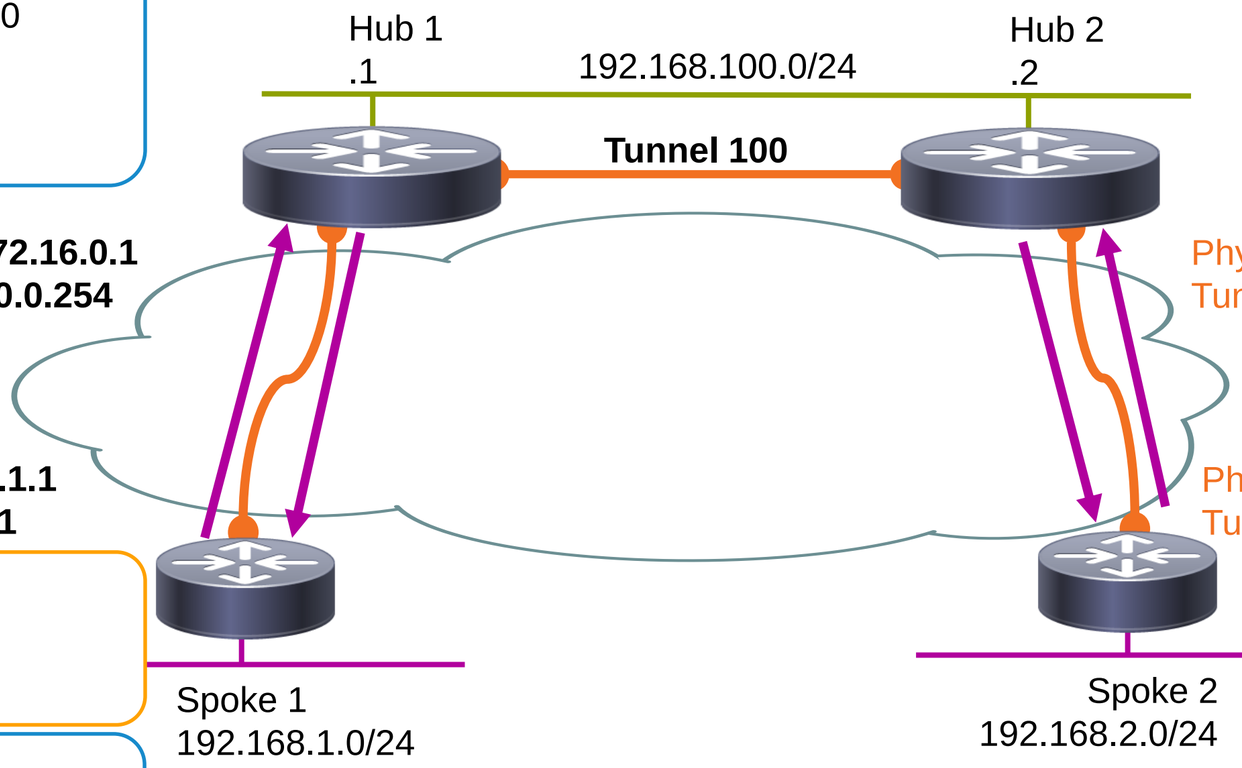

Routing Table

C 10.0.0.254 → Loopback0
C 192.168.100.0/24 → Eth0
S 192.168.0.0/16 → Tunnel100
S 10.0.0.0/8 → Tunnel100
S 10.0.0.1 → V-Access1
B 192.168.1.0/24 → 10.0.0.1

Routing Table

C 10.0.0.253 → Loopback0
C 192.168.100.0/24 → Eth0
S 192.168.0.0/16 → Tunnel100
S 10.0.0.0/8 → Tunnel100
S 10.0.0.2 → V-Access1
B 192.168.2.0/24 → 10.0.0.2

NHRP Table

-

NHRP Table

-

Routing Table

C 192.168.1.0/24 → Eth0
C 10.0.0.1 → Tunnel0
S 0.0.0.0/0 → Dialer0
S 10.0.0.254/32 → Tunnel0
B 192.168.0.0/16 → 10.0.0.254

Routing Table

C 192.168.2.0/24 → Eth0
C 10.0.0.2 → Tunnel1
S 0.0.0.0/0 → Dialer0
S 10.0.0.253/32 → Tunnel1
B 192.168.0.0/16 → 10.0.0.253

BGP complex ? Not really…

Spoke Configuration

```
router bgp 1
 bgp log-neighbor-changes
 neighbor 10.0.0.253 remote-as 1
 neighbor 10.0.0.254 remote-as 1

 address-family ipv4
  network 192.168.2.0
  neighbor 10.0.0.253 activate
  neighbor 10.0.0.254 activate
  maximum-paths ibgp 2
```

Spoke prefix to advertise

Any other routing protocol will do but mind the scalability and resiliency against packet losses.

All protocols were not created equal.

BGP shown in this presentation as a “no brainer”.

Hub Configuration

Summary prefixes to advertise to all spokes

```
ip route 10.0.0.0 255.0.0.0 Tunnel100 tag 2
ip route 192.168.0.0 255.255.0.0 Tunnel100 tag 2

router bgp 1
 bgp log-neighbor-changes
 bgp listen range 10.0.0.0/24 peer-group Flex

 address-family ipv4
  neighbor Flex peer-group
  neighbor Flex remote-as 1
  redistribute static route-map rm
 exit-address-family
!
route-map rm permit 10
 match tag 2
```

Dynamically accept spoke BGP peering!

route-map filters static routes to redistribute in BGP

Routing IPv6 with IKEv2 or BGP

With IKEv2 routing

```
crypto ikev2 profile default
 match identity remote fqdn domain cisco.com
 identity local fqdn Hub1.cisco.com
 authentication remote rsa-sig
 authentication local rsa-sig
 pki trustpoint TP
 dpd 10 2 on-demand
 …
 virtual-template 1
 !
crypto ikev2 authorization policy default
 route set remote ipv6 2001::/64
 route set remote ipv6 2002::/64
```

Same as v4… just specify ipv6 ☺

With BGP

One peering, for both IPv4 and IPv6

```
router bgp 1
 bgp log-neighbor-changes
 bgp listen range 10.0.0.0/16 peer-group Flex
 neighbor Flex peer-group
 neighbor Flex remote-as 1
 !
 address-family ipv4
  redistribute static route-map rm
  neighbor Flex activate
 exit-address-family
 !
 address-family ipv6
  redistribute static route-map rm
  neighbor Flex activate
 exit-address-family
```

Spokes Zero Touch Provisioning

SUDI Certificates

Installed at factory in ACT-2 chip (when applicable)

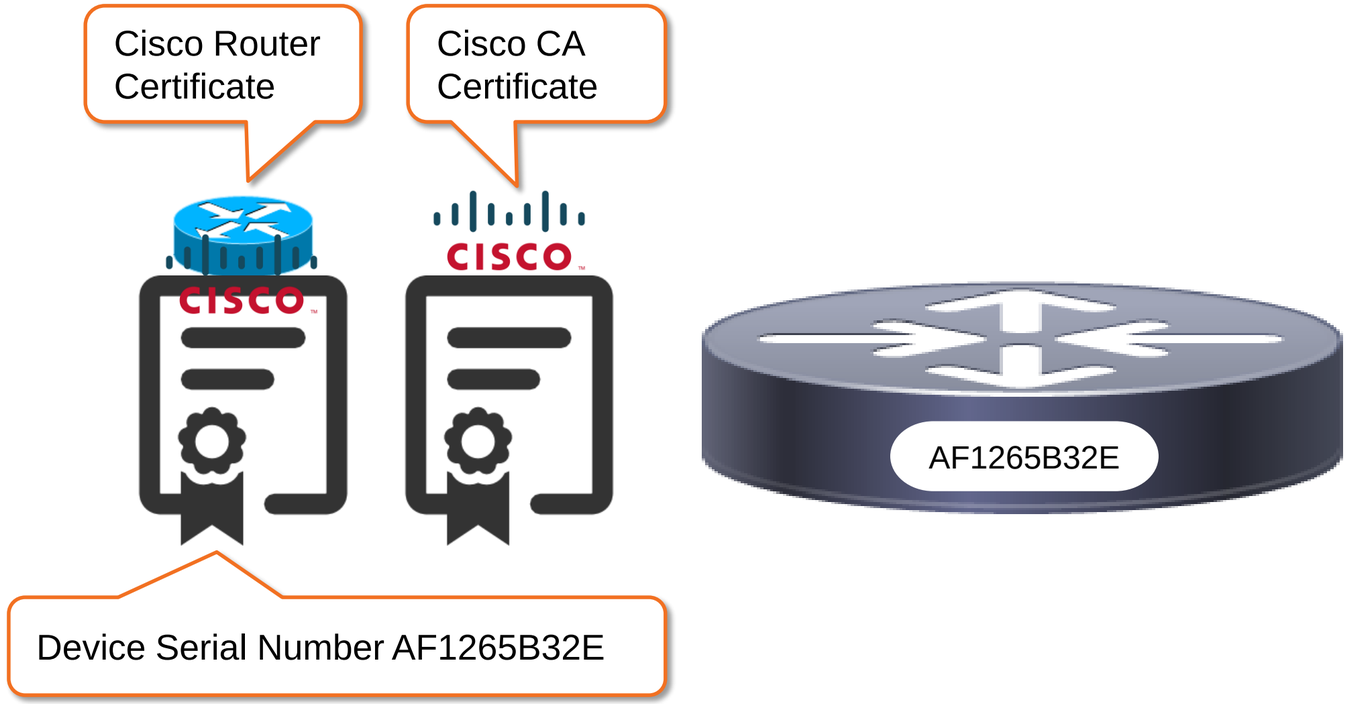

Zero Touch using IOS PKI Server

Initial State

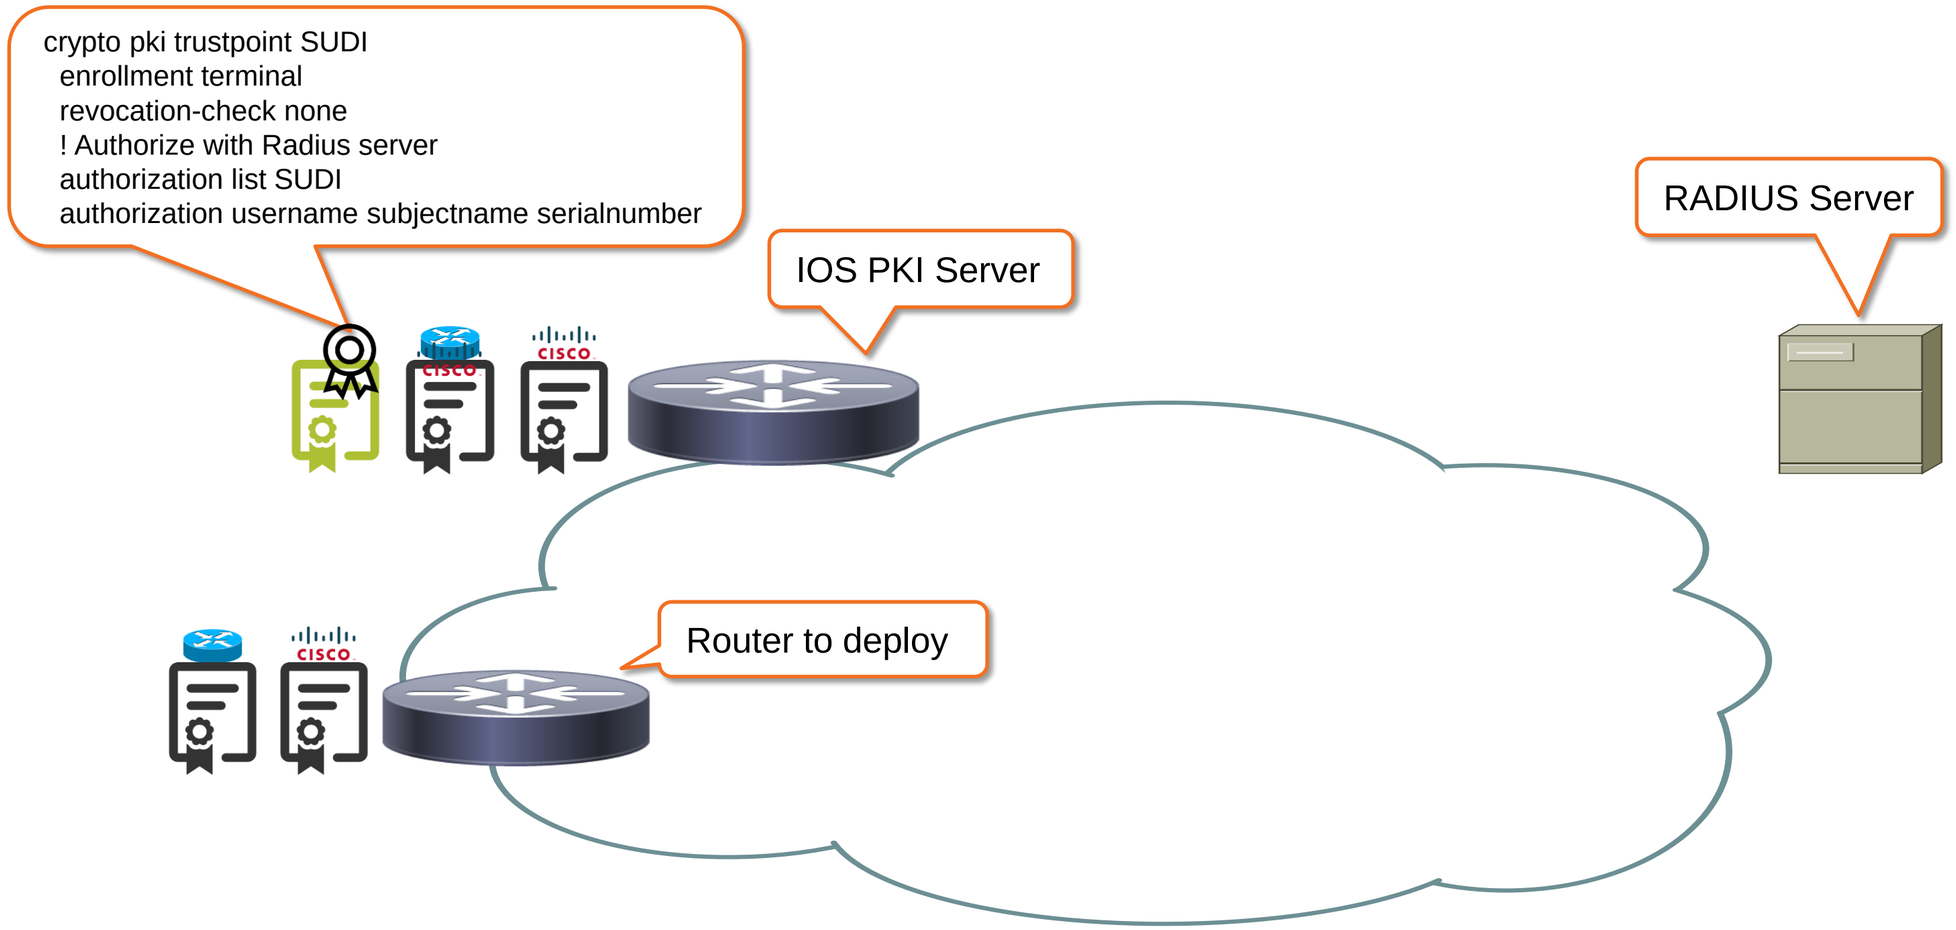

Zero Touch using IOS PKI Server

USB to bootstrap branch

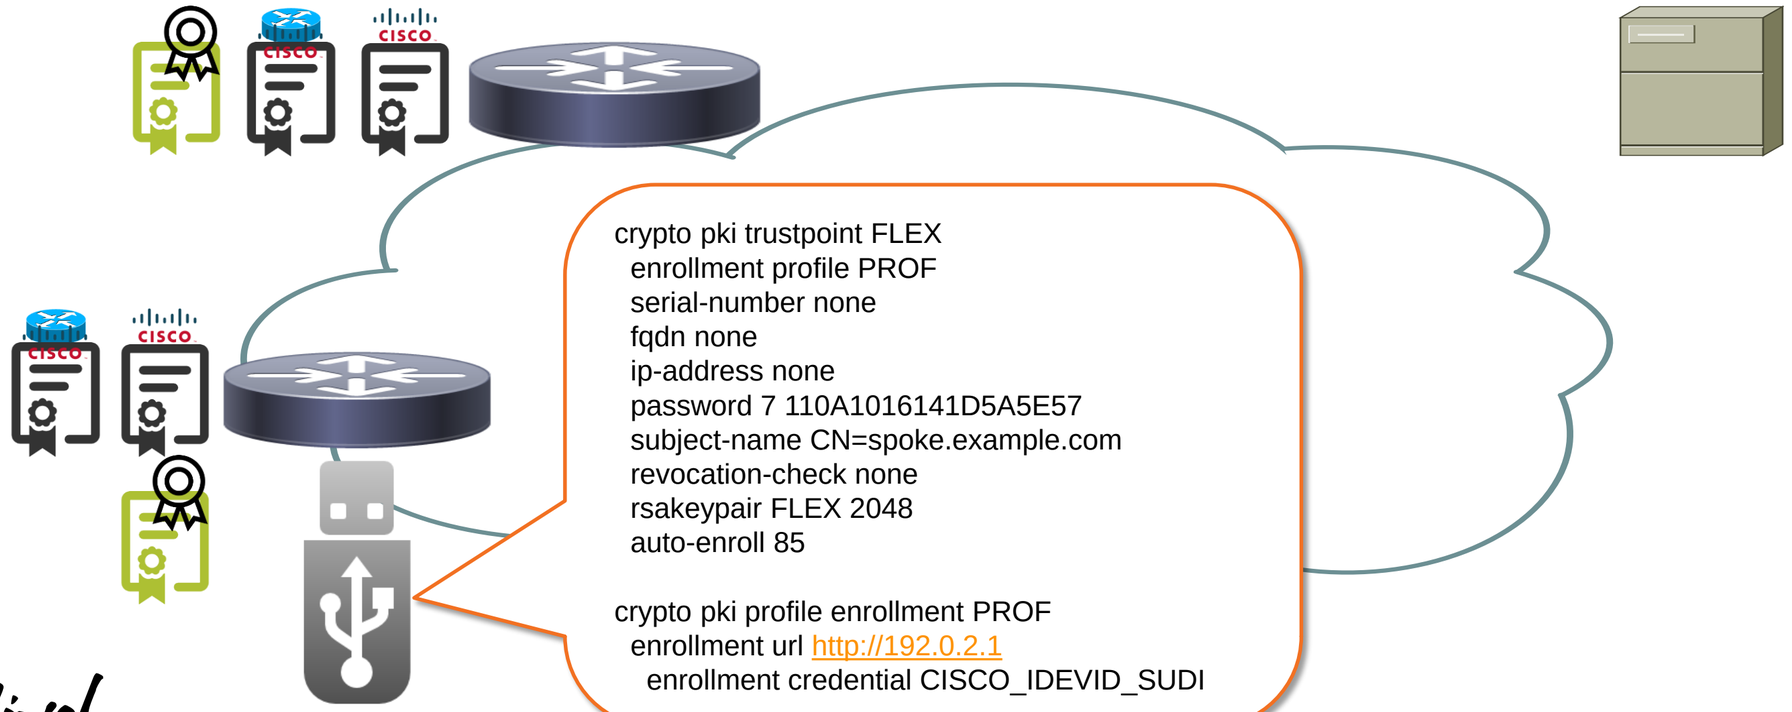

Zero Touch using IOS PKI Server

USB to bootstrap branch

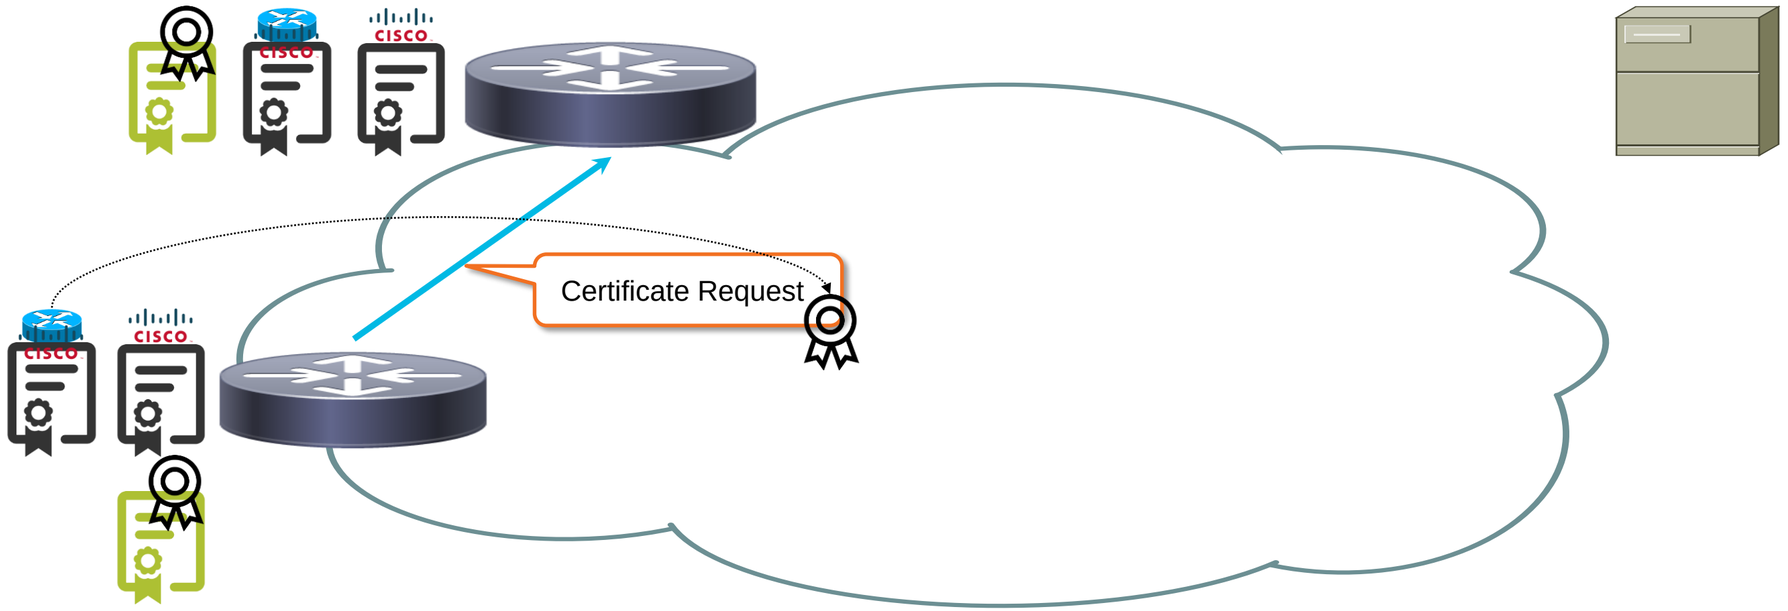

Zero Touch using IOS PKI Server

Authorizing Certificate Grant

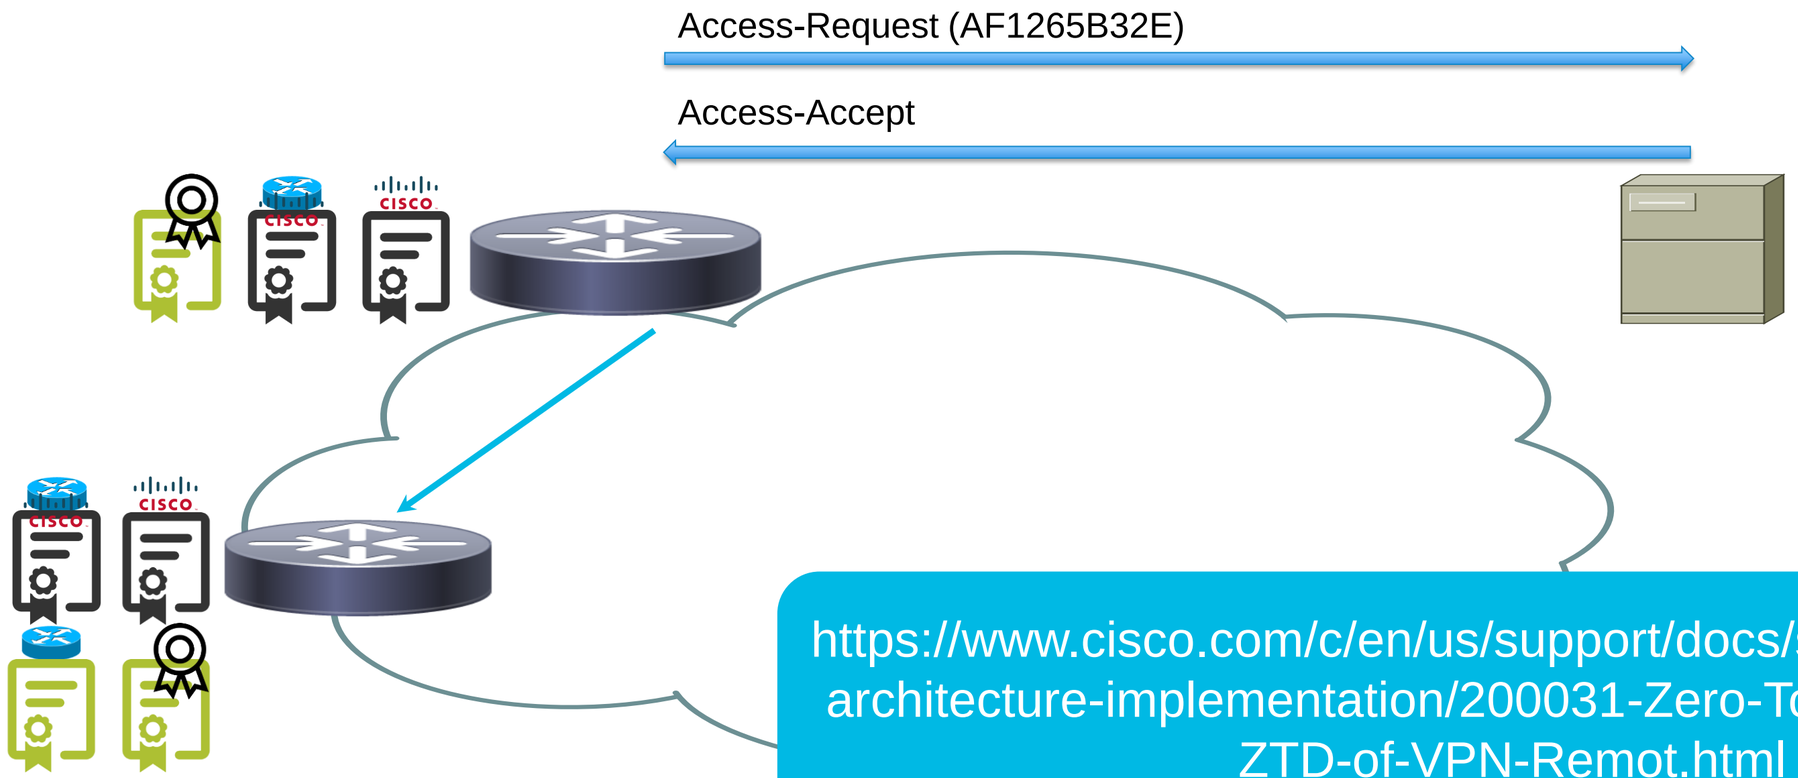

https://www.cisco.com/c/en/us/support/docs/security-vpn/ipsec-architecture-implementation/200031-Zero-Touch-Deployment-ZTD-of-VPN-Remot.html

MPLS over FlexVPN with Shortcut Switching

MPLS VPN o Flex

Objective: end-to-end VRF separation

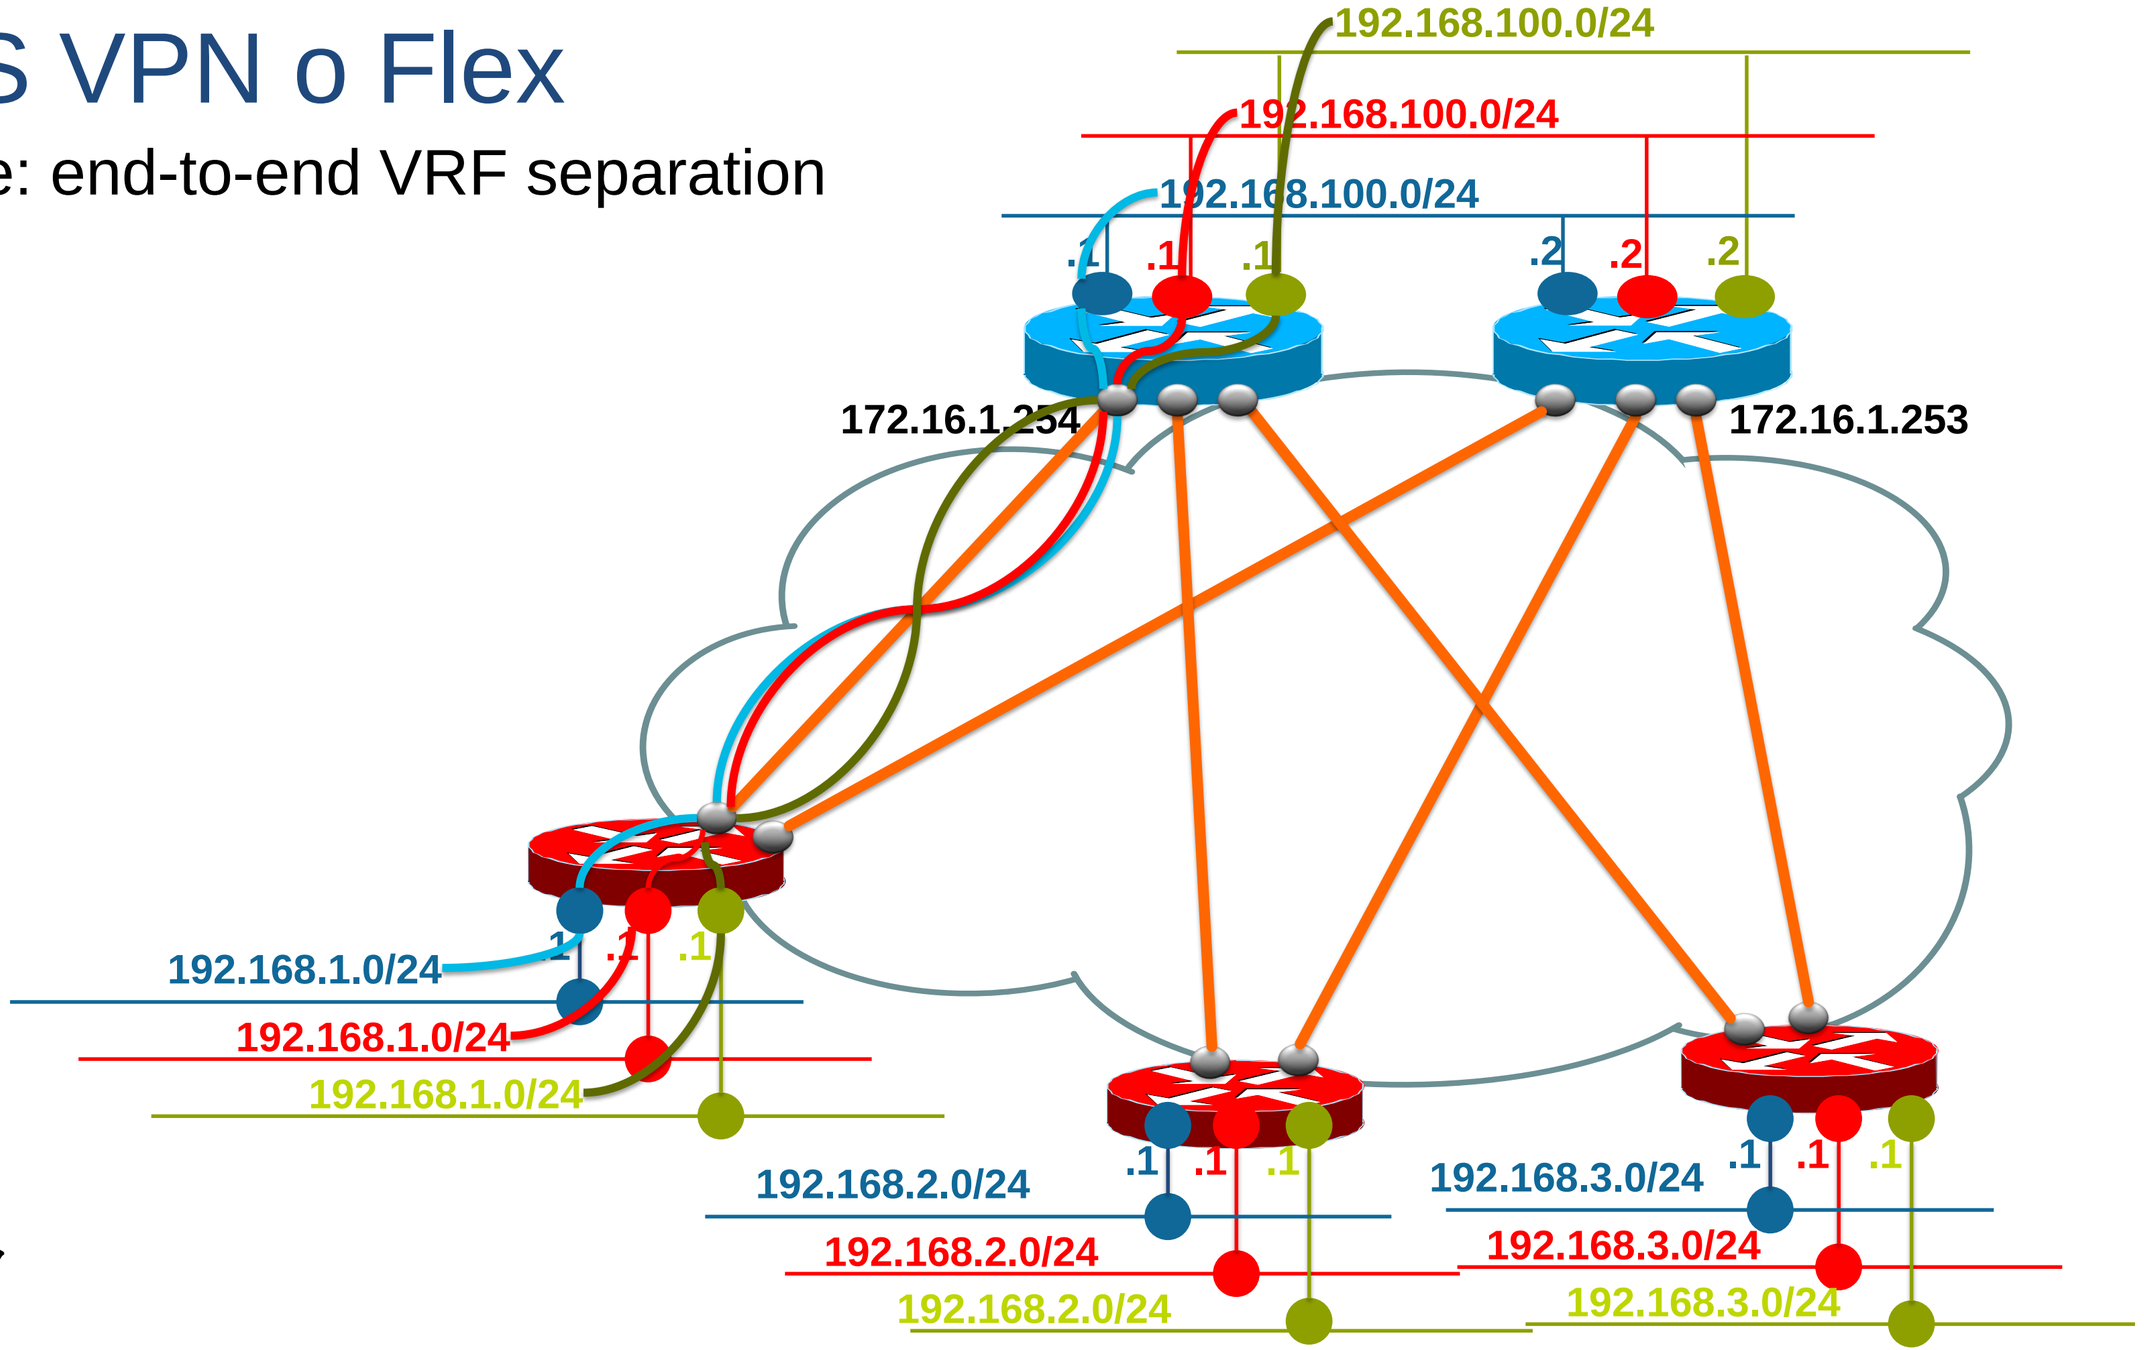

MPLS VPN o Flex

Going LDP Free

Hub private interface(s) in inside VRF or MPLS

Virtual-Access' in GRT, run MPLS

Tunnels create "back-to-back" links

➔ LDP not needed !!

Spoke tunnels run MPLS

Private interfaces in VRF's

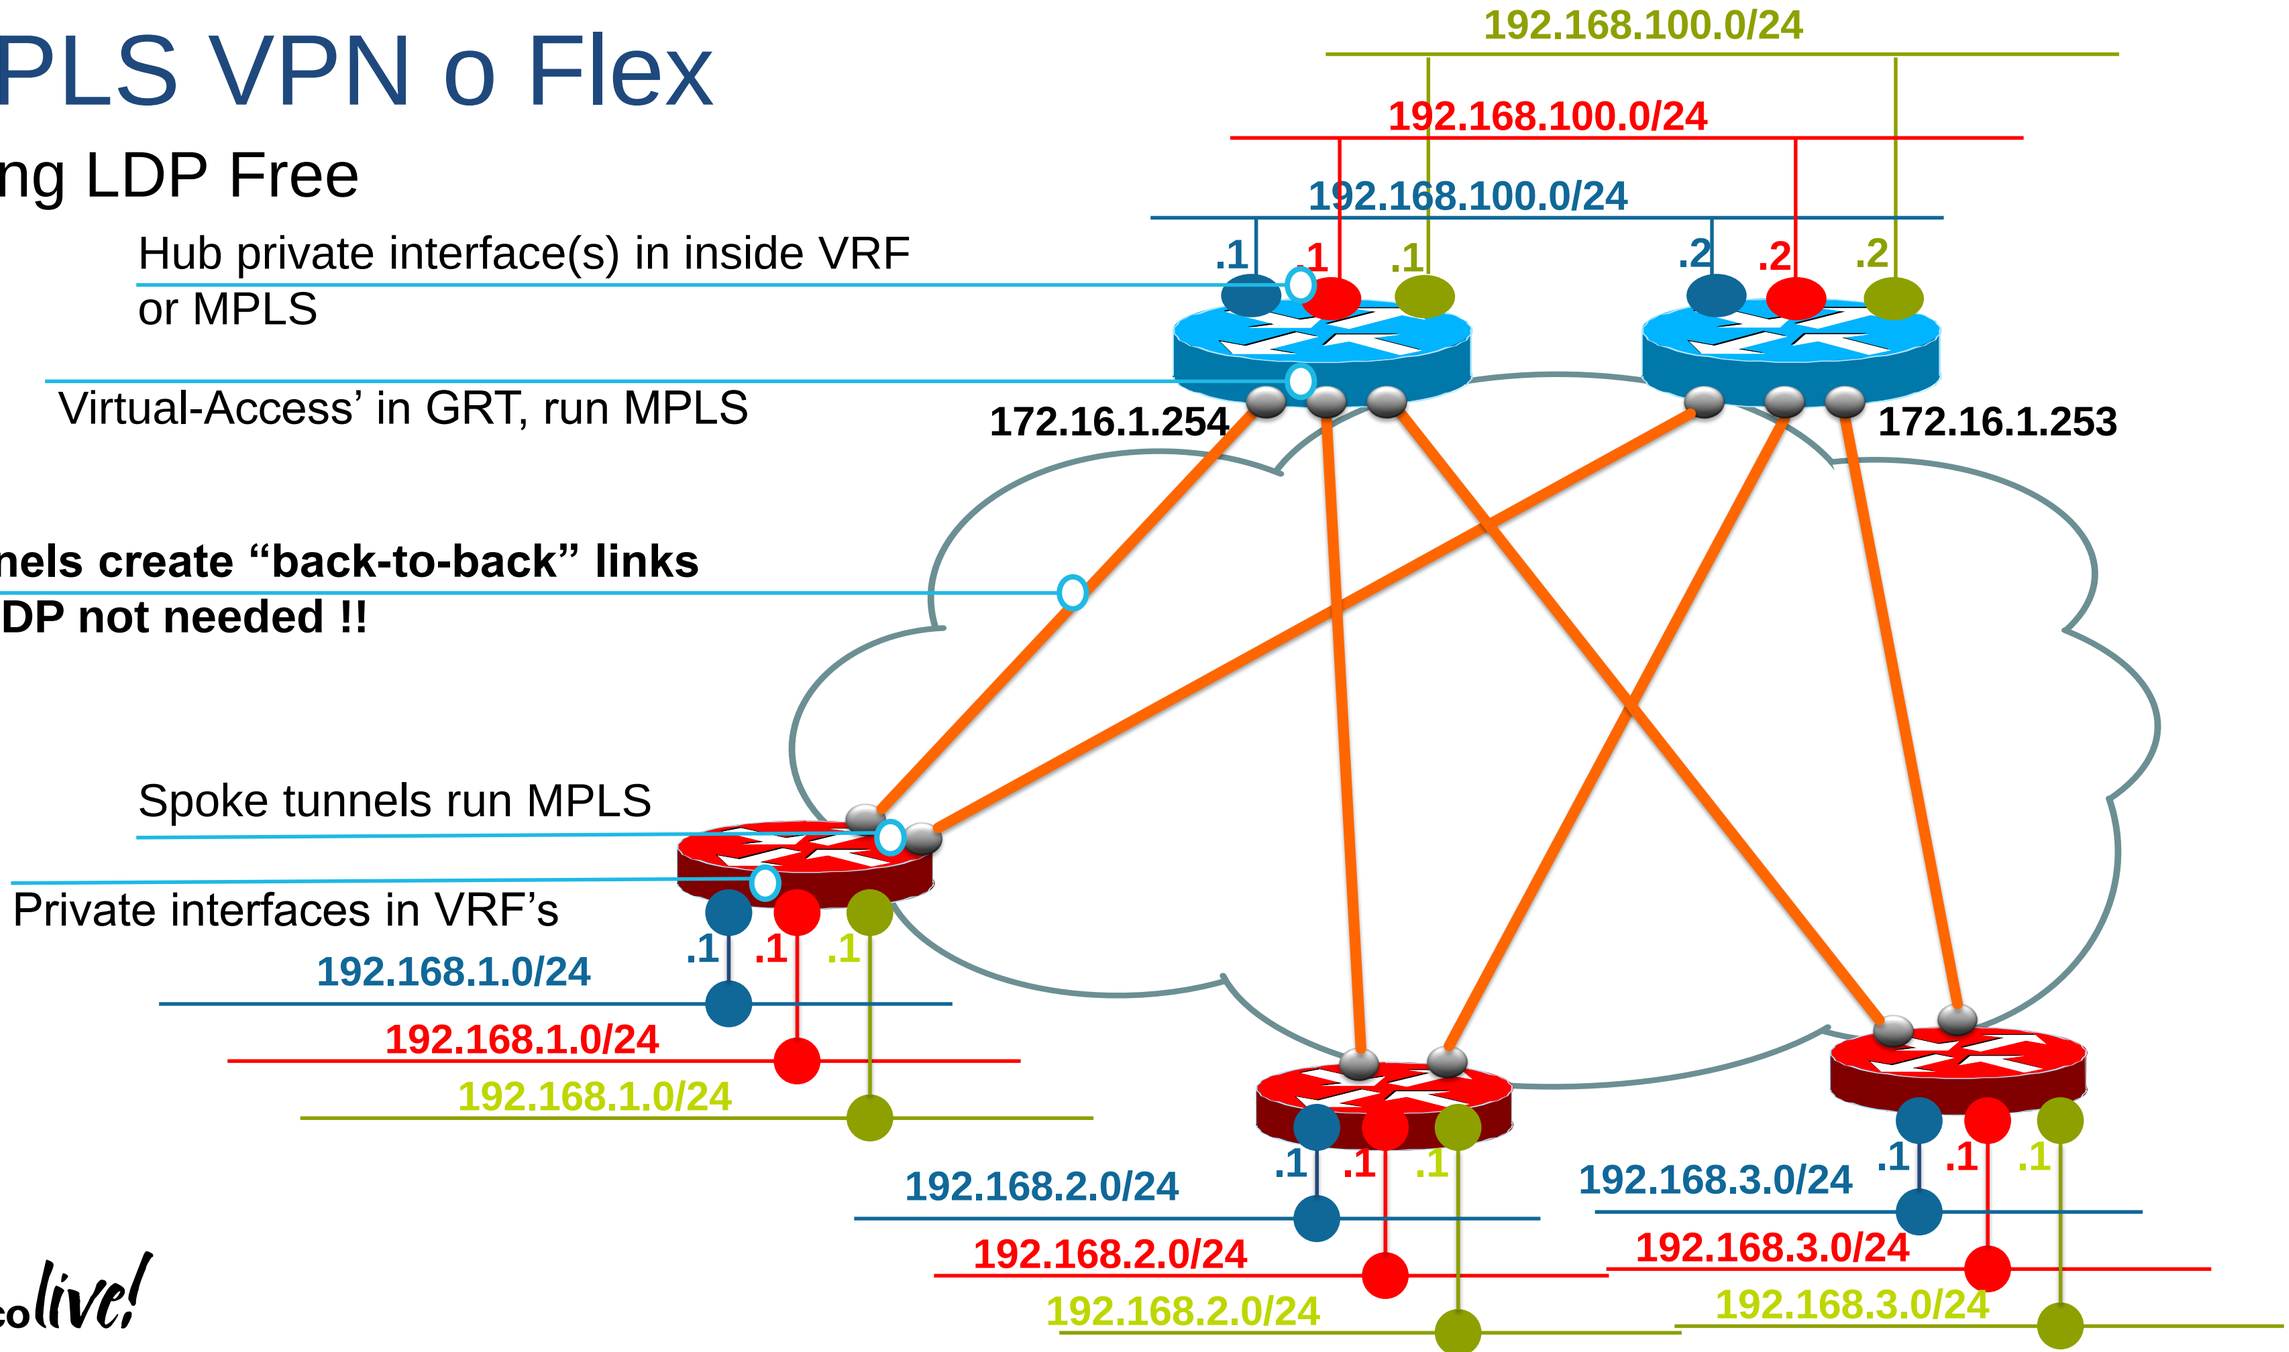

MPLS VPN o Flex

Extreme Summarization

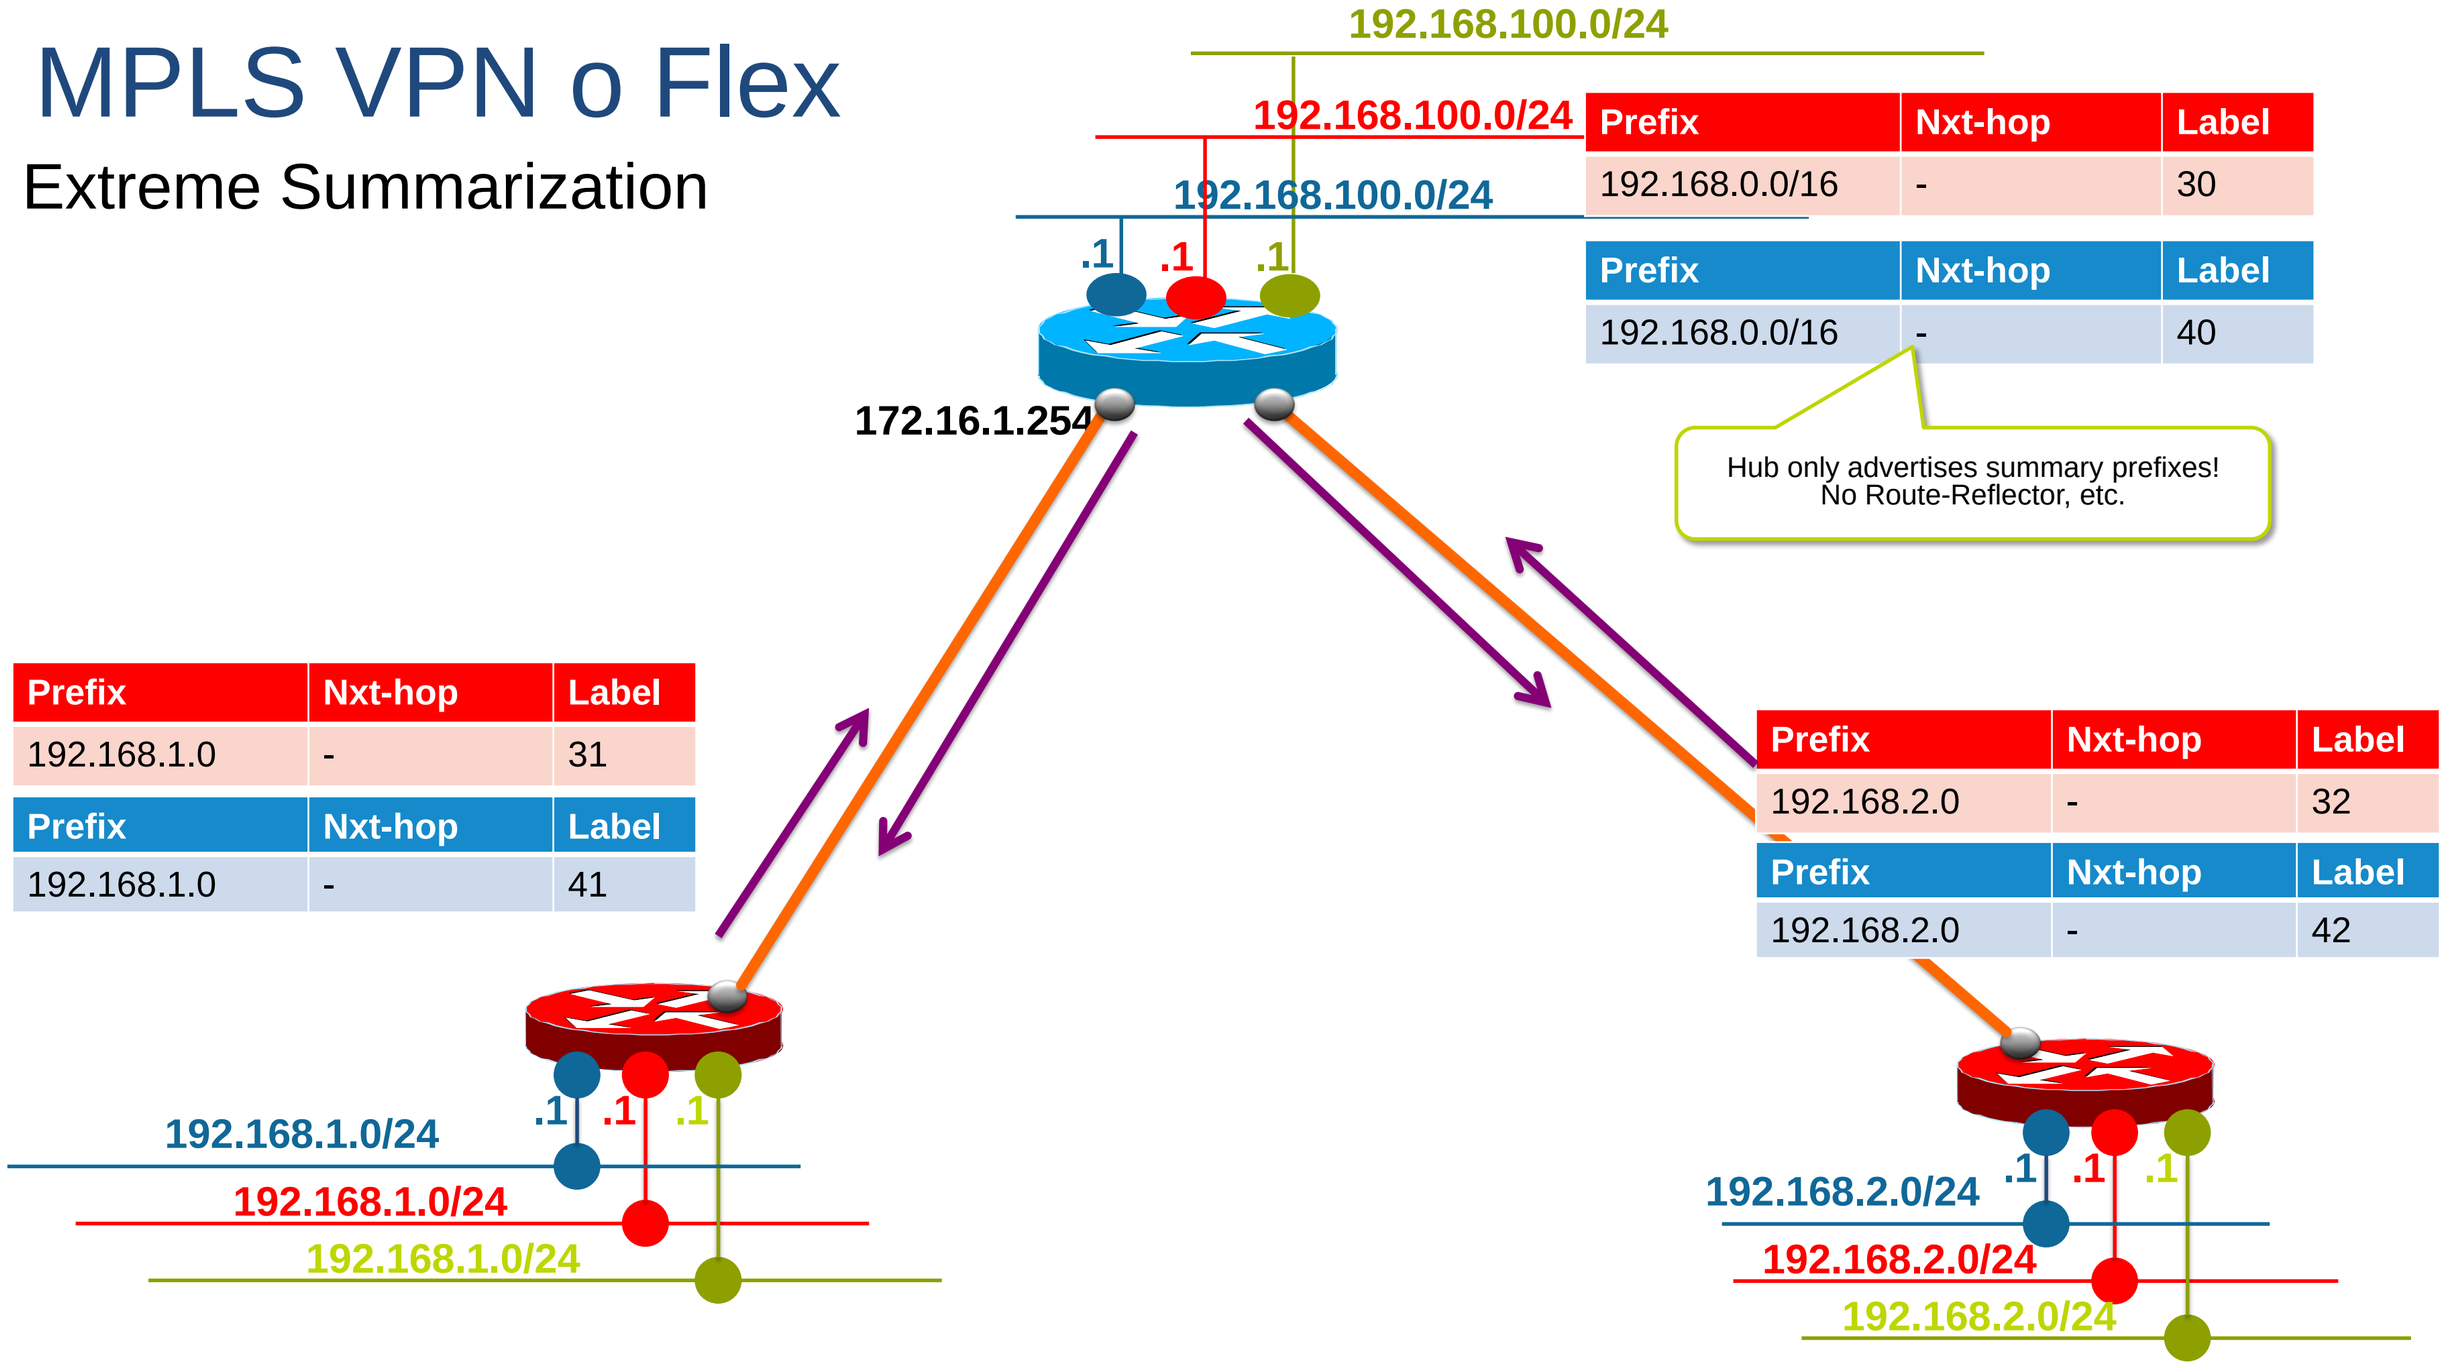

MPLS VPN o Flex

Summary Label Exchange

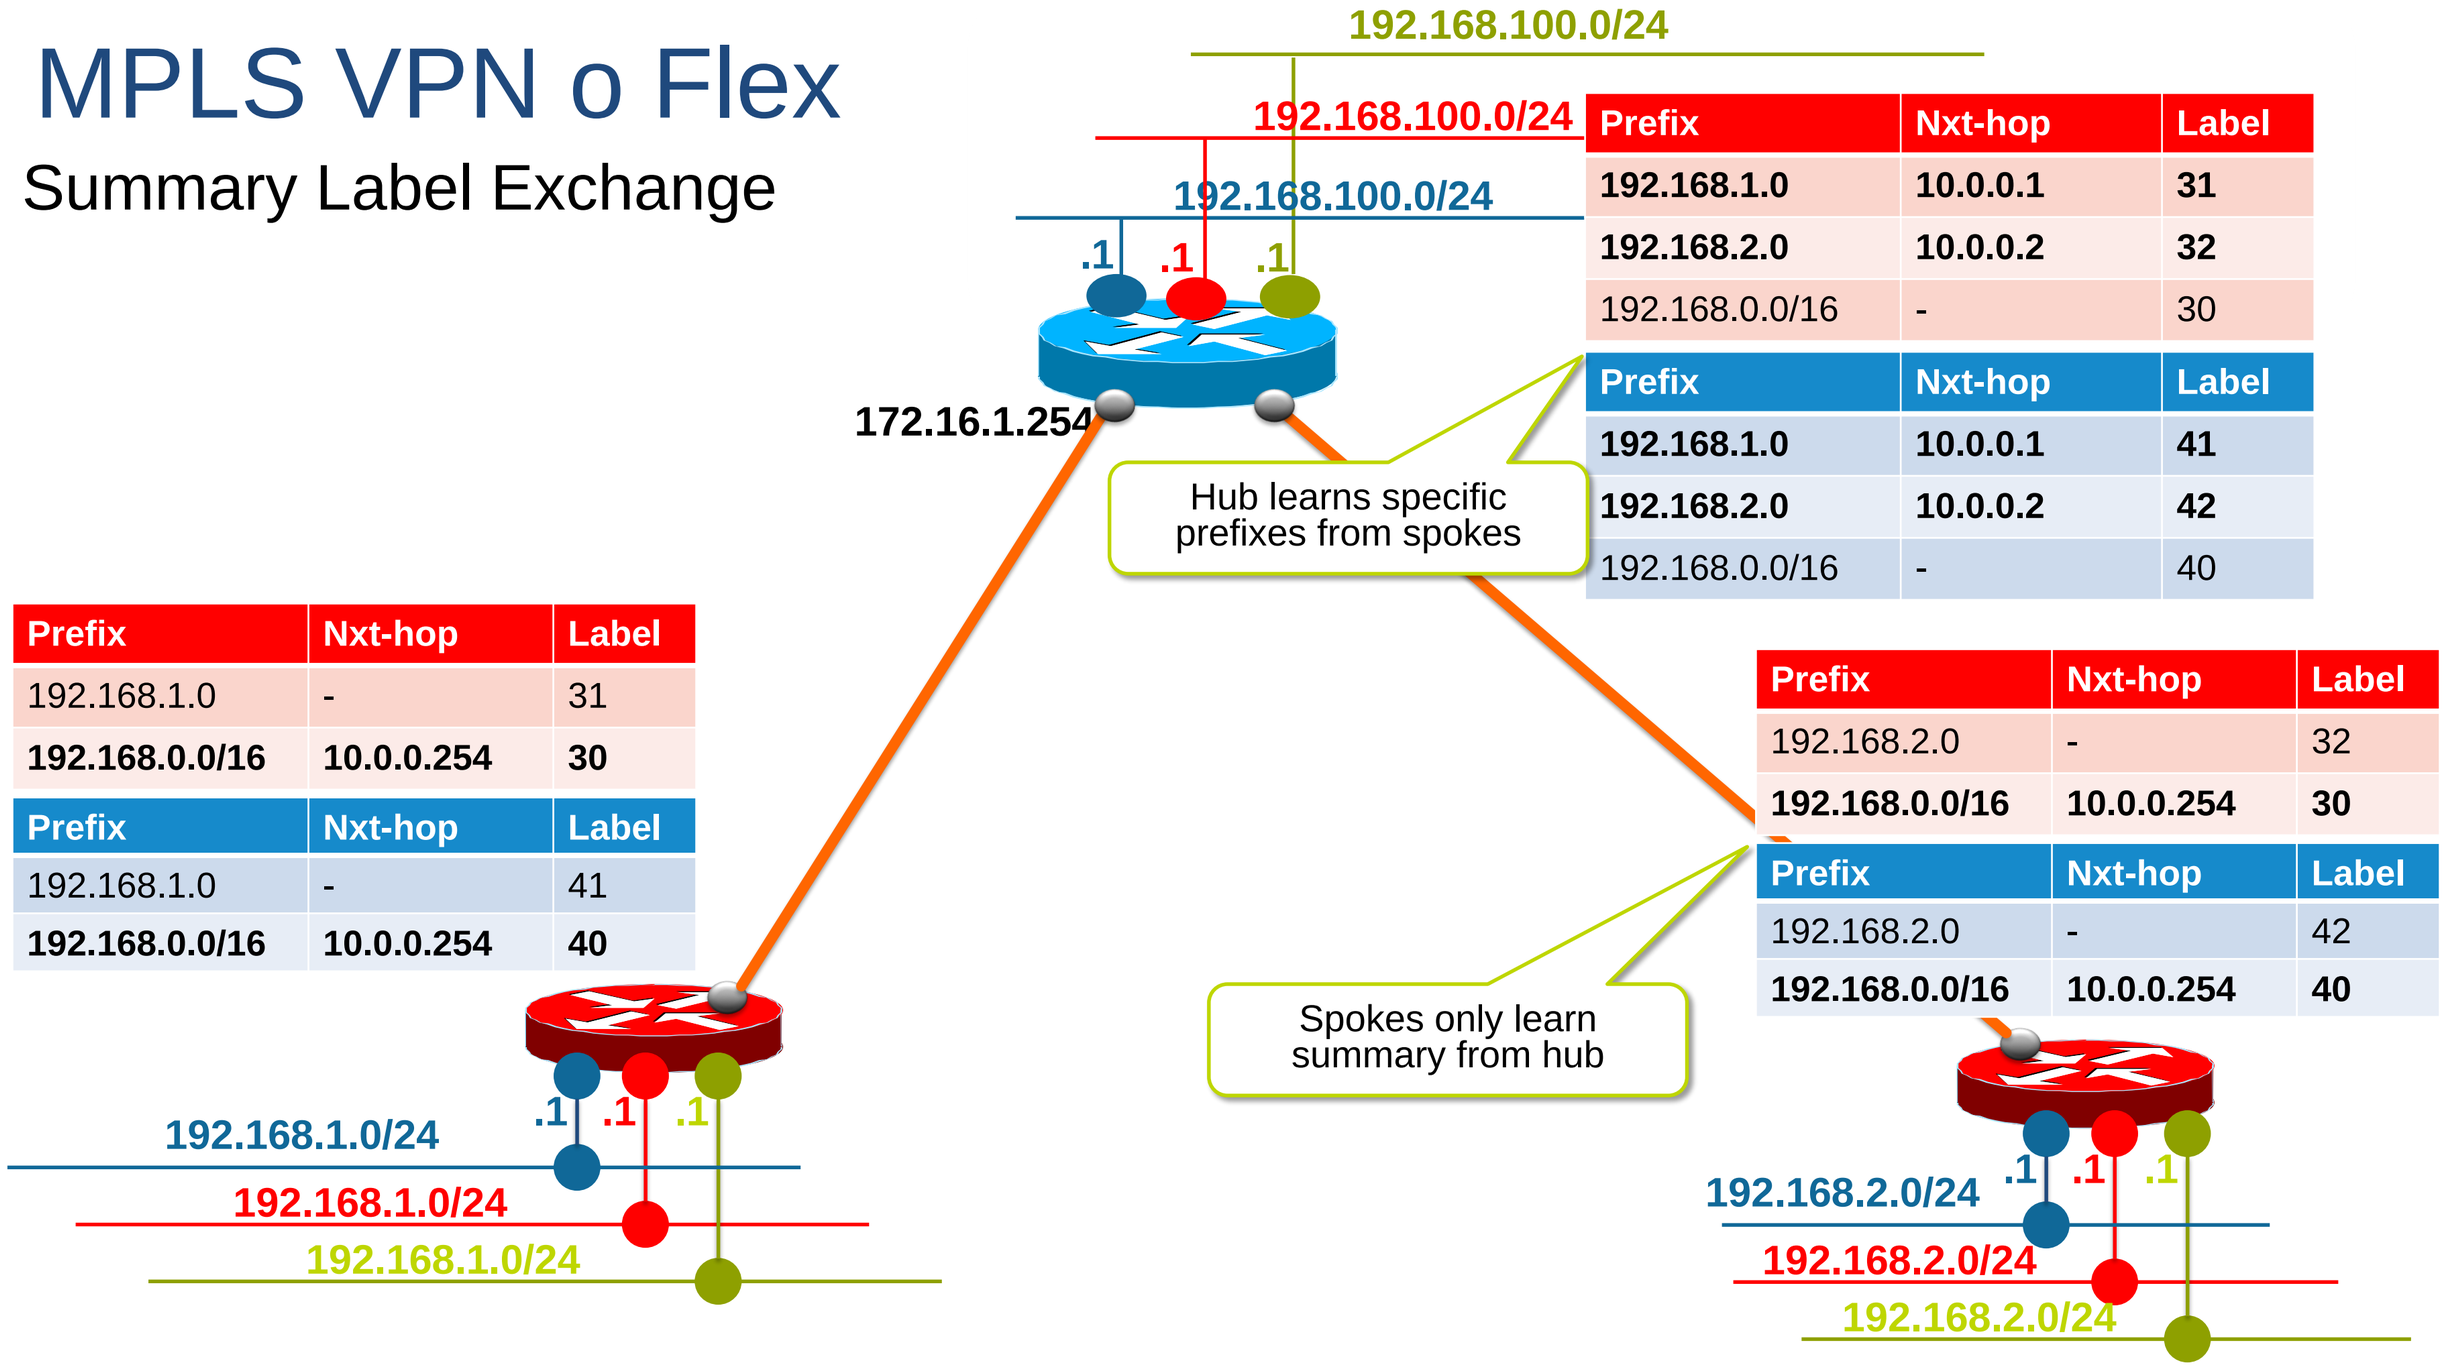

Prefix	Nxt-hop	Label
192.168.1.0	10.0.0.1	31
192.168.2.0	10.0.0.2	32
192.168.0.0/16	-	30

Prefix	Nxt-hop	Label
192.168.1.0	10.0.0.1	41
192.168.2.0	10.0.0.2	42
192.168.0.0/16	-	40

Prefix	Nxt-hop	Label
192.168.1.0	-	31
192.168.0.0/16	10.0.0.254	30

Prefix	Nxt-hop	Label
192.168.1.0	-	41
192.168.0.0/16	10.0.0.254	40

Prefix	Nxt-hop	Label
192.168.2.0	-	32
192.168.0.0/16	10.0.0.254	30

Prefix	Nxt-hop	Label
192.168.2.0	-	42
192.168.0.0/16	10.0.0.254	40

MPLS VPN o Flex

Hub & Spoke FIB's and LFIB's

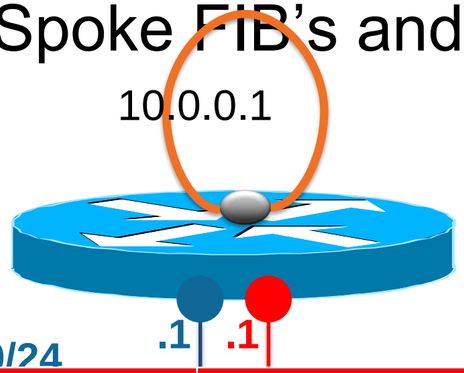

192.168.1.0/24

VRF FIBs

Prefix	Adjacency
192.168.1.0/24	Glean (e0)
192.168.0.0/16	10.0.0.254 **30**

Prefix	Adjacency
192.168.1.0/24	Glean (e1)
192.168.0.0/16	10.0.0.254 **40**

FIB

Prefix	Adjacency
10.0.0.254	Tunnel0 (**Null**)
0.0.0.0/0	Dialer0

LFIB

Loc.	Out	I/F
31	POP	VRF RED
41	POP	VRF BLUE

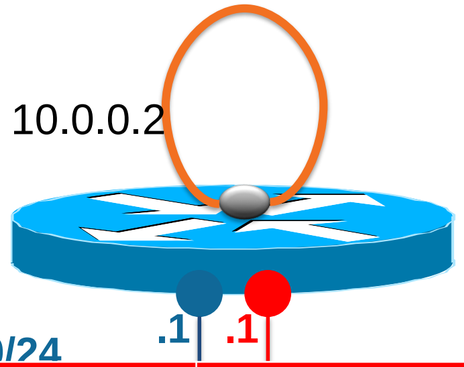

192.168.2.0/24

VRF FIBs

Prefix	Adjacency
192.168.2.0/24	Glean (e0)
192.168.0.0/16	10.0.0.254 **30**

Prefix	Adjacency
192.168.2.0/24	Glean (e1)
192.168.0.0/16	10.0.0.254 **40**

FIB

Prefix	Adjacency
10.0.0.254	Tunnel0 (**Null**)
0.0.0.0/0	Dialer0

LFIB

Loc.	Out	I/F
32	POP	VRF RED
42	POP	VRF BLUE

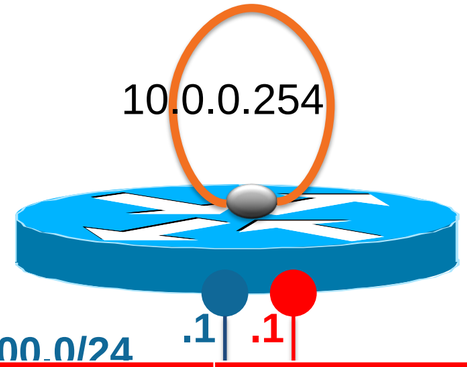

192.168.100.0/24

VRF FIBs

Prefix	Adjacency
192.168.1.0/24	10.0.0.1 **31**
192.168.2.0/24	10.0.0.2 **32**

Prefix	Adjacency
192.168.1.0/24	10.0.0.1 **41**
192.168.2.0/24	10.0.0.2 **42**

FIB

Prefix	Adjacency
10.0.0.1	VA-1 (**Null**)
10.0.0.2	VA-2 (**Null**)
0.0.0.0	Dialer0

LFIB

Loc.	Out	I/F
30	POP	VRF RED
40	POP	VRF BLUE

MPLS VPN o Flex

Spoke → Hub

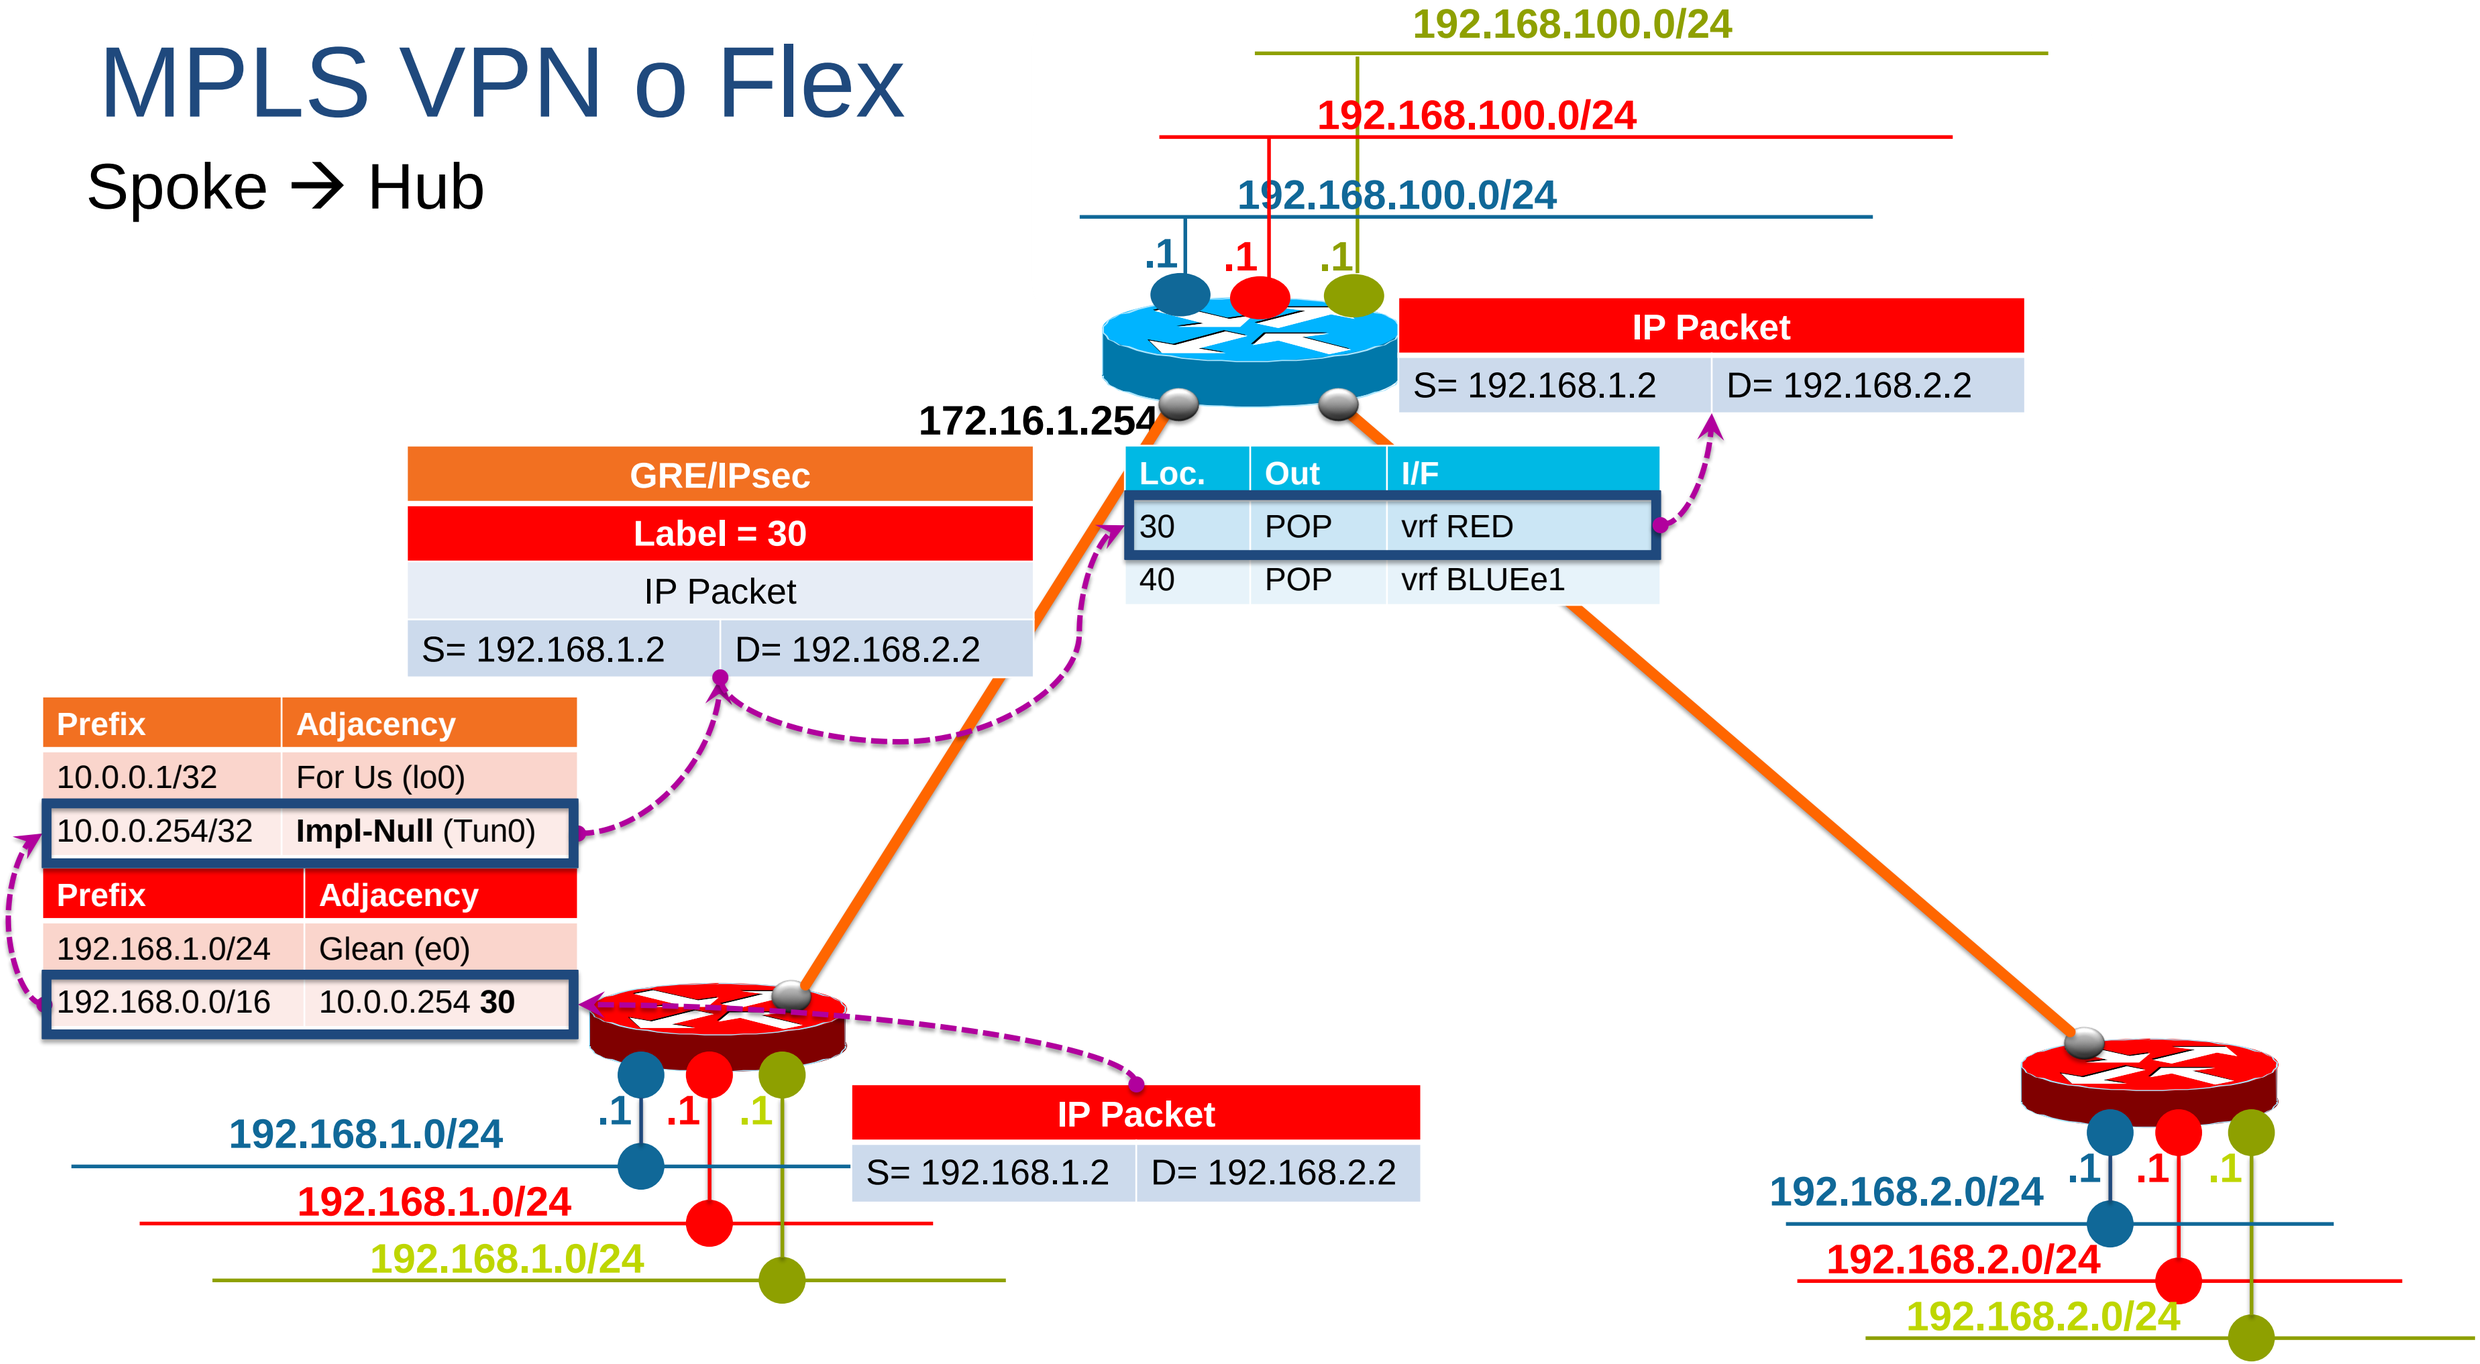

MPLS VPN o Flex

Hub → Spoke

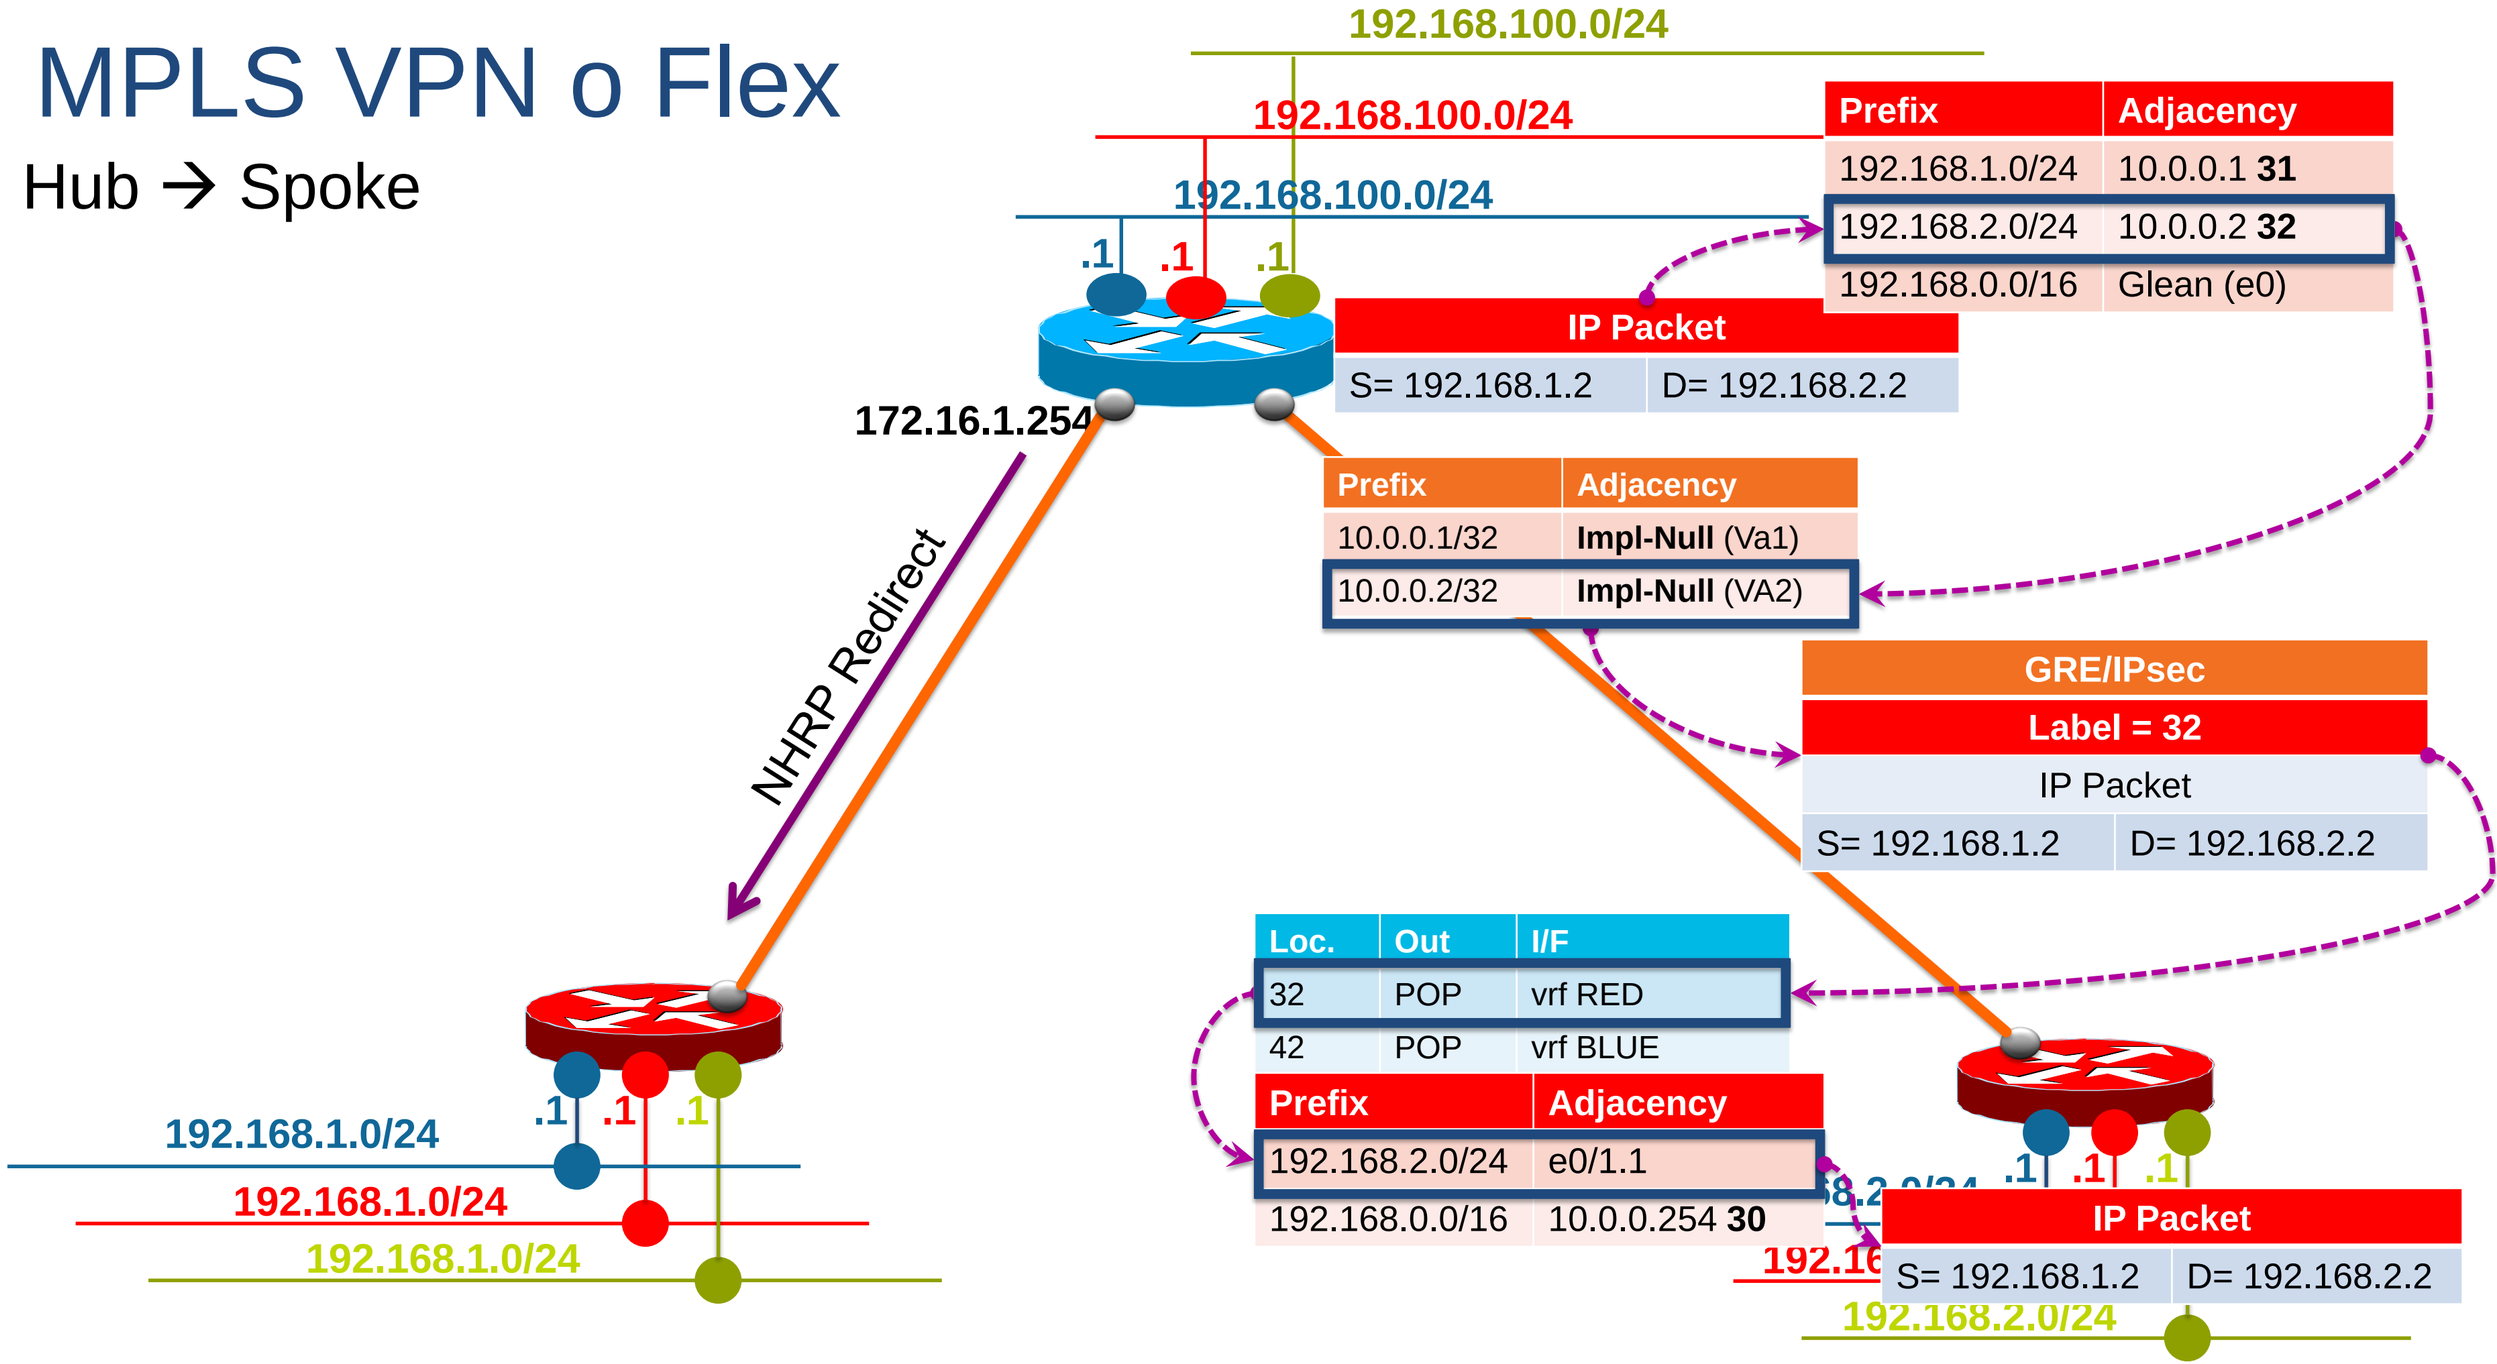

MPLS VPN o Flex

Enhanced NHRP Notifications

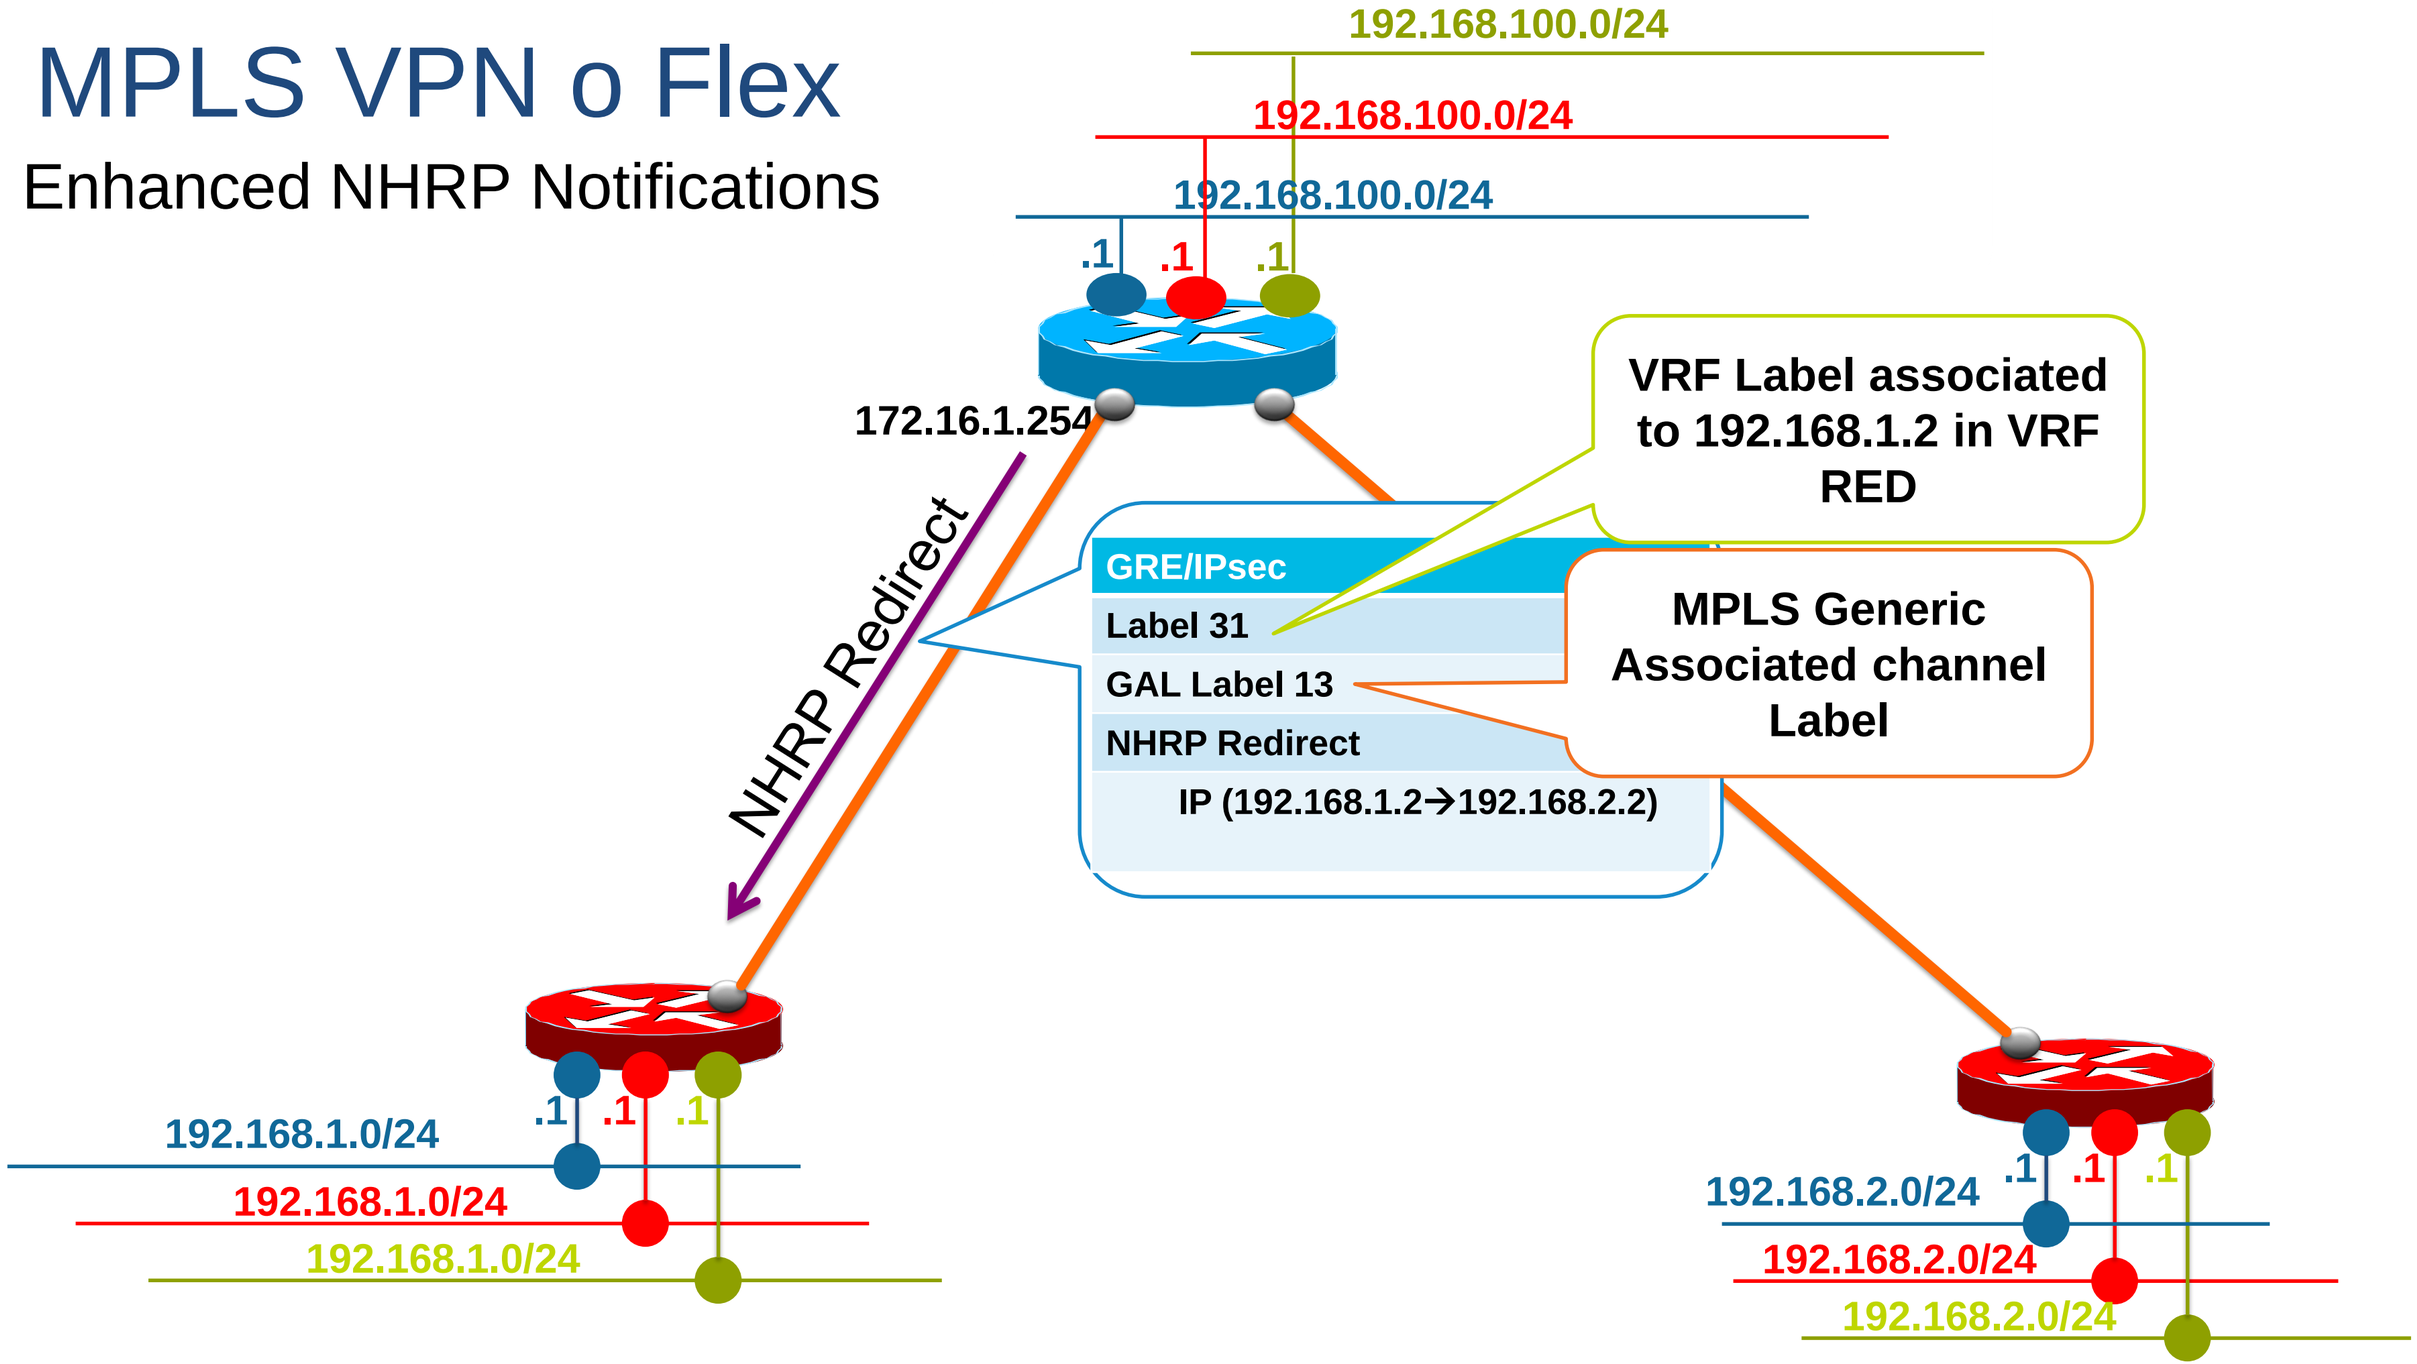

Hub Routing

BGP and Interfaces

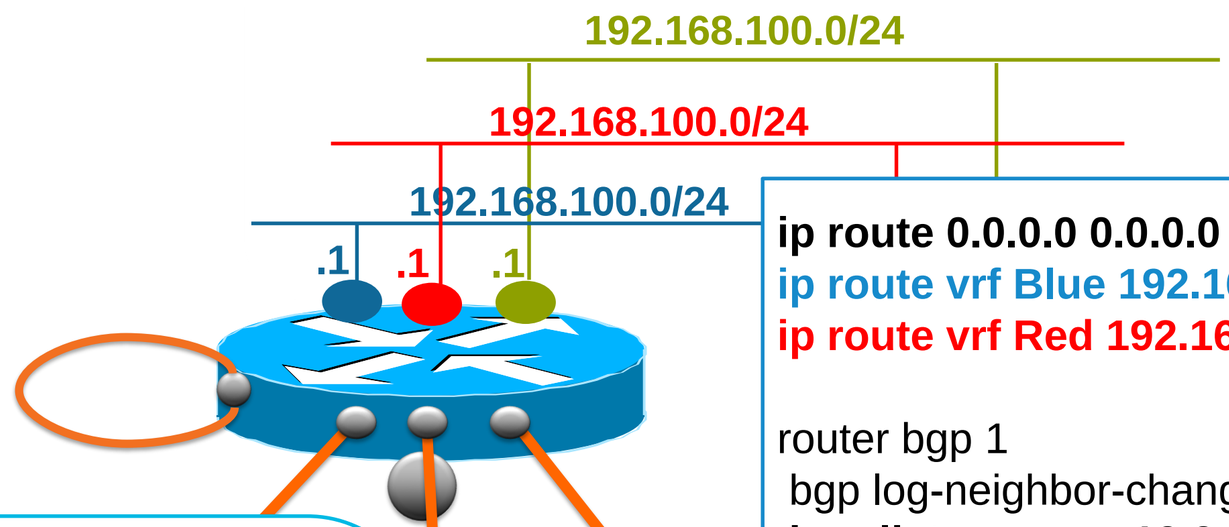


```
interface Virtual-Template1 type tunnel
 ip unnumbered Loopback0
 ip nhrp network-id 1
 ip nhrp redirect
 mpls nhrp
 tunnel protection ipsec profile default

interface Ethernet0/0
 ip address 172.16.1.254 255.255.255.0

interface Loopback0
 ip address 10.0.0.254 255.255.255.255
```

Activate NHRP redirects and give NHRP control over MPLS

WAN Interface can be in Front VRF

Tunnels and Loopback in Global Routing Table

```
ip route 0.0.0.0 0.0.0.0 172.16.1.2
ip route vrf Blue 192.168.0.0 255.255.0.0 Null0
ip route vrf Red 192.168.0.0 255.255.0.0 Null0

router bgp 1
 bgp log-neighbor-changes
 bgp listen range 10.0.0.0/16 peer-group Flex
 neighbor Flex peer-group
 neighbor Flex remote-as 1
 neighbor Flex timers 5 15

 address-family vpnv4
  neighbor Flex activate
  neighbor Flex send-community extended

 address-family ipv4 vrf Blue
  network 192.168.0.0 mask 255.255.0.0

 address-family ipv4 vrf Red
  network 192.168.0.0 mask 255.255.0.0
```

Activate VPNv4

Advertise each VRF

Spoke Routing Configuration

BGP and Static Routes

```
interface Tunnel0
 ip unnumbered Loopback0
 ip nhrp network-id 1
 ip nhrp shortcut virtual-template 1
 mpls nhrp
 tunnel source Ethernet0/0
 tunnel destination 172.16.1.254
 tunnel protection ipsec profile default

interface Tunnel1
 (same as Tunnel0 – points to hub 2)

interface Virtual-Template 1 type tunnel
 ip unnumbered Loopback0
 ip nhrp network-id 1
 ip nhrp shortcut virtual-template 1
 mpls nhrp
 tunnel protection ipsec profile default

interface Ethernet0/0
 ip address 172.16.2.1 255.255.255.0

interface Loopback0
 ip address 10.0.0.2 255.255.255.255
```

Same as on Hub: Start MPLS forwarding without LDP

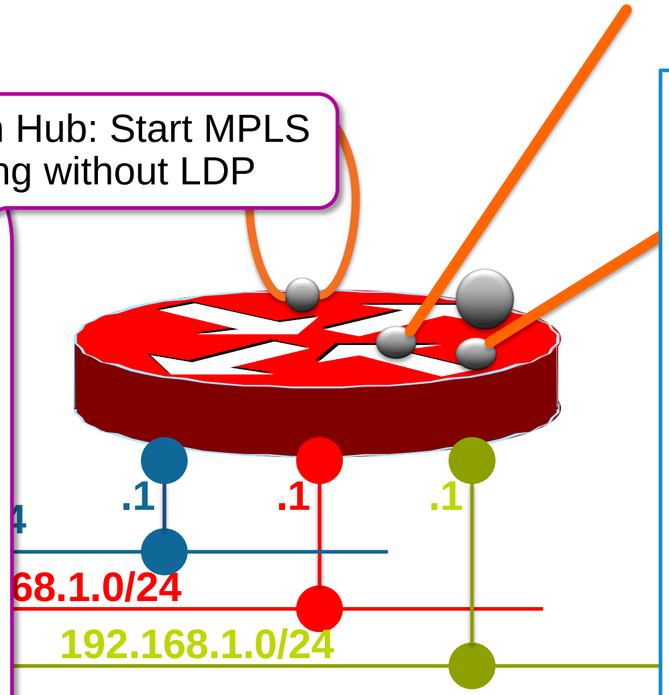

WAN Interface can be in Front VRF

Tunnels and Loopback in Global Routing Table

```
ip route 0.0.0.0 0.0.0.0 172.16.2.2

router bgp 1
 bgp log-neighbor-changes
 neighbor Flex peer-group
 neighbor Flex remote-as 1
 neighbor Flex timers 5 15
 neighbor 10.0.0.253 peer-group Flex
 neighbor 10.0.0.254 peer-group Flex

 address-family vpnv4
  neighbor Flex send-community extended
  neighbor 10.0.0.253 activate
  neighbor 10.0.0.254 activate

 address-family ipv4 vrf Blue
  redistribute connected
  maximum-paths ibgp 2

 address-family ipv4 vrf Red
  redistribute connected
  maximum-paths ibgp 2
```

Activate VPNv4

Advertise each VRF

Multicast over FlexVPN

Multicast Source location

Hub private interface(s) in Inside VRF (light)

Virtual-Access in iVRF

No VRF on spokes
(for this example)

Source behind spoke
is our focus now

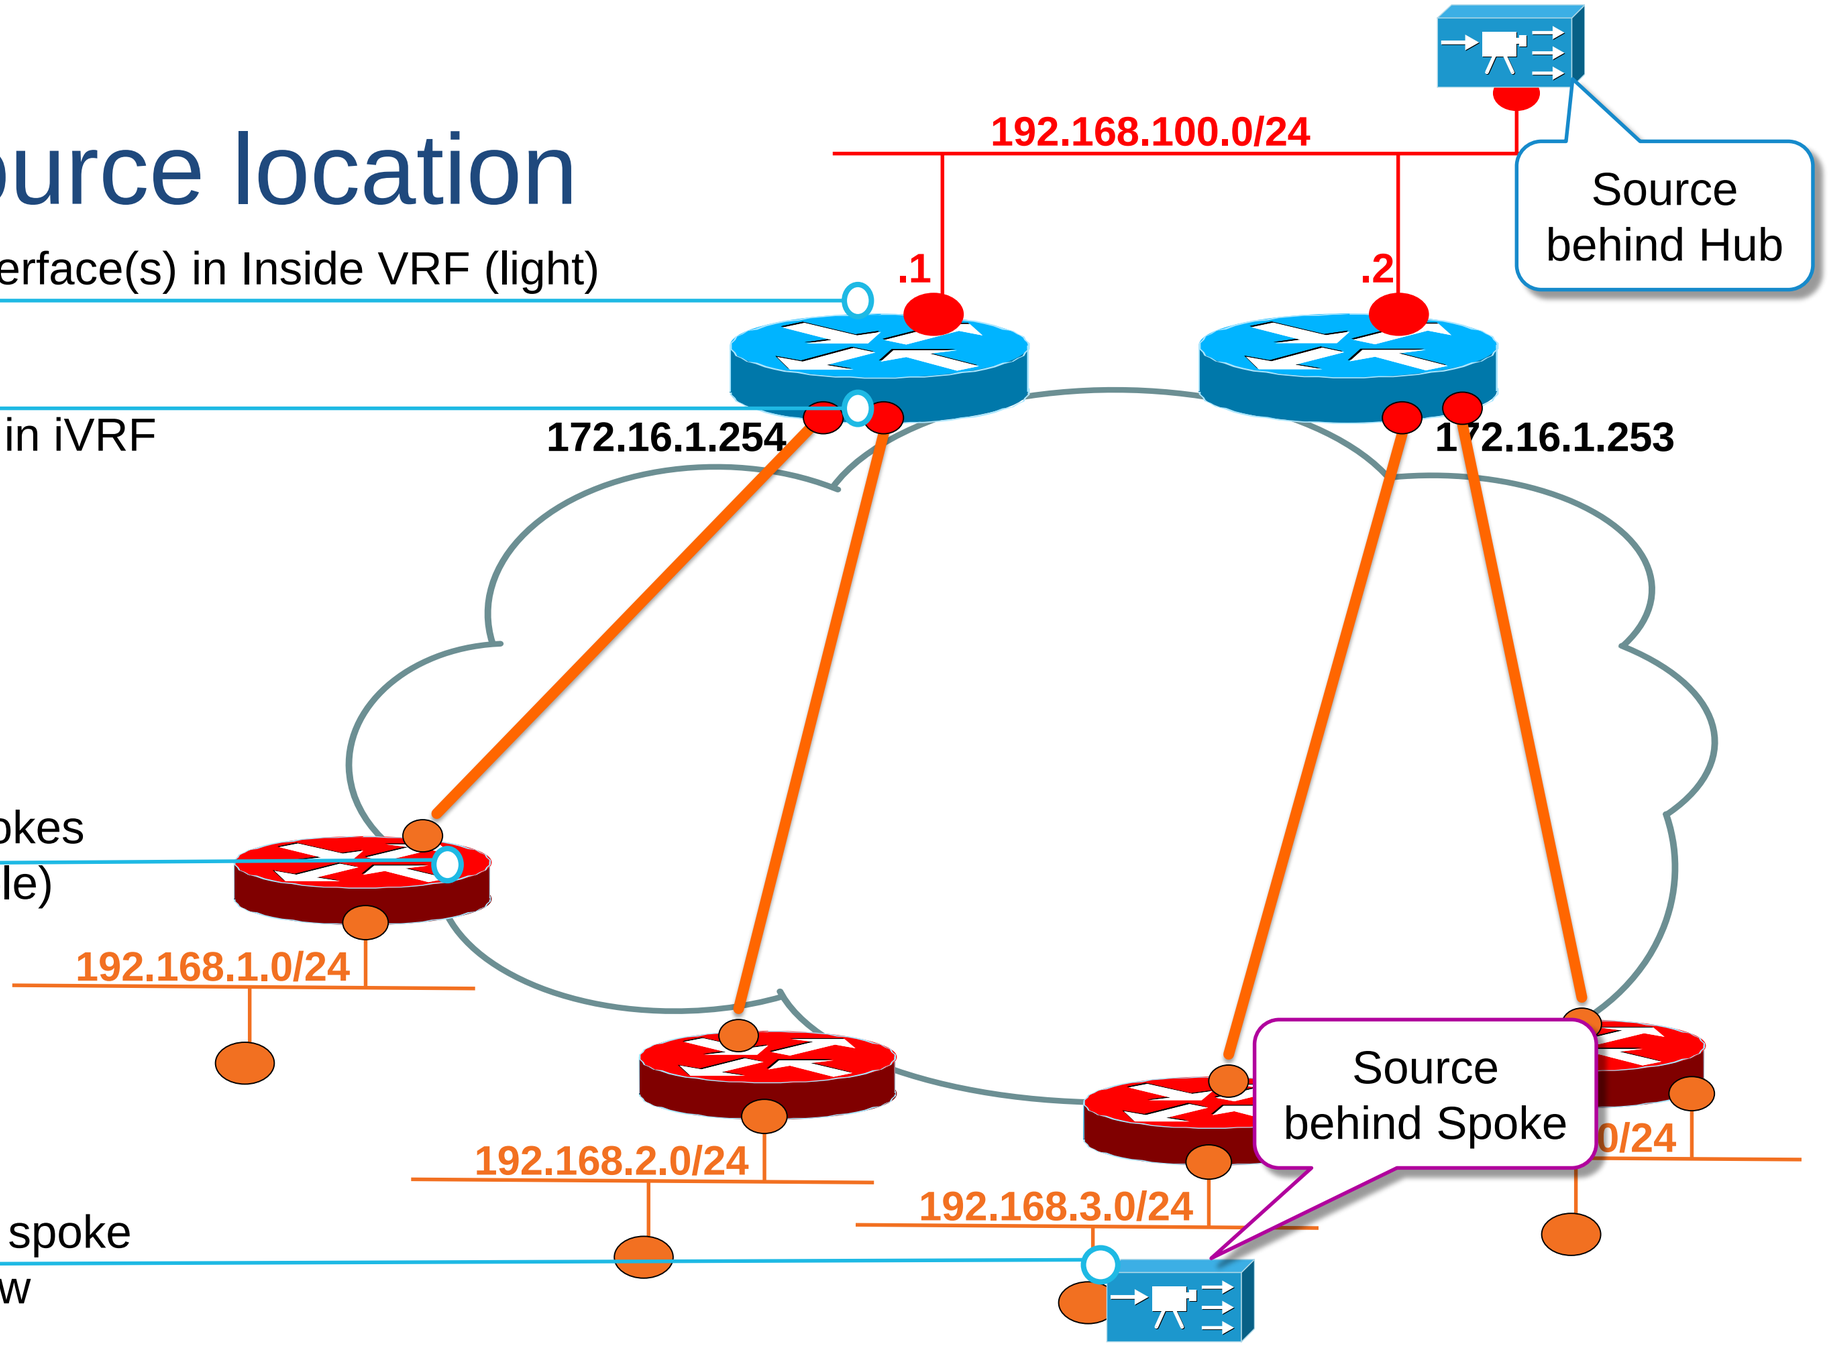

PIM Sparse-Mode Overview

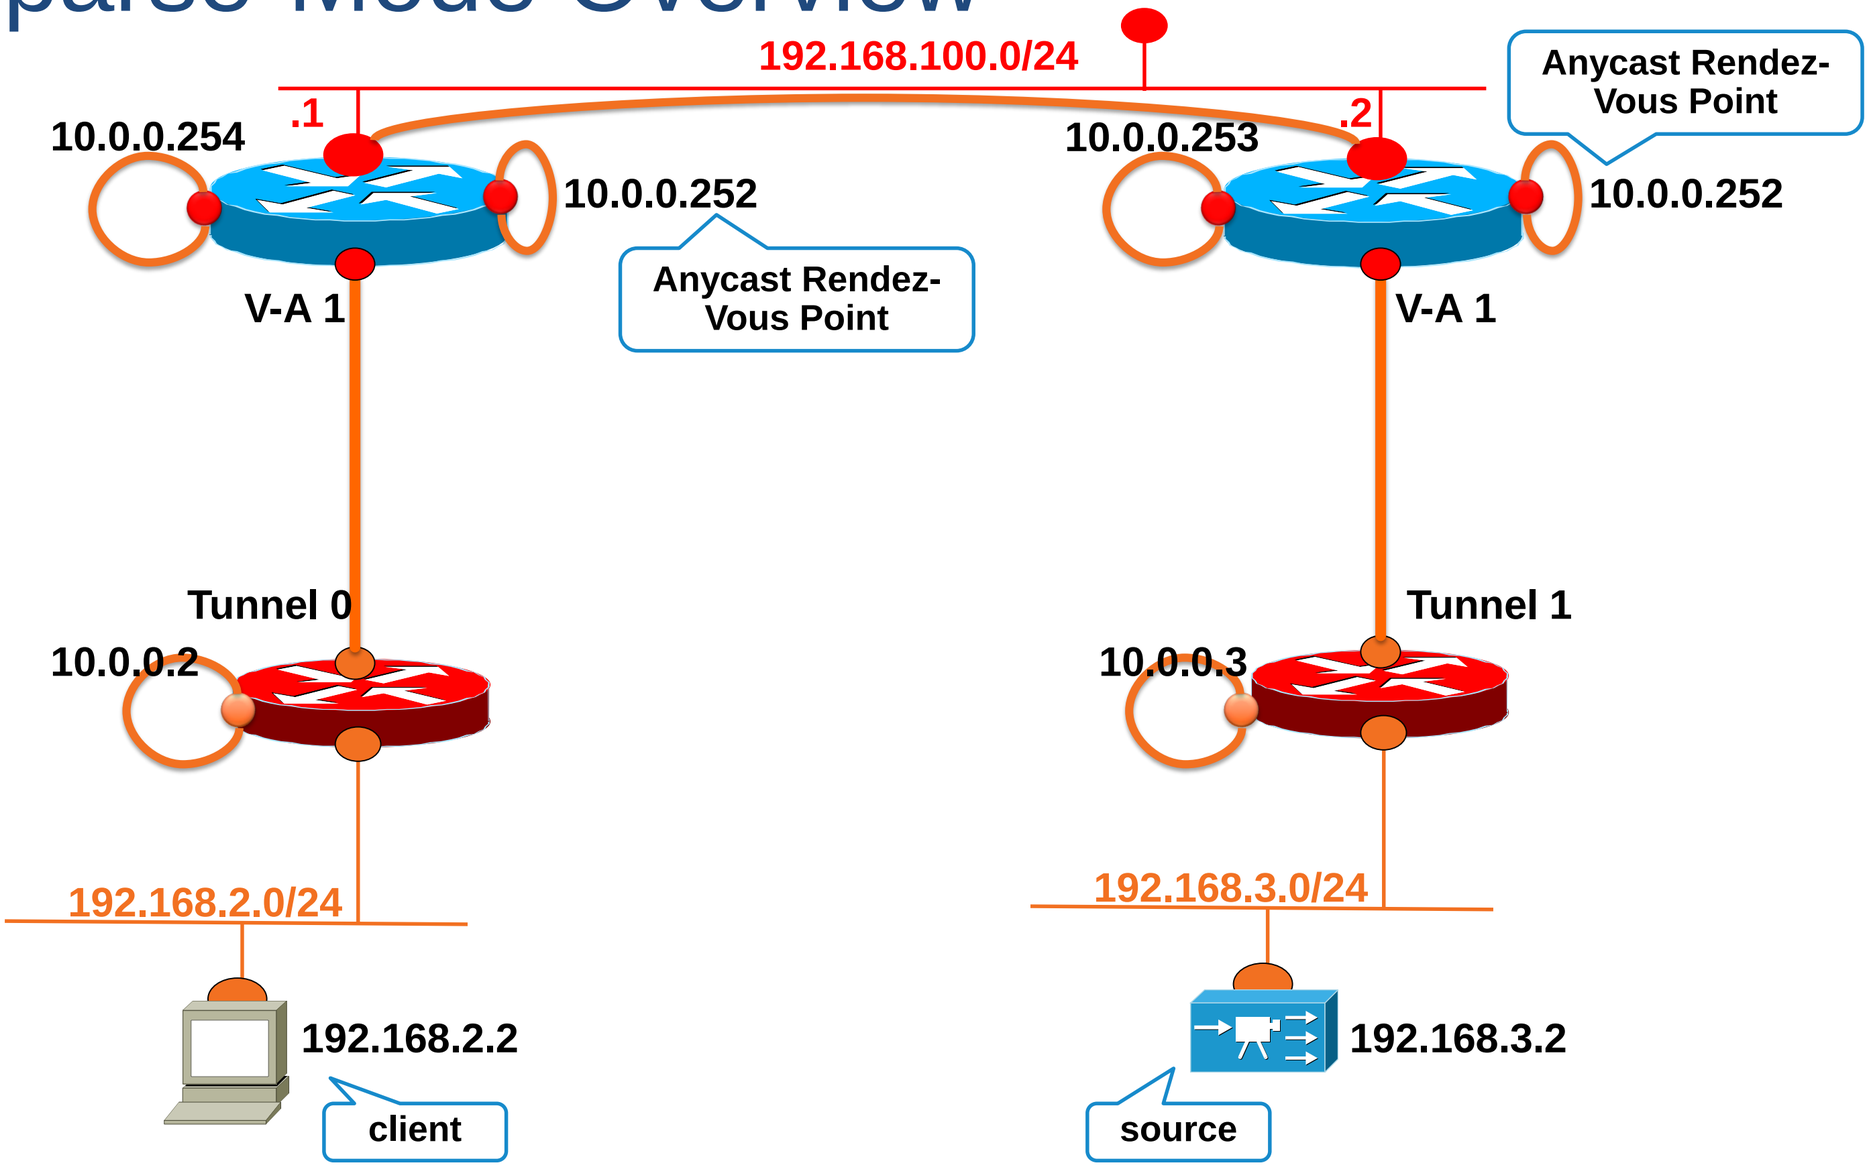

Recommendation – SAFI 2 & SAFI 129

Hub will remote-control Spoke's MuRIB via BGP

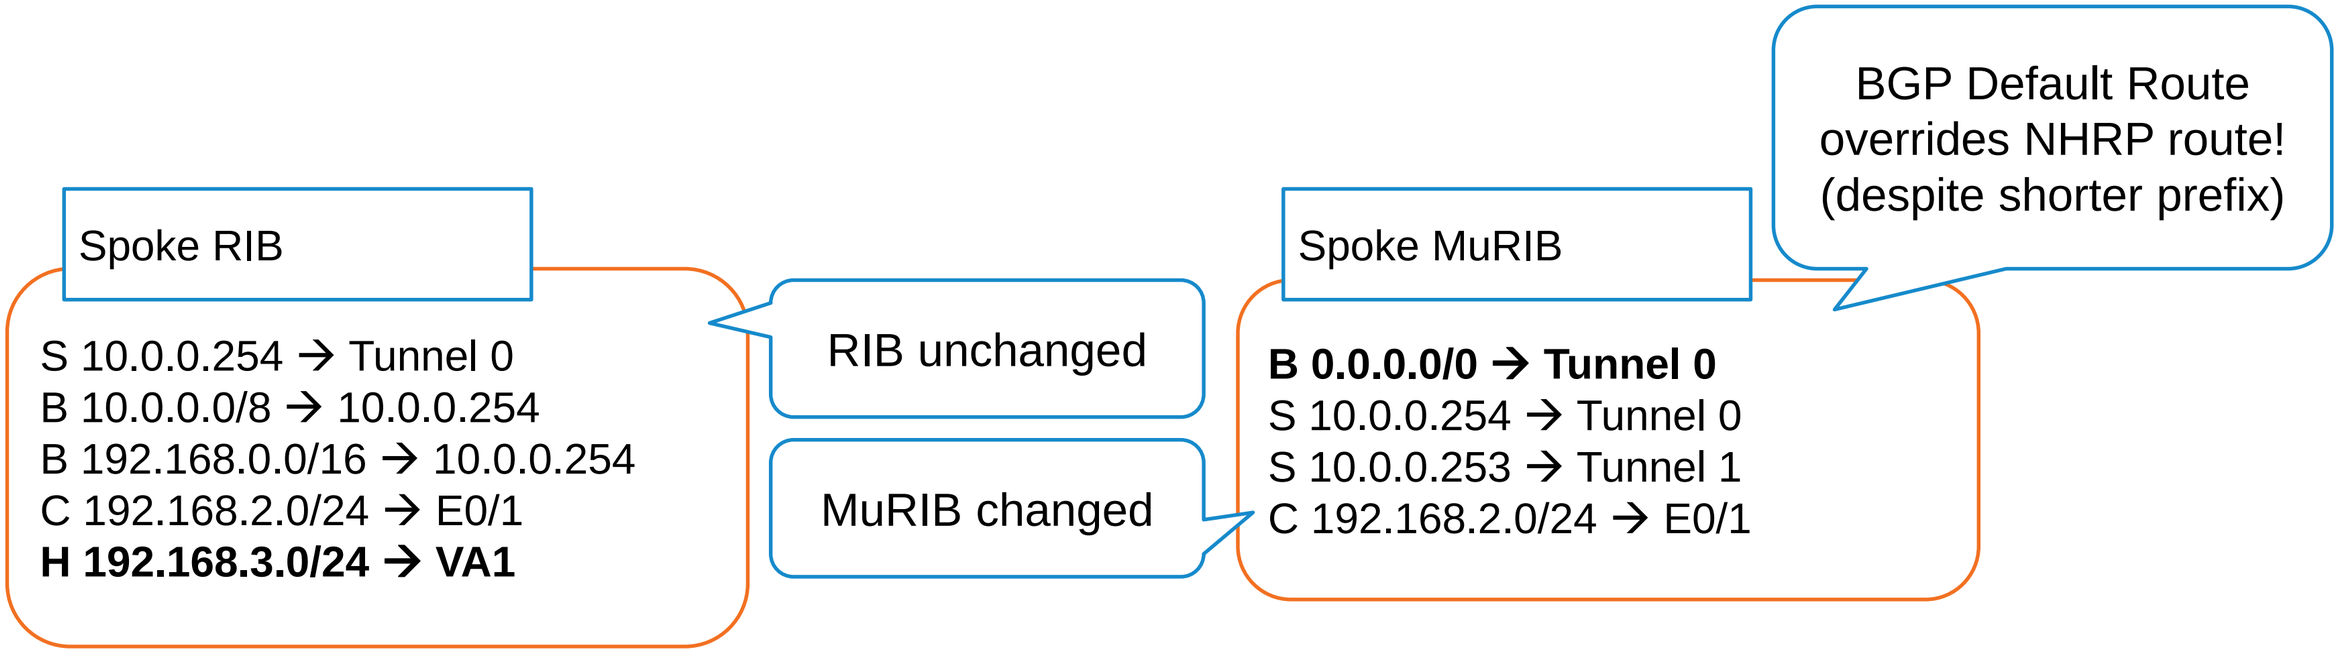

PIM Sparse-Mode – Unicast RIB and MuRIB

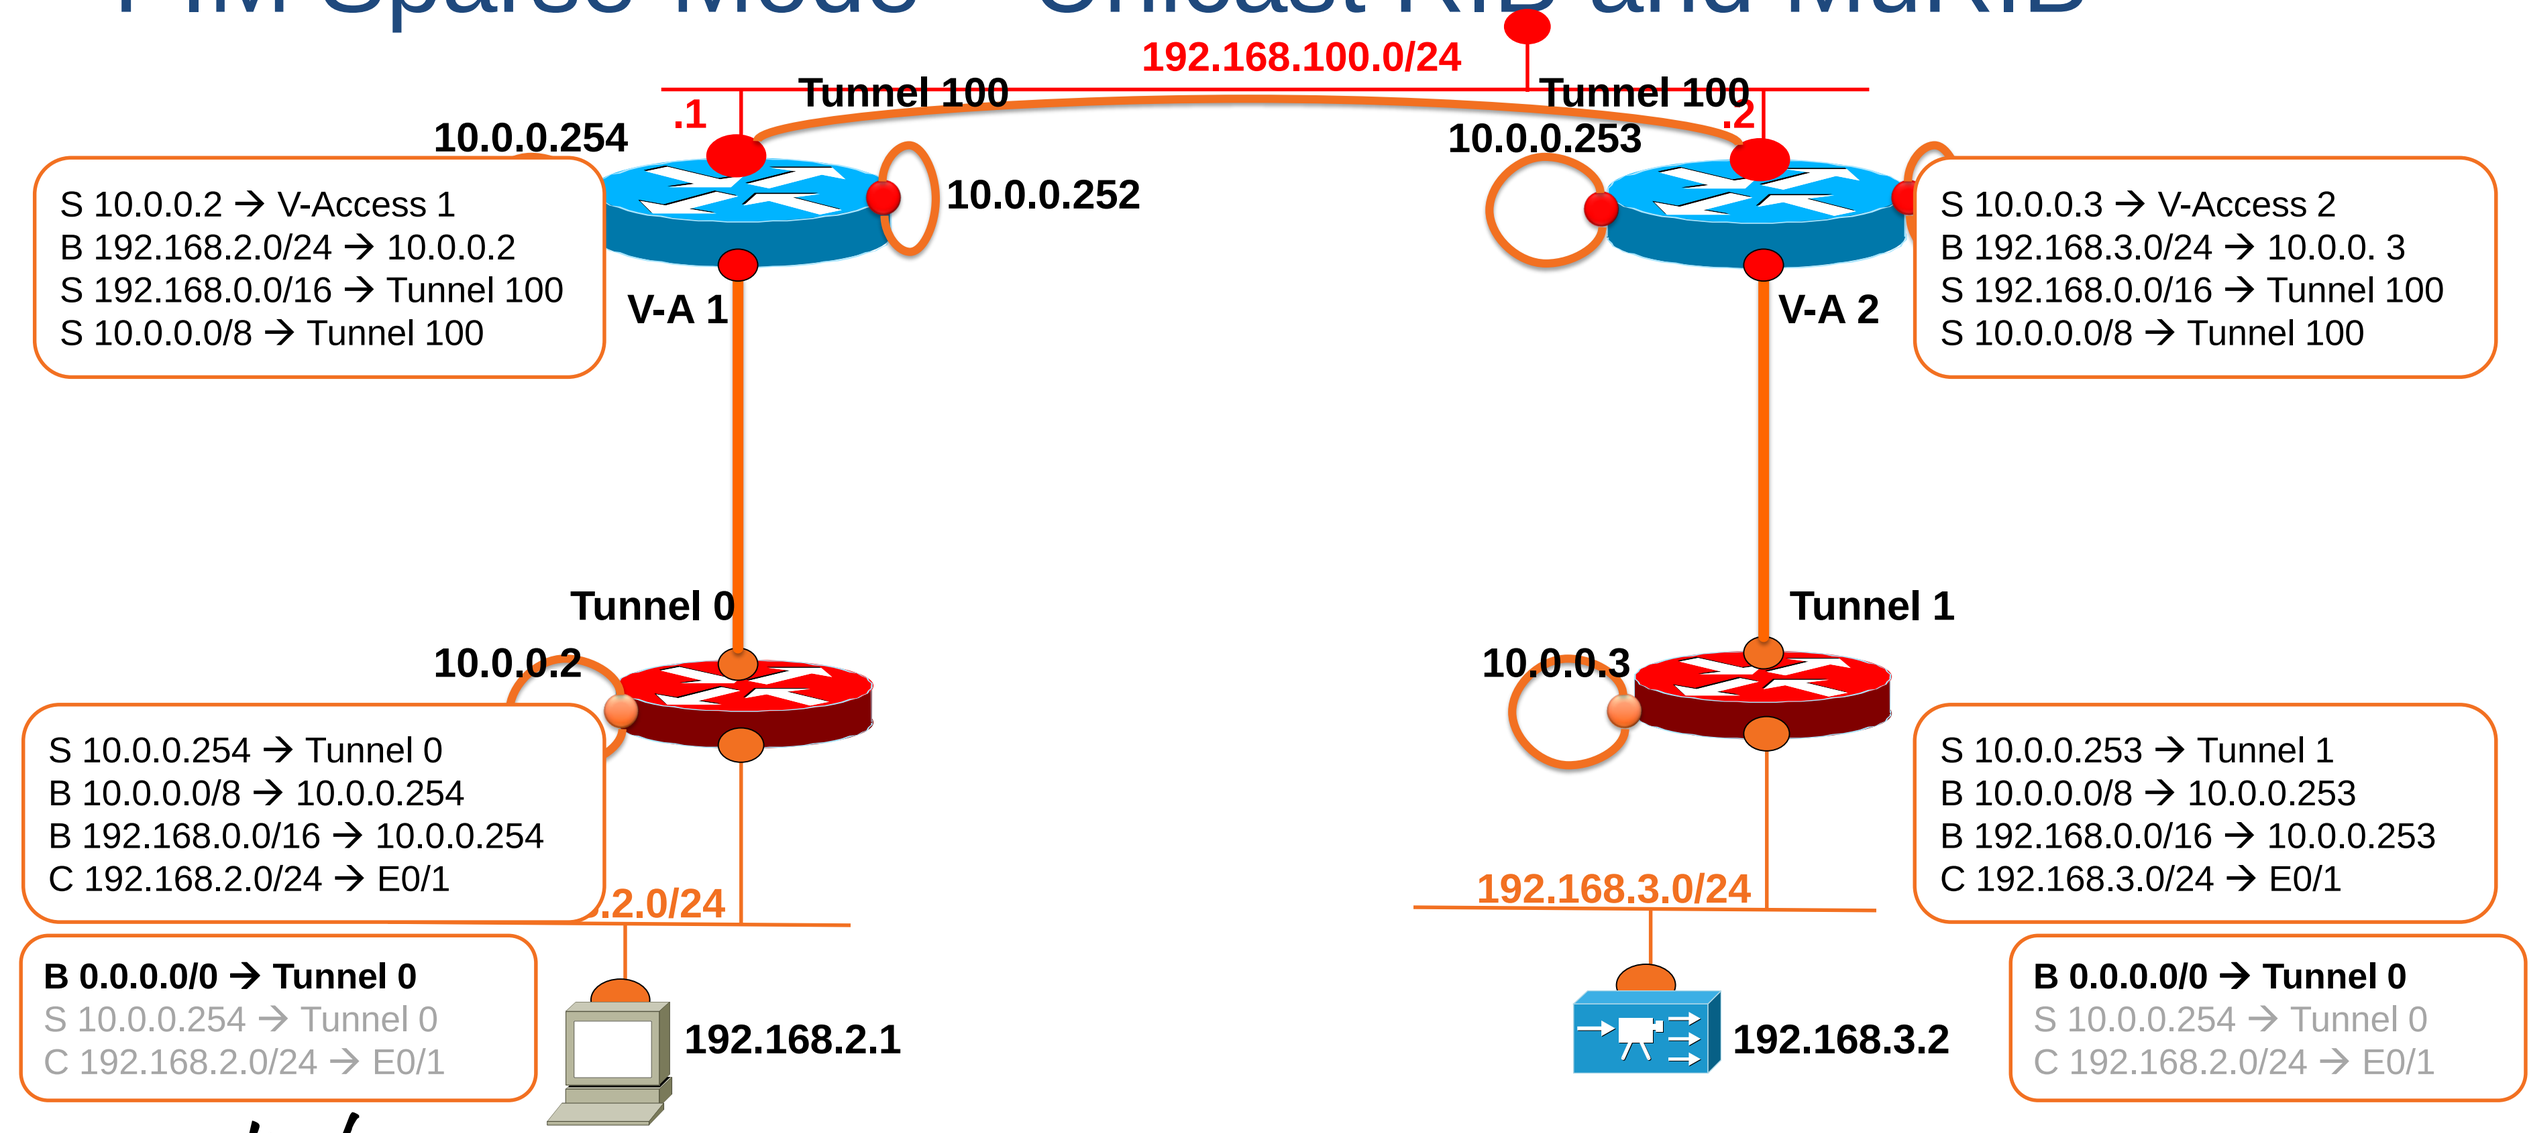

PIM Sparse-Mode – Mcast Flows End-to-End

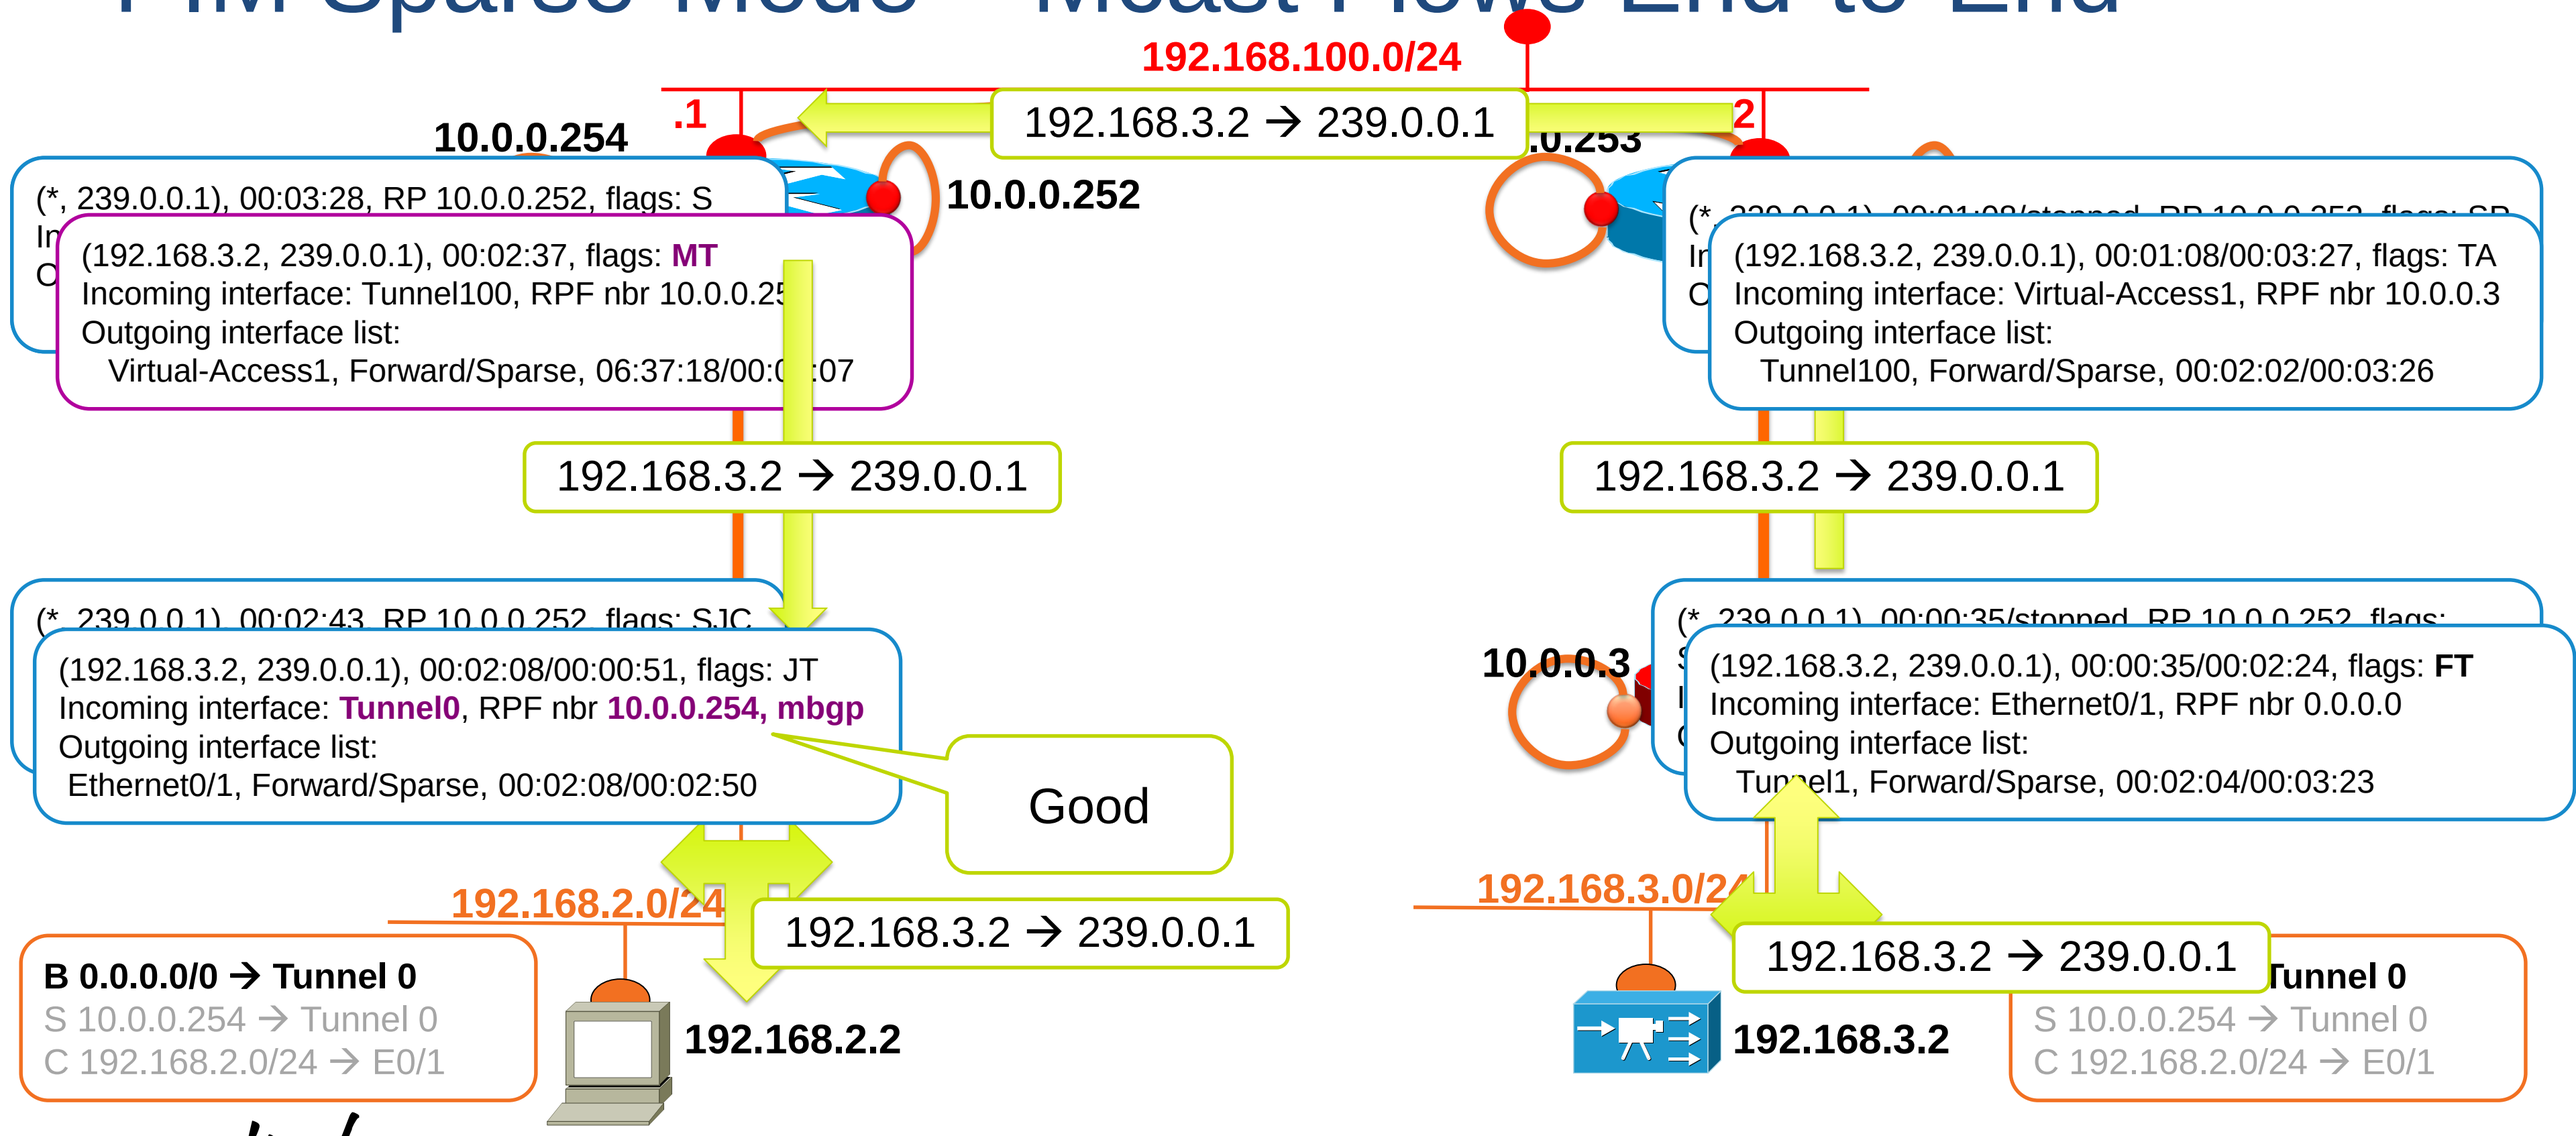

(S,G) Entry Now Usable w/ SAFI 129 & 2

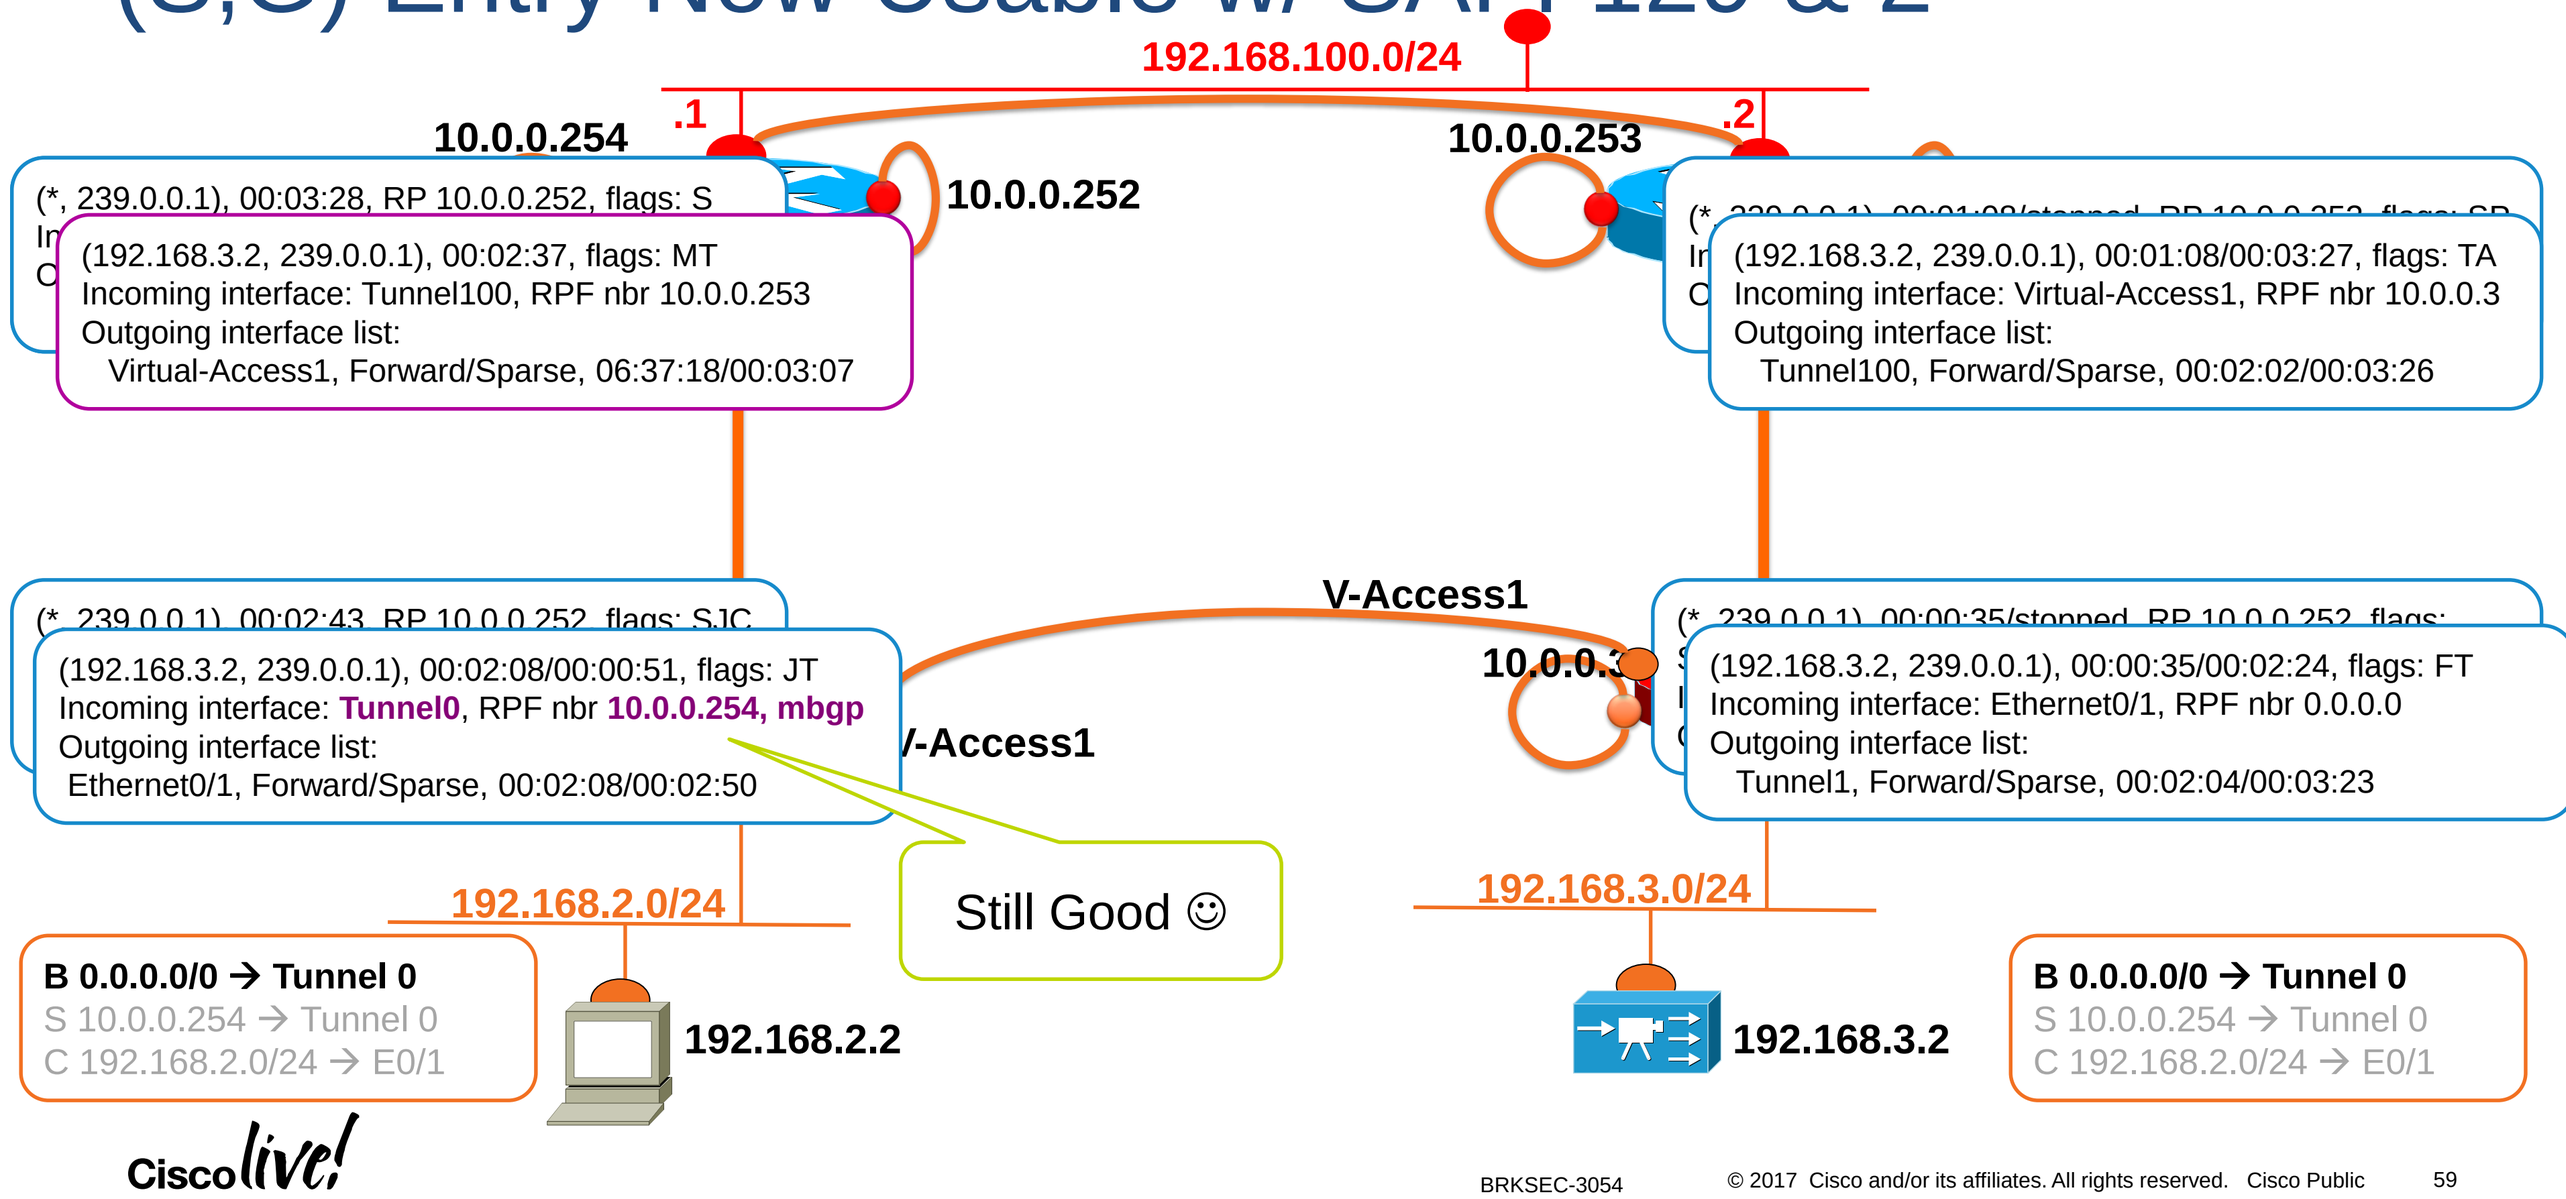

Flex & Sparse-Mode – Hub Configuration

Hub 1 – Flex, Multicast and Interfaces

```
crypto ikev2 profile default
 match identity remote fqdn domain cisco.com
 identity local fqdn Hub1.cisco.com
 authentication remote rsa-sig
 authentication local rsa-sig
 pki trustpoint TP
 dpd 10 2 on-demand
 aaa authorization group cert list default default
 virtual-template 1

interface Virtual-Template1 type tunnel
 vrf forwarding RED
 ip unnumbered Loopback0
 ip access-group AllowMyBGP in
 ip pim sparse-mode
 ip nhrp network-id 1
 ip nhrp redirect
 tunnel protection ipsec profile default
```

Activate Sparse-Mode

```
ip pim vrf RED rp-address 10.0.0.252

ip msdp vrf RED peer 10.0.0.253 connect-source Loopback0
ip msdp vrf RED cache-sa-state
```

Rendez-Vous Point Definition

MSDP for RP synchronization

```
vrf definition RED
 rd 2:2
 address-family ipv4
 address-family ipv6

interface Loopback0
 vrf forwarding RED
 ip address 10.0.0.254 255.255.255.255
!
interface Loopback1
 vrf forwarding RED
 ip address 10.0.0.252 255.255.255.255
!
interface Tunnel100
 vrf forwarding RED
 ip unnumbered Loopback0
 ip pim sparse-mode
 ip nhrp network-id 1
 ip nhrp redirect
 tunnel source Ethernet0/1
 tunnel destination 192.168.100.2
 tunnel vrf ivrf
```

Anycast Rendez-Vous Point Loopback

Activate Sparse-Mode

Flex & Sparse-Mode – Hub Configuration

Hubs Common BGP Configuration – With SAFI 2 and SAFI 129

```
ip route vrf RED 10.0.0.0 255.0.0.0 Tunnel100 tag 2
ip route vrf RED 192.168.0.0 255.255.0.0 Tunnel100 tag
2

router bgp 1
 bgp log-neighbor-changes
 bgp listen range 10.0.0.0/24 peer-group Flex
 !
 address-family ipv4 vrf RED
  redistribute static route-map rm
  neighbor Flex peer-group
  neighbor Flex remote-as 1
  neighbor Flex next-hop-self all
 exit-address-family
!
address-family ipv4 multicast vrf ivrf
  neighbor Flex peer-group
  neighbor Flex remote-as 1
  neighbor Flex default-originate
 exit-address-family

route-map rm permit 10
 match tag 2
```

Vanilla BGP Configuration

Use mBGP to advertise a prefix in the MuRIB only

Default Originate advertises 0.0.0.0/0

Flex & Sparse-Mode – Spoke Configuration

Client/Receiver and Source Spoke

```
crypto ikev2 profile default
 match identity remote fqdn domain cisco.com
 identity local fqdn Spoke2.cisco.com
 authentication remote rsa-sig
 authentication local rsa-sig
 pki trustpoint TP
 dpd 10 2 on-demand
 aaa authorization group cert list default default
 virtual-template 1
```

```
ip pim rp-address 10.0.0.252

router bgp 1
 bgp log-neighbor-changes
 neighbor 10.0.0.253 remote-as 1
 neighbor 10.0.0.254 remote-as 1
address-family ipv4
  network 192.168.2.0
  neighbor 10.0.0.253 activate
  neighbor 10.0.0.254 activate
  maximum-paths ibgp 2
address-family ipv4 multicast
  neighbor 10.0.0.253 activate
  neighbor 10.0.0.254 activate
```

Rendez-Vous Point Definition

Receive SAFI2 & SAFI129

```
interface Loopback0
 ip address 10.0.0.2 255.255.255.255

interface Tunnel0
 ip unnumbered Loopback0
 ip pim sparse-mode
 ip nhrp network-id 1
 ip nhrp shortcut virtual-template 1
 tunnel source Ethernet0/0
 tunnel destination 172.16.1.1
 tunnel protection ipsec profile default
!
interface Tunnel1
 ip unnumbered Loopback0
 ip pim sparse-mode
 ip nhrp network-id 1
 ip nhrp shortcut virtual-template 1
 tunnel source Ethernet0/0
 tunnel destination 172.16.4.1
 tunnel protection ipsec profile default

interface Virtual-Template1 type tunnel
 ip unnumbered Loopback0
 ip nhrp network-id 1
 ip nhrp shortcut virtual-template 1
 tunnel protection ipsec profile default
```

Activate PIM on Tunnel Interfaces

Activate PIM on Tunnel Interfaces

No PIM on V-Template!

Routing Based Resiliency

FlexVPN Backup
Routing Based Multi-Hub Resiliency (1)

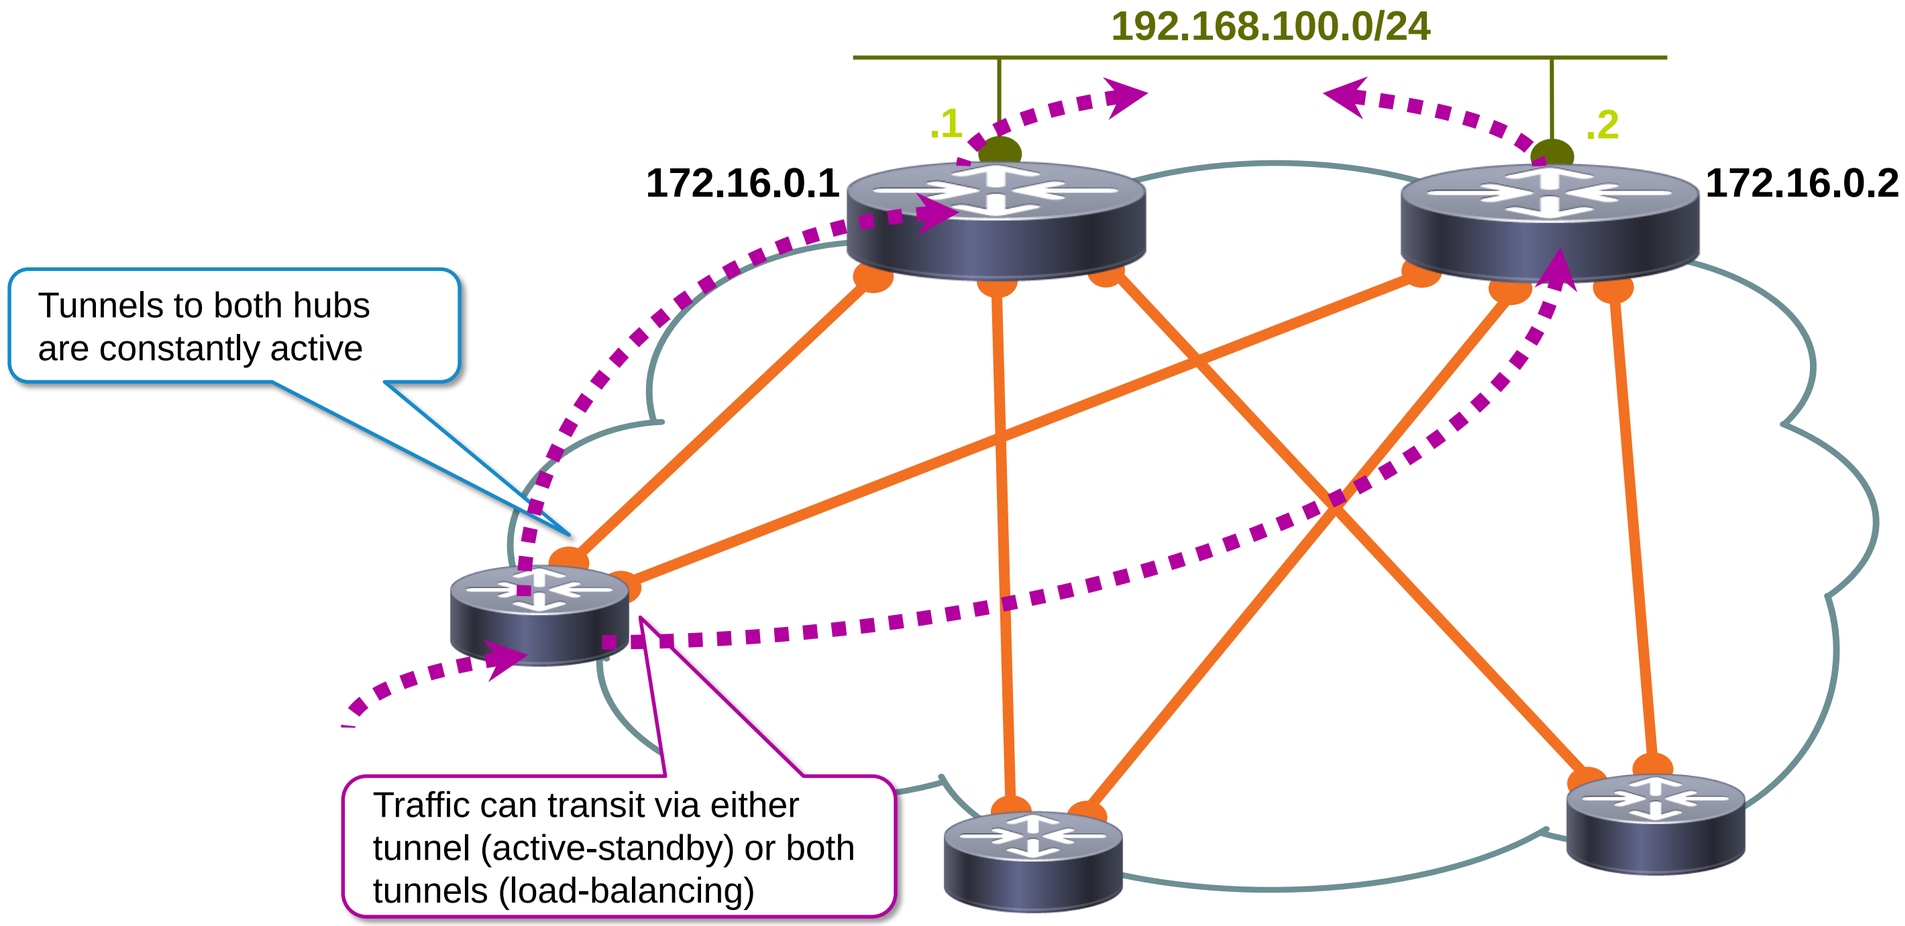

FlexVPN Backup
Routing Based Multi-Hub Resiliency (2)

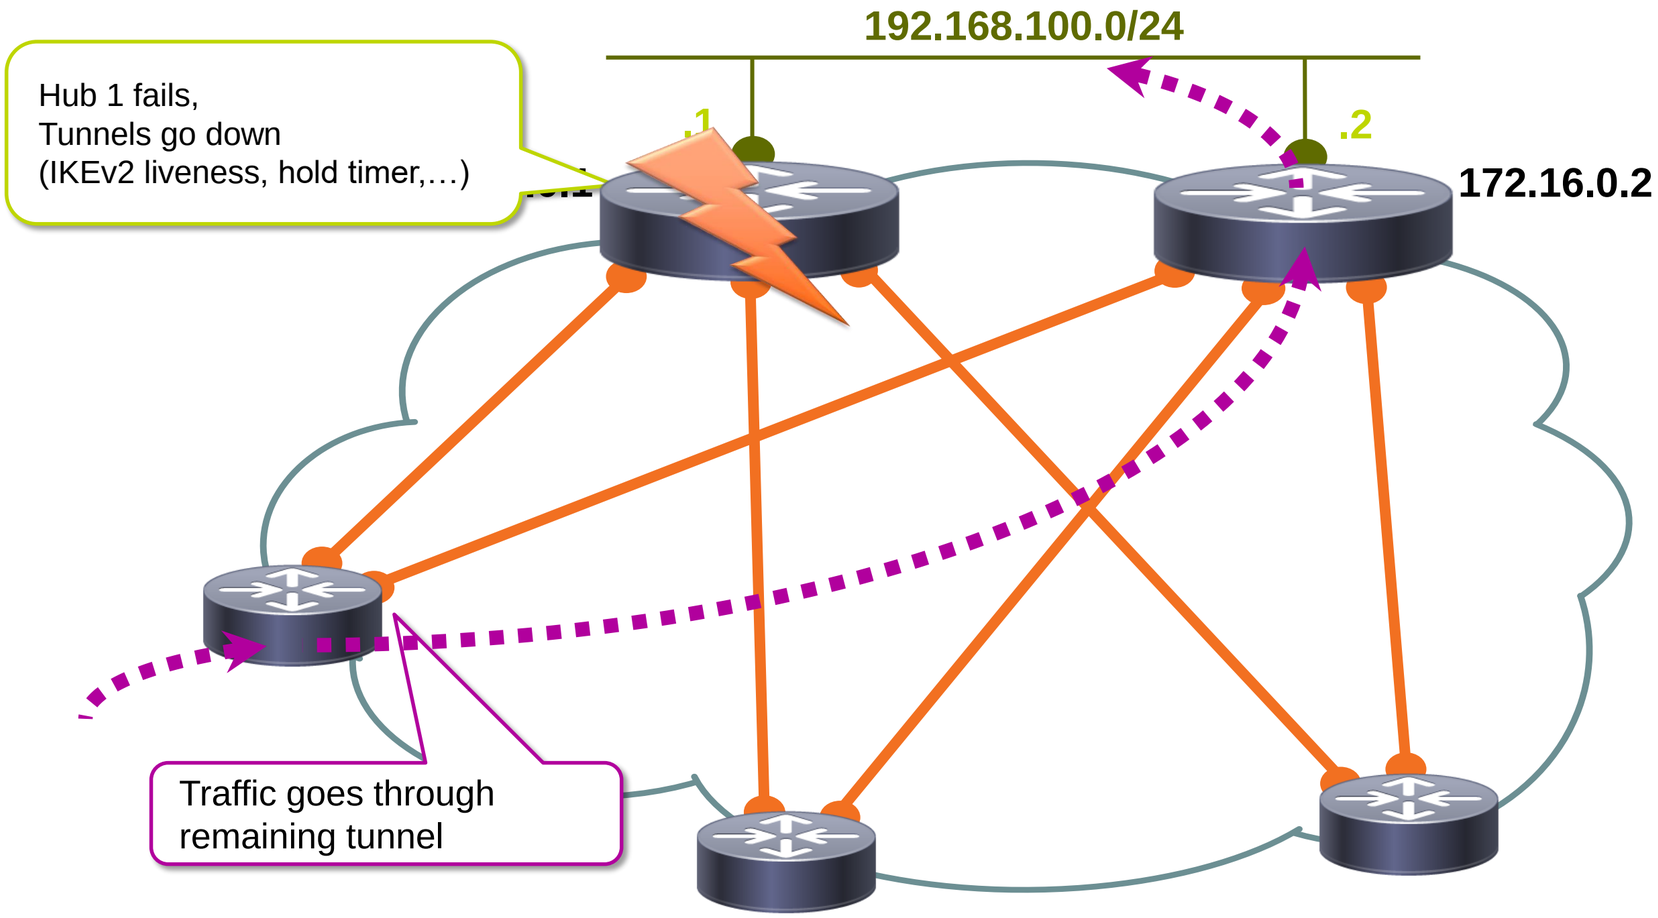

Routing Based Resiliency Probing over WAN

A simple setup…

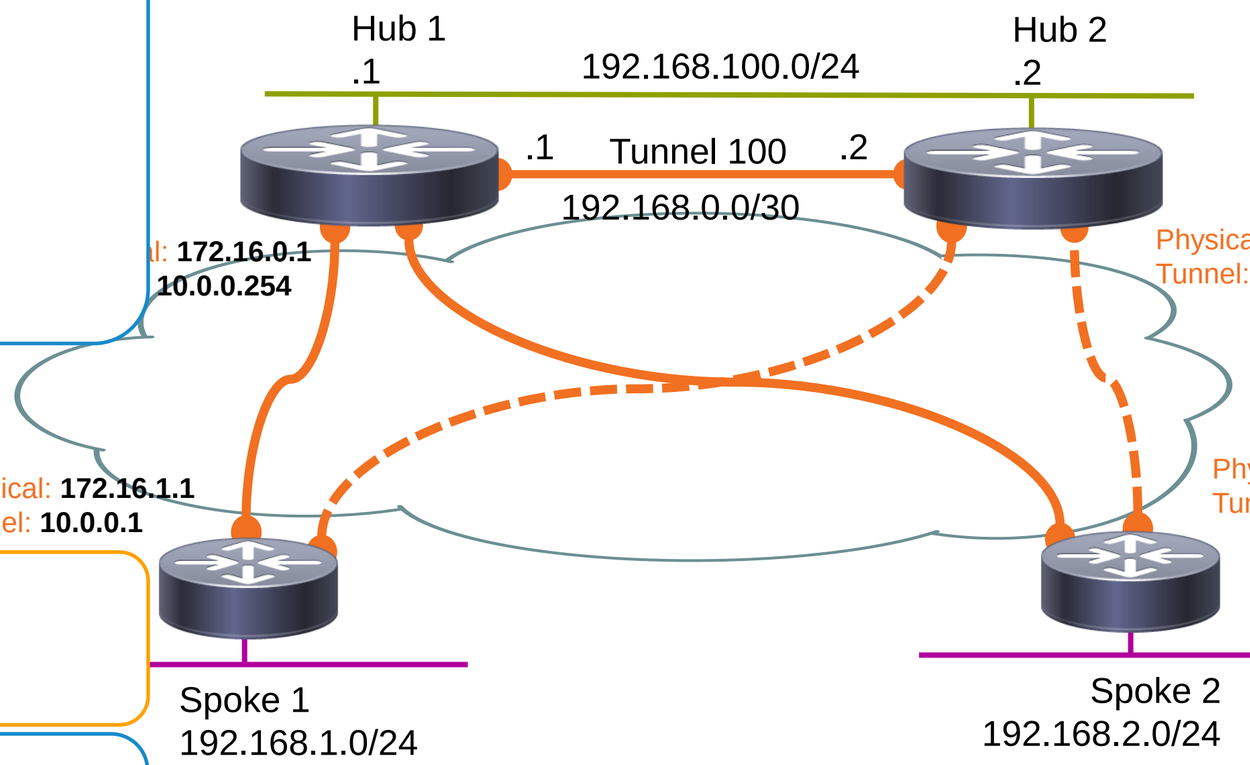

Routing Table

C 10.0.0.254 → Loopback0
C 192.168.100.0/24 → Eth0
S 192.168.0.0/16 → Null0 tag 2
S 10.0.0.0/8 → Tunnel100
S 1.0.0.1/32 → Null0 tag 2
B 1.0.0.2/32 → Tunnel100
S 10.0.0.1 → V-Access1
S 10.0.0.2 → V-Access2
B 192.168.1.0/24 → 10.0.0.1
B 192.168.2.0/24 → 10.0.0.2

Routing Table

C 10.0.0.253 → Loopback0
C 192.168.100.0/24 → Eth0
S 192.168.0.0/16 → Null0 tag 2
S 10.0.0.0/8 → Tunnel100
S 1.0.0.2/32 → Null0 tag 2
B 1.0.0.1/32 → Tunnel100
S 10.0.0.1 → V-Access1
S 10.0.0.2 → V-Access2
B 192.168.1.0/24 → 10.0.0.1
B 192.168.2.0/24 → 10.0.0.2

NHRP Table

-

NHRP Table

-

Routing Table

C 192.168.1.0/24 → Eth0
C 10.0.0.1 → Tunnel0
S 0.0.0.0/0 → Dialer0
S 10.0.0.254/32 → Tunnel0
S 10.0.0.253/32 → Tunnel1
B 192.168.0.0/16 → 10.0.0.254
[B 192.168.0.0/16 → 10.0.0.253]

Routing Table

C 192.168.2.0/24 → Eth0
C 10.0.0.2 → Tunnel1
S 0.0.0.0/0 → Dialer0
S 10.0.0.254/32 → Tunnel0
S 10.0.0.253/32 → Tunnel1
B 192.168.0.0/16 → 10.0.0.254
[B 192.168.0.0/16 → 10.0.0.253]

Method #1: Faster Hello's

Hub Configuration

```
router bgp 1
  bgp log-neighbor-changes
  bgp listen range 10.0.0.0/8 peer-group SPOKES
  neighbor SPOKES peer-group
  neighbor SPOKES remote-as 1
  neighbor SPOKES timers 1 3
  address-family ipv4
    neighbor SPOKES activate
```

Spoke Configuration

```
router bgp 1
  bgp log-neighbor-changes
  neighbor 10.0.0.254 remote-as 1
  neighbor 10.0.0.254 timers 1 3
  neighbor 10.0.0.253 remote-as 1
  neighbor 10.0.0.253 timers 1 3
```

BGP can go as fast as 1 second hello's with a 3 seconds Hold Timer ➔ Failover in 3 seconds

Monitor IOS CPU level – expect about 10% CPU background load at 500 spokes (RP2)

Convergence (massive reconnect) may be affected by process starvation. Test –test –test –test.

Method #2: BFD between hub and spokes

Hub Configuration

```
bfd map ipv4  10.0.0.0/8  10.0.0.0/8 mh1
bfd-template multi-hop mh1
  interval min-tx 200 min-rx 200 multiplier 3

router bgp 1
  bgp log-neighbor-changes
  bgp listen range 10.0.0.0/8 peer-group SPOKES
  neighbor SPOKES peer-group
  neighbor SPOKES remote-as 1
  neighbor SPOKES ebgp-multihop 2
  neighbor SPOKES fall-over bfd multi-hop
  address-family ipv4
  neighbor SPOKES activate
```

Spoke Configuration

```
bfd map ipv4  10.0.0.0/8  10.0.0.0/8 mh1
bfd-template multi-hop mh1
  interval min-tx 200 min-rx 200 multiplier 3

router bgp 1
  bgp log-neighbor-changes
  neighbor 10.0.0.254 remote-as 1
  neighbor 10.0.0.254 fall-over bfd multi-hop
  neighbor 10.0.0.253 remote-as 1
  neighbor 10.0.0.253 fall-over bfd multi-hop
```

On ASR1K, ESP CPU will offload BFD; IOS unaffected
problem moved, not fully solved

Microbursts can cause false positives; very hard to monitor

Fast probing – Mind the microbursts (I)

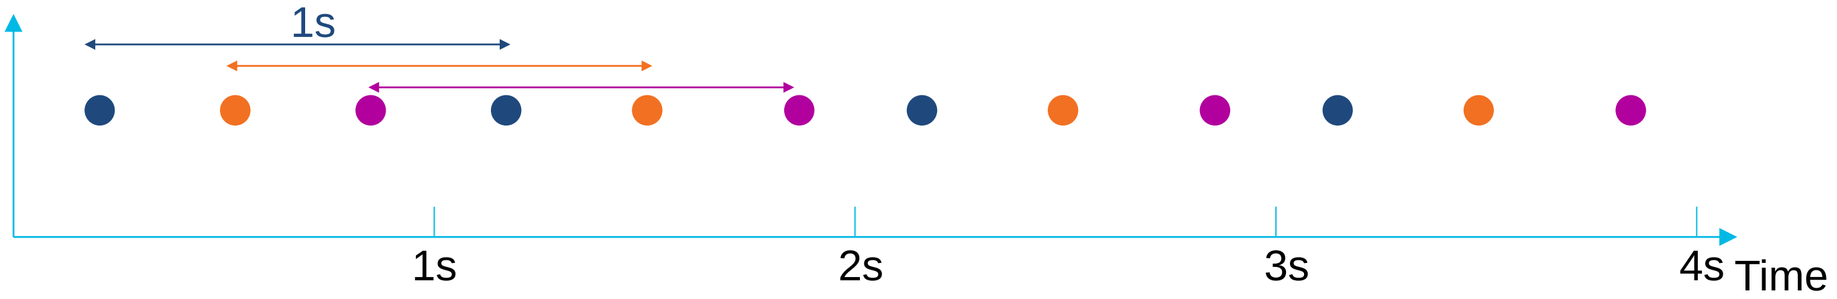

Even probe distribution

Fast probing – Mind the microbursts (II)

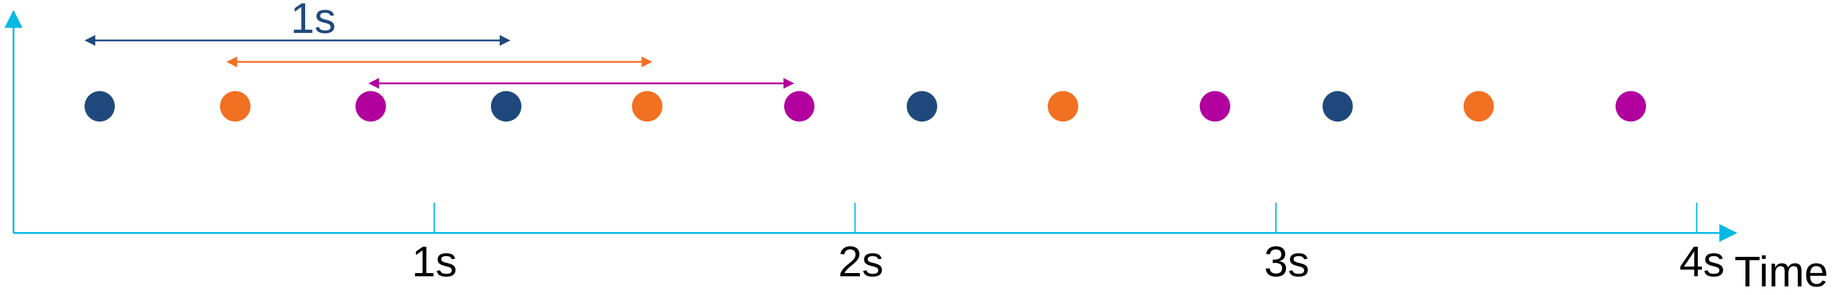

Even probe distribution

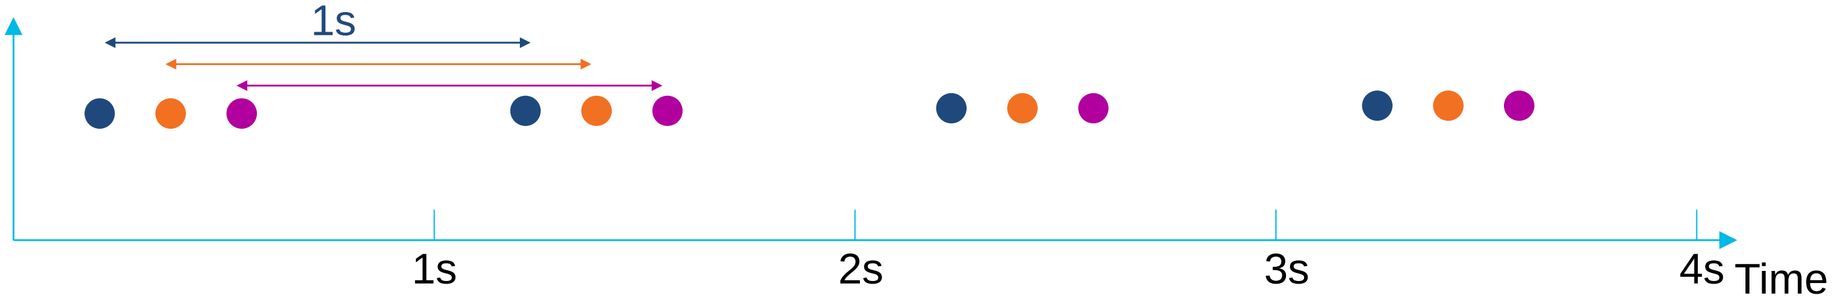

Spokes clock drifts affect distribution

Not so even anymore

Fast probing – Mind the microbursts (III)

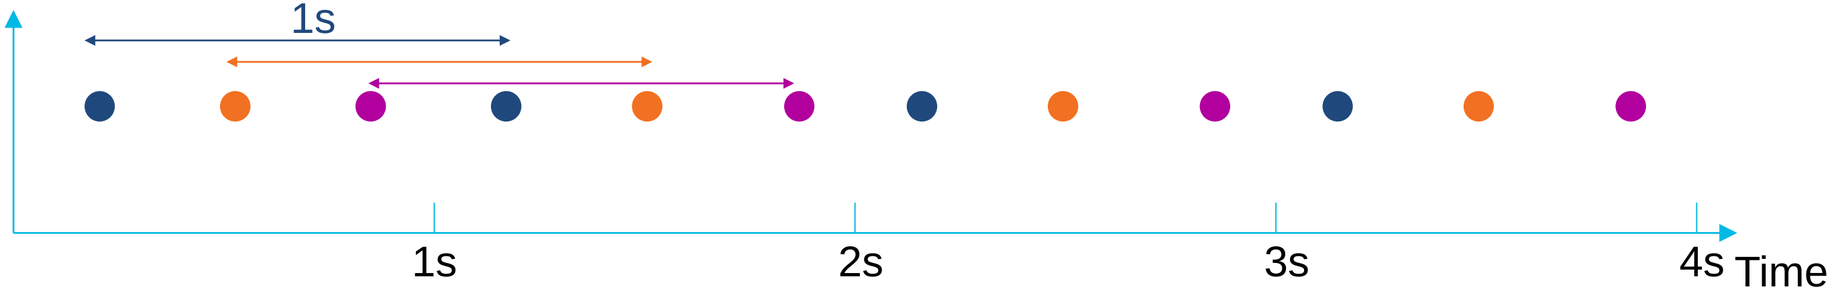

Even probe distribution

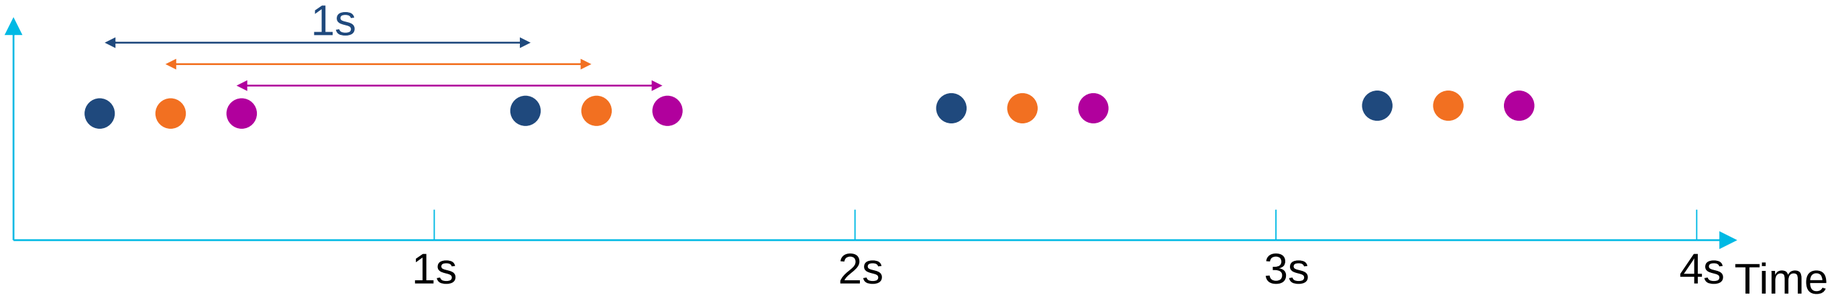

Spokes clock drifts affect distribution

Not so even anymore

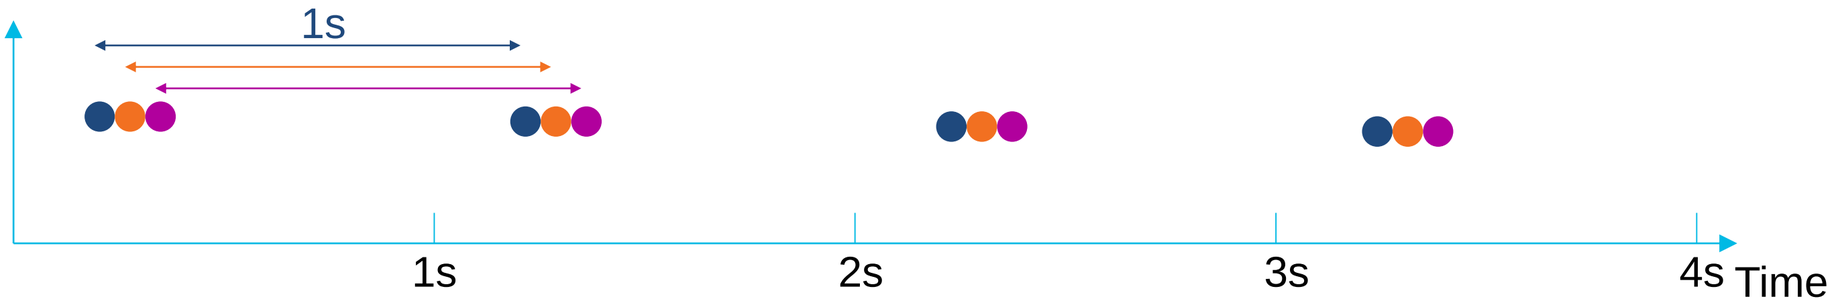

Further clock drift yields very uneven distribution and microbursts

Routing Based Resiliency WAN Friendly

Inter-hub BGP – BFD keepalives

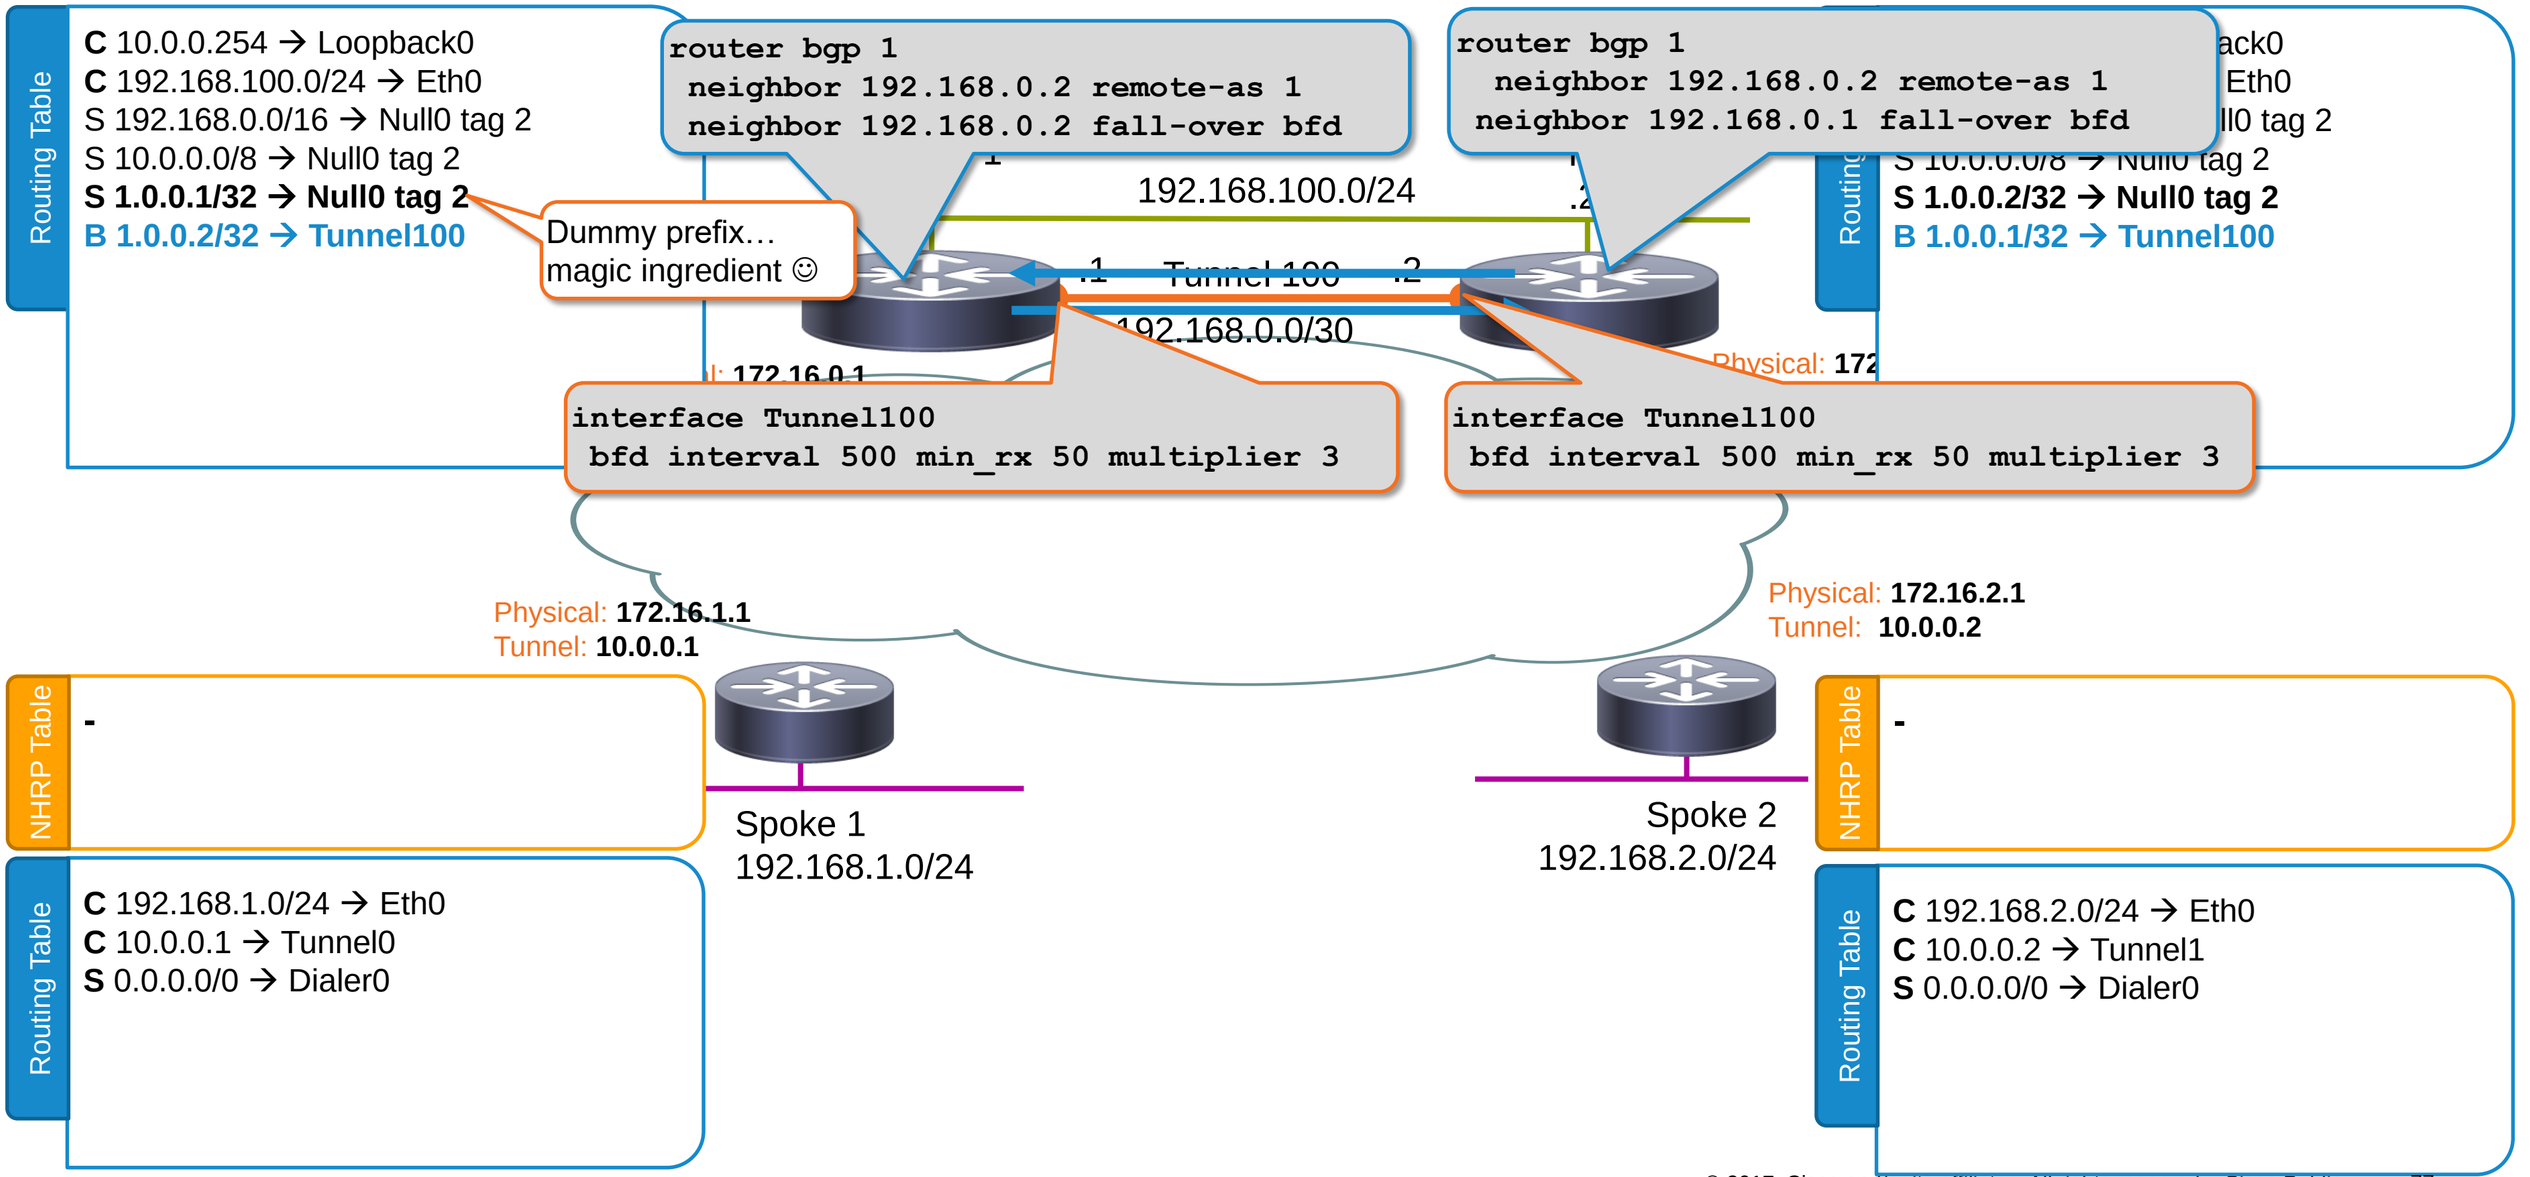

Spokes Connect – Next-Hop w/ High Distance

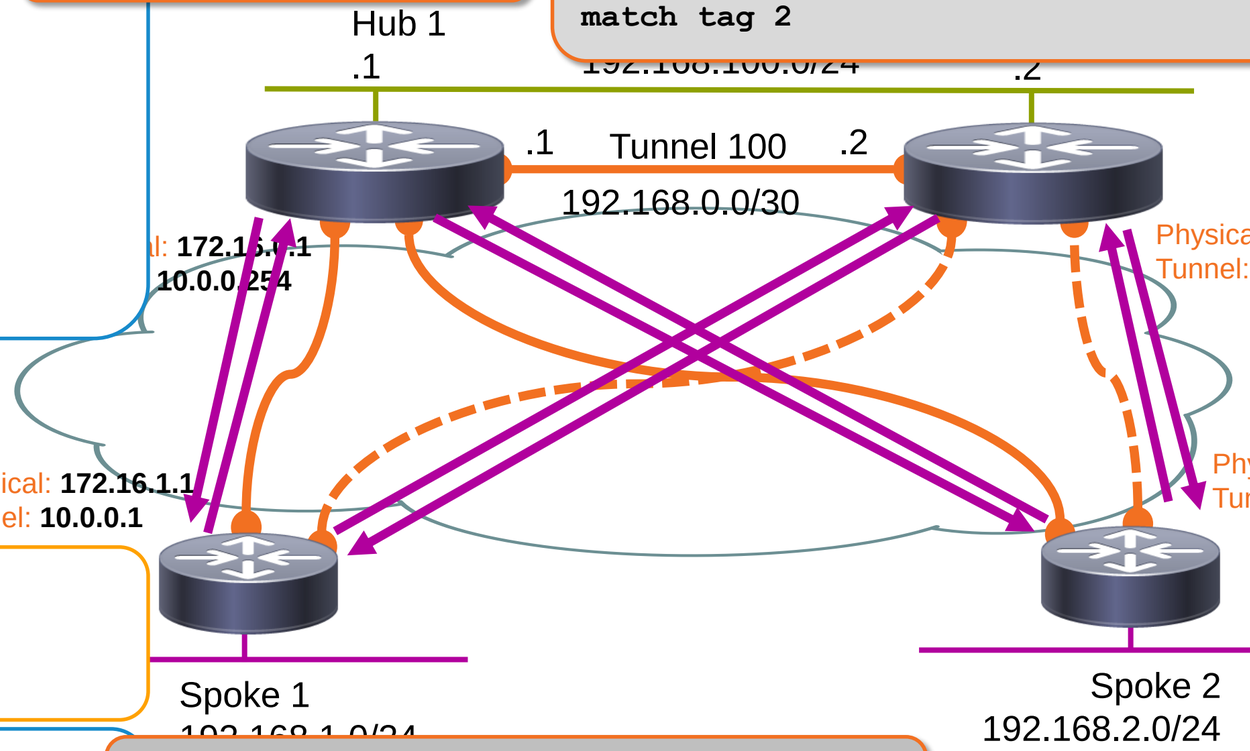

Routing Table

C 10.0.0.254 → Loopback0
C 192.168.100.0/24 → Eth0
S 192.168.0.0/16 → Null0 tag 2
S 10.0.0.0/8 → Null0 tag 2
S 1.0.0.1/32 → Null0 tag 2
B 1.0.0.2/32 → Tunnel100
S 10.0.0.1 → V-Access1
S 10.0.0.2 → V-Access2
B 192.168.1.0/24 → 10.0.0.1
B 192.168.2.0/24 → 10.0.0.2

tag 2 routes got redistributed in BGP and advertised to the spokes

```
router bgp 1
 redistribute static route-map TAG2
route-map TAG2
 match tag 2
```

Routing Table

C 10.0.0.253 → Loopback0
C 192.168.100.0/24 → Eth0
S 192.168.0.0/16 → Null0 tag 2
S 10.0.0.0/8 → Null0 tag 2
S 1.0.0.2/32 → Null0 tag 2
B 1.0.0.1/32 → Tunnel100
S 10.0.0.1 → V-Access1
S 10.0.0.2 → V-Access2
B 192.168.1.0/24 → 10.0.0.1
B 192.168.2.0/24 → 10.0.0.2

NHRP Table

-

NHRP Table

-

Routing Table

C 192.168.1.0/24 → Eth0
C 10.0.0.1 → Tunnel0
S 0.0.0.0/0 → Dialer0
S 10.0.0.254/32 → Tunnel0 dist. 210
S 10.0.0.253/32 → Tunnel1 dist. 210
B 192.168.0.0/16 → 10.0.0.254 dist. 200

```
crypto ikev2 authorization policy default
 route accept any distance 210
```

Routing Table

C 192.168.2.0/24 → Eth0
C 10.0.0.2 → Tunnel1
S 0.0.0.0/0 → Dialer0
S 10.0.0.254/32 → Tunnel0 dist. 210
S 10.0.0.253/32 → Tunnel1 dist. 210
B 192.168.0.0/16 → 10.0.0.254 dist. 200

Traffic Flows – recursive routing applies

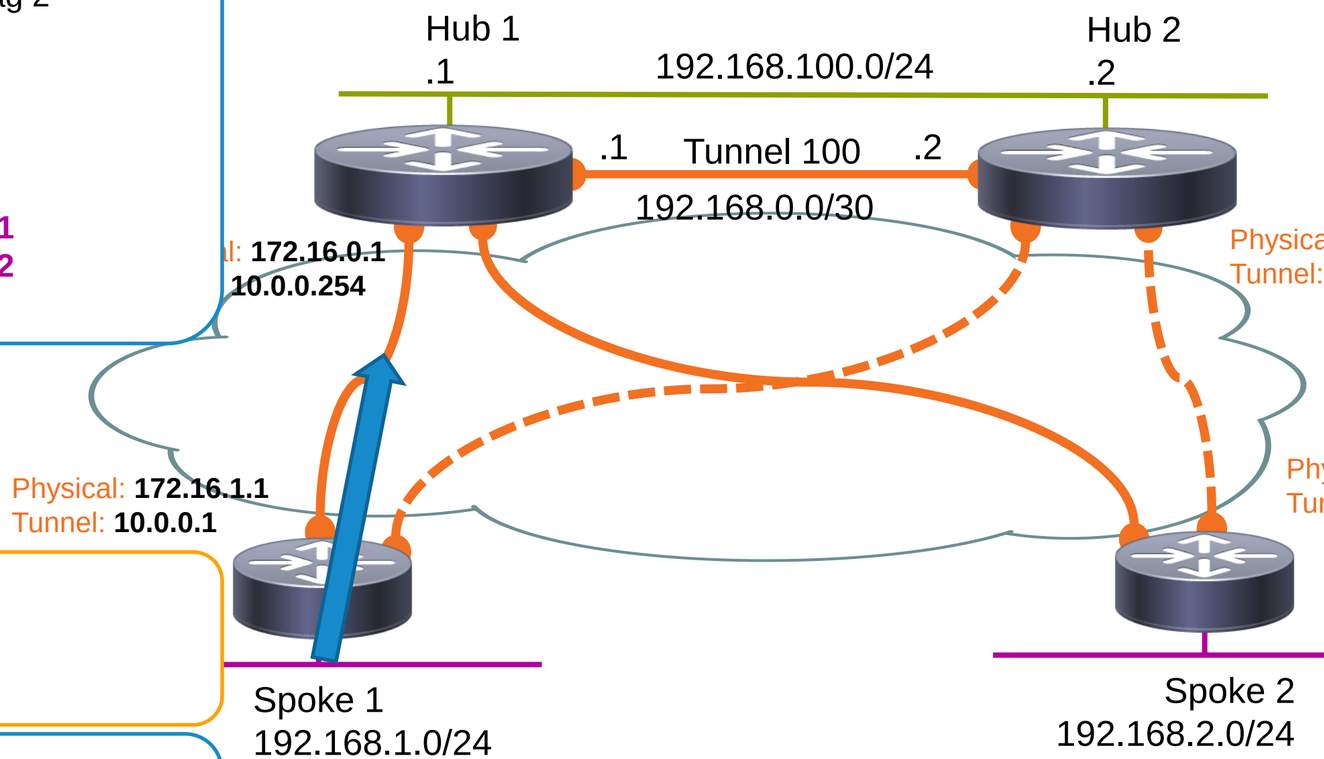

Routing Table

C 10.0.0.254 → Loopback0
C 192.168.100.0/24 → Eth0
S 192.168.0.0/16 → Null0 tag 2
S 10.0.0.0/8 → Tunnel100
S 1.0.0.1/32 → Null0 tag 2
B 1.0.0.2/32 → Tunnel100
S 10.0.0.1 → V-Access1
S 10.0.0.2 → V-Access2
B 192.168.1.0/24 → 10.0.0.1
B 192.168.2.0/24 → 10.0.0.2

Routing Table

C 10.0.0.253 → Loopback0
C 192.168.100.0/24 → Eth0
S 192.168.0.0/16 → Null0 tag 2
S 10.0.0.0/8 → Tunnel100
S 1.0.0.2/32 → Null0 tag 2
B 1.0.0.1/32 → Tunnel100
S 10.0.0.1 → V-Access1
S 10.0.0.2 → V-Access2
B 192.168.1.0/24 → 10.0.0.1
B 192.168.2.0/24 → 10.0.0.2

NHRP Table

-

NHRP Table

-

Routing Table

C 192.168.1.0/24 → Eth0
C 10.0.0.1 → Tunnel0
S 0.0.0.0/0 → Dialer0
S 10.0.0.254/32 → Tunnel0 dist. 210
S 10.0.0.253/32 → Tunnel1 dist. 210
B 192.168.0.0/16 → 10.0.0.254 dist. 200

Routing Table

C 192.168.2.0/24 → Eth0
C 10.0.0.2 → Tunnel1
S 0.0.0.0/0 → Dialer0
S 10.0.0.254/32 → Tunnel0 dist. 210
S 10.0.0.253/32 → Tunnel1 dist. 210
B 192.168.0.0/16 → 10.0.0.254 dist. 200

Say Hub 1 Crashed…

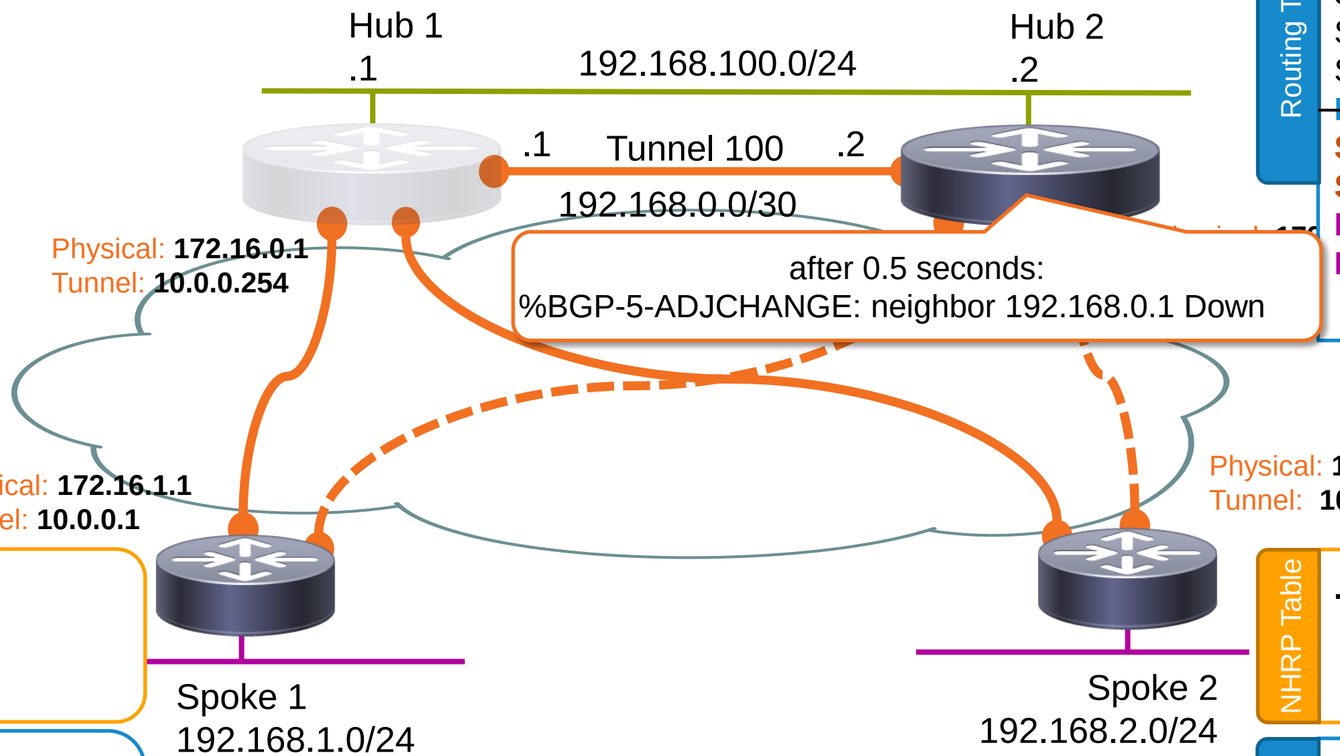

Routing Table

C 10.0.0.253 → Loopback0
C 192.168.100.0/24 → Eth0
S 192.168.0.0/16 → Tunnel100 tag 2
S 10.0.0.0/8 → Null0 tag 2
S 1.0.0.2/32 → Null0 tag 2
~~**B 1.0.0.1/32 → Tunnel100**~~
S 10.0.0.1 → V-Access1
S 10.0.0.2 → V-Access2
B 192.168.1.0/24 → 10.0.0.1
B 192.168.2.0/24 → 10.0.0.2

NHRP Table

-

Routing Table

C 192.168.1.0/24 → Eth0
C 10.0.0.1 → Tunnel0
S 0.0.0.0/0 → Dialer0
S 10.0.0.254/32 → Tunnel0 dist. 210
S 10.0.0.253/32 → Tunnel1 dist. 210
B 192.168.0.0/16 → 10.0.0.254 dist. 200

NHRP Table

-

Routing Table

C 192.168.2.0/24 → Eth0
C 10.0.0.2 → Tunnel1
S 0.0.0.0/0 → Dialer0
S 10.0.0.254/32 → Tunnel0 dist. 210
S 10.0.0.253/32 → Tunnel1 dist. 210
B 192.168.0.0/16 → 10.0.0.254 dist. 200

We have achieved High-Availability

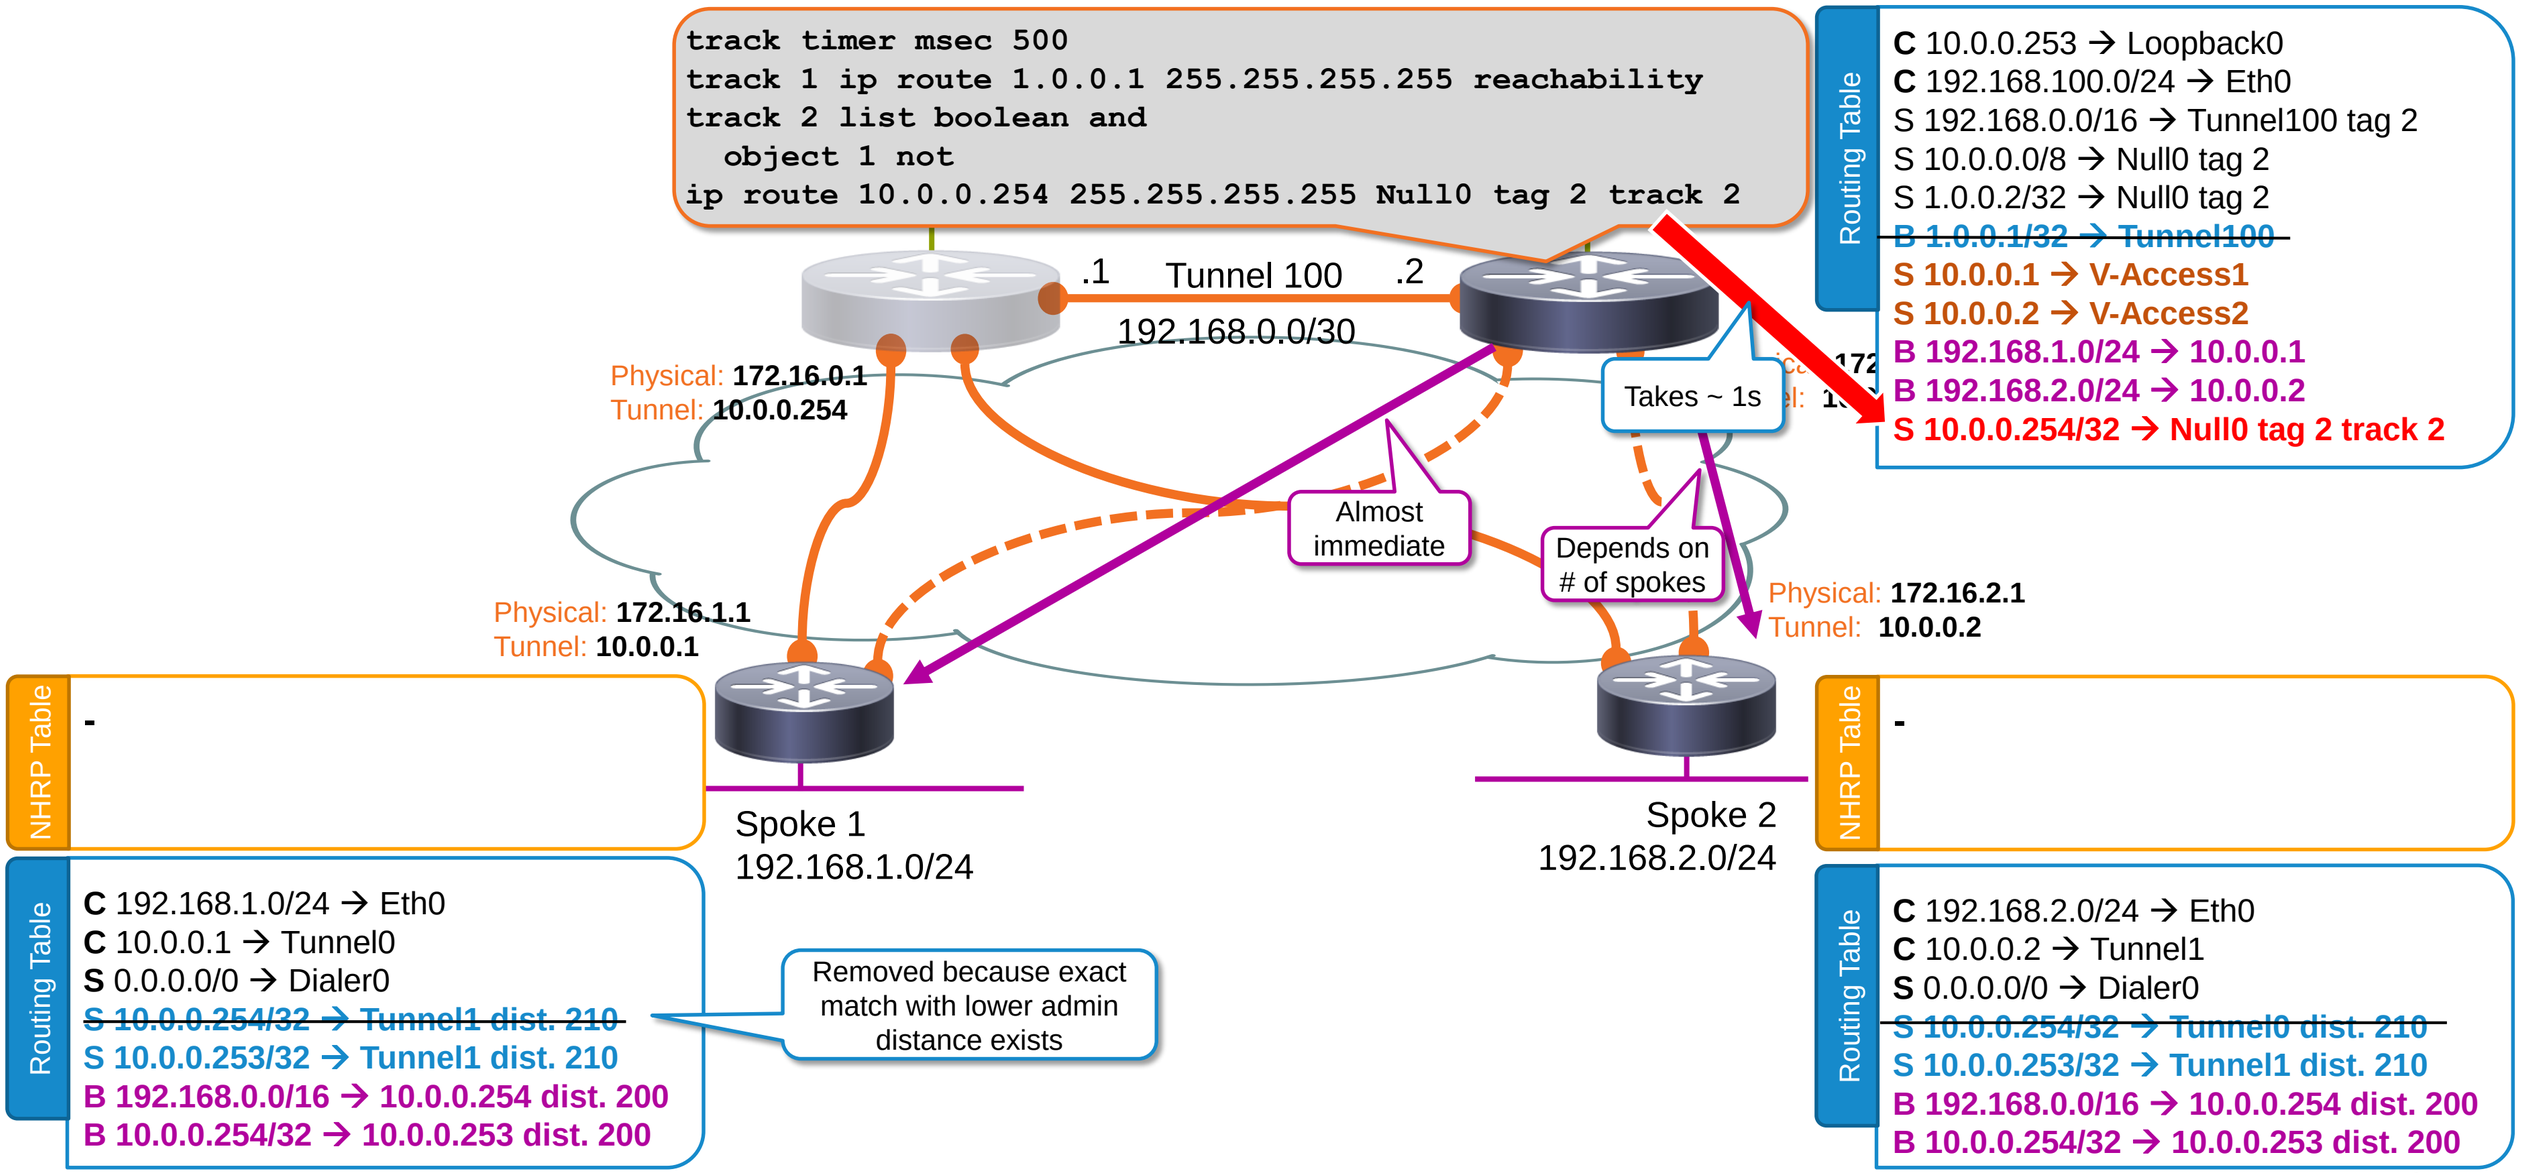

We have achieved High-Availability

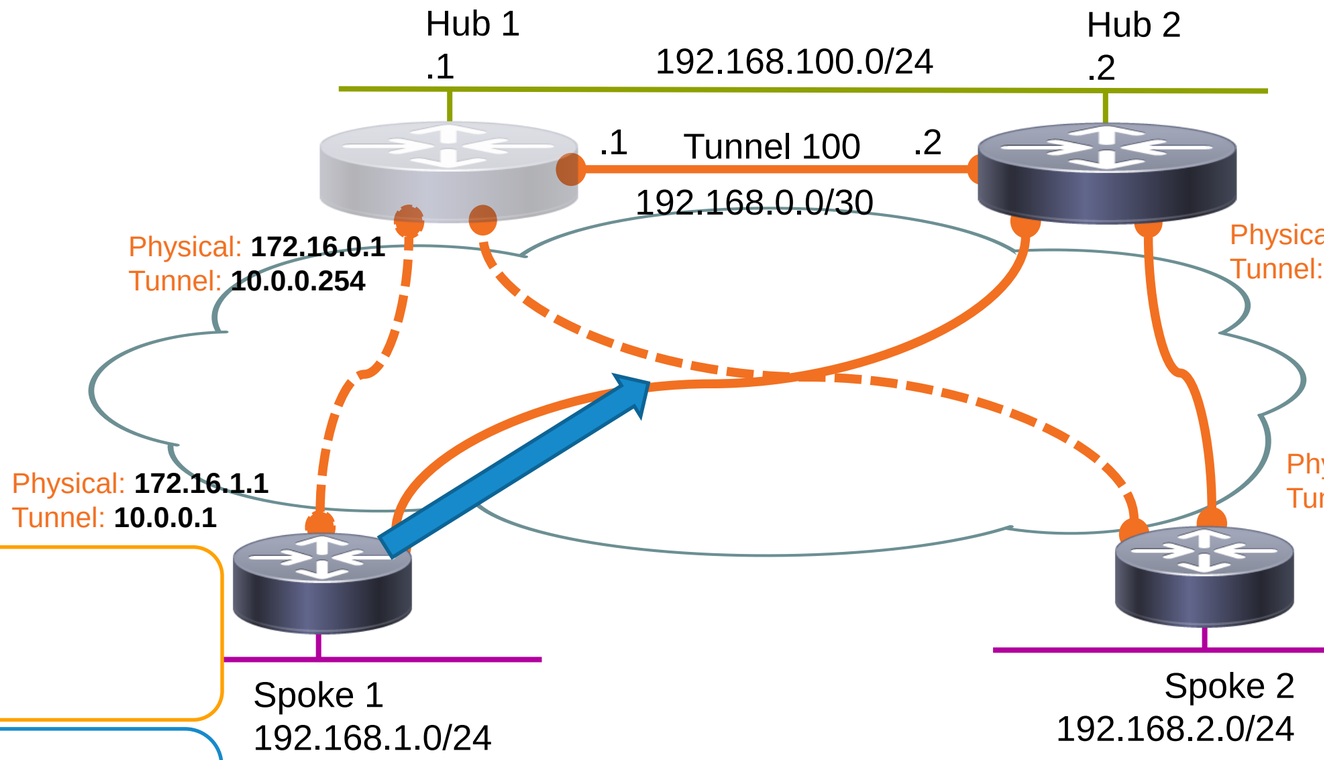

Routing Table

C 10.0.0.253 → Loopback0
C 192.168.100.0/24 → Eth0
S 192.168.0.0/16 → Tunnel100 tag 2
S 10.0.0.0/8 → Null 0 tag 2
S 1.0.0.2/32 → Null0
~~**B 1.0.0.1/32 → Tunnel100**~~
S 10.0.0.1 → V-Access1
S 10.0.0.2 → V-Access2
B 192.168.1.0/24 → 10.0.0.1
B 192.168.2.0/24 → 10.0.0.2
S 10.0.0.254/32 → Null0 tag 2 track 2

NHRP Table

-

Routing Table

C 192.168.1.0/24 → Eth0
C 10.0.0.1 → Tunnel0
S 0.0.0.0/0 → Dialer0

S 10.0.0.253/32 → Tunnel1 dist. 210
B 192.168.0.0/16 → 10.0.0.254 dist. 200
B 10.0.0.254/32 → 10.0.0.253 dist. 200

NHRP Table

-

Routing Table

C 192.168.2.0/24 → Eth0
C 10.0.0.2 → Tunnel1
S 0.0.0.0/0 → Dialer0

S 10.0.0.253/32 → Tunnel1 dist. 210
B 192.168.0.0/16 → 10.0.0.254 dist. 200
B 10.0.0.254/32 → 10.0.0.253 dist. 200

Non-Routed Backup Mechanisms

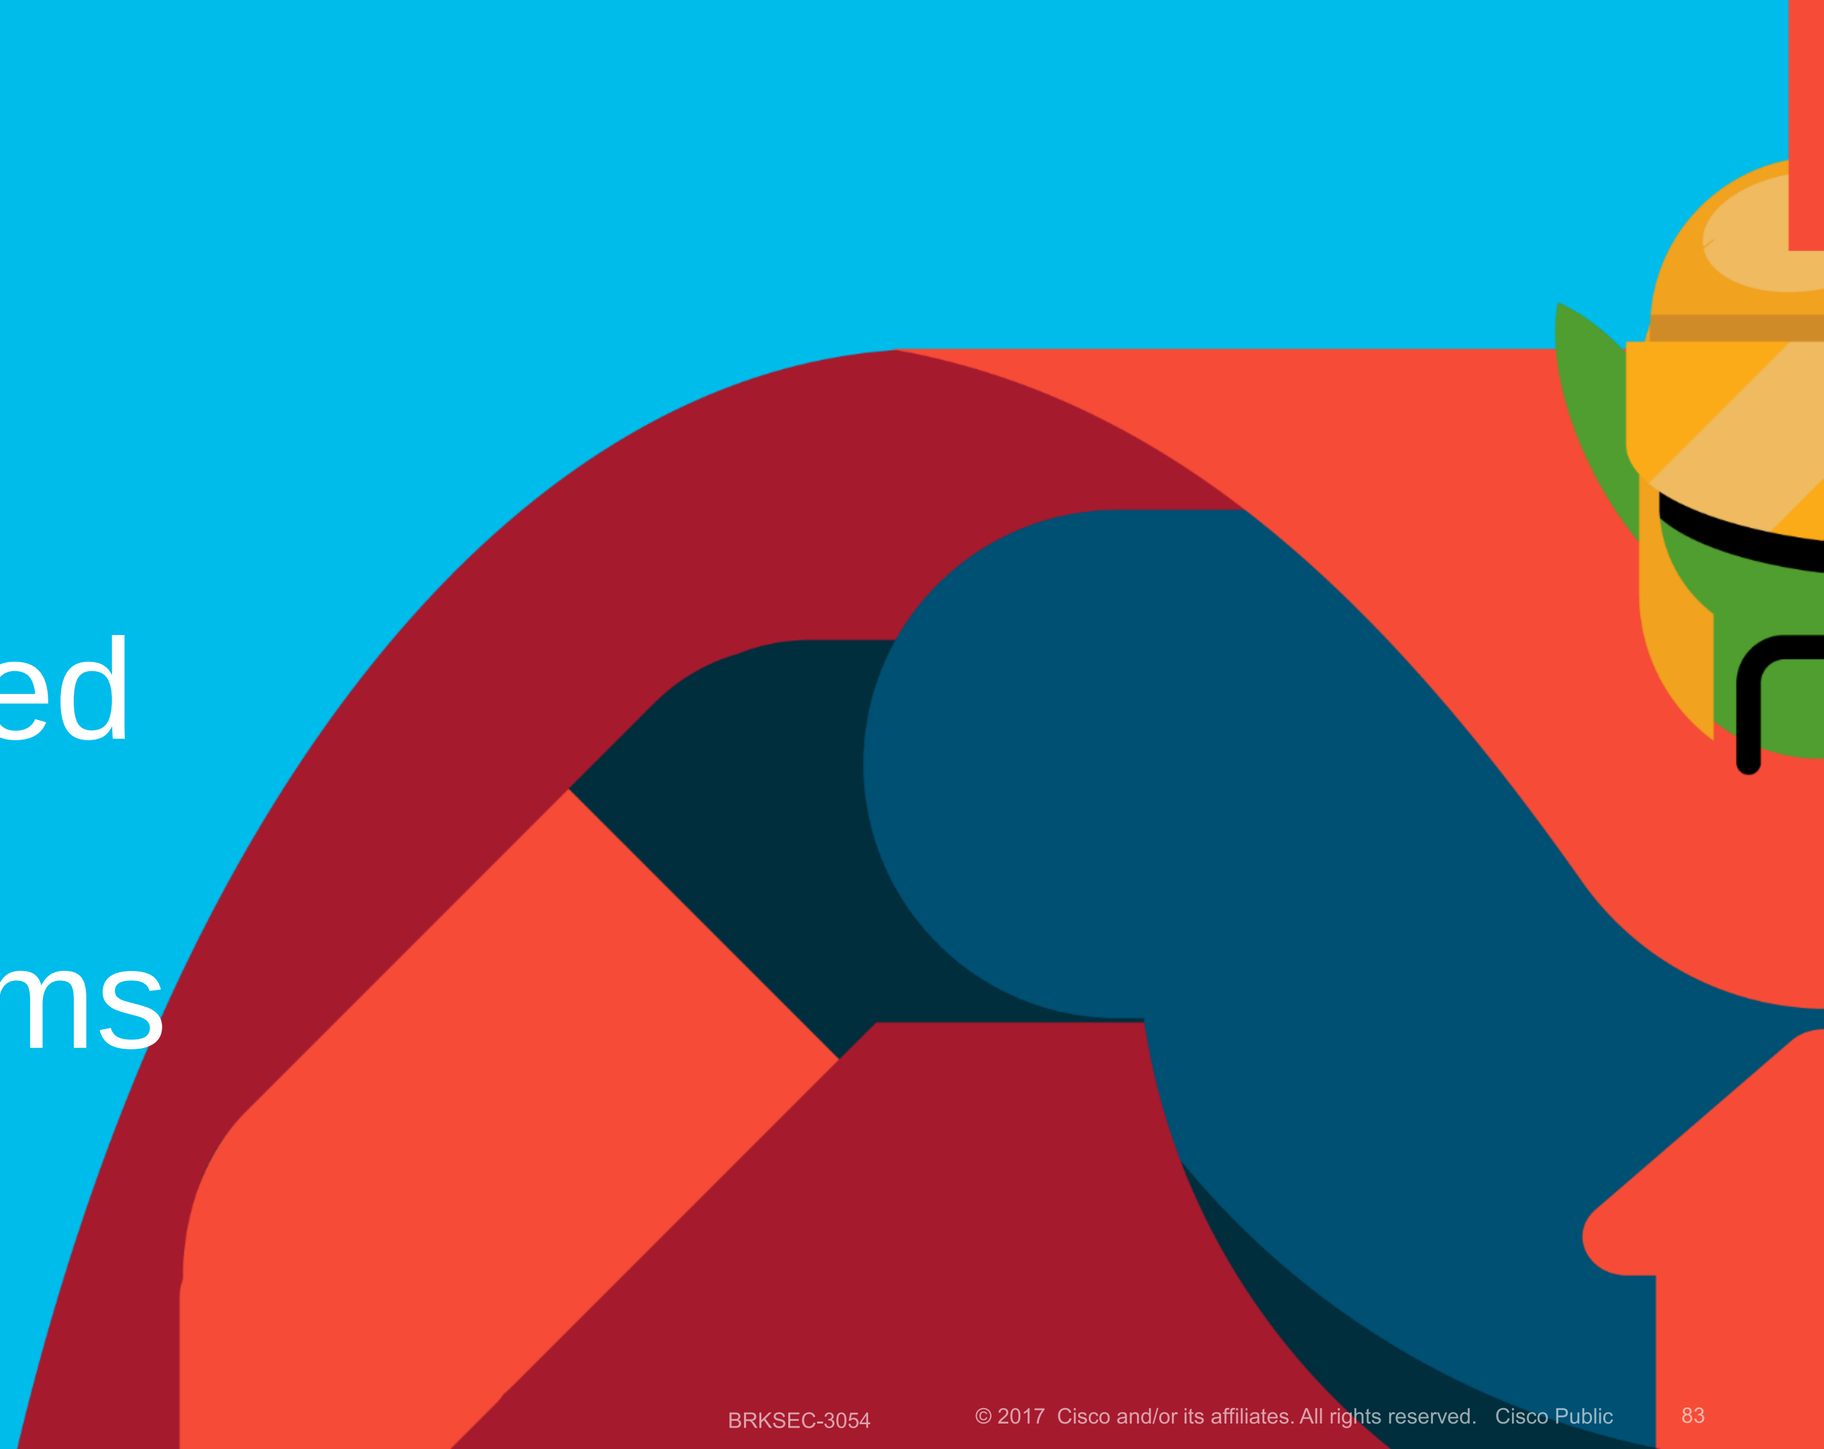

FlexVPN Backup Peers (1)

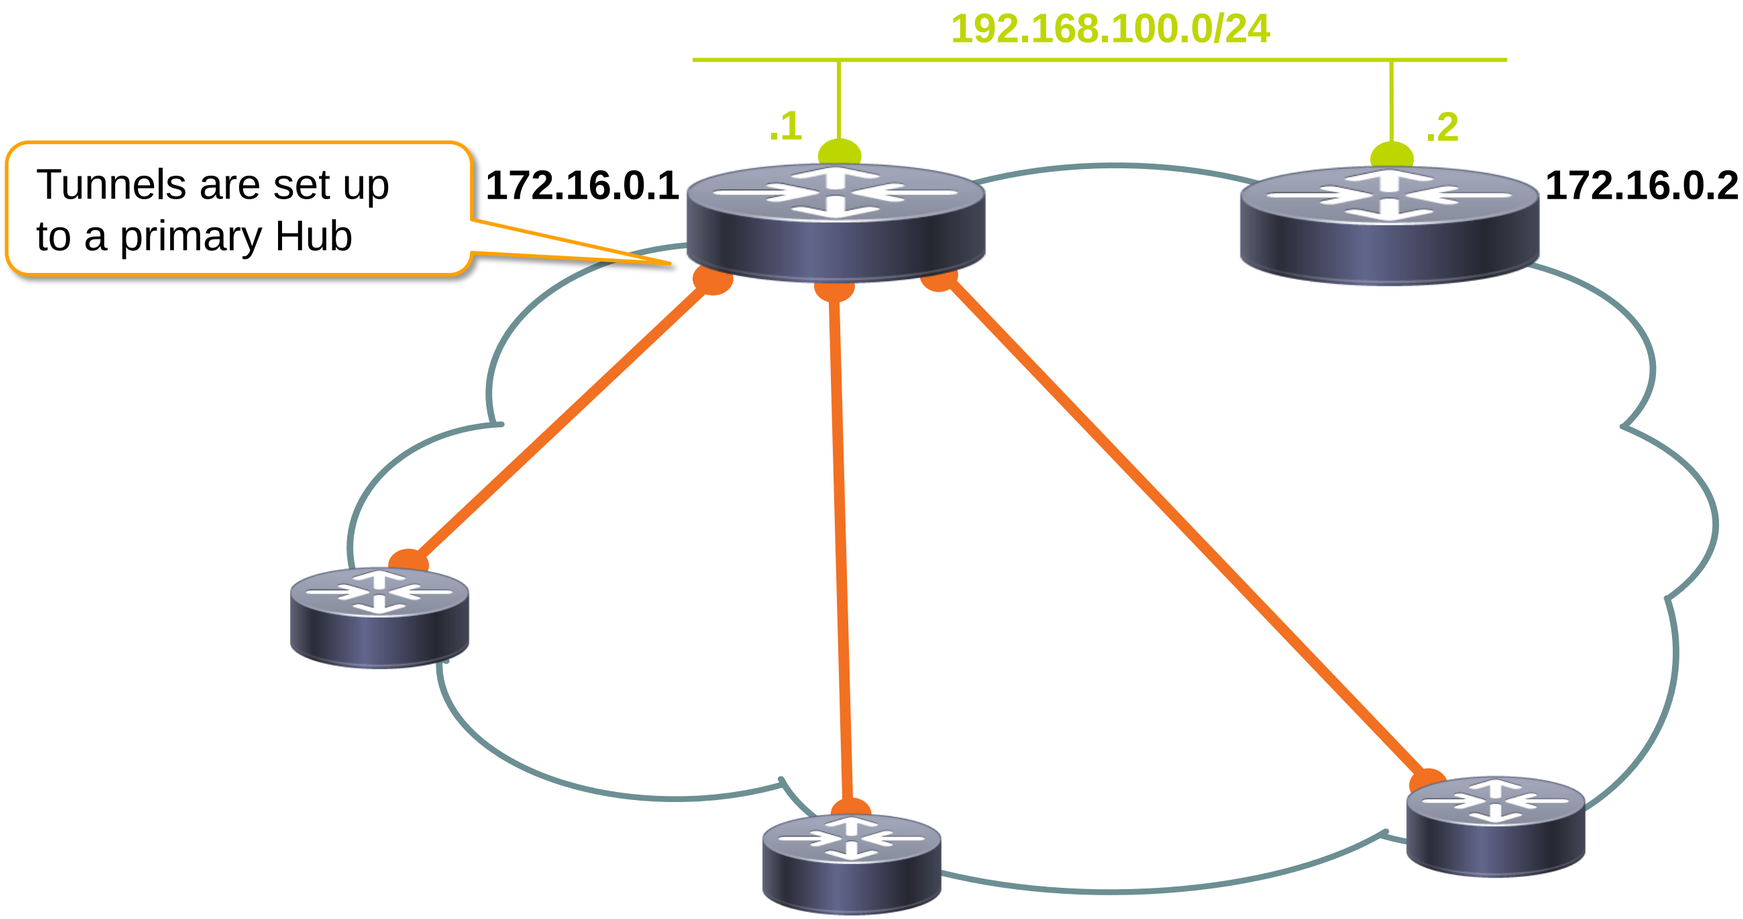

FlexVPN Backup Peers (2)

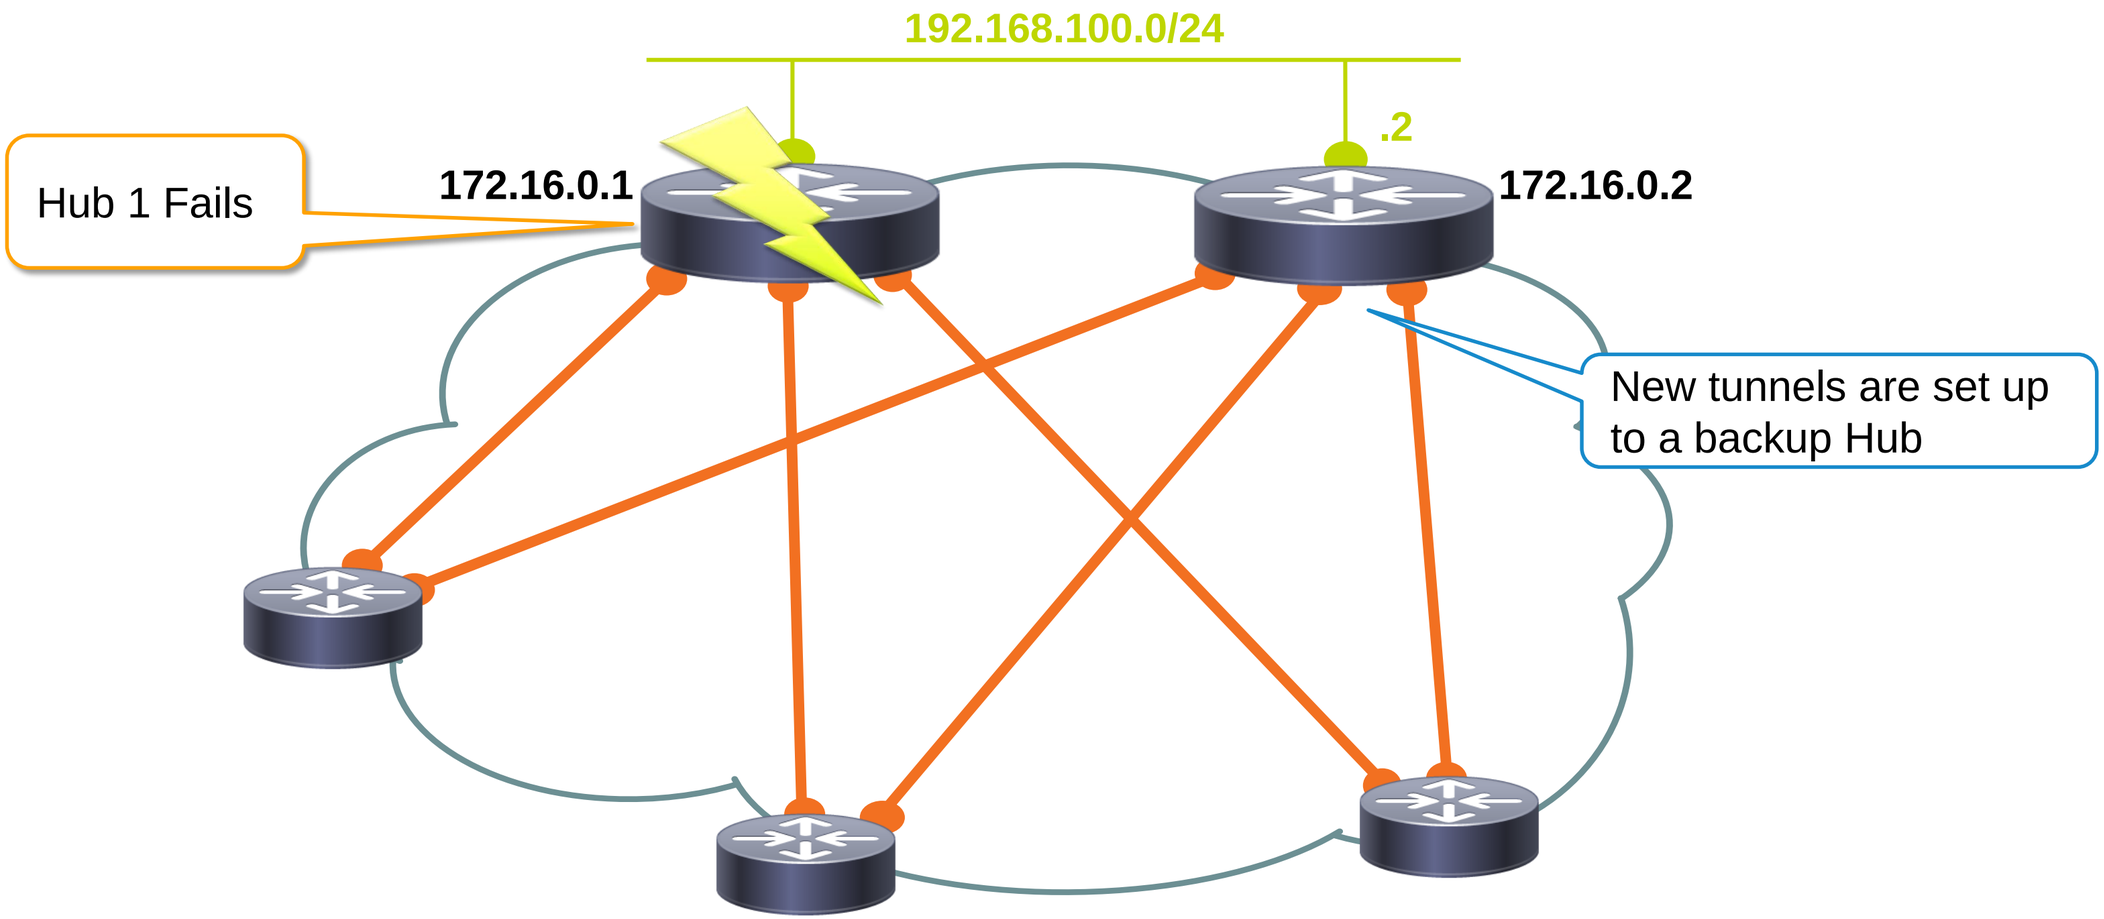

FlexVPN Backup Peers (3) – Spoke Config.

Also works with Routing Protocol

```
aaa authorization network default local

crypto ikev2 profile default
 match certificate HUBMAP
 identity local fqdn Spoke1.cisco.com
 authentication remote rsa-sig
 authentication local pre-shared
 keyring local
 pki trustpoint CA
 aaa authorization group cert list default default
 dpd 30 2 on-demand

crypto ikev2 client flexvpn default
 client connect tunnel 0
 peer 1 172.16.1.254
 peer 2 172.16.1.253

interface Tunnel0
 ip address negotiated
 tunnel source FastEthernet0/0
 tunnel destination dynamic
 tunnel protection ipsec profile default
```

Detect Hub Failure

To Primary Hub

To Secondary Hub

Destination managed by FlexVPN

Powerful Peer Syntax

```
peer reactivate
peer <n> <ip>
peer <n> <ip> track <x>
peer <n> <fqdn> [dynamic [ipv6]]
peer <n> <fqdn> [dynamic …] track <x>
```

Switch back

Nth source selected only if corresponding track object is up

RADIUS Backup List Attribute

```
ipsec:ipsec-backup-gateway
```

Up to 10 backup gateways pushed by config-exchange

```
crypto ikev2 authorization policy default
 route set interface
 route set access-list 99
```

FlexVPN Backup Groups

Warrant that a peer, belonging to different peer-lists in the same backup group, is never active in multiple peer-list at a given time

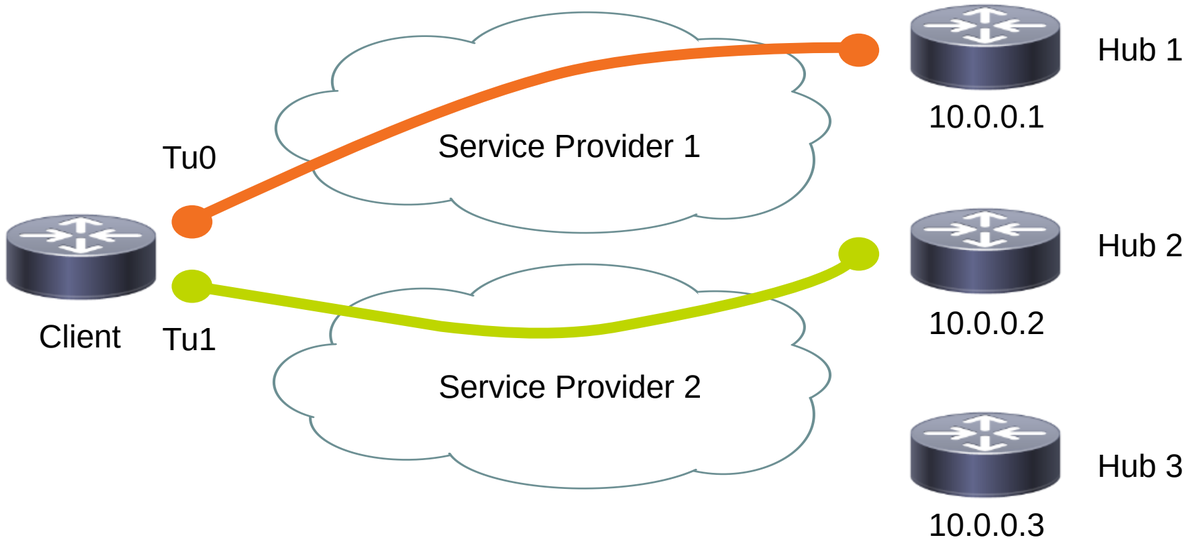


```
crypto ikev2 flexvpn client remote1
 peer 1 10.0.0.1
 peer 2 10.0.0.2
 peer 3 10.0.0.3
 backup group 1
 client connect Tunnel0
crypto ikev2 flexvpn client remote2
 peer 1 10.0.0.1
 peer 2 10.0.0.2
 peer 3 10.0.0.3
 backup group 1
 client connect Tunnel1
!
interface Tunnel0
 ip address negotiated
 …
 tunnel destination dynamic
 …
interface Tunnel1
 ip address negotiated
 …
 tunnel destination dynamic
 …
```

10.0.0.1 cannot be used as already active in remote1 peer-list from same group

FlexVPN Tunnel Pivot

- Use when different Service Providers are used to connect to remote host

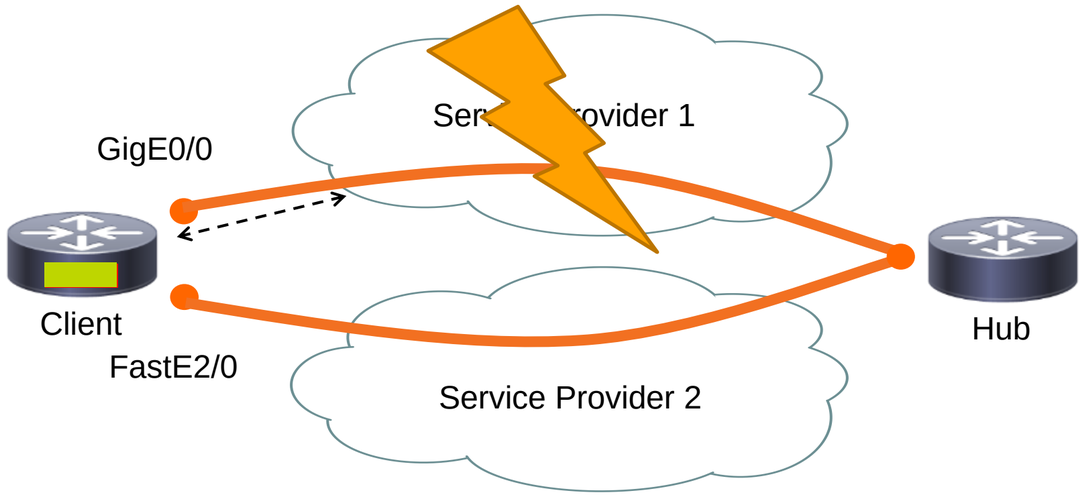

Tracker state (Up/Down)

ICMP-echo IP SLA probe

IPsec Tunnel

```
track 1 ip sla 1 reachability

crypto ikev2 flexvpn client remote1
  peer 10.0.0.1
  source 1 interface GigabitEthernet0/0 track 1
  source 2 interface FastEthernet2/0
  client connect tunnel 0

interface Tunnel0
   ip address negotiated
   …
   tunnel source dynamic
   tunnel destination dynamic
   …
```

Example: Multi-tenant Hybrid Access

Use Case: Mixed Client & Branch Access

- Requirements:
  - Single responder for software clients & remote branches (spokes)
  - Spoke-to-spoke tunnels enabled on a per-branch basis
  - VRF enforced per user/branch
  - Branches use IKE certificates, clients use EAP (password or TLS certificates)
- Proposed solution:
  - Single IKEv2 profile & V-Template
  - Differentiated AAA authorization depending on authentication method

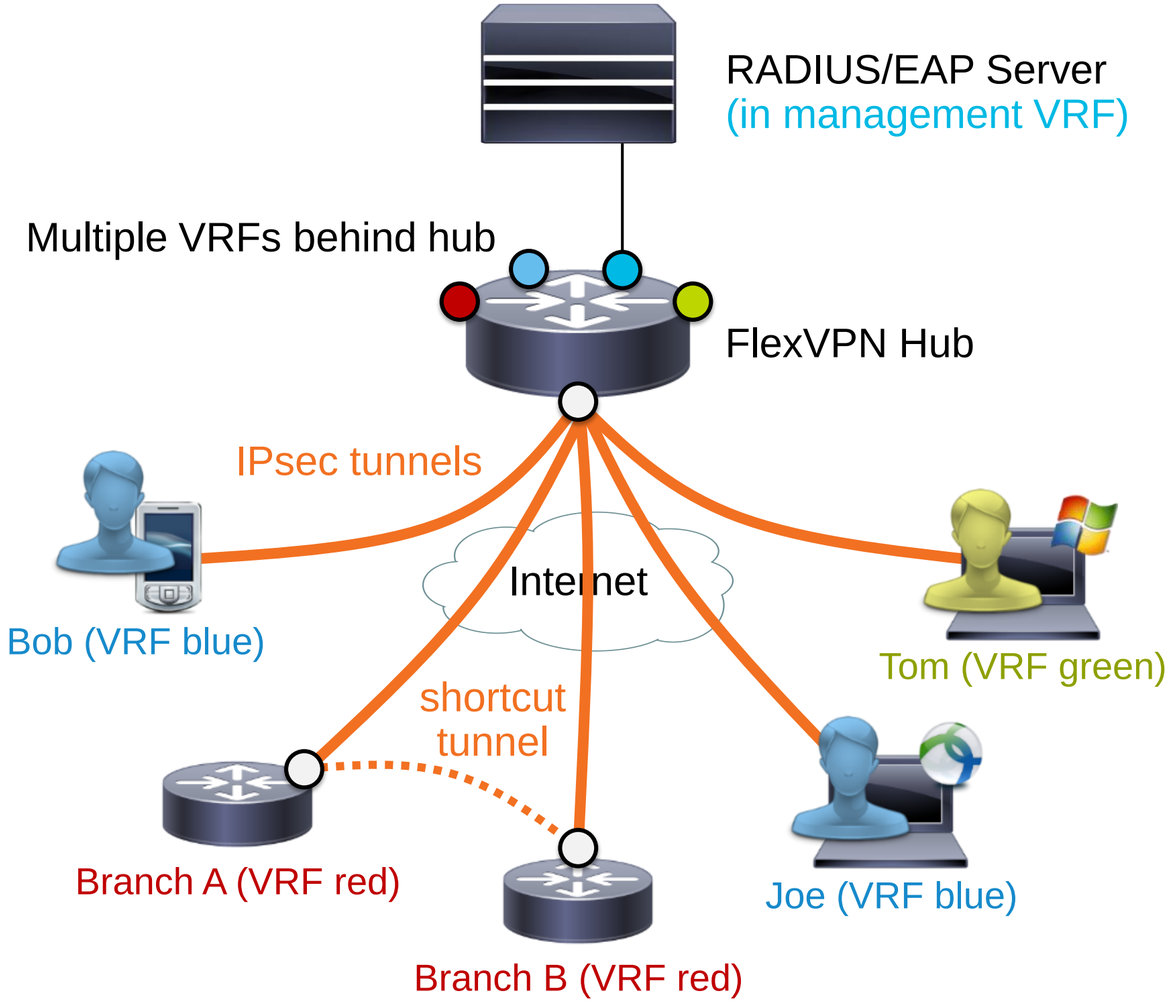

Tunnel modes made easy

```
crypto ikev2 profile prof1
 …
 virtual template 1
interface virtual-template 1
 …
 tunnel mode ipsec ipv4
 tunnel protection ipsec profile default
```

```
crypto ikev2 profile prof2
 …
 virtual template 2
interface virtual-template 2
 …
 tunnel mode ipsec ipv6
 tunnel protection ipsec profile default
```

```
crypto ikev2 profile prof3
 …
 virtual template 3
interface virtual-template 3
 …
 tunnel mode gre ip
 tunnel protection ipsec profile default
```

```
crypto ikev2 profile prof4
 …
 virtual template 4
interface virtual-template 4
 …
 tunnel mode gre ipv6
 tunnel protection ipsec profile default
```

```
crypto ikev2 profile default
 …
 virtual template 1 mode auto

interface virtual-template 1
 …
```

FlexVPN Server Configuration

RADIUS-based EAP authentication and AAA authorization

Match on FQDN domain for branches

Match statements for clients
(depending on allowed client types)

Allow peers to authenticate using either EAP or certificates

User authorization (using attributes returned during EAP authentication)

Branch authorization using RADIUS

Automatic detection of tunnel mode[1]
(pure IPsec tunnel mode for clients, GRE/IPsec for branches/spokes)

```
aaa new-model
aaa authentication login my-rad group my-rad
aaa authorization network my-rad group my-rad
!
crypto ikev2 profile default
 match identity remote fqdn domain example.com
 match identity remote {key-id | email | address} ...
 identity local dn
 authentication remote rsa-sig
 authentication remote eap query-identity
 authentication local rsa-sig
 pki trustpoint my-ca
 aaa authentication eap my-rad
 aaa authorization user eap cached
 aaa authorization user cert list my-rad
 virtual-template 1 mode auto
!
interface Virtual-Template1 type tunnel
 no ip address
 [no need to specify tunnel mode]
 tunnel protection ipsec profile default
```

[1] Starting with IOS-XE 3.12S

RADIUS Server Configuration

Clients can perform password-based or TLS-based EAP authentication
(TLS: RADIUS account = CN or UPN)

User attributes returned by RADIUS upon successful EAP authentication

Branch attributes returned by RADIUS during AAA authorization step

Add/remove NHRP to enable/disable spoke-to-spoke tunnels per branch

Exchange prefixes via IKEv2 routing, branch prefix(es) controlled by branch

Branch prefix controlled by AAA server
(installed as local static route)

```
joe
   cleartext-password=c1sc0!
   ipsec:addr-pool=blue
   ip:interface-config=vrf forwarding blue
   ip:interface-config=ip unnumbered Loopback1
   ip:interface-config=service-policy output blue-pol
   ip:interface-config=...

branch1.example.com
   ip:interface-config=vrf forwarding red
   ip:interface-config=ip unnumbered Loopback3
   ip:interface-config=ip nhrp network-id 3
   ip:interface-config=ip nhrp redirect
   ipsec:route-set=prefix 192.168.0.0 255.255.0.0
   ipsec:route-accept=any

branch2.example.com
   ip:interface-config=vrf forwarding green
   ip:interface-config=ip unnumbered Loopback2
   ipsec:route-set=prefix 192.168.0.0 255.255.0.0
   ipsec:route-set=local 192.168.1.0 255.255.255.0
```

Software Client Management

Mobile Device Management

Clients › Piotr's iPad

Client details | Refresh details | Edit details

Name: Piotr's iPad
Model: iPad mini Retina
Serial: F9FMWHSFFCM8
Warranty: Apple
Tags: recently-added
Auto tags: iOS devices
Charge: 23%
Owner: Set an owner

OS

Version: iOS 10.3.1

Security

Encryption: Both file-level and block-level capable
Passcode: Not present
Jailbroken?: No

Management

Settings: up-to-date
Apps: up-to-date
Supervised: No
Kiosk application: -
Managed Profile: No
Device Owner: No
Enrollment date: 23:39 Apr 28 2017

Storage

Device storage: 9.0 GB / 11.6 GB 77%

Approximate location | Refresh location

Warszawa, Poland (via IP, updated 10 minutes ago)

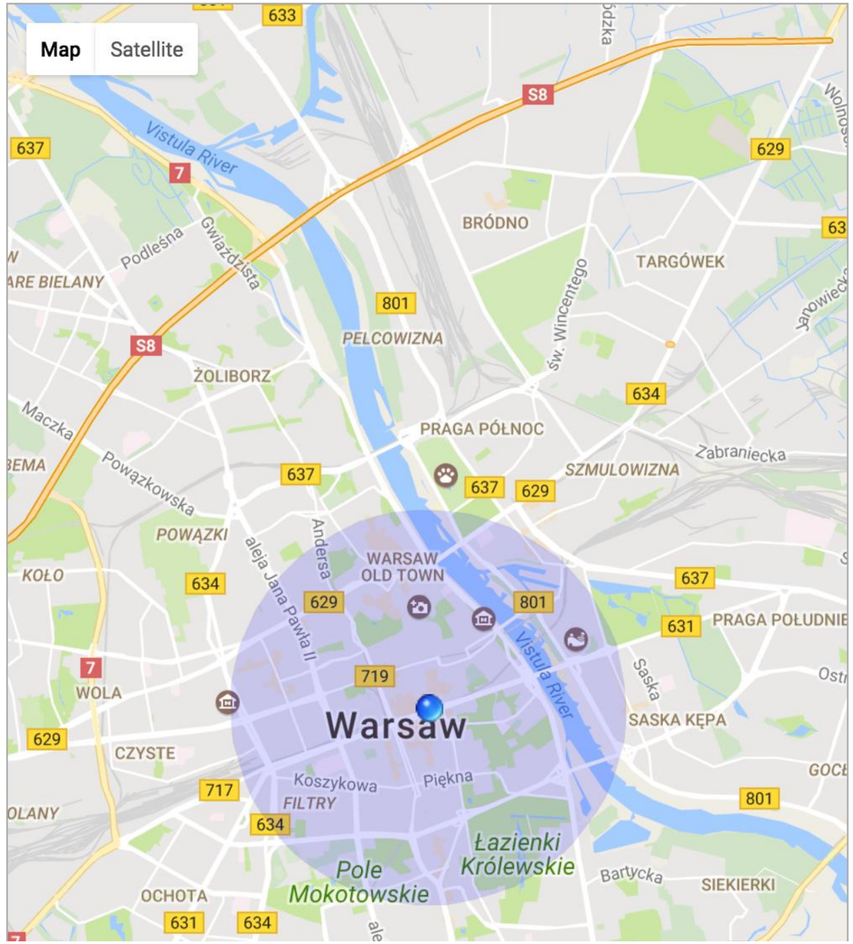

Creating an IPsec AnyConnect profile

[**ipsec**://][<**AUTHENTICATION**>[":"<**IKE-IDENTITY**>"@"]] <**HOST**>[":"<**PORT**>]["/"<**GROUP-URL**>]

Parameter	Description
ipsec	: Indicates that this is an IPsec connection. If omitted, SSL is assumed.
AUTHENTICATION	Specifies the authentication method for an IPsec connection. If omitted, EAP-AnyConnect is assumed. Valid values are: • EAP-AnyConnect • EAP-GTC • EAP-MD5 • EAP-MSCHAPv2 • IKE-RSA
IKE-IDENTITY	Specifies the IKE identify when AUTHENTICATION is set to EAP-GTC, EAP-MD5, or EAP-MSCHAPv2. This parameter is invalid when used for other authentication settings.
HOST	Specifies the server address. The hostname or IP address to be used.
PORT	Currently ignored, included for consistency with the HTTP URI scheme.
GROUP=URL	Tunnel group name appended to the server name.

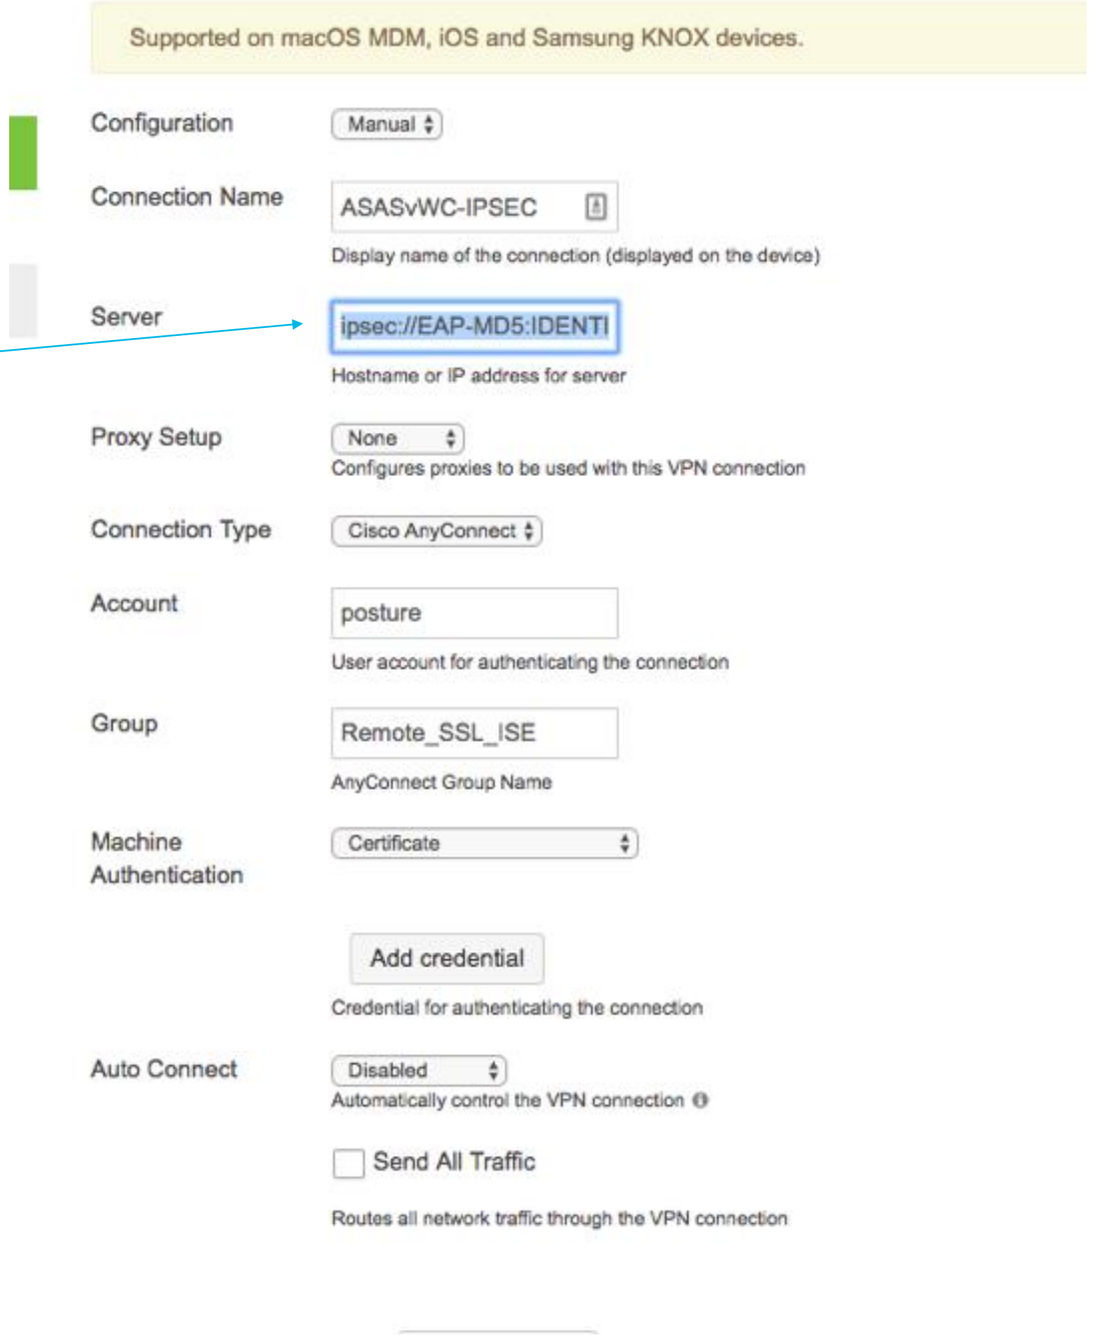

AAA & Augmented Security

Generic Profile Derivation

Full Example (seldom this complex)

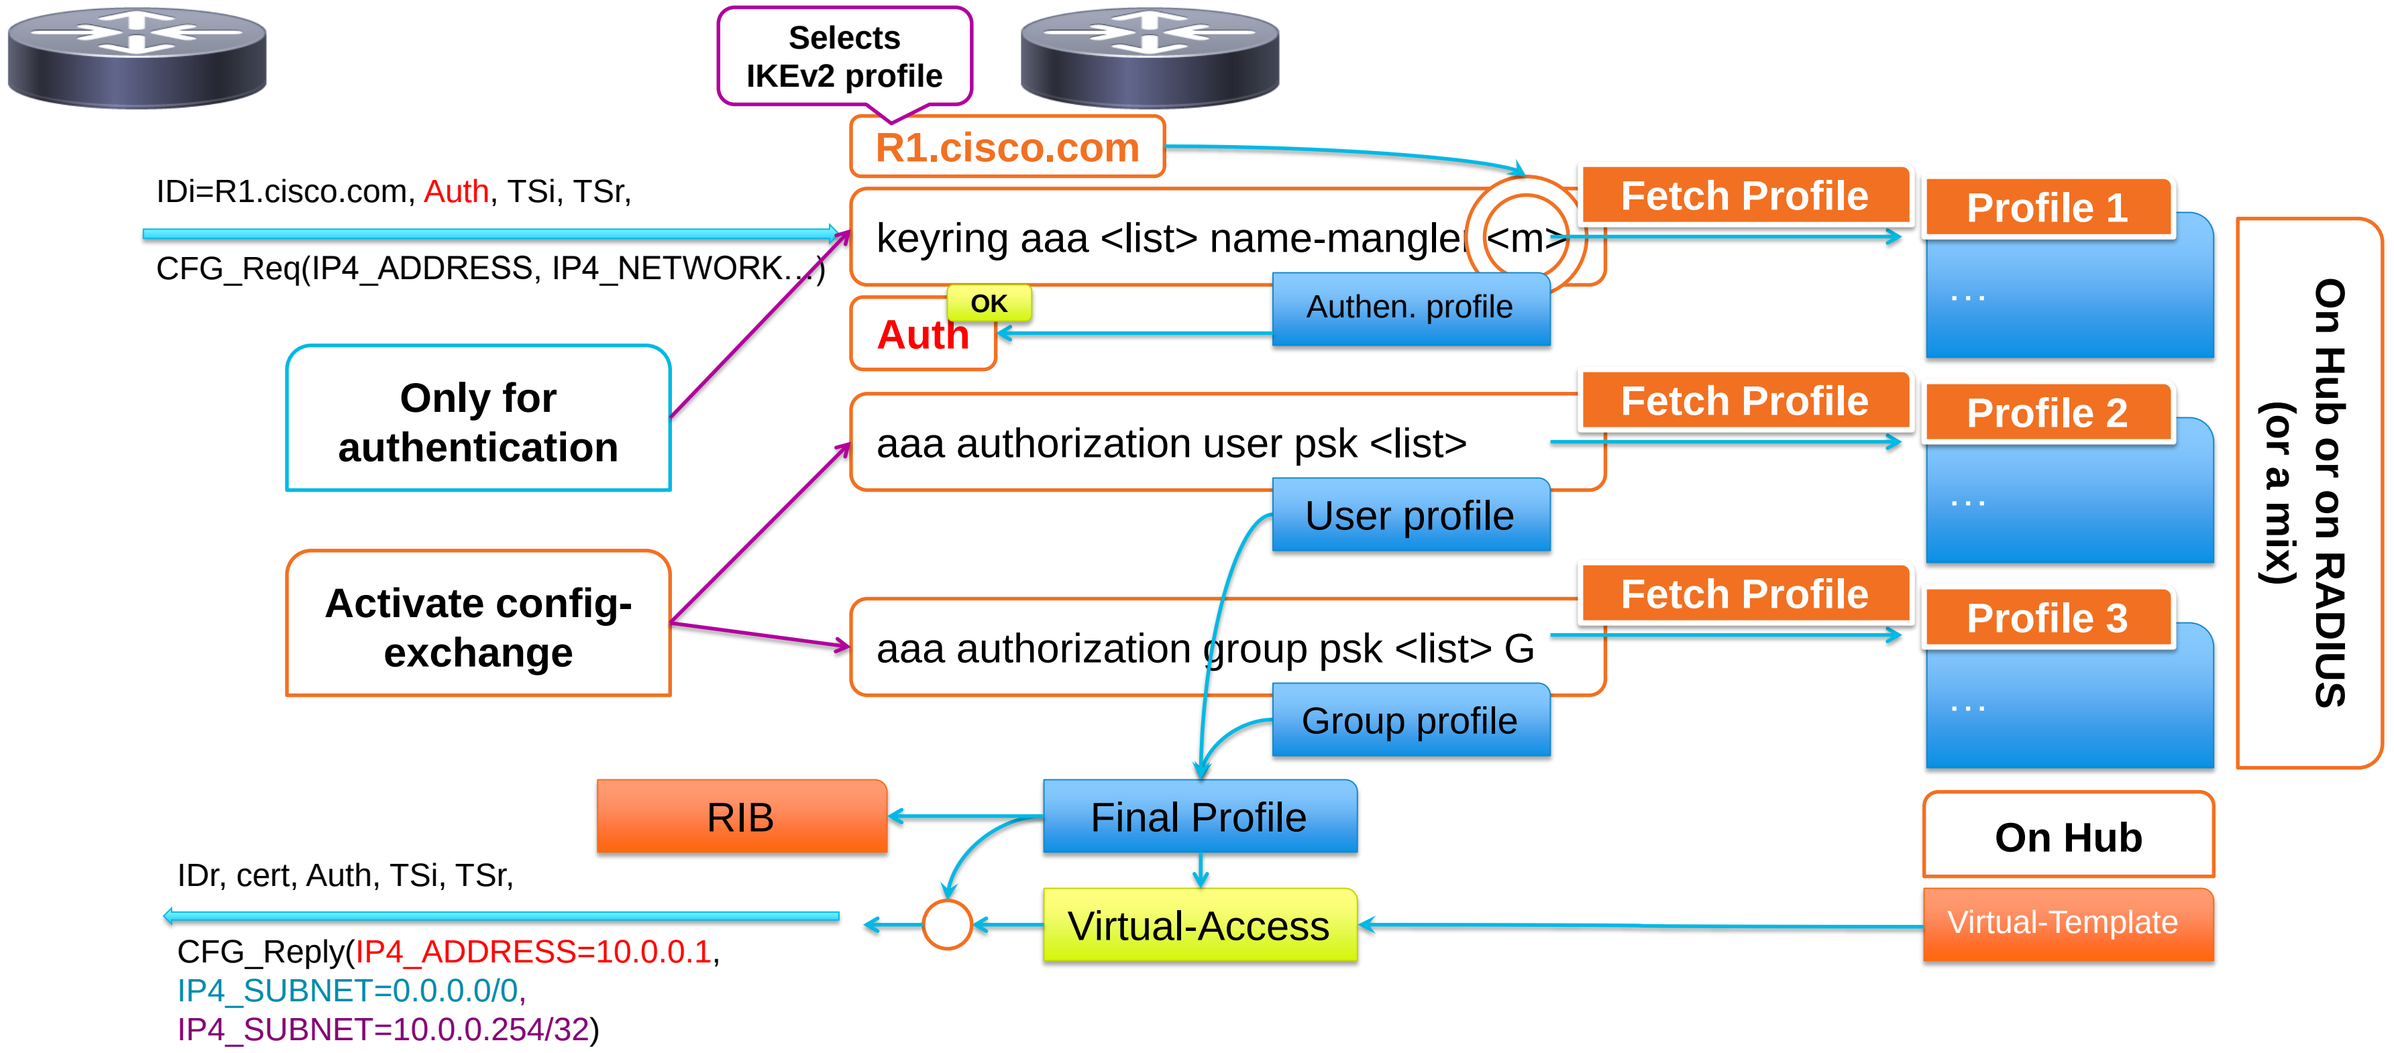

Risk: Route Injection

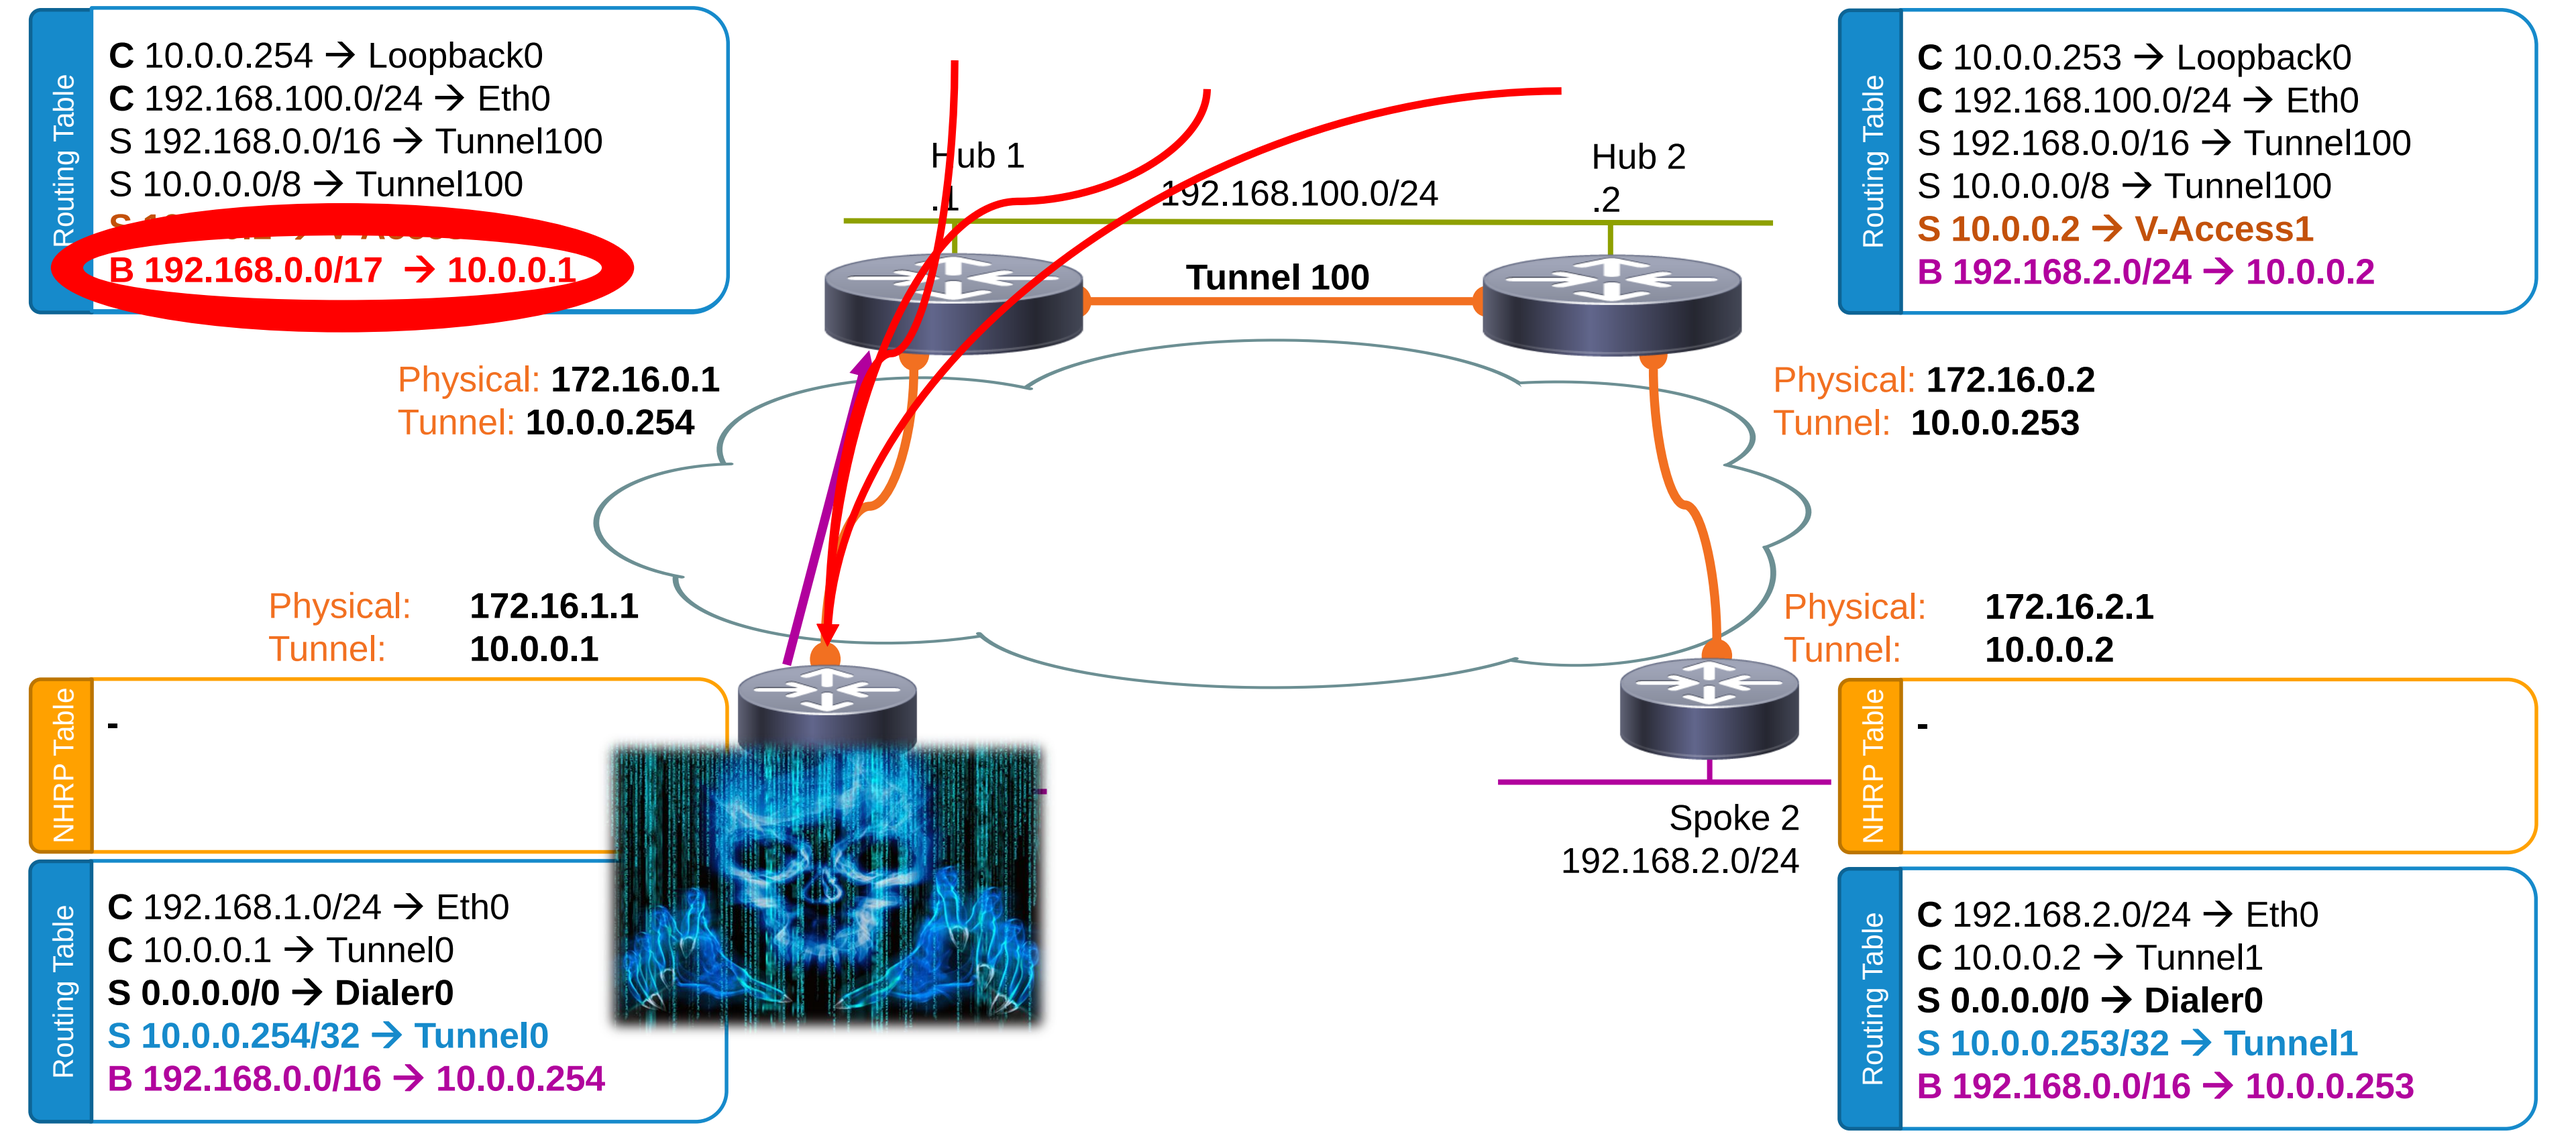

Risk: Traffic Injection

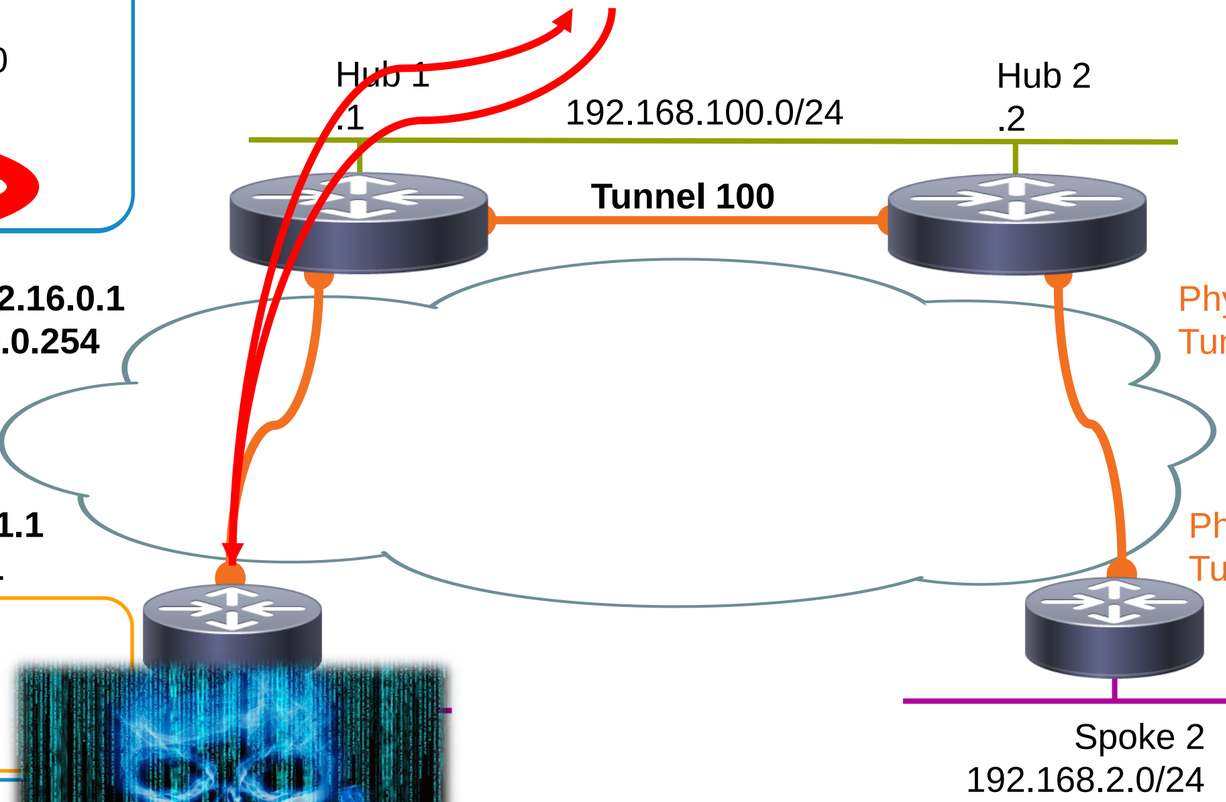

Routing Table

C 10.0.0.254 → Loopback0
C 192.168.100.0/24 → Eth0
S 192.168.0.0/16 → Tunnel100
S 10.0.0.0/8 → Tunnel100
B 192.168.0.0/17 → 10.0.0.1

Routing Table

C 10.0.0.253 → Loopback0
C 192.168.100.0/24 → Eth0
S 192.168.0.0/16 → Tunnel100
S 10.0.0.0/8 → Tunnel100
S 10.0.0.2 → V-Access1
B 192.168.2.0/24 → 10.0.0.2

NHRP Table

-

NHRP Table

-

Routing Table

C 192.168.1.0/24 → Eth0
C 10.0.0.1 → Tunnel0
S 0.0.0.0/0 → Dialer0
S 10.0.0.254/32 → Tunnel0
B 192.168.0.0/16 → 10.0.0.254

Routing Table

C 192.168.2.0/24 → Eth0
C 10.0.0.2 → Tunnel1
S 0.0.0.0/0 → Dialer0
S 10.0.0.253/32 → Tunnel1
B 192.168.0.0/16 → 10.0.0.253

AAA local & remote routes

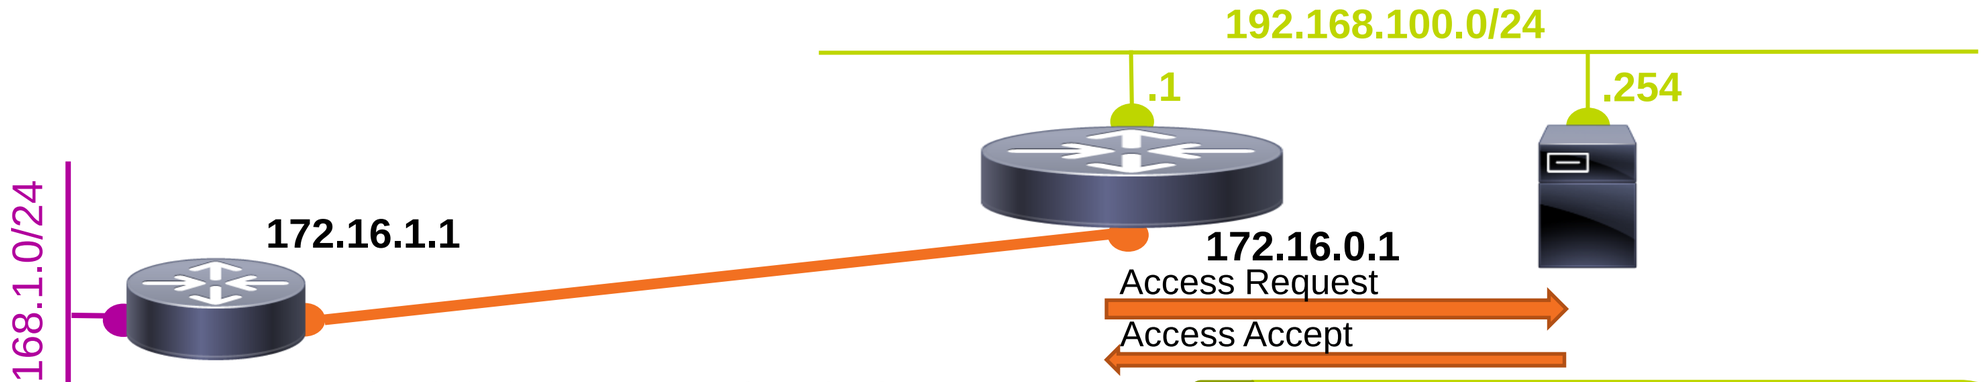

SA Prop (AES-256, SHA-1, DH 5), KEi, Ni

SA Prop (AES-256, SHA-1, DH 5), KEr, Nr

IDi=Spoke1.cisco.com, Auth, TSi, TSr,

CFG_Req(IP4_SUBNET…)

IDr, cert, Auth, TSi, TSr,

CFG_Reply(IP4_SUBNET=10.0.0.254/32, 192.168.0.0/16; IP4_ADDRESS=10.0.0.1)

CFG_set(IP4_SUBNET=~~10.0.0.1/32~~, ~~192.168.0.0/17~~)

CFG_ack()

Interfaces

Ethernet0/0: 172.16.1.1
Ethernet0/1: 192.168.1.1
Tunnel0: **10.0.0.1**

Routing Table

172.16.0.1/32 → 172.16.1.254 (E0/0)
192.168.1.0/24 → Ethernet 0/1
10.0.0.254/32 → Tunnel 0
192.168.0.0/16 → Tunnel 0

PROFILE

Spoke1.cisco.com
ipsec:route-set=interface
ipsec:route-set=ipv4 remote 192.168.0.0 255.255.0.0
ipsec:route-set=ipv4 remote 10.0.0.0 255.0.0.0
ipsec:route-set=ipv4 local 192.168.1.0 255.255.255.0
ipsec:route-accept=deny *(default)*
Framed-address=10.0.0.1
ip:interface-config=ip verify unicast reverse path

Interfaces

Ethernet0/0: 172.16.0.1
Ethernet0/1: 192.168.100.1
Loopback0: **10.0.0.254/32**
VirtualAccess1: 10.0.0.254/32

Routing Table

0.0.0.0/0 → 172.16.0.254 (E0/0)
192.168.100.0/24 → Ethernet 0/1
10.0.0.1/32 → VirtualAccess1
192.168.1.0/24 → VirtualAccess1

AAA local & remote routes

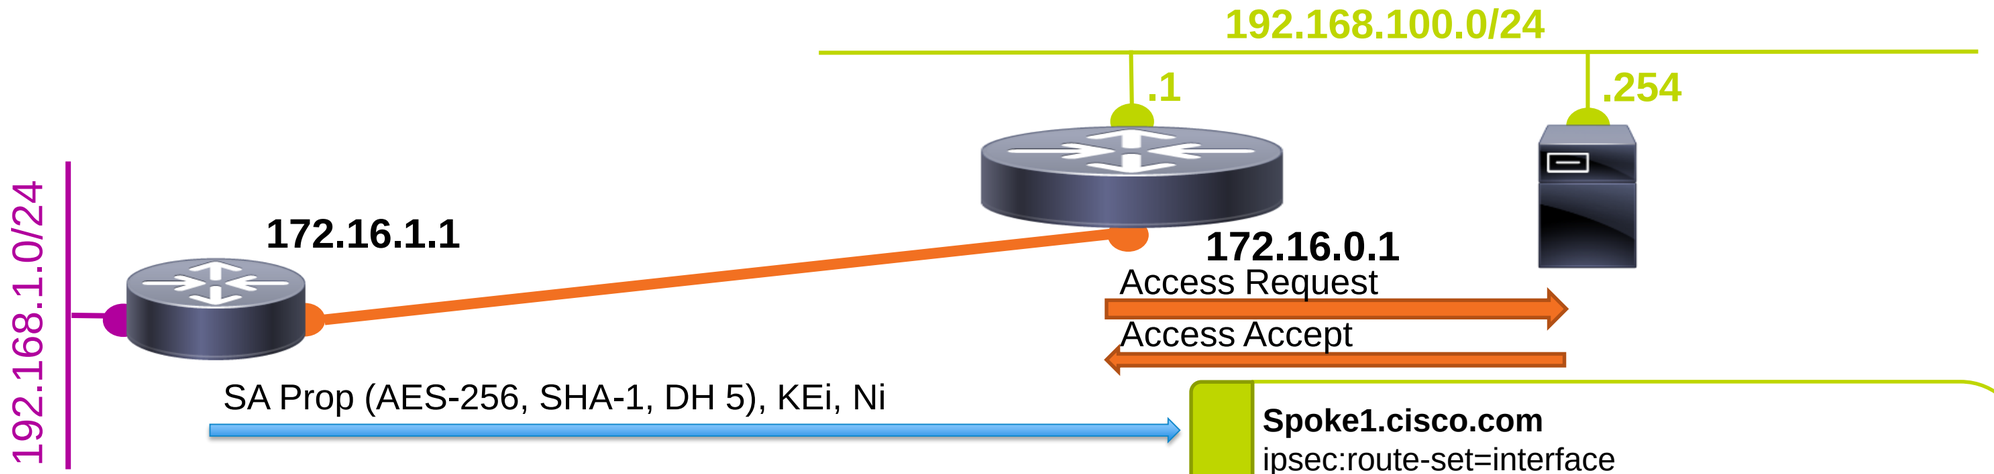

SA Prop (AES-256, SHA-1, DH 5), KEi, Ni

SA Prop (AES-256, SHA-1, DH 5), KEr, Nr

IDi=Spoke1.cisco.com, Auth, TSi, TSr,

CFG_Req(IP4_SUBNET…)

IDr, cert, Auth, TSi, TSr,

CFG_Reply(IP4_SUBNET=10.0.0.254/32, 192.168.0.0/16; IP4_ADDRESS=10.0.0.1)

CFG_set(IP4_SUBNET=10.0.0.1/32, **192.168.0.0/17**)

CFG_ack()

Interfaces

Ethernet0/0: 172.16.1.1
Ethernet0/1: 192.168.1.1
Tunnel0: **10.0.0.1**

Routing Table

172.16.0.1/32 → 172.16.1.254 (E0/0)
192.168.1.0/24 → Ethernet 0/1
10.0.0.254/32 → Tunnel 0
192.168.0.0/16 → Tunnel 0

PROFILE

Spoke1.cisco.com
ipsec:route-set=interface
ipsec:route-set=ipv4 remote 192.168.0.0 255.255.0.0
ipsec:route-set=ipv4 remote 10.0.0.0 255.0.0.0
ipsec:route-set=ipv4 local 192.168.1.0 255.255.255.0
ipsec:route-accept=deny *(default)*
Framed-address=10.0.0.1
ip:interface-config=ip verify unicast reverse path

Interfaces

Ethernet0/0: 172.16.0.1
Ethernet0/1: 192.168.100.1
Loopback0: **10.0.0.254/32**
VirtualAccess1: 10.0.0.254/32

Routing Table

0.0.0.0/0 → 172.16.0.254 (E0/0)
192.168.100.0/24 → Ethernet 0/1
10.0.0.1/32 → VirtualAccess1
192.168.1.0/24 → VirtualAccess1

For Your Reference

ISE - Custom Attribute

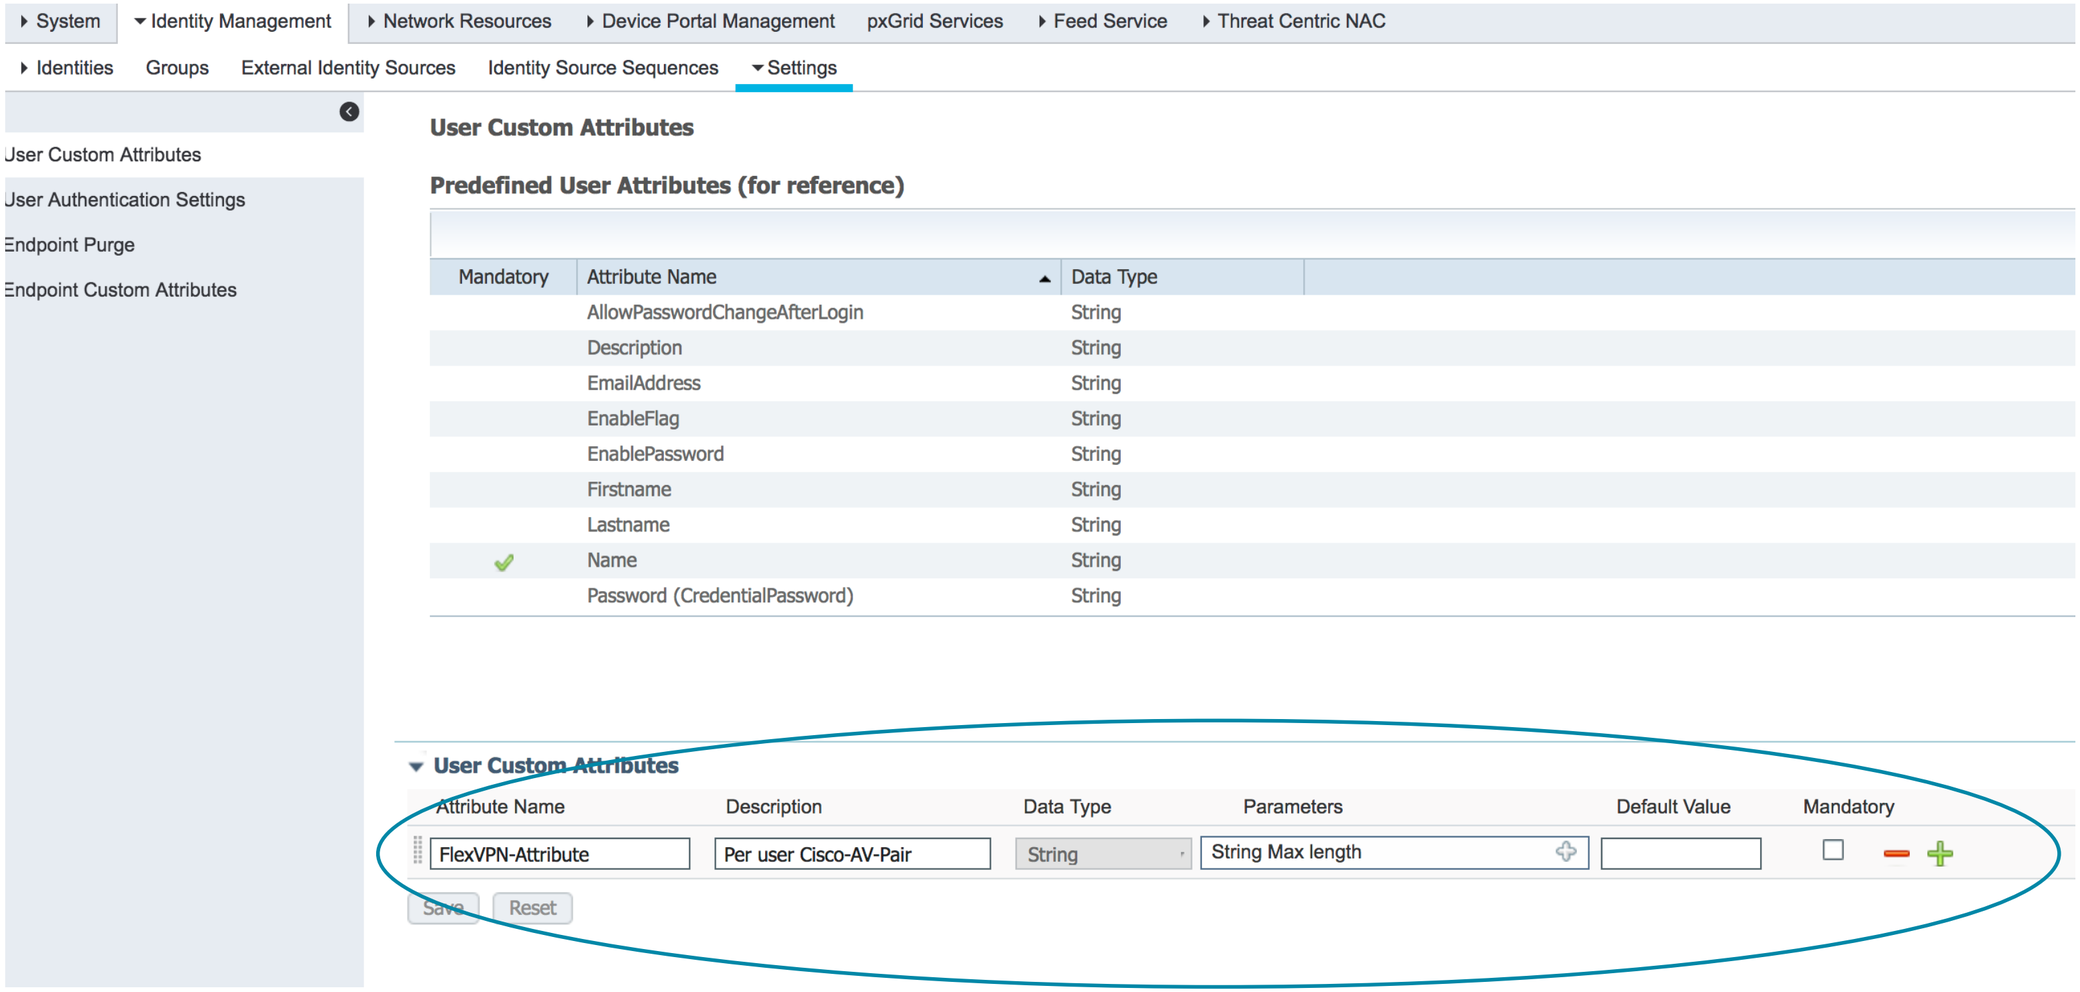

ISE – Configuring custom-attribute per Spoke

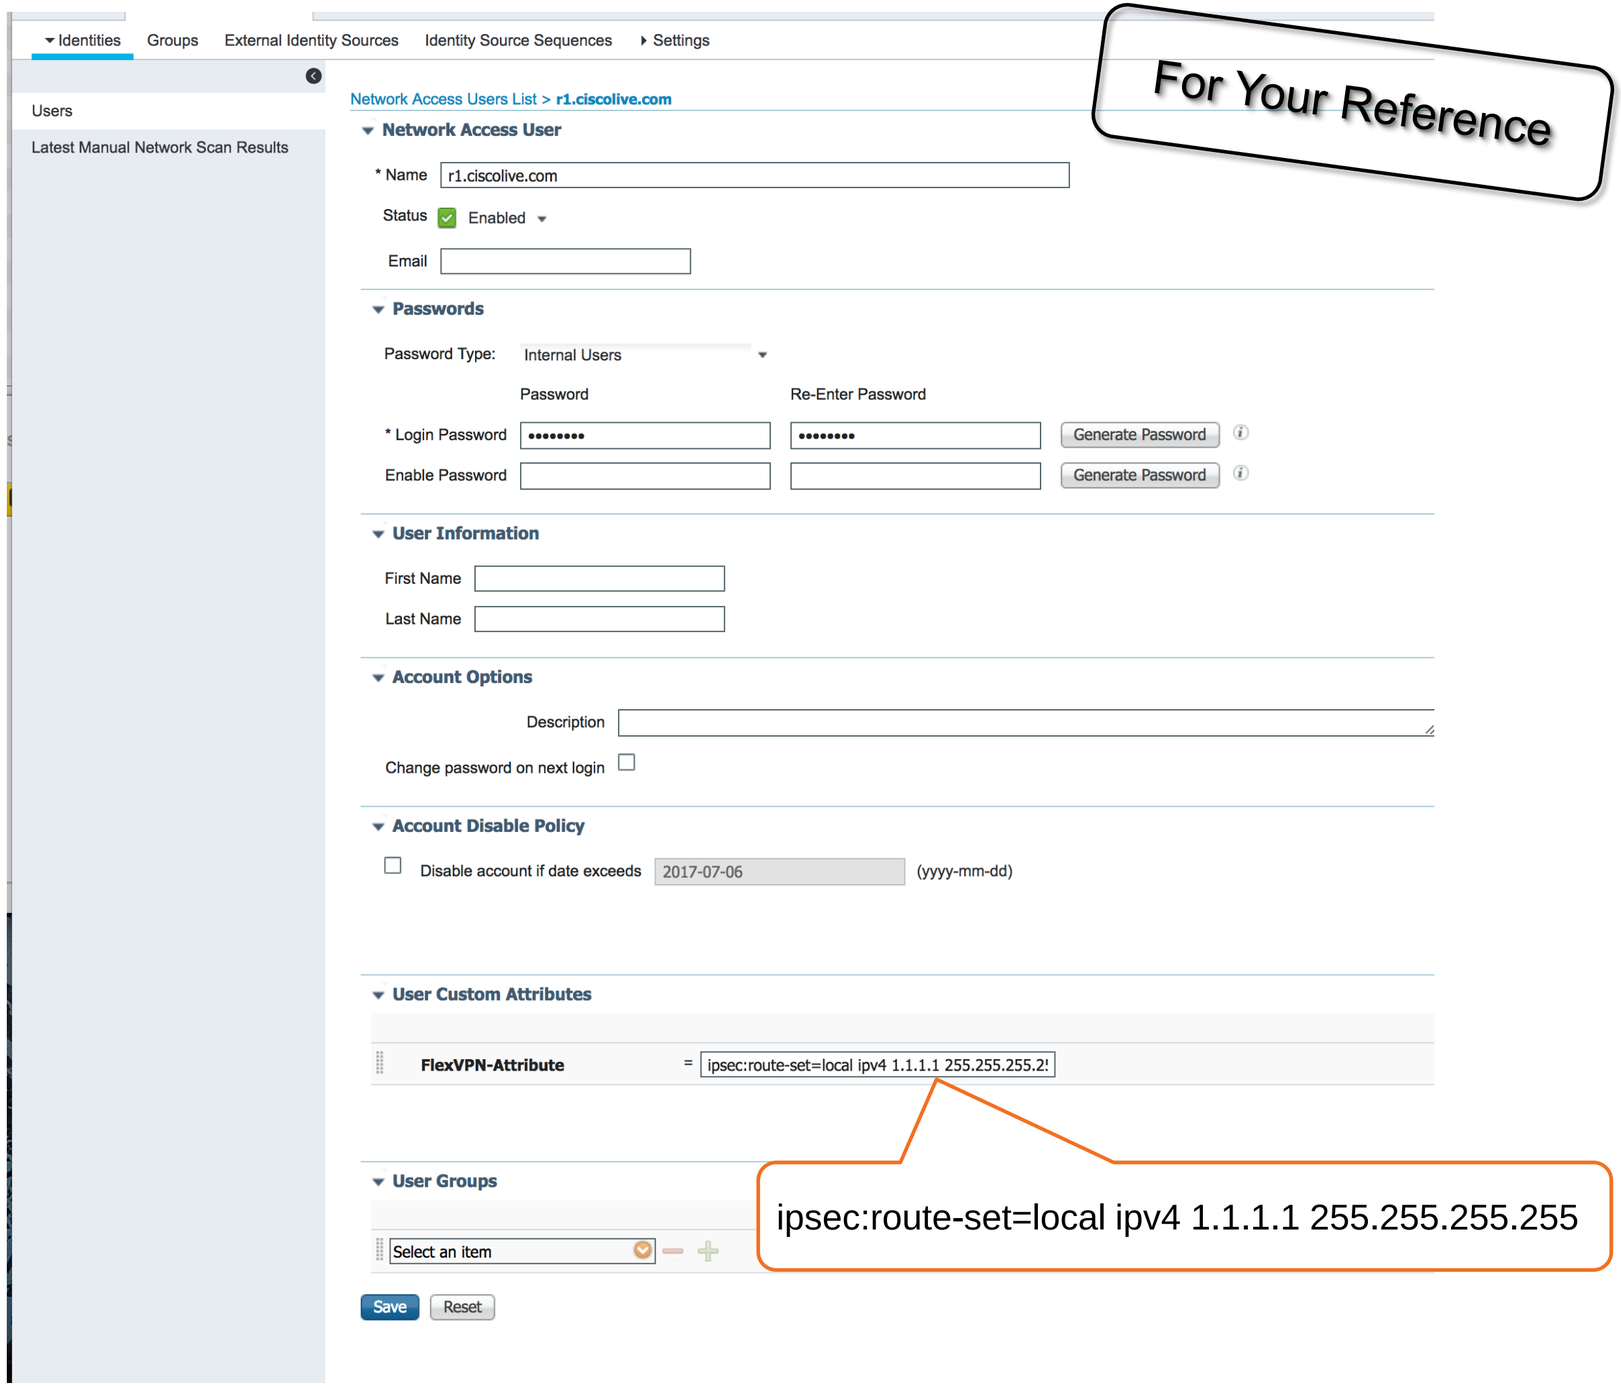

ISE – Authorization profile assigning customer attribute + additional attributes

For Your Reference

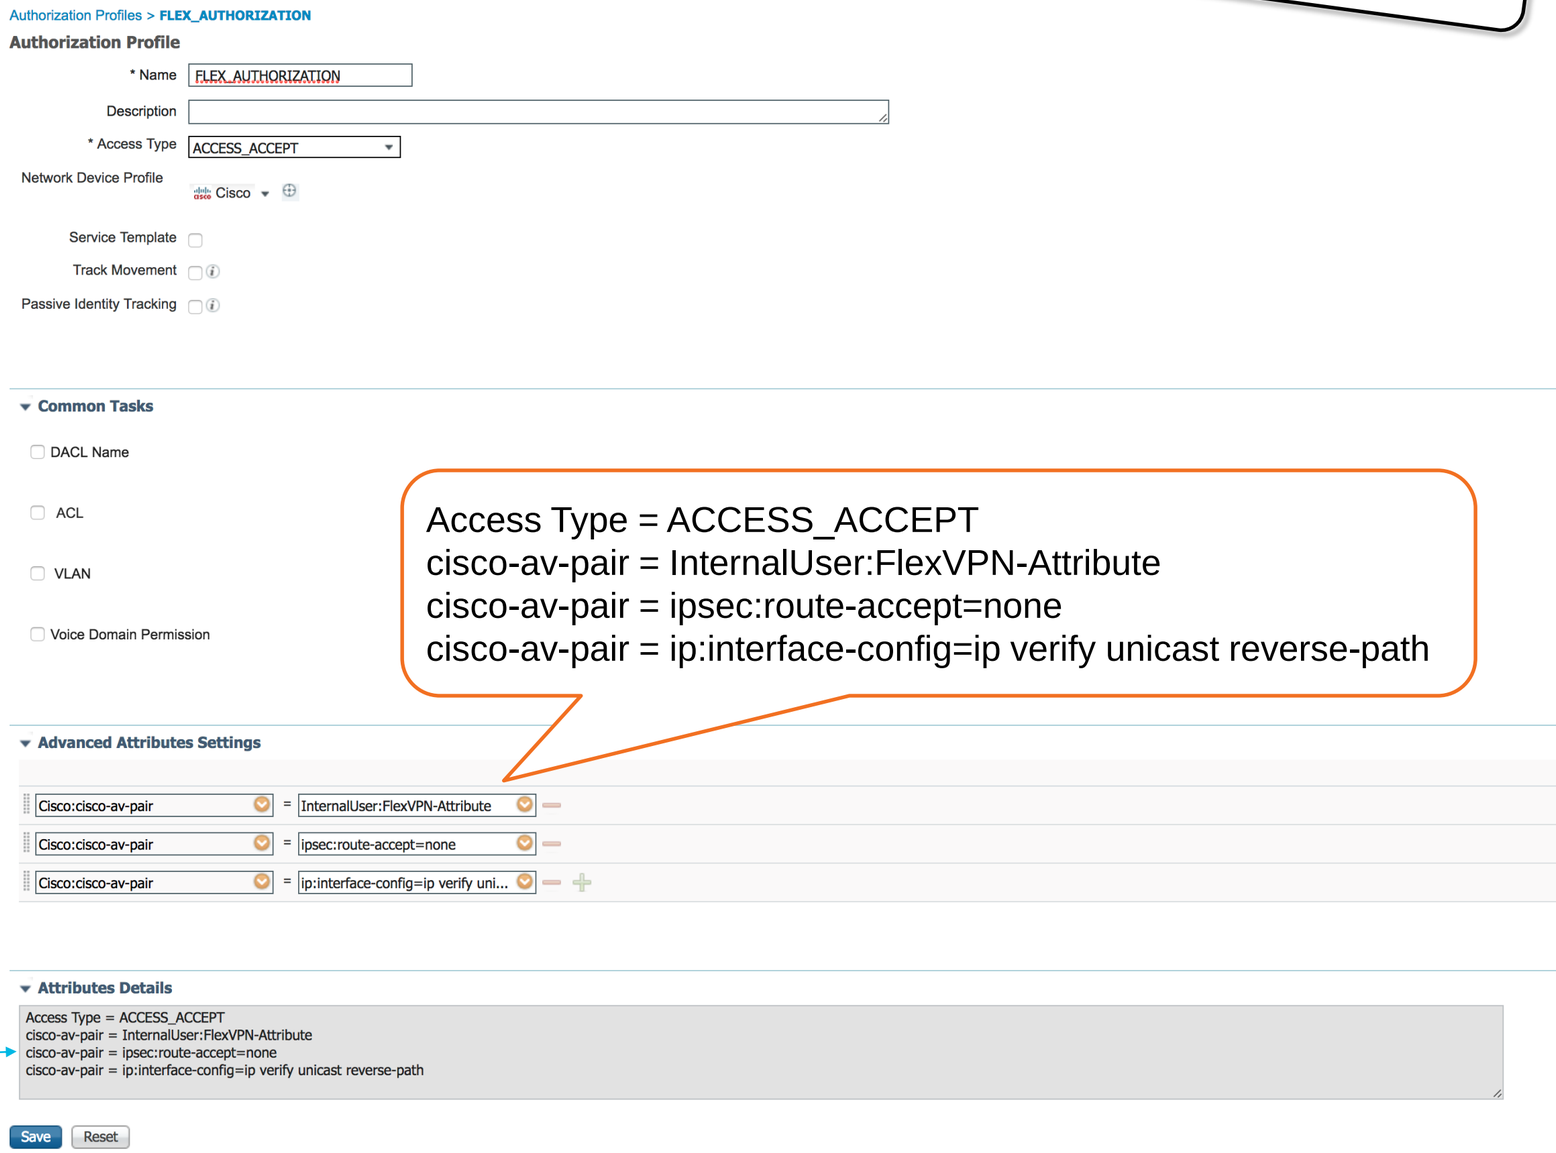

For Your Reference

ISE – Authorization Policy rule

Status	Rule Name	Conditions (identity groups and other conditions)	Permissions
✓	AUTHORIZATION	if Radius:Service-Type EQUALS Outbound	then FLEX_AUTHORIZATION
✓	RA_NOT_COMPLIANT	if Session:PostureStatus NOT_EQUALS Compliant	then RA_NOT_COMPLIANT
✓	RA_COMPLIANT	if Session:PostureStatus EQUALS Compliant	then RA_COMPLIANT
✓	Default	if no matches, then DenyAccess	

Matching done based on Service-Type == 5 (Outbound). Used by IOS only for authorization (never for authentication)

Hub Configuration

For Your Reference

```
radius server ISE
 address ipv4 172.16.140.212 auth-port 1645 acct-port 1646

aaa group server radius ISE
 server name ISE
aaa authorization network FLEX group ISE



aaa server radius dynamic-author
 client 172.16.140.212 server-key CISCO

crypto ikev2 name-mangler MANGLER
 dn common-name

crypto ikev2 profile default
 match identity remote any
 identity local dn
 authentication remote rsa-sig
 authentication local rsa-sig
 pki trustpoint MY_CERT
 aaa authorization user cert list FLEX name-mangler MANGLER password Cisco123
 virtual-template 1
```

Password set as User's password in ISE

TrustSec Concepts

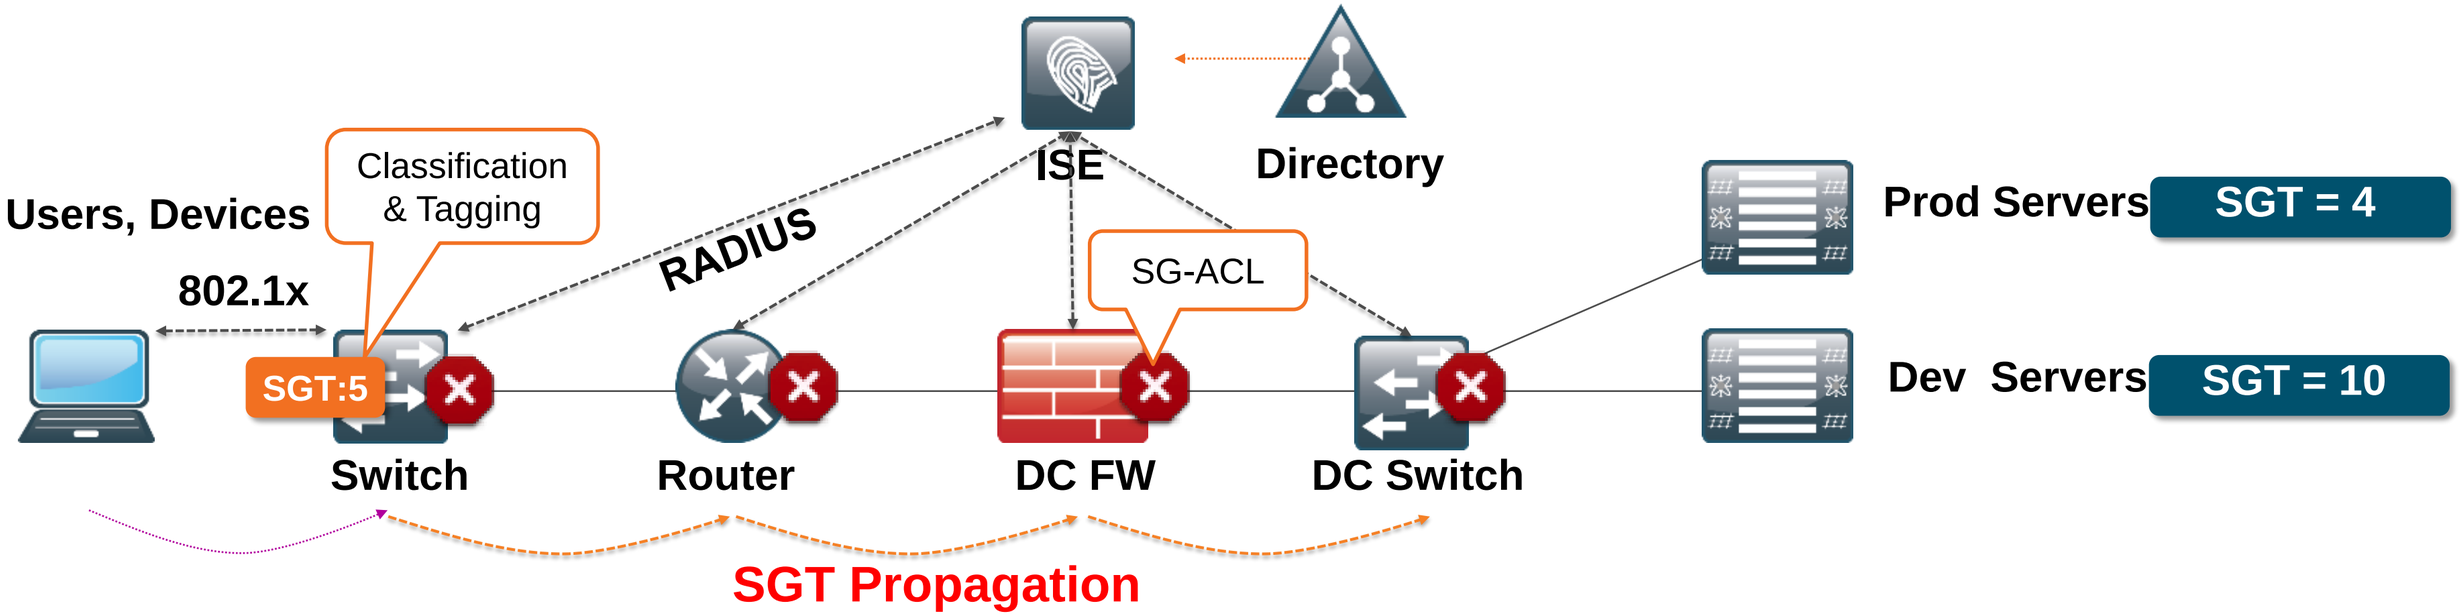

- Classification of systems/users based on **context** (user role, device, location etc.)
- Context (role) expressed as Security Group Tag (SGT)
- Firewalls, routers and switches use SGT to make filtering decisions
- Classify once – reuse SGT multiple times anywhere on network, or….

TrustSec Matrix – source of the SG-ACL

Source \ Destination	BYOD 15/000F	Guests 6/0006	INTERNET 5/0005	ISE 4/0004	Quarantined_Sys... 255/00FF	RA_OK 3/0003	RA_QUARANTINE 2/0002	Routers 16/0010	Unknown
5/0005									
ISE 4/0004							PERMIT_TCP, Deny IP		
Quarantined_Sys... 255/00FF									
RA_OK 3/0003	PERMIT_IP		PERMIT_IP	PERMIT_IP		PERMIT_IP	DENY_IP	PERMIT_IP	
RA_QUARANTINE 2/0002	DENY_IP	DENY_IP	DENY_IP	PERMIT_TCP, Deny IP	DENY_IP	DENY_IP	DENY_IP	PERMIT_DNS, DENY_IP	PERMIT_DNS, Deny IP
Routers 16/0010							DENY_IP	PERMIT_IP	

FlexVPN and SGT

SGT can be imposed at the branch (FlexVPN client side) or at the hub (FlexVPN server side).

The choice depends on capabilities and security level of the remote site.

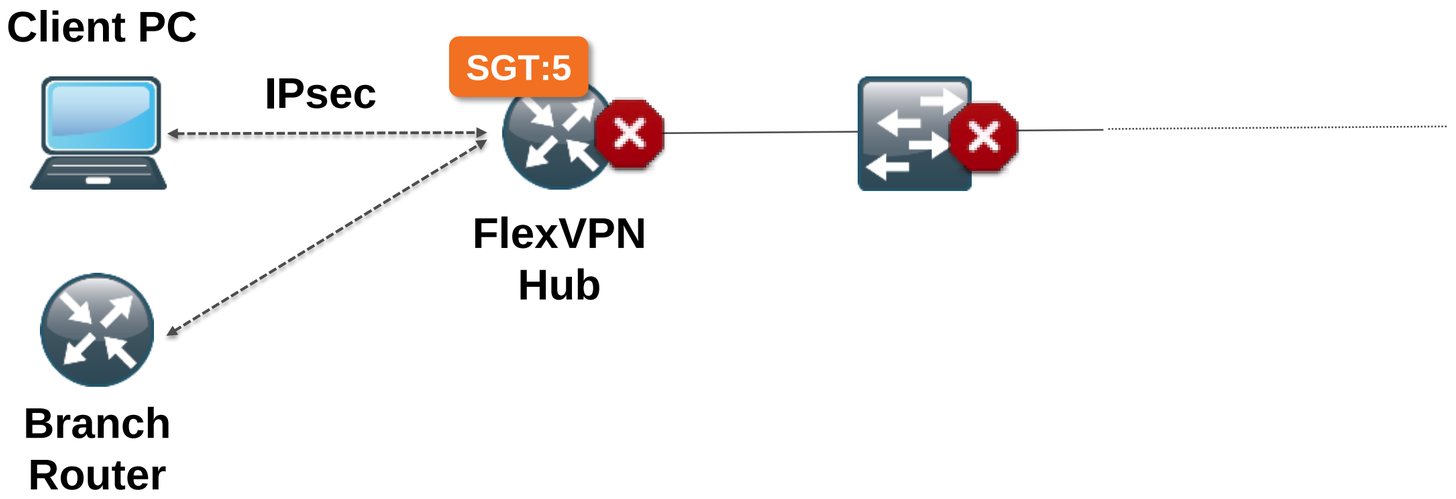

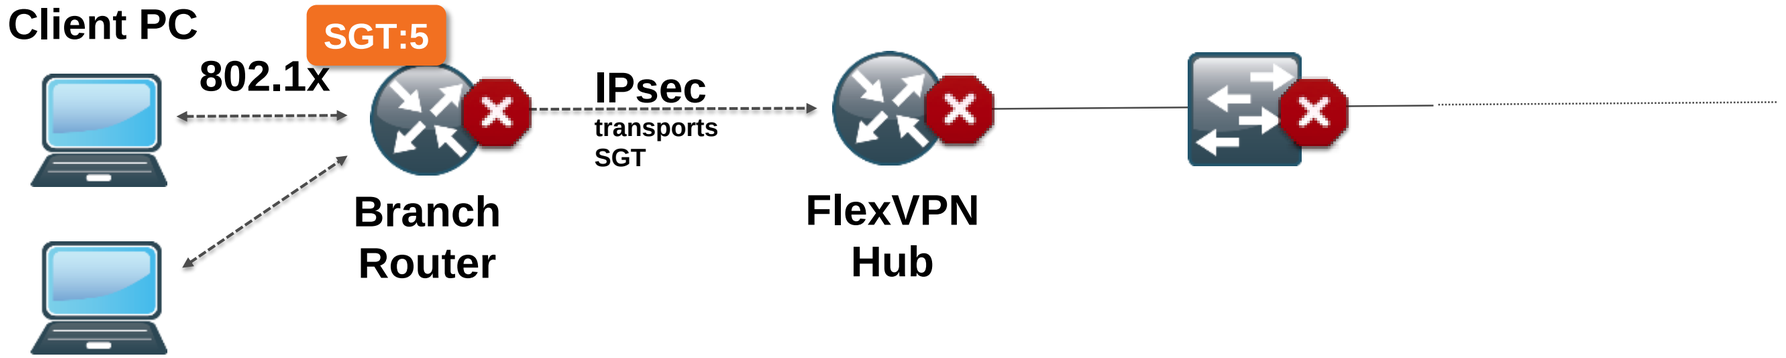

Change of Authorization

How CoA works

Session is set up – V-Access is populated

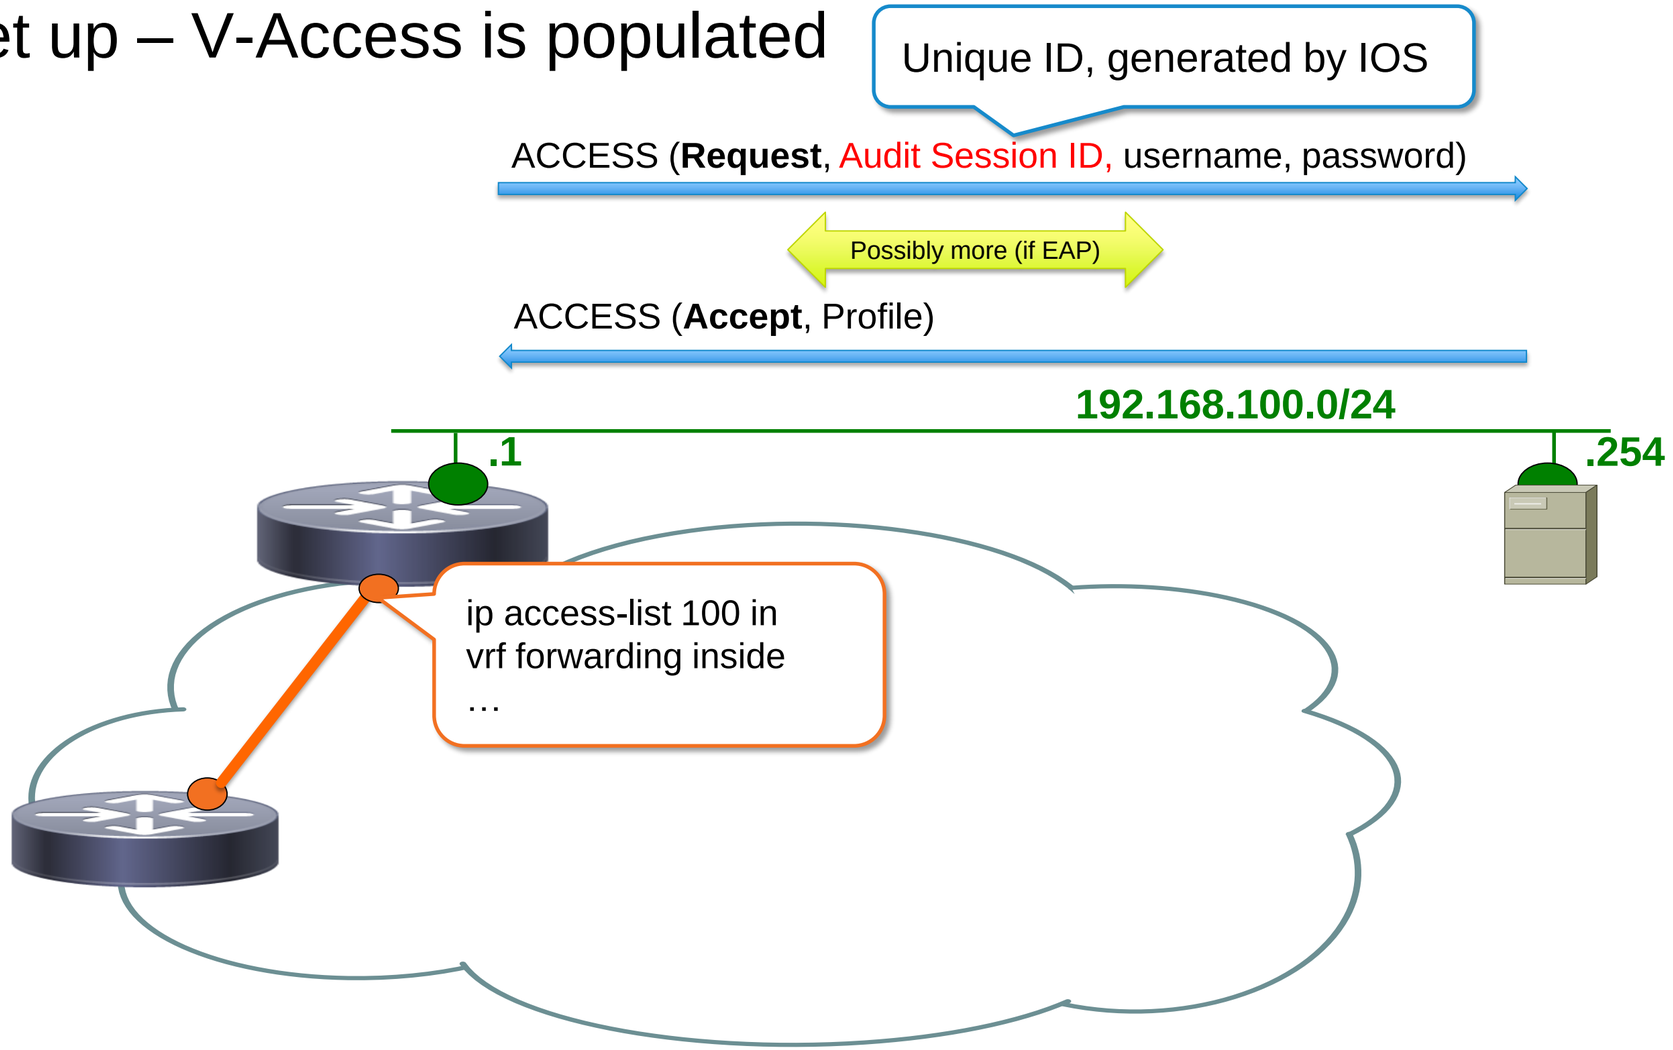

CoA – Packet of Disconnect

Remote clearing of a session

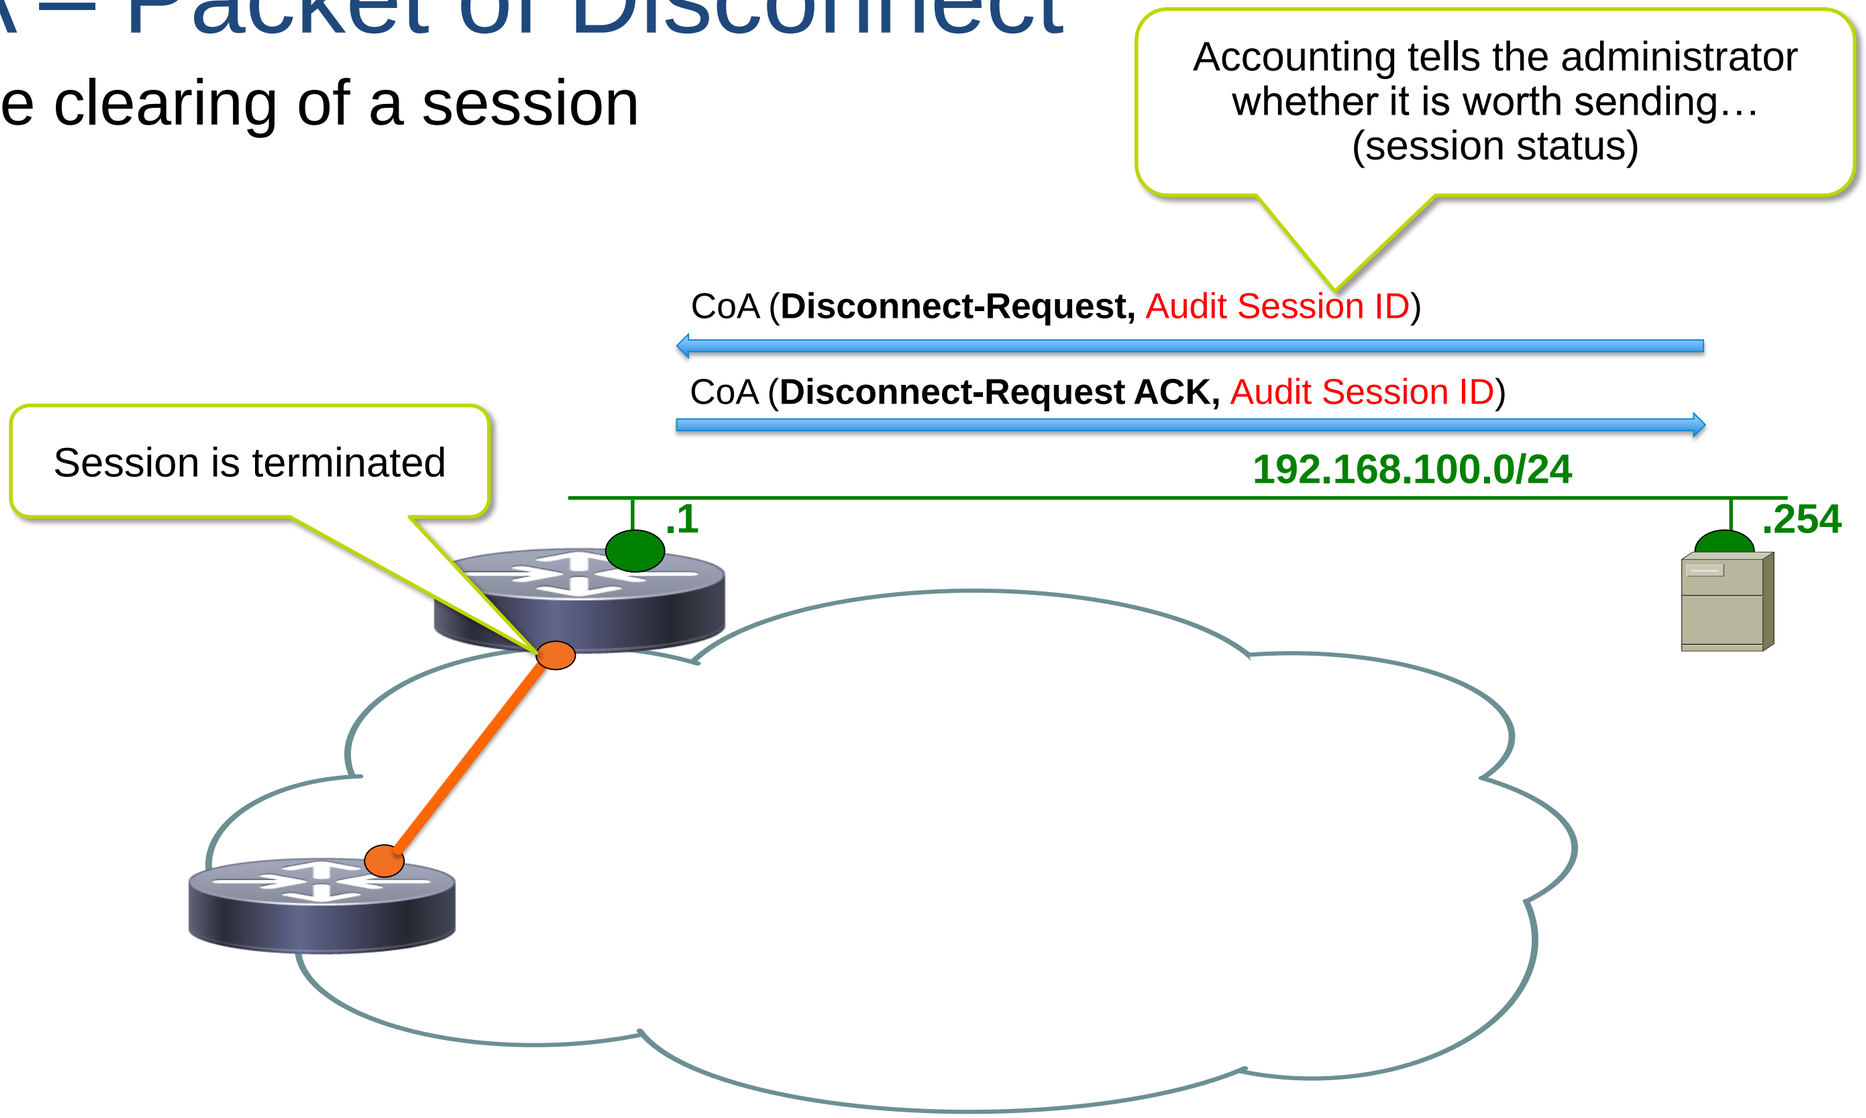

CoA – Change of Authorization

The Real Thing ™

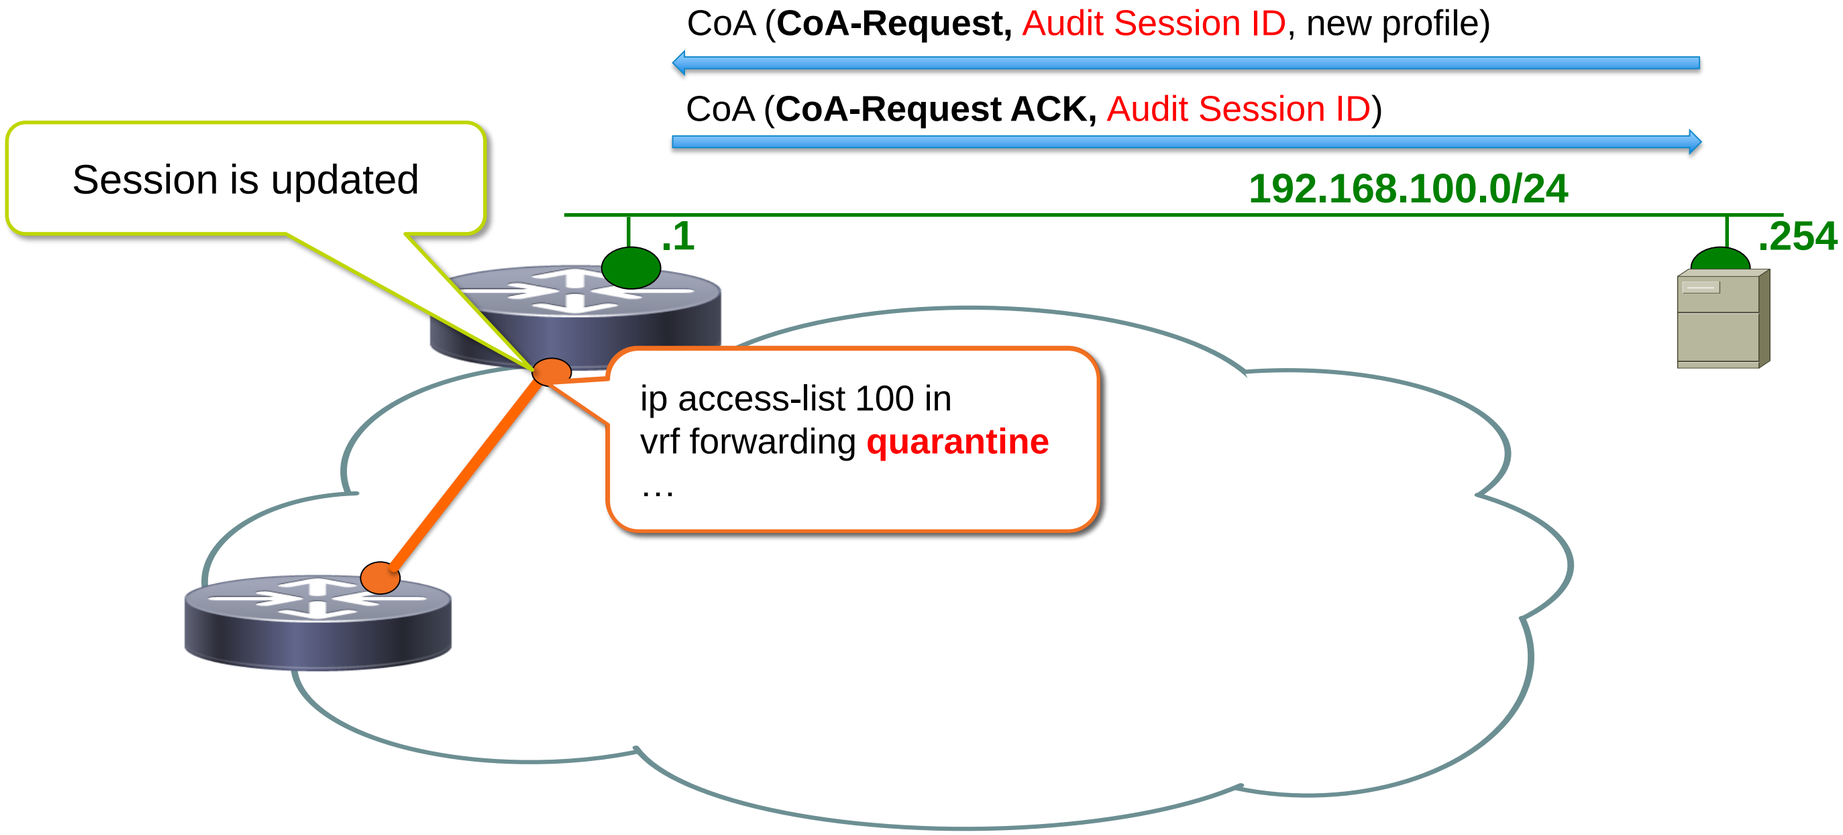

Putting it all together – Posture Flow

Client

IPsec

FlexVPN Hub

DNS Server

ISE

SG-TAG 5 and SG-ACL are applied to v-access

Data traffic is marked and rejected by SG-ACL

SGT:5

DNS Request (enroll.cisco.com)

DNS Response (10.0.0.1)

AnyConnect Posture Module connects to ISE and performs Posture

If Posture Status is compliant,
a new SG-TAG is assigned to the VPN Session
via a CoA-Push message

Data traffic is marked, passes SG-ACL (and unmarked if unneeded)

SGT:6

Accounting for Monitoring

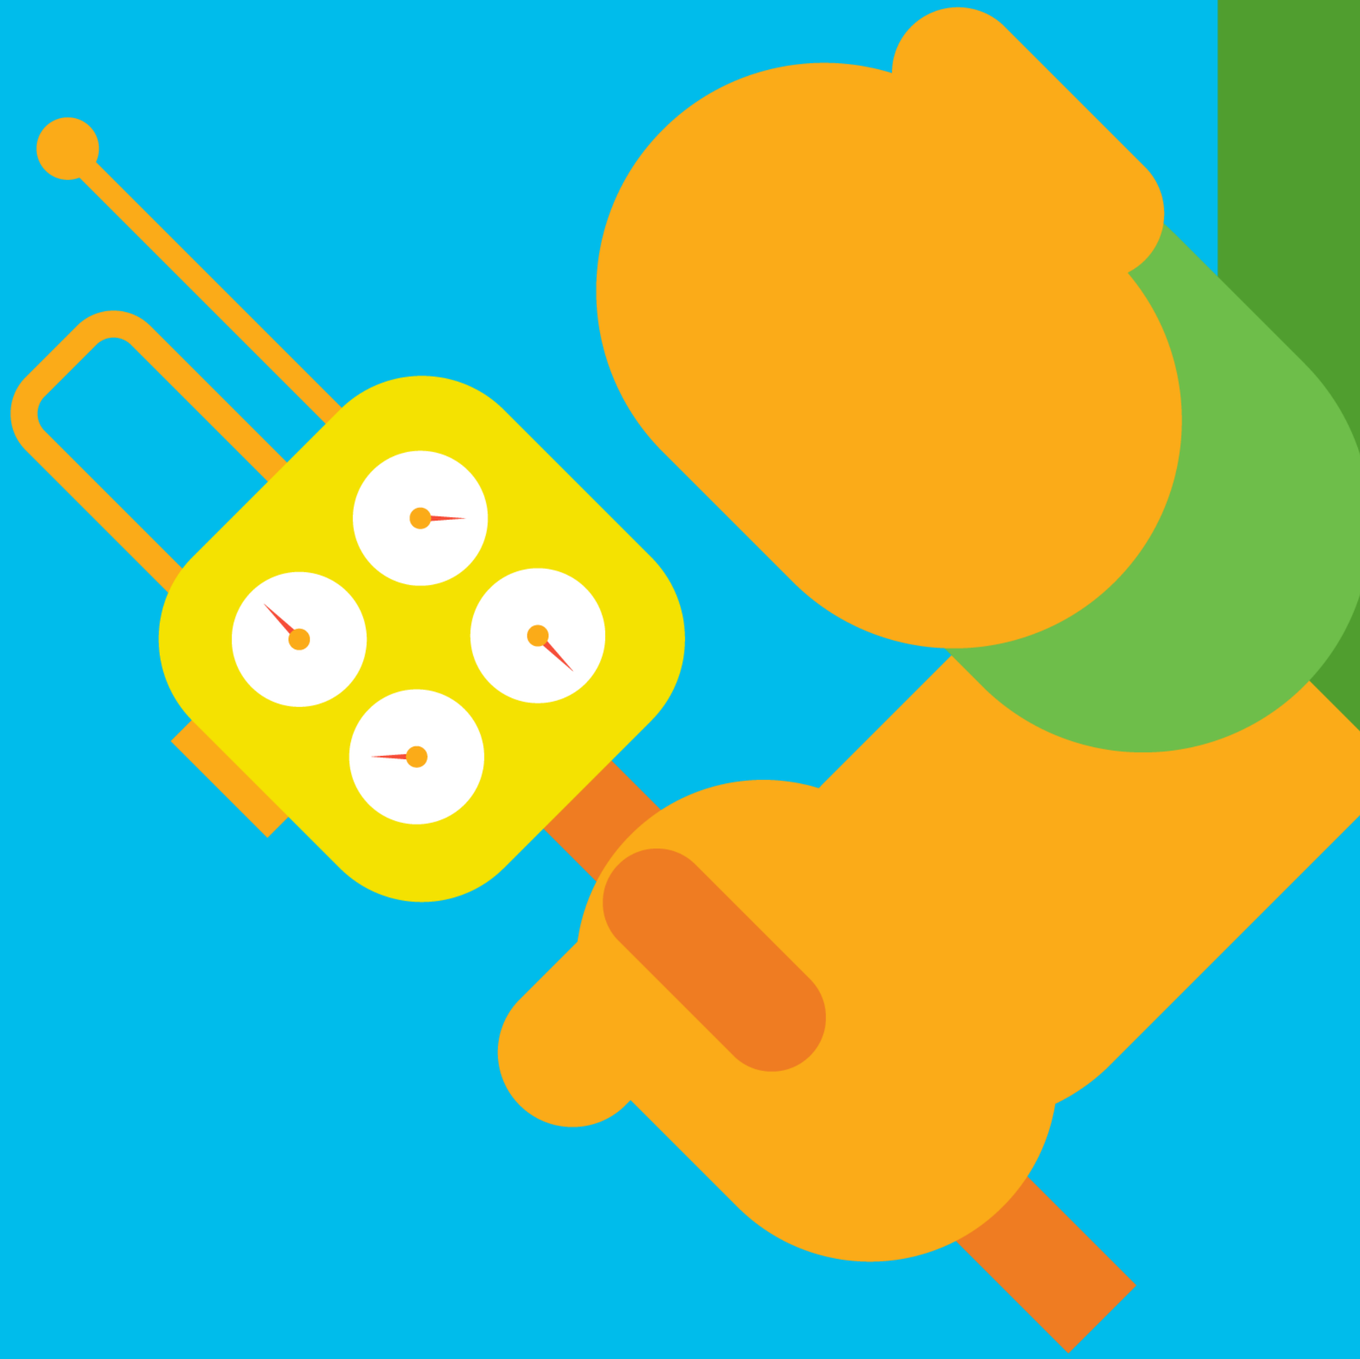

AAA Accounting

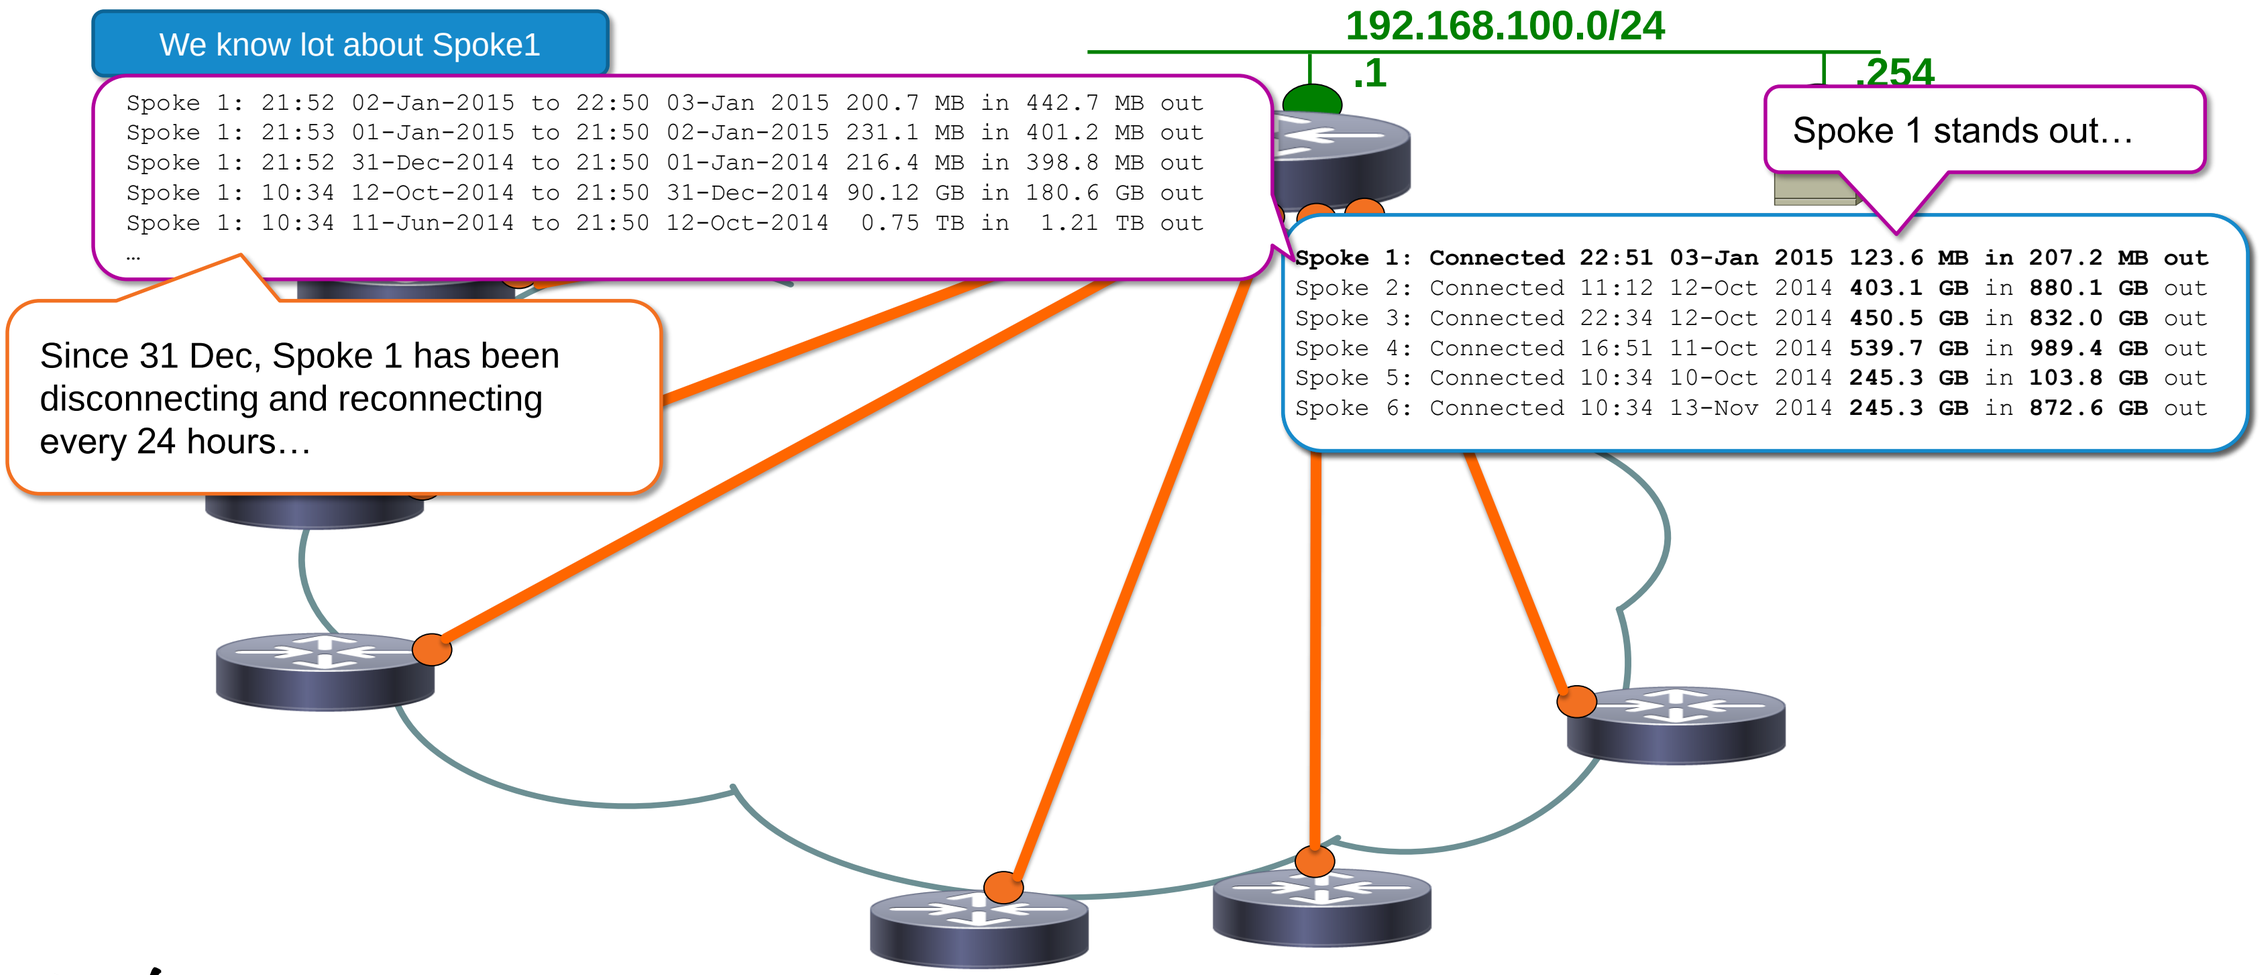

Activating AAA Accounting

And why it is a good idea too…

```
aaa group server radius MyRADIUS
  server-private 192.168.104.101 key cisco

aaa accounting network ACCT start-stop group MyRADIUS

crypto ikev2 profile default
 match identity fqdn domain mycompany.com
 authentication local rsa-sig
 authentication remote rsa-sig
 pki trustpoint TP
 aaa authorization group cert list default default
 aaa accounting cert ACCT
 virtual-template 1
```

Tell IKEv2 to report session status

A Good Idea ?

- Because it is simple!
- Captures even short lived sessions
  ➔ event driven vs. polling (e.g. SNMP)
- Reliable protocol (acknowledged)
  ➔ more reliable than SNMP traps
- Maps the identity to the statistics
  ➔ no more crossing tables (IP→ID)
- You may need it anyway
  - Authorization, IP pool…

A simplistic configuration

RADIUS based Authentication, Authorization and Accounting

```
aaa group server radius ISE
 server-private 192.168.104.101 key CISCO
!
aaa authentication login ISE group ISE
aaa authorization network ISE group ISE
aaa accounting network ISE start-stop group ISE
!
aaa server radius dynamic-author
 client 192.168.104.101 server-key CISCO
 auth-type all
!

crypto ikev2 profile default
 match identity remote any
 identity local dn
 authentication remote eap query-identity
 authentication local rsa-sig
 pki trustpoint TRUSTPOINT
 aaa authentication eap ISE
 aaa authorization user eap cached
 aaa accounting eap ISE
 virtual-template 1
```

EAP Authentication

Authorization

Accounting (optional but recommended)

FlexVPN Load Balancer

FlexVPN Load-Balancer Bootstrap

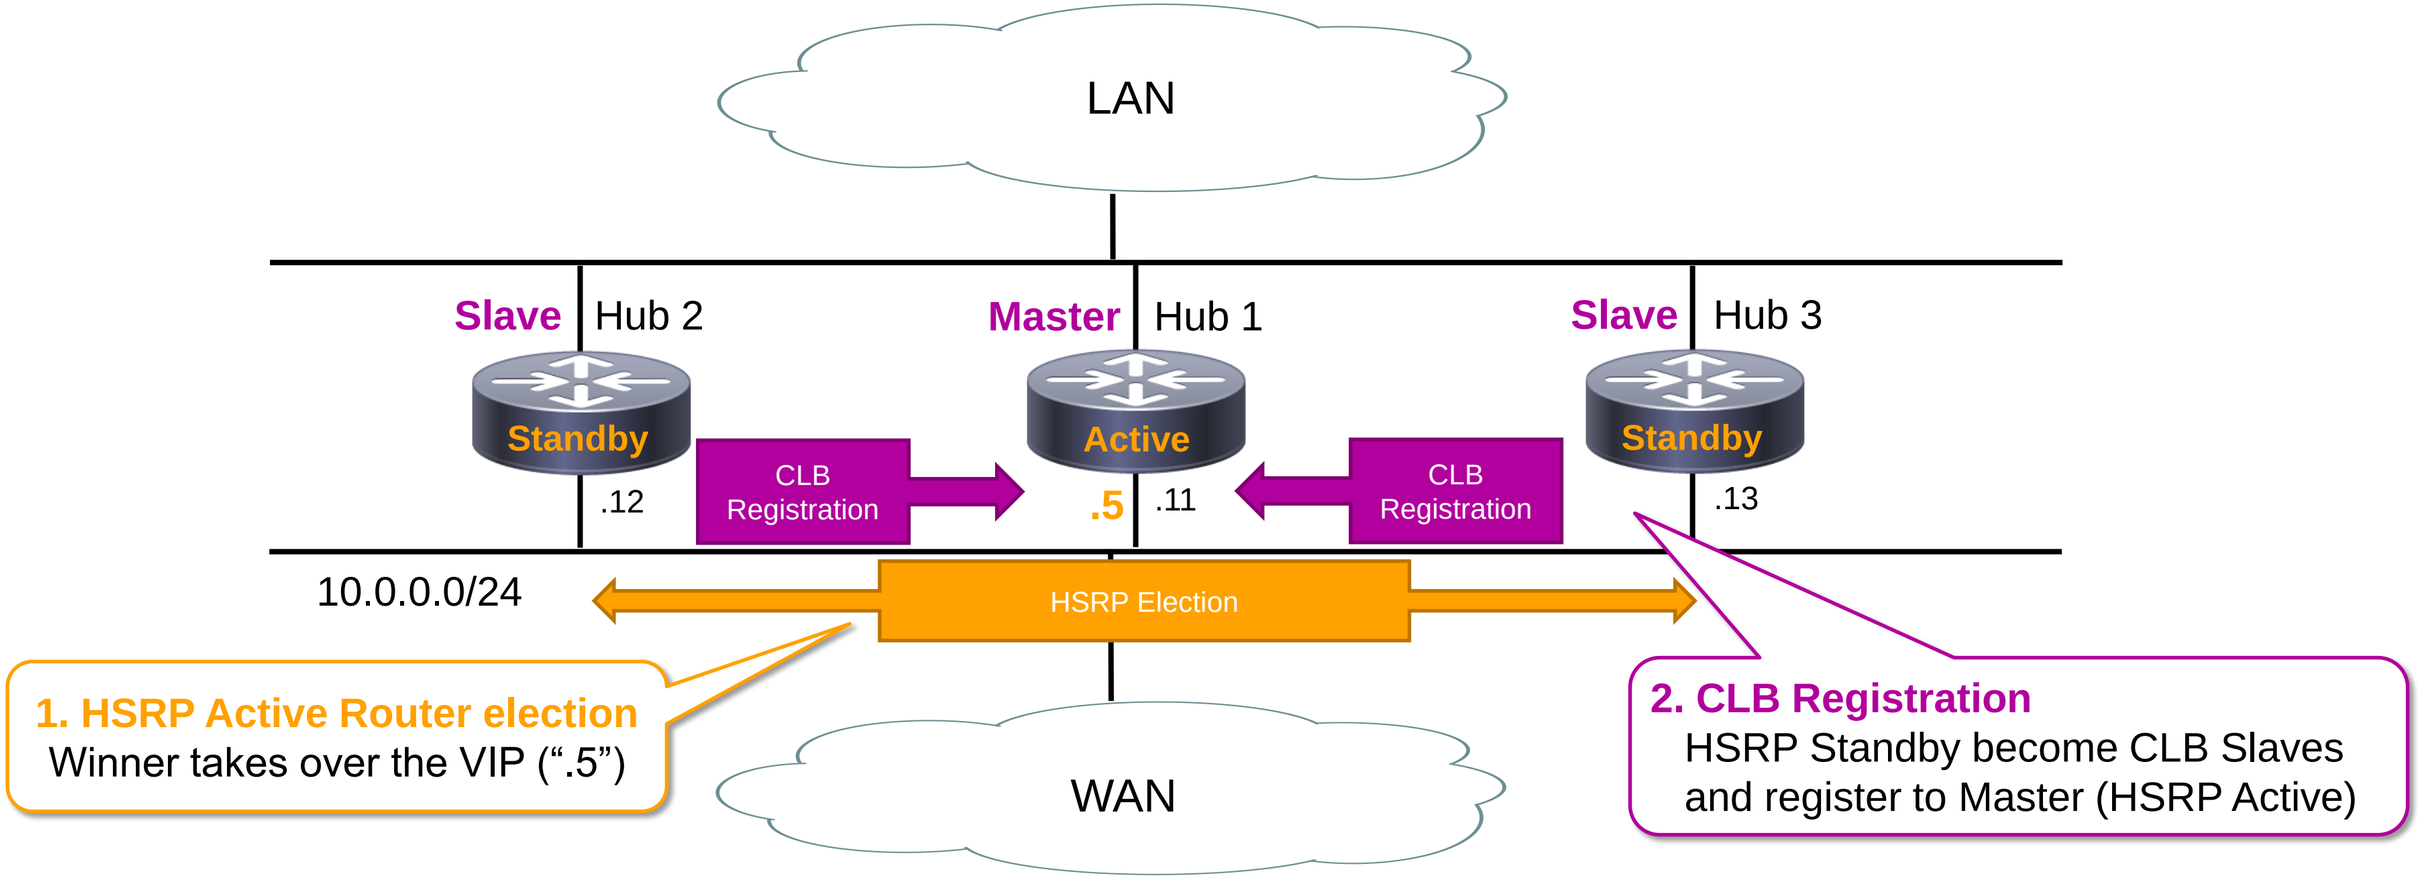

1. HSRP Active Router election
Winner takes over the VIP (".5")

2. CLB Registration
HSRP Standby become CLB Slaves and register to Master (HSRP Active)

```
On Hub 1:
*Nov 20 12:43:58.488: %CLB-6-CLB_SLAVE_CONNECTED: Slave 10.0.0.13 connected.
*Nov 20 12:43:58.493: %CLB-6-CLB_SLAVE_CONNECTED: Slave 10.0.0.12 connected.
```

FlexVPN Load-Balancer Client Connection

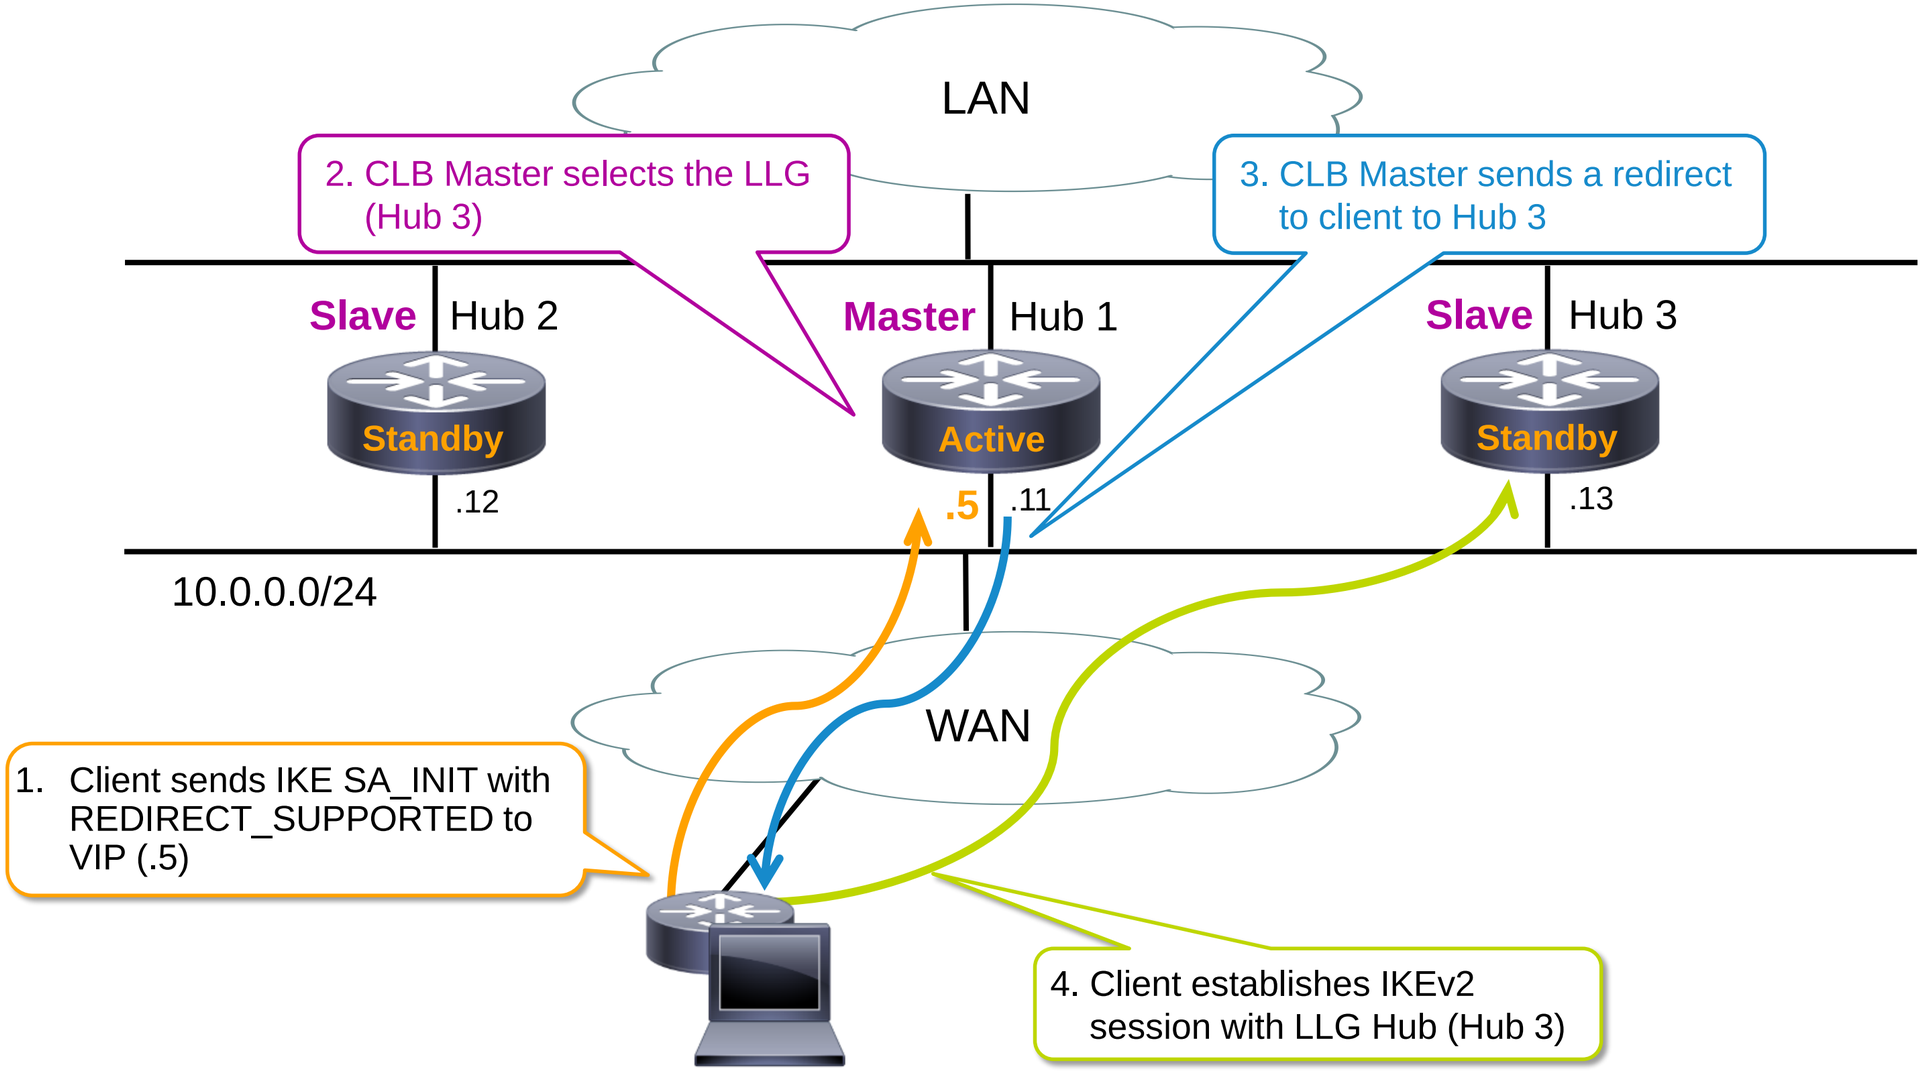

For Your Reference

FlexVPN Load-Balancer – Hub 1 Configuration

```
crypto ikev2 redirect gateway init
!
crypto ikev2 profile default
 match identity remote fqdn domain cisco.com
 identity local fqdn Hub1.cisco.com
 authentication remote rsa-sig
 authentication local rsa-sig
 pki trustpoint TP
 dpd 10 2 on-demand
 aaa authorization group cert list default default
 virtual-template 1
!
crypto ikev2 authorization policy default
 route set interface
!
crypto ikev2 cluster
 standby-group vpngw
 slave max-session 10
 no shutdown
```

Activates the sending of IKEv2 redirects during SA_INIT

```
!
interface Ethernet0/0
 ip address 10.0.0.11 255.255.255.0
 standby 1 ip 10.0.0.5
 standby 1 name vpngw
!
interface Loopback0
 ip address 172.16.1.11 255.255.255.0
!
interface Virtual-Template1 type tunnel
ip unnumbered Loopback0
ip mtu 1400
tunnel source Ethernet1/0
tunnel protection ipsec profile default
```

HSRP Group Name must match IKEv2 Cluster configuration

For Your Reference

FlexVPN Load-Balancer – Client Configuration

```
crypto ikev2 authorization policy default
 route set interface
!
crypto ikev2 redirect client max-redirects 10
!
crypto ikev2 profile default
 match identity remote fqdn domain cisco.com
 identity local fqdn Spoke2.cisco.com
 authentication remote rsa-sig
 authentication local rsa-sig
 pki trustpoint TP
 dpd 10 2 on-demand
 aaa authorization group cert list default default
 virtual-template 1
!
crypto ikev2 client flexvpn VPN_LB
 peer 1 10.0.0.5
 client connect Tunnel0
```

Activates IKEv2 redirection support and limit redirect count (DoS prevention)

```
interface Tunnel0
  ip address 172.16.1.100 255.255.255.0
  ip mtu 1400
  tunnel source Ethernet0/0
  tunnel destination dynamic
  tunnel protection ipsec profile default
```

FlexVPN Peer configured with the VIP address **only**

Before we part…

Route Exchange Protocol Selection

Branch-Hub		Use case				
IKEv2 Recommended	Simple, large scale	Static (No redistribution IGP→IKE)	Simple branches (< 20 prefixes)	Identity-based route filtering	Lossy networks	High density hubs
BGP Recommended	Simple to complex, large scale	Dynamic (Redistribution IGP → BGP)	Complex branches (> 20 prefixes)	Powerful route filtering – not identity based	Lossy networks	High density hubs up to 350K routes
EIGRP not recommended at large scale	Simple to complex	Dynamic (Redistribution IGP → IGP)	Semi-complex branches (> 20 prefixes)	Intermediate route filtering – not identity based	Lossless networks (very rare)	< 5000 prefixes at hub

Hub-Hub	Use case		
BGP Recommended	Large amount of prefixes (up to 1M)	Road to scalability	Powerful route filtering
IGP (EIGRP, OSPF)	< 5000 prefixes total	Perceived simplicity	

Ideal for M2M, IoT, Field, B2B, Managed Svc,…

Connected Vehicle

ISP & Managed Service

ATM’s

Fleet Connectivity

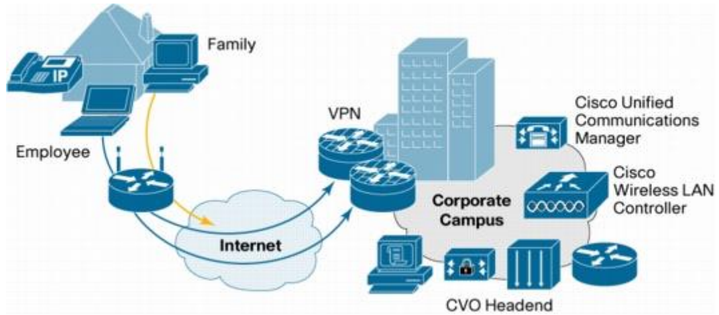

Virtual Office, Retails,…

SmartGrid and all Utilities

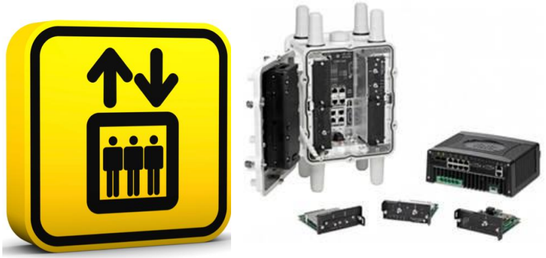

Elevators and IoT via Field Area Routers

Construction, Oil, and field deployments

Spotlight on ESR & IR platforms

- ESR – Embedder Services Routers
- Regular IOS
- Mobile networks in vehicles, mobile users, harsh environments
- 3 ESR models – 5915, **5921** (runs on Linux!) and 5940
- 3 IR models – IR 809, IR819 & IR 829 – ruggedized fog-computing platforms

ESR5921 Bring Your Own Hardware

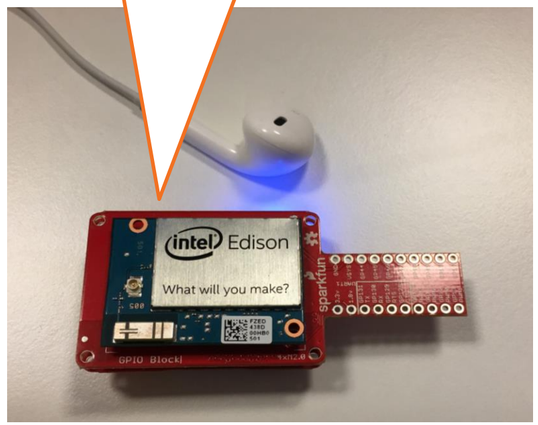

IR 809

IR 829

ESR5915

Complete Your Online Session Evaluation

- Give us your feedback to be entered into a Daily Survey Drawing. A daily winner will receive a $750 gift card.
- Complete your session surveys through the Cisco Live mobile app or on www.CiscoLive.com/us.

Don’t forget: Cisco Live sessions will be available for viewing on demand after the event at www.CiscoLive.com/Online.

Cisco live!

Continue Your Education

- Demos in the Cisco campus
- Walk-in Self-Paced Labs
- Lunch & Learn
- Meet the Engineer 1:1 meetings
- Related sessions

Thank you

You're

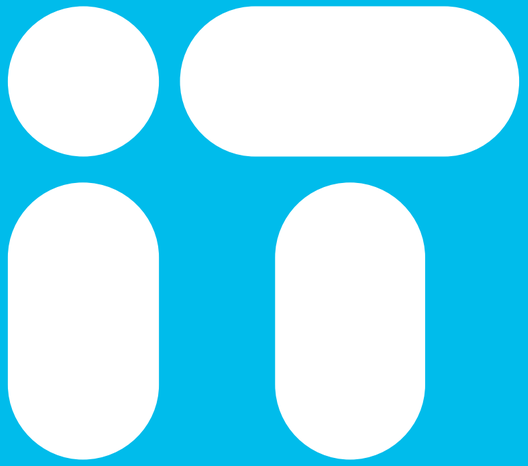

Cisco live!